SERDP

FINAL REPORT

Simulation, Signal Extraction and Augmented Visualization for 3D BOSS Data

Timothy Marston
The Applied Physics Laboratory, University of Washington

Daniel Plotnick
Penn State University, Applied Research Laboratory

May 2022

SERDP Project MR18-1051

DISTRIBUTION STATEMENT A
This document has been cleared for public release.

This report was prepared under contract to the Department of Defense Strategic Environmental Research and Development Program (SERDP). The publication of this report does not indicate endorsement by the Department of Defense, nor should the contents be construed as reflecting the official policy or position of the Department of Defense. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the Department of Defense.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
18-05-2022	Final report	July 2018 - June 2022

4. TITLE AND SUBTITLE
Signal Extraction, Image Correction and Augmented Visualization for 3D BOSS data

5a. CONTRACT NUMBER
W912HQ-18-C-0066

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Timothy M. Marston

5d. PROJECT NUMBER
MR18-1051

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
University of Washington, Applied Physics Laboratory
1013 NE 40th ST
Seattle, WA, 98105

8. PERFORMING ORGANIZATION REPORT NUMBER
A130433

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Strategic Environmental Research and Development Program
4800 Mark Center Drive, Suite 16F16,
Alexandria, VA 22350-3605

10. SPONSOR/MONITOR'S ACRONYM(S)
SERDP

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
MR18-1051

12. DISTRIBUTION / AVAILABILITY STATEMENT
Distribution A
Approved for public release, distribution is unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT
The final report for the project "Simulation, Signal Extraction, and Augmented Visualization of Buried Object Scanning Sonar Data" gives a detailed rundown of the different data processing advancements achieved during the course of the project. Highlights include the development of a very efficient 3D synthetic aperture sonar beamformer, new structural acoustic filtering techniques, and the development of a software library for live and post-processing data captured by buried object scanning sonar (BOSS) systems.

15. SUBJECT TERMS
Buried Object Scanning Sonar, synthetic aperture, volumetric imaging

16. SECURITY CLASSIFICATION OF: A			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Timothy M. Marston
Unclassified	Unclassified	Unclassified	UU	333 337	19b. TELEPHONE NUMBER *(include area code)* 206 543 1300

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

Table of Contents

Abstract 1

Executive Summary 2

Introduction 2

Objectives 3

Technical approach 4

Results and Discussion 6

Data simulation and beamformer testing 6

Field test results: AUV data 8

Field test results: Multi-Sensor Towbody 9

Implications for Future Research and Benefits 11

Objective 12

Background 12

Materials and Methods 17

I. The Buried Object Scanning Sonar (BOSS) 18

Design overview 18

AUV BOSS 21

The Multi-Sensor Towbody (MuST) 21

MuST eBOSS Calibration 23

II. BOSS data simulation 37

Overview 37

Scattering model 38

Simulation parameters 38

Simulation example 40

III. Beamformer development 44

Fast Factorized Back-projection (FFBP) 44

Extension of FFBP to 3D downward looking sonars 49

Asymmetry correction 53

Stage 1 sub-aperture origin and orientation definitions 54

Sampling constraints in θ 56

Frame to sub-aperture conversion 57

Hierarchical sub-aperture combinations 58

Cartesian image generation 59
Acoustic camera images and additional data products 62
Efficient sub-aperture shifting 65
Beamforming summary 66
IV. The MuST Toolbox 67
Overview 67
MuSTLive 68
Renavigation 84
File association and linking 93
SAS post-processing 94
Feature-based renavigation 117
Post-Mission Analysis 129
V. Structural Acoustic Feature Extraction 146
Overview 146
Glint filtering 147
Volume noise reduction 149
VI. Additional visualization considerations and tools 156
MIP vs Alpha renderings 156
Unincorporated data visualization tools 159
Results 163
I. Overview 163
II. BOSSim and 3D FFBP beamformer results 163
Overview 163
Simulation results: data slices 164
Simulation results: SNR, BP vs FFBP 166
Simulation results: Resolution 168
Simulation results: spectral aliasing 169
Simulation results: interpolation comparisons 170
Simulation results: computational efficiency 171
Simulation results: summary of FFBP performance 171
III. Field test results: ONR AUV-mounted BOSS 172
Overview 172

3D FFBP beamforming results 172
Mosaic results 174
Autofocusing 177
AUV BOSS summary 178
IV. Field test results: MuST First Deployment in Lake Washington 178
Overview 178
Data Collection, 8/2/2012 179
Data collection notes, Aug. 5^{th} – 8^{th} 183
Data sets 184
Inertial Labs INS and DVL navigation processing 193
V. Field test results: Sequim Bay 2019 (MuST) 193
Overview 193
Live-streaming Data Visualization 195
Beamforming Results 198
BOSS Structural Acoustic Data: First Look 205
Summary 208
VI. Field test results: GS4 INS Integration 208
Overview 208
Targets and cable 209
Lake Washington PB4Y 209
Range-alias imaging demonstration 212
VII. Field test results: 2020 Sequim Bay UXO trial testbed survey 212
Overview 212
Renavigation results 217
SAS post-processing results 222
Technical Evaluation 239
VIII. Field test results: 2021 INS integration tests 240
IX. Field test results: 2021 Sequim Bay UXO Testbed Survey 243
Overview 243
Intra-survey BOSS data products 244
Mosaic results 251
Target-of-interest ROI processing results 262

Detection results 289
Technical Evaluation 295
Conclusions and Implications for Future Research 296
Literature Cited 299
Appendix A: Buried Cable Tracking 302
Appendix B: list of publications 307
Conference papers: 307
Journal articles: 307

List of Tables

Table 1: Table of calibration measurements performed on the MuST eBOSS array at the APL-UW calibration facility in the spring of 2020. 25
Table 2: *BOSSim* parameters for a point-target data simulation demonstration. 42
Table 3: Back-projection and FFBP completion time tables for different interpolants. Each operation was performed six times and the means and standard deviations are shown in the table. 171
Table 4: Distribution statistics for image alignment using the navigation solution with the three stages of refinement. For raw navigation data with no refinement, the statistics in the top row will all be biased low due to the lack of adequate sampling of the actual displacement distribution. The MAE value of zero for feature-based navigation is due to discretization of the displacement measurements (5cm resolution), resulting in a large number of 0m displacement measurements. 219

List of Figures

Figure 1: Example BOSS images gathered from the literature. Fig. A. is from one of the early papers discussing synthetic aperture processing on BOSS systems and shows two images generated using a 30-ping integration length [2]. The plots in Fig. B were gathered from SERDP project MM-1533 [3] and the plots in Fig. C. are from SERDP project MR-2103. 13
Figure 2: A 120-degree field of view low-frequency acoustics image of a solid 30 by 60 cm. aluminum cylinder, produced using simulated data [7], and viewed to catch the broadside and end-on glints (A). Figure B. shows narrow-aperture low-frequency “looks” at the same target but using real data from a field experiment, and Figure C. shows a three-dimensional low-frequency acoustic image of the same target as in (A) and (B). All figures are from [6]. 13
Figure 3: Plot (A) shows the three-dimensional acoustic color measurement extracted from a solid 2-1 aluminum cylinder using a low-frequency sonar mounted on an AUV. Dimensions are azimuth (theta), elevation (phi), and absolute spatial wavenumber |K|. Plot (B) shows the same acoustic color plot but viewed as two maximum intensity projections through the elevation dimension (top) and wavenumber dimension (bottom). Different physical rigid and elastic scattering mechanisms are labeled in (B). These figures were lifted from the final report of [6]. 14
Figure 4: Acoustic color measurement for a target without (left) and with (right) a clutter reduction filter implemented in the wavenumber domain [6]..... 15

Figure 5: The original BOSS system designed by S. G. Schock et al at Florida Atlantic University. The figures, borrowed from [1], show A) the layout of the projectors (circular elements) and hydrophones, B) a photograph of the projector and hydrophone array on the underside of the BOSS platform, and C) a picture of the BOSS platform from above. 18
Figure 6: A subsequent BOSS design from FAU, which used a linear array of 40 receiver elements (20 elements on each wing) which, when combined with SAS processing, improved resolution over previous designs. Image taken from Appendix A of the final report of SERDP project MR-2103... 19
Figure 7: the BOSS concept of operation. A down-looking source with a very broad beam and low center frequency ensonifies the subsediment volume (sediment surface = flat brown rectangle), and a ribbon-like synthetic aperture array is used to backproject the data into the volume, creating a 3D image. 20
Figure 8: Concept art for the MuST system (left) and a photo of the actual system (right). The leading edge of the eBOSS array is the yellow strip visible on the bottom of the towfish on the figure on the right. Note that in the concept image the eBOSS array was located on the aft end of the towfish, whereas it is on the fore-side for the final product. 21
Figure 9: The eBOSS array dismounted from the towfish, viewed from the bottom. The three sources are visible in the image: the omnidirectional Neptune sources are the black spheres near the sides of the array. The VLA array is the brown grid located near the middle of the system. 22
Figure 10: The MuST eBOSS array removed from the towfish and mounted on the ram used for suspending and rotating the array in the water column during calibration. 24
Figure 11: Illustration of the calibration setup used for calibrating the MuST eBOSS array showing the calibrated reference transducer on the left and the eBOSS array mounted in the configuration used to calibrate the center Neptune source on the right. 24
Figure 12: Mounting configurations 1 (left) and 2 (right) used to measure the aft-to-fore beam patterns of the receiver array (configuration 1) and transducers (configuration 2). The ram is attached to the port side of the array in both configurations. The only difference in these configurations is that the panel is shifted more directly over the transducers in configuration 2. 25
Figure 13: Mounting configurations 3 (left) and 4 (right) used to measure the port-to-starboard beam patterns of the central (left) and port-side (right) transducer. Configuration 3 was also used to measure the port-to-starboard sensitivity of the central receiver array panels. 26
Figure 14: The MuST eBOSS array mounted in configuration 3 showing the eBOSS panels and Neptune sources (left), and the port and starboard sides of the array (right). 26
Figure 15: The array being lowered into the moon pole in configuration 4 (left), and the submerged array mounted in configuration 3 (right). The ram stabilizer, a metal collar which fits around the ram, is visible in the photo on the right. 27
Figure 16: The broadside free-field voltage sensitivity of the eBOSS array elements 1 to 64 over the frequency range 3 to 35 kHz. 28
Figure 17: Starboard panel receiver array element directivity plots in counts / decaPascals at each of the tested CW pulse frequencies. 29
Figure 18: The individual beampattern sensitivities measured for elements 1 – 16 for the starboard side panel of the MuST eBOSS receiver array. Sensitivities are shown in dB referenced to V/uPa. The total dynamic range of the sensitivity plot is 40dB. In several transmitters a null near 20 degrees at broadside is visible. 30
Figure 19: The composite starboard-side receiver panel beam pattern shown in polar coordinates (left) and Cartesian coordinates (right). Note the sensitivity drop near 20 kHz. 30
Figure 20: Port-to-starboard sensitivities at discrete frequencies for a set of elements near the center of the array. Note the dip in sensitivity near broadside at 21 kHz. 31

Figure 21: Port-to-starboard sensitivities in V/uPa for elements 28 – 35. 31
Figure 22: The composite port-to-starboard beam pattern for the two central panels of the array (i.e. elements 17 – 48). Note the very different directivity characteristics vs. the aft-to-fore sensitivity. In both cases, however, there is a dip in sensitivity near 20 kHz. 32
Figure 23: The aft-to-fore directivity measured for the central Neptune transducer. Note the strong interference fringes caused by interference with reflection off of the eBOSS array mount. 33
Figure 24: The port-to-starboard directivity measured for the central Neptune transducer. The interference pattern created by the array is almost identical to that shown in Fig. 23. 33
Figure 25: Port side transmitter port-to-starboard beam pattern in uPa /V at 1 meter. Unlike the central transmitter the port side transmitter shows a high degree of asymmetry in the response, however the reflection nulls are still a dominant feature. 34
Figure 26: A volume rendering created in the calibration facility by moving the array vertically in configuration #3 to create a synthetic aperture. The target is a hollow 60 x 30 cm aluminum cylinder. The dynamic range of the image is 30dB referenced to the brightest point. The direction toward the array is illustrated in both figures which show the volume rotating around different axes. 35
Figure 27: An image showing both the cylindrical shell and the ¼ nylon coated steel cable suspended next to it. The cable is not as strong of a scatterer as the cylinder, but it is still readily visible, with a positive SNR. 35
Figure 28: The 1.5 meter ¼ inch nylon coated steel cable suspended by itself. The 30dB dynamic range, now referenced to the brightest point on the cable, clearly shows reverberation noise within the volume. 36
Figure 29: A suspended coil of ¼ inch nylon coated steel cable. The coil is 0.9 meters in diameter and contained several loops of the steel cable. The cable is noticeably brighter than the individual segment, as evidenced by the decreased relative level of the background noise. 36
Figure 30: A single 0.9m diameter loop of cable. In several of the rendered orientations a faint ring is visible behind the direct return; it is possible that this faint ring may be the visible manifestation of the array mount reflection responsible for creating the nulls in the transmitter beam patterns (i.e. the plots shown in Fig.'s 23 – 25). 37
Figure 31: Three views of the distribution of simulated scatterers, array elements (dark strip down middle) and sources (red dots). 42
Figure 32: A MIP through the channel-dimension of the raw data matrix, showing the profile of the signal data in range and ping. Note that the undulating minimum range values match the vertical motion of the array shown in Fig. 31. 43
Figure 33: Slices of the volume corresponding to the raw data captured by the array on pings 51 (left) and 52 (right). Notice the convergence of the target delay arcs on different sides of the image. This is due to the shifting location of the source: on the figure on the left the source is on the starboard side of the array, and in the figure at right the source is on the port side of the array. 43
Figure 34: A volumetric rendering of the output of the 3D back-projection demonstration code from the previous page. The simulated scatterer located at the center of the volume is visible. 46
Figure 35: Beamforming execution times for the back-projection routine as computed on a RAZER Blade15 laptop using the following values of N: 8, 16, 32, 64, and 128. The red curve is a plot of the function $y = 2.1*10^{-8}x^5$. 47
Figure 36: An illustration of the FFBP process for 2D SAS or SAR systems. The left-hand side, (column A.) shows an aperture that has been divided into small segments, each with two contributing channels. These are beamformed into polar aperture images that have only two "angles" (*A.*, top figure). In subsequent columns, *B.* and *C.*, the sub-aperture images are combined in adjacent pairs to halve the number of images but double the resolution of each image. In column *C*, for example, each

sub-aperture image is comprised of 8 contributing pings and has 8 samples in the angle dimension. This process is repeated recursively through several un-illustrated stages (column *D.*) until the final stage in which all pings for the aperture or block of pings have been combined to form a single high resolution polar image (column *E.*). A polar-to-Cartesian transformation is performed to generate a final image in Cartesian coordinates and perhaps coherently combine with images from other blocks. 48
Figure 37: The figure on the left was generated using standard back-projection and the figure on the right was computed using a fast-factorized back-projection algorithm tuned to reduce interpolation artifacts. The FFBP algorithm executed over 40 times faster than the standard back-projection algorithm. 49
Figure 38: Illustration of the coordinate system of the down-looking synthetic aperture sonar, used for subsequent illustrations of the 3D FFBP beamforming process. Note that the coordinate system is in the North-East-Down (NED) convention, and that *x* is the direction of array motion. 50
Figure 39: An example of a series of real-aperture frames generated in vehicle body coordinates. 54
Figure 40: Example of a bounding-box constructed around the extrema of the visible sonar beam, used to determine the region of the 3D image that will be generated. This box should bound the region "visible" to the sonar. (Snipped is from an animation showing the vehicle, inset in the figure, traversing through the sonar path). 60
Figure 41: A flowchart illustrating the 3D FFBP algorithm. Two dimensional raw data matrices captured every ping in subplot A) are beamformed to real aperture in polar coordinates *B*). The axi-symmetry correction $\boldsymbol{r} \rightarrow \boldsymbol{r}$ allows these frames to be represented across θ without multiple back projection operations C). Adjacent frames are combined in a common spherical coordinate frame to form the first sub-aperture stage D), equivalent to column *A* in Fig. 36. As in Fig. 36, this process is iterated until the specified set of pings (e.g. 256) forming a full chunk are combined (G). A three-dimensional spherical-to-Cartesian interpolation is performed to bring the chunk into the Cartesian image domain, where it can be combined with other chunks. 62
Figure 42: A series of *acoustic camera* frames from the MuST system as it travels over a pipe and solid aluminum cylinder. The vertical axis *Theta* is angle in the synthetic aperture dimension, and the horizontal axis *Phi* is angle in the real-aperture dimension. 63
Figure 43: An example of a MIP through the θ dimension of a final-stage spherical-coordinate sub-aperture. This looks along the synthetic aperture and visualizes the scattering strength vs. angle of incidence for the different sediment layers. 64
Figure 44: An example of a set of MIP's through the φ (real-aperture) dimension, captured frame-by-frame showing a set of floats moving through the field of view of the sonar. Some details of the sediment layering can also be observed. 65
Figure 45: Efficiency vs. shift size for a 256 sample sub-aperture showing the relative speedup vs. complete re-computation using the stage combination method described previously (left), and the fractional efficiency, showing product of the speed gain and the fractional aperture shift (right). The red line is Q=1, the full-aperture reconstruction case. 66
Figure 46: A flowchart showing the MuST eBOSS data workflow used for live and post-mission analysis. 67
Figure 47: The main screen of MuSTLive. 69
Figure 48: SAS parameters screen. This screen enables the user to define the length of the synthetic aperture used for beamforming (in number of pings), along with various options that can affect focus quality and image contrast. 70
Figure 49: MuSTLive's data visualization options screen. A large variety of options are available to the user. 71

Figure 50: The two MIP projections that can be visualized in MuST Live: a projection of the SAS 3D image (represented by the translucent gray cuboid) along the X-axis onto the Z-Y plane, and a projection along the Z-axis onto the X-Y plane. The Z-X projection was not found to be as operationally useful as these two. Note the items are not to scale (the actual dimensions of the MuST are smaller, and actual dimensions of the SAS volume are larger)..... 72
Figure 51: The SAS MIPs data visualization window. The top window represents the Z-projection of the 3D volume onto the X-Y plane, *limited to the region bounded by the red lines shown in the Z-Y projection on the bottom*. This region shows a random segment of sediment and sub-sediment layering from Lake Washington. 72
Figure 52: An X-Z slice through a fishing boat scanned by the MuST. The field-of-view of the sonar (+/- 60 degrees across-track) and range (16 meters, limited by the 40 Hz pulse repetition frequency) is clearly visible in this window. The parallel red lines represent the vertical boundaries of the range for which a MIP is computed and projected in the X-Y plane in the MIPs view. 73
Figure 53: RAS data visualizations: the figure on the left shows the current "real aperture image," a polar-format single-ping image displayed in range vs. angle, and the figure on the right shows a projection through angle of a ~3 minute span of RAS images, resulting in a range-vs-ping plot that is very similar to the output of a standard sub-bottom profiling sonar. 74
Figure 54: Moving average of pulse-compressed envelope data, shown on the left, and raw data on the right which has not been pulse-compressed. No dynamic range compression is being applied in the image so the color-scale has been pre-set to enable visualization of lower amplitude signals at the cost of clipping the direct-path signals, which occur in the first 400 samples. The vertical axes of both figures correspond to range bins, and the horizontal samples correspond to elements of the eBOSS array, which contains 64 channels total..... 75
Figure 55: Connection diagnostics window showing ping difference values and frame-per-second (FPS) reception rate..... 75
Figure 56: Vehicle pose window. The MuST towbody and cardinal directions are visible. The orientation of the towbody updates in realtime based on the feedback from the inertial navigation sensor. When the vehicle is surfaced, as in the example, the position of the waterline is shown as a blue disk..... 76
Figure 57: The navigation map window, annotated to show features..... 77
Figure 58: A zoomed out navigation map window, with the outer ring radius set to 1 km. This is the zoomed-out version of Fig. 57. The map shows the location of the Lake Washington shoreline. Map markers shown the location of shipwrecks and a cable (green line) are visible. 78
Figure 59: A waypoint text file imported into MuSTLive. This waypoint text file corresponds to the detections found on day 5 during the 2021 Sequim Bay testbed survey..... 79
Figure 60: Waypoint update window: by clicking "Select waypoints from saved contacts" in the MuSTLive data visualization panel and selecting from a contact, this window can be opened. New detections can be added to the contact list from this page. Metadata for the contact in the form of location and local imagery is also saved as an external file that is reloadable in a visualization tool that is part of the MuST software library. 80
Figure 61: The waypoint selection window after the colormap has been modified and several objects have been tagged. 81
Figure 62: The navigation map for Sequim Bay (left) showing two passes through a UXO testbed, and the corresponding coverage map produced by the scans (right). 82
Figure 63: Examples of different colormaps that can be selected from the colormap popup menu under "Miscellaneous" in the MuSTLive data visualization screen. 83

Figure 64: Default data visualization layout created by clicking the “Arrange” button. Visual preference is given to the scrolling SAS display; the map and orientation features are present and the processor can be monitored for overburdening using the connection diagnostics display. 83
Figure 65: A GPS re-acquisition jump. The path of the towfish is plotted in this figure with depth color coded using the scale on the right. When the vehicle is less than one meter deep a white ring is placed around the marker, indicating the antennae are exposed. The black ring is a 9-meter radius circle placed just prior to the GPS update being incorporated into the position estimated by the INS. 84
Figure 66: The renavigation solution (colored) vs. the original INS navigation position estimate (black). The reacquisition discontinuity has been completely removed. 86
Figure 67: The user interface that appears when the RENAV program is run. RENAV is used to construct a renavigation solution from MuST INS data. 87
Figure 68: Specifying the roots for importing navigation data from *.ejsf* files into memory 88
Figure 69: INS navigation data displayed in the ReNAV program following loading data from *.ejsf* files. 88
Figure 70: The renavigation solution (cyan) overlaid on the INS navigation solution (pink and white).... 90
Figure 71: A zoomed in view of the track data shown in Fig. 70. Locations in which GPS reacquisition has caused discontinuities are highlighted in the yellow ovals. The renavigation solution smoothly transitions between the submerged and surfaced states with no position lurch. 91
Figure 72: Map contact data loaded into RENAV via the “Load waypoints” option. Note the oval shape of the circular testbed region. The image is being displayed on an equal-axis scale which, for latitude and longitude, will cause a stretching based on the latitude of the map. 92
Figure 73: The *JSF_table* user interface. Users input the relevant sonar data directories, link a renavigation file and specify a path and name for the new table file. 93
Figure 74: The MuST SAS post-processing tool user interface. 95
Figure 75: The SAS processing toolbox showing the options available when the Batching tab is selected (right side). 97
Figure 76: An illustration of the positioning of layers relative to the sediment interface. Rather than representing constant range sways relative to the sonar, the layers represent depth slices relative to the water/sediment interface. The layers above are truncated to improve interpretability, but they span the entire length of the imaged region. 97
Figure 77: Example real aperture and SAS maximum intensity projection (MIP) profiles, projected through real-aperture angle phi (across-track). A scatterer before and after SAS processing is highlighted by the red arrow: upon SAS beamforming the range-migration arcs collapse into a well-focused scatterer. 99
Figure 78: An X-Y MIP layer rasterized into UTM Northings and Eastings coordinates. Note that sediment texture is visible in the imagery and that the nadir glint, which ordinarily manifests as a bright stripe down the middle of the swath, is completely absent due to the application of the de-glinting algorithm. 99
Figure 79: A flowchart of the de-glinting and normalization process. Some relationships (such as the application of the coherence-based weighting functions to the altitude-compensated and high-pass filtered data) have been omitted to preserve flowchart clarity. 100
Figure 80: The data driven altitude estimation process - the profile (top) is filtered and masked to accentuate edges near the sediment transition using the DVL altitude estimate (center). The maxima of this masked region are fitted with a robust curve fitting solution that extrapolates outside of the measurement region using the derivative of the DVL altitude estimate (bottom). 101

Figure 81: Data (MIP through theta) after altitude and phase screen corrections: the specular glint from the direct path to the sediment lines up across all pings. 102
Figure 82: Moving from left to right: 1) a slice of the filtered volume used to determine general trends and discontinuities in range; 2) the normalization function which ranges from 0 (full attenuation) to 1 (no attenuation), 3) the magnitude of the de-glinted slice, still showing elevated levels of backscattering at the surface near nadir, and 4) the slice after applying the normalization function. 103
Figure 83: The three stages of de-glinting: a range-MIP through the original volume (Top), a range-MIP through the de-glinted volume (Center), and a range-MIP through the normalized volume (Bottom). 103
Figure 84: An example of a mosaic generated without the de-glinting algorithm (left) and with the de-glinting algorithm (right) applied. Note that though there are other differences (the figure on the right has additional measures for set continuity and registration applied), the effect of the specular glints is overwhelming. 104
Figure 85: A MuST eBOSS image of the Lake Washington PB4Y captured at a 40 Hz pulse repetition frequency while moving at 2 m/s. The image is clear and well-focused, but some aliasing artifacts are present, resulting from undersampling the synthetic aperture array. 105
Figure 86: Range zones for alias imaging. Standard beamforming occurs in zone 0, and the maximum range is determined by the round-trip distance travelled by sound as limited by the pulse-repetition frequency. Zone 1 corresponds to the swath of ranges ensonified by the previous ping, and zone 2 from two-pings previous. 106
Figure 87: The SAS_VIEW program main screen with some of the features annotated. 107
Figure 88: Demonstration of the linking between the projection in the top window and the spacing and location of the projection limit lines in the bottom window. In the figure on the left the projection lines (annotated with the two black arrows on the left) bound the exposed portion of the shipwreck. On the right, the projection lines bound the bottom of the keel and the sediment below the ship. The acoustic shadow produced by the ship is visible in the projection. 108
Figure 89: Modifying the axis of projection to be in the along-track dimension by selecting the central view-point button in the left-hand column. The projection-limit lines now become vertical dashed lines and are visible mid-ship. In the upper projection window different decks of the ship and ladder-like features become visible. 109
Figure 90: The third (lowest) viewpoint button changes the upper projection axis to be in the across-track dimension. To accommodate adjustment of the projection limits the bottom projection axis is altered to the z (height) dimension. The projection limit lines are visible as horizontal dashed lines in this viewpoint. The limit lines are currently set to isolate the port-aft section of the wreck highlighting a deck rail and some lower-deck features. 110
Figure 91: A 3D maximum intensity projection rendering created from bounds specified by the user after selecting the “Grab 3D MIP” option in the upper right corner of SAS_view. The blue annotation arrows show the relationship between the user selections and the dimensions of the rendered region. 111
Figure 92: Two volume renderings of a UXO generated in SAS_view using the “Grab 3D VOX” button (left) and the “COLOR VOX” button (right). The colorized option rotates color vs. height. 112
Figure 93: The COLOR MIP volume rendering tools opened by clicking the “COLOR MIP” button in SAS_view and selecting a region of interest. The box in the left-most window, which can be dragged by the user’s cursor, modifies the viewpoint perspective. The three color-channel images used to create the RGB composite are visible in the background. 113

Figure 94: An example of the visualization produced by selecting “Grab ORTHO” from the bottom left tool panel and selecting the region outlined by the cyan box. Three orthographic projections are plotted with units on uniformly scaled images. 114
Figure 95: An example of the visualization produced by “Grab Upper” with the Axis mode set to equivalent scaling. Note that de-glinting was applied to this dataset and the distracting Nadir-glint has been removed. 114
Figure 96: An example of the Williams detector results overlaid on the geo-rectified scan from which the detections were made. The detector operates on the volumetric data rather than the projected data and can be used to detect objects that are not easily detectable in a MIP projection that spans the entire scanned sediment volume. 116
Figure 97: A sampling of some looks at a target captured in the 2021 Sequim Bay UXO test bed (left) and the result of fusing the set in a color-by-aspect imaging scheme (right). The taper at the nose of the UXO is visible in the fused data-product shown on the right. 117
Figure 98: A multi-look fusion example showing the results of combining looks with progressively better navigation accuracy 118
Figure 99: An outline of the renavigation algorithm developed for working with BOSS data. 120
Figure 100: Layer data stored for a JSF file (top) and following projection onto target field grid (bottom). 121
Figure 101: the effects of pre-processing on gridded MIP layer data. 122
Figure 102: Gridded sonar data (left) and the binary mask representing the valid data region (middle). The white region in the figure on the left represents the valid region used for overlap detection after the blue bounding area has been eroded. 123
Figure 103: Total displacement vs correlation quality determined via equation (5). Quality values below 6 are randomly distributed over the search area, indicating the noise floor of the quality measurement. 124
Figure 104: A comparison of correlation quality measurements and the distance measurement weightings arrived at via an IRLS solver, for which weights were initialized using the correlation qualities mapped using Eq. 59. 125
Figure 105: Displacement-to-timestamp mapping procedure: the overlap centroid (magenta star) is associated with the timestamp for the path locations of closest approach (magenta lines) for the contributing scan paths (green and red lines). 126
Figure 106: Mosaics of a survey region created using the *Max* operation (left), *mean* (middle), and *contraharmonic mean* (right). The maximum has the fewest image artifacts created by scan density and beam pattern trends, followed very closely by the contraharmonic mean. The arithmetic mean (center) exhibits many artifacts based on density of coverage and beam-pattern effects. 128
Figure 107: A zoomed view of some targets in the region of interest showing the effects of the different combination methods. The *max* operator produces images with highly speckled, almost feathery appearance. The mean and contraharmonic means create nicely de-speckled images. 128
Figure 108: The BOSS post-mission analysis tool user interface. 129
Figure 109: The PMA user interface after making several modifications to the mosaic visualization: the combination method has been set to “mean,” the black and white points have been altered to improve contrast, the layer straddling the interface has been selected, the colormap has been altered and the markers from the last day of survey operations have been included. 130
Figure 110: An example of target selection in the PMA tool. Selections are denoted by red stars. Four locations have been selected. 131
Figure 111: The beamforming options selection window 132

Figure 112: The ROI review screen created upon beamforming the selections shown in Fig. 110. This screen will be the launching pad for most of the user's ROI observation and manipulation and classification activities. 133
Figure 113: The look-data display window showing projections through the spatial domain (top) and filtered data (bottom). An inset view of the mosaic with the location of the current target denoted with a yellow star is shown in the upper right corner. 134
Figure 114: The 3D set-view visualization tool. Clicking and dragging the viewpoint box on the left simultaneously changes the viewpoint of the rendered looks on the right. The image on the bottom right is a magnitude-sum of all of the contributing target looks. 135
Figure 115: The Look-Editor user interface, containing options for look selection, visualization, denoising and length-scale measuring. 136
Figure 116: The MIP viewer opened upon selecting "Launch MIP viewer." 137
Figure 117: The alpha-blending rendering opened by selecting "Launch VOX viewer." 137
Figure 118: The acoustic-color visualization tools opened by selecting "View AC data." 138
Figure 119: The original acoustic color plot prior to filtering, selected by toggling the "Dataset" switch to "Original." 138
Figure 120: The custom denoising window pane. The user defines a boundary for each window which is combined to create a three-dimensional spatial mask around the target. 139
Figure 121: Mask polynomials created in cylindrical coordinates. Note the labels above the images show the dimensions through which the maximum is being found. The "Render & Test" button has also been activated. 140
Figure 122: The acoustic-color spectra for the masked object (top) vs. the original, un-masked object (bottom). The spatial mask is shown as a transparent 3D polygon on the right. 141
Figure 123: The custom solution (top) vs. the previously saved solution (bottom). The figure perspective has also been altered on the right. 141
Figure 124: The length measurement tool opened by clicking "Measure" (left panel). The interface is very similar to the 3D masking interface. 142
Figure 125: 3D length scale measurement example: A line is defined via the Z-MIP (left) and Northings-MIP (right). The Z-MIP line was first selected by the user and defines the line endpoints in Eastings and Northings. The Northings-MIP provided the Z-dimension values for the endpoints. The length of the target is approximately 42 centimeters. 143
Figure 126: The data export panel in the ROI review tool. 143
Figure 127: An example of the field update caused by exporting data corresponding to a particular target. 144
Figure 128: The processed ROI points are now displayed as colored circles. Note that the "Query" button is now active. 146
Figure 129: A flowchart of the structural-acoustic feature enhancement algorithm developed in this project. Stage 1 (left) removes the coherent portion of specular glints in the spatial domain, and stage 2 transforms the 3D acoustic color data created from the de-glinted image into the frequency - k_θ domain, where filter boundaries are auto-detected and applied. 147
Figure 130: Visual depiction of the de-glinting process used for structural acoustic feature enhancement. The sediment glint feature is detected and isolated via iterative reweighted least-squares smoothing and subtracted from the original dataset. Two views of the original dataset are shown on the left, the extracted glint feature is shown in the middle, and the filtered dataset is shown on the right. 148
Figure 131: MIP through cylindrical φ (phi) axis showing a UXO target before (left) and after the origin shift + padding operation (right). 150

Figure 132: Mapping from a 3D image spatial spectrum to CSAS-format data. The left-most image shows a MIP of the 3D FFT of the cylindrical-format image cube (projected through real-aperture angle φ) showing the sampled segment of the wavenumber spectrum resembling an annulus. This is mapped to the frequency vs. angle domain via a Cartesian-to-polar transformation (center). The FFT can be computed in the look angle (theta) dimension to arrive at the ω-k_θ representation (where a wavenumber filter is generated and applied) on the top right, or in the CSAS-like time-angle representation (bottom right), which shows the impulse response vs. look angle for the target..... 151
Figure 133: Knee-point detection. The weighted spectral energy is integrated outwards from both side of the estimated center-point and the knee-points of the integration are used as the filter boundaries. 152
Figure 134: Application of the wavenumber filter. The left figure shows a MIP through the raw data, $S(k_\varphi,\omega,k_\theta)$, which is used to define the coverage of the wavenumber filter via the knee-points shown in Fig. 32. The filter (center) is then directly multiplied with the data. 153
Figure 135: A maximum intensity projection of the filtered "CSAS format" data $\boldsymbol{S\varphi, t, \theta}$ (top) is performed along θ and φ to create the time series on the bottom, used to determine the temporal gating thresholds for the data. 154
Figure 136: The projected data (solid black line, top) and robust quadratic fit (red dashed line top), and the compensated projection (solid black line, bottom) and the smoothed compensated line (red dashed line, bottom). The red arrow points to the peak of the smoothed curve. This point determines the center point for outward accumulation..... 155
Figure 137: The 3D acoustic color dataset used to create the projection shown in Fig. 135 before (left) and after (right) the application of the structural acoustic enhancement filter described in this section. 156
Figure 138: Augmented MIP views through a volume containing a UXO facsimile target. The MIP volume is rendered from the perspective of a user-defined camera location. Attenuation as a function of range combined with perspective scaling and multiple looks creates a greatly improve sense of the 3D spatial relationships between scatterers..... 157
Figure 139: An illustration of the tradeoff between feature preservation and spatial relationship preservation for spheres embedded in non-uniform noise. The MIP rendering on the left clearly shows the key features of the volume, including the three spheres and spatial variance of the noise, but all spatial scale is lost. In contrast, alpha-blending (right 3 figures) preserves the spatial relationships of the spheres, but no single threshold captures the major features of the volume. 157
Figure 140: Side-by-side renderings of a multi-scan volume mosaic containing a fishing vessel. The MIP rendering on the left shows structural details inside of the hull, whereas the alpha-blended rendering on the right is more useful for visualizing the external features of the vessel. 158
Figure 141: Augmented MIP's (top) of a multi-layer sediment volume containing a target, juxtaposed with color-by-depth alpha renderings generated in VOXview. 158
Figure 142: The wavenumber coverage displays created using the *k_coverage*() tool for an example ROI. 159
Figure 143: Two ROI rendered using CBA coloration. The left target is using a power mean exponent of 1, the volume on the right is rendered using a power mean exponent of 2..... 160
Figure 144: Color-by-depth ROI volume visualizations rendered using color-MIP's. Objects are detections found during the 2021 Sequim Bay UXO testbed survey..... 161
Figure 145: RenderSpect showing spectral data for 2, 1, 1, and 3 separate scans (top row) and the ROI aggregate (bottom). Note that the aggregate creates a piece-wise reconstruction of features such as the broad-side specular and interference fringes from corner scattering that are not well captured by any single scan within the set..... 161
Figure 146: MIP views of the denoised aggregate spectrum (left) and the original noisy aggregate (right). Note the bright vertical spectral line extending upwards from the center of the spatial spectrum: this

spike corresponds to the sediment specular glint. This feature is completely absent in the spectrum on the left. 162
Figure 147: The ROI_TriView visualization showing an alpha-blended rendering (left), MIP rendering (middle) and iso-surface rendering (right) viewed from the same perspective. 162
Figure 148: Visualization of the simulated system track and point scatterers, along with the slicing planes sampled from the volume that will be used to compare the performance of standard back-projection with 3D FFBP. 164
Figure 149: A perspective-MIP rendering of the 3D image generated by applying the 3D FFBP algorithm to the data generated using BOSSim and the simulation parameters listed in Table 1. The points are well behaved with low side-lobes and the distribution follow the expected radial distribution around the sensor trajectory. The amplitude is plotted in a linear scale normalized by the value of the brightest point. 165
Figure 150: slices through the image cube using FFBP (top row) and back-projection (bottom row), imaged using a log-scale color axis and 50dB dynamic range to accentuate interpolation artifacts. Color-scale is 50dB relative to the brightest point in the image. 166
Figure 151: A data slice through the 3D FFBP image plotted in Fig. 150.A at the line z = 5.51 meters. The peaks are highlighted with a star, the median signal level is shown with a dashed line and the noise level is shown with a bold line. The signal-to-noise level is measured to be 35.3 dB. 167
Figure 152: A data slice through the 3D back-projection image plotted in Fig. 150.D at the line z = 5.51 meters. The peaks are highlighted with a star, the median signal level is shown with a dashed line and the noise level is shown with a bold line. The signal-to-noise level is measured to be 37.9 dB, about 2.5 dB greater than for the 3D FFBP case 167
Figure 153: BP (red) and FFBP (black) profiles plotted on top of each other, showing how similar the general structure and levels of the noise and point spread functions are. 167
Figure 154: SNR estimates for back-projection (red, circles) vs. FFBP (black, stars). Vertical axis limited to exaggerate the differences. BP tends to have ~2 to 3dB higher SNR relative to interpolation-artifact noise floor. 168
Figure 155: Peak widths vs. range, back-projection (red, circles) vs. FFBP (black, stars). The decreasing resolution is due to the fact that the resolution is aperture limited not beam-width limited in the simulation. 168
Figure 156: Peak width ratios, FFBP / BP, vs. range. FFBP resolution is, at worst, only about 15% lower than back-projection. Qualitatively, this difference is nearly impossible to detect in images. 169
Figure 157: Spatial spectra of 150.A (FFBP, left) and 150.D (BP, right). The spatial spectrum of the image beamformed by FFBP stays within the bounds specified by the integration limits (denoted by the white dashed lines), and the signal bandwidth (solid white curves). In contrast, the BP routine, which had no integration angle limits imposed beyond that of the aperture itself, shows energy distributed beyond the integration angle limits. 169
Figure 158: SNR comparisons between FFBP implementations using spline (red, circles) interpolation and linear (black, stars) interpolation. The interpolation can make a 4 to 5 dB difference in SNR. 170
Figure 159: Peak width comparison in the along-track dimension for FFBP implementations using spline (red, circles) and linear (black, stars) interpolation. Unlike with SNR, there doesn't appear to be a huge difference in the resolution results for the different interpolation methods. 170
Figure 160: Frames generated from one block of pings. These frames will be used to make the first stage spherical coordinate sub-apertures in the 3D FFBP algorithm. 172
Figure 161: Stage-1 low resolution spherical coordinate sub-apertures generated from the frames visualized in Fig. 160. 173

Figure 162: Stage-2 spherical coordinate sub-aperture images. The pair-wise coherent combination of the images in Fig. 32 has resulted in half the number of apertures having twice the resolution. 173
Figure 163: Stage-3 spherical coordinate sub-aperture images. 173
Figure 164: Final stage (stage-4) spherical coordinate sub-aperture. Depending on the block-size and value for *K* the 3D FFBP algorithm may require fewer, or more, stages to generate the final sub-aperture corresponding to a single ping block. 173
Figure 165: Cartesian data projection. Note that the sediment no longer appears to lie on a curved manifold. 174
Figure 166: Z-MIP mosaics of two scans from the AUV-mounted BOSS system. The target, a buried pipe, is distinctly visible, along with the end-markers, at ~2 meters east, -18 meters north. 175
Figure 167: Zoom-in of the buried pipe target and the end-markers. A ripple field is visible on the left side of the image 175
Figure 168: Frame 1 of an acoustic camera / Z-MIP generation demo. The Acoustic camera view is on the left, with angle in the synthetic aperture dimension being in the vertical axis, and angle in the real-aperture dimension being the horizontal axis. The Z-MIP mosaic imagery (as it is generated) is shown on the right with a white box surrounding the most recent addition to the mosaic. 176
Figure 169: A subsequent frame of the mosaic, showing the progression of the mosaic. 176
Figure 170: A frame showing the pipe target in the acoustic camera image (left) and its contribution to the mosaic, on the right. 177
Figure 171: Data from channel 1 (left) and channel 25 (right) showing signal degradation due to panel noise. 179
Figure 172: An example of the raw data captured from the BOSS system showing sensor noise (horizontal striations). The starboard-side array was malfunctioning and thus shows no signal (channels 49-64). The bright signal dominating the first couple milliseconds is the direct path signal from the source to the array. Pulse compression artifacts have wrapped around to the time series beyond 25 milliseconds. The specular glint signal from the sediment arrives at about 24 milliseconds. 180
Figure 173: Noisy sonar data before (left) and after (right) de-noising. This ping was acquired on 8/8/19 (day 4) but is used here because the sensor is at low elevation and the masking of the sediment layers is easy to see 181
Figure 174: Noisy sediment profile (left) and de-noised profile (right) found by taking an average of the mean across the channels for the pulse-compressed data. Images are plotted using the same 40dB log scale color axis. 181
Figure 175: The MuST towfish position in UTM coordinates for a scan as estimated by the INS' internal Kalman filter. The bounding box required to enclose the path estimate is greater than 100 *kilometers* on edge, despite the track in reality only being on the order of a kilometer. The red points near the bottom left show the locations at which a GPS signal is received. The INS estimate is good when the GPS signal is received but rapidly diverges from having physical meaning when the GPS signal is lost. 182
Figure 176: An estimate of roll angle using the effects of gravity on the raw linear-accelerometers in the INS, and the roll estimate actually output by the INS as estimated by its internal Kalman filter. Over the timespan depicted the towfish was not noticeably rolling, and the INS filter is giving clearly erroneous results. 182
Figure 177: The difference between the INS' estimate for roll and the simple post-processing estimate found using the orientation of the gravity vector measured via the linear accelerometers. 183
Figure 178: Sensor response averaged over several thousand pings (left) and a relative channel sensitivity calibration in dB relative to the median response across the array. 185

Figure 179: A single ping real aperture image visualized in polar (left) and Cartesian (right) coordinates. 185
Figure 180: RAS images without de-noising, showing the spoke (left), and with de-noising, showing the null beam (right). Images are envelope averages of 50 RAS images, and plotted on the same 40dB dynamic range log-scale axis. Note that the null-beam isn't a perfect null because the broadside has phase curvature that is not modelled by a coherent sum across panels. Hence, a small amount of residual signal remains in the broadside direction. 186
Figure 181: Averaged real-aperture polar images for the alternating Neptune configuration (left) and the VLA configuration (right). Linear scale. 186
Figure 182: VLA profile, 8/8/2019, 8:14 AM (PST). Target range = 20m, not visible in scan. 187
Figure 183: VLA profile, 8/8/2019, 8:50 AM (PST). Directly over pipe (target arcs visible in far right panel)..... 187
Figure 184: VLA profile, 8/8/2019, 9:20 AM (PST). Slug target +5 m. west. 188
Figure 185: VLA profile, 8/8/2019, 9:20 AM (PST). Figure on the right shows a MIP through the angle-span containing the target (slug), making the target arcs visible. Note that the "layers" over this angle range are actually side-lobe artifacts. 188
Figure 186: VLA profile, 8/8/2019, 9:20 AM (PST). A comparison of the profile scan showing the target with an average of the real-aperture images containing the target. The target is visible around +45 degrees (starboard). The high sidelobe level is an artifact of the window function used to weight the array during beamforming, poor uniformity of sensor response across the array, and de-noising process. Note that, due to VLA directivity the signal level from the target is low enough that poor placement could have put the target in a sidelobe that would completely drowned the target signal in noise. The wider directivity of the Neptune sources would prevent this type of problem from occurring as frequently..... 189
Figure 187: VLA profile, 8/8/2019, 10:10 AM (PST). Directly over target (slug). The target arcs corresponding to the slug are enlarged in the panel at right. Note that the target appears to have sunk down to the second sediment layer, and is also casting a shadow. Range aliasing (reception of backscattering data from previous pings) is visible in the top left and top right of the left and middle panels as well. This type of problem exists due to the rapid ping rate. The ping rate could be decreased, however this would cause major undersampling in the along-track direction, hampering SAS processing that will be applied to the data in the future. 189
Figure 188: VLA profile, 8/8/2019, 10:56 AM (PST). Nearly directly over target (pipe). 190
Figure 189: Alternating Neptunes, 8/8/2019, 1:15 PM (PST). This scan performed at different angle through target field. The scan is directly over the pipe and close to the slug. 190
Figure 190: Alternative view of the data in Fig. 189. An adaptive filter was used to remove specular glints from sediment layers, leaving the rapidly evolving target arcs as a residual. A range-MIP through the real-aperture beamformed data was performed, showing the relative location of target and clutter target arcs relative to the track of the MuST. SAS processing, when applied in the future, should focus these arcs into targets and clutter. Vertical axes are angle, horizontal axes are ping number. 190
Figure 191: Alternating Neptunes, 8/8/2019, 1:40 PM (PST). Directly over pipe. Target is enlarged in the far right panel. 191
Figure 192: Alternating Neptunes, 8/8/2019, 3:07 PM (PST). Slug off to starboard. 191
Figure 193: Alternating Neptunes, 8/8/2019, 4:09 PM (PST). Directly over slug. Channels 39-42 failed, in panel #3. The plot on the far right is an excellent example of range-aliasing from previous pings: as the vehicle climbs in altitude the backscatter data from the previous pings begin to fill in the region of the watercolumn. 192
Figure 194: Alternating Neptunes, 8/8/2019, 4:36 PM (PST). Over slug. Final scan of the day. 192

Figure 195: This figure shows the re-navigation solution as an inset to the INS estimate shown on left. The re-navigation position estimate is in blue, with blue dots spaced every 100 meters of track. The star marks the initial position of the towfish in the plot. 193
Figure 196: Track lines over the muddy target field and ship coordinates from day 1, 9/9/2019 (A) and day 2, 10/9/2019 (B). Scans were also conducted over the sandy region on day 2, visible as the set of race-track loops in the upper left of the figure. The region of the map shown in (B) is shown in its greater context in (C). 194
Figure 197: Ship coordinates for day 3, 11/9/2019 (left) and day 4, 12/9/2019 (right). The bold rectangle in the figure at right is approximately 500 by 1700 meters. 194
Figure 198: Target field item position maps for the sandy site (left) and muddy site (right). Targets were positioned 5 meters east / west of the center line for each field, and had 10 meters spacing in the north / south direction between other objects on the same side of the line. The total length of the fields were approximately 60 meters. 195
Figure 199: Live pulse-compressed channel data visualization panel. This window would refresh at approximately the ping repetition rate. Channel number is the horizontal axis. The port-side element is channel #1 and the starboard-side element is channel #64. Temporal sample bin is in the vertical dimension. The bright signal near the top of the matrix is the direct path, which peaks in amplitude near the port side, indicating that the transmitter configuration is the alternating Neptune configuration, and that this measurement is from a port-side transmission. Two channels, #20 and #39, are inactive, and two more, 3 & 4, have reduced sensitivity. On subsequent days of the experiment (Sept. 10 – 12), these panels were replaced with panels from EdgeTech that had an improved design and performance was improved. 196
Figure 200: A screenshot of the real-aperture, single ping visualization panel. The bright artifacts near the image origin are beamforming artifacts from the direct-path signal to the array. In this image the sediment is approximately 10 meters below the sonar. The image is in Cartesian coordinates and approximately 20 x 35 meters in size. Only the highlighted region in the window contains beamformed acoustic data, however. This image is refreshed at approximately the pulse repetition rate..... 196
Figure 201: The scrolling sediment profile panel. This image is an updating queue of profiles captured using the central angles of the real-aperture frames. The horizontal axis is range from the sensor, and the vertical axis is ping (the numbers are reversed, new data appears at sample #500 and works towards sample #1). The profile scrolls from right to left as data is acquired A target near the surface is highlighted in the figure. The red strip near the top is the direct-path signal. The signal level is clipped to reduce the dynamic range of the image and make the lower level target signatures visible. 197
Figure 202: Vertical MIP view. This is a scrolling window that shows the azimuthal displacement in meters from the track of the vehicle, and scrolls from top to bottom. The horizontal axis is port/starboard displacement relative to the towfish, and the vertical axis is sample. The image is generated by computing a MIP in the vertical dimension through the real-aperture frames The target circled in Fig. 74 is visible in this figure at (350, -2). 197
Figure 203: A MIP in range through the raw data (top) and through the data-cube generated by applying SAS beamforming. The ellipses in the top figure are the un-compressed target arcs corresponding to two targets in the bottom figure, denoted with the arrows. The target field is the field in the muddy sediment. 198
Figure 204: Z-MIPS through the entire beamformed volume (top) and just the top half-meter through the sediment, including the surface. When the range of the MIP is truncated, the target field becomes immediately obvious due to a reduction in clutter. 199

Figure 205: A MIP through the eastings dimension of mud field dataset shown in Fig. 203 and 204. A strongly backscattering sediment layer is visible ~3 meters below the water/mud interface that masks the target response when a MIP is calculated through the vertical dimension. Targets are the discrete bright scatterers clustered within a meter of the water/mud interface. 199
Figure 206: Crab-pot patch showing the Z-MIP and zoomed patch (top and bottom left), and iteratively changing the height plane of the slice through the volume. It can be seen that the crab-pot makes a disturbance in the sediment (middle plot), and is also visible in the watercolumn (bottom right two figures). Note that the dynamic range of the bottom right two figures has been altered to make the crab-pots visible. The low scattering strength of the crab-pot is why the object is invisible in the Z-MIP (bottom left). 200
Figure 207: Target field end-marking floats. The bottom three images are MIP's through depth, eastings and northings, showing the targets and their anchors about half a meter beneath the sediment. The focus for these targets is very good despite being in the watercolumn. 201
Figure 208: Target field in sand, showing a relatively strong specular glint compared to the mud patch shown in Fig. 204. 202
Figure 209: Single transmitter configuration (top) vs. dual transmitter configuration (bottom) using the Neptune transmitters. The figure on the top was generated from a dataset in which the data corresponding to port-side Neptune transmission was withheld, effectively making the data single-source. 203
Figure 210: Amplitude values of two rows of the Z-MIP in Fig. 209 (rotated), showing the amplitude value of the row for the dual and single source configurations (top) and the location of the profile shown in the Z-MIP (bottom row plots). 204
Figure 211: Profiles of two more targets of opportunity for the dual and single transmitter configurations (top), and the context of the slices (bottom). Note that in the profile on the top right there are two closely spaced scatterers that are resolved in the dual transmitter configuration but cannot be resolved in the single transmitter configuration. 204
Figure 212: Transmitter configuration comparisons. The scan plotted on the top shows imaging results for the muddy target field using the dual Neptune configuration, and the bottom image shows the image generated using the VLA transmitter. Lines are drawn showing some of the common targets between the scans. 205
Figure 213: ROI's detected in the 2019 Sequim muddy target field. The top figure shows a MIP through a thin layer near the sediment / water interface and many targets are visible (compare with Fig. 207 where *all* of the targets are very prominent except the crab pot). The bottom two figures show a Z-MIP and Eastings-MIP with outlines around some regions of interest. Regions of interest are labelled as follows: 1 = Stake, 2 = Cylinder, 3 = Tank, 4 = 105mm HEAT, 5 = Cinder Block, 6 = buried unknown, 7 = Howitzer Shell Replica. 206
Figure 214: Rotations of MIP volume renderings corresponding to the ROI's outlined in Fig. 213. 207
Figure 215: Acoustic color projections through the real-aperture dimension created from the hollow 30 x 60 cm aluminum cylinder (left), scuba tank (second from left), "Pinocchio" mortar (2nd from right) and cinder-block (right). 207
Figure 216: An image of the target and cable array deployed in Lake Washington for the 6/2020 INS integration test. Remarkably, in addition to the deployed targets, the nylon recovery rope is also visible. 209
Figure 217: A top and 3D view of the Lake Washington PB4Y bomber rendered with alpha-blending and a color-by-depth colorization scheme. The path of the MuST went directly aft the port and starboard gun nacelles. The wings, tail, tail-fin, nacelles, and three of the engines (the fourth was lost during the failed salvage attempt) are readily visible. 210

Figure 218: Cross sections of the plane in the along-track dimension moving from aft to fore, 1 – 6. 210
Figure 219: Cross-sections through the wing and forward. The cross-section of the wing and the landing gear beneath are visible. Most of the detail of the front of the plane is lost due to forward scattering. 211
Figure 220: An across-track slice directly through the fuselage of the airplane. Resolution is considerably lower in this dimension yet many features of the plane are still very interpretable. 211
Figure 221: The first range-aliasing image created with the MuST BOSS sonar: echoes from the ship were in Zone 2 (see Fig. 86), i.e. 2 pings previous to the current. Several decks of the ship are reconstructed nicely: decks are annotated in the across-track MIP on the top and corresponding numbered horizontal slices of these levels are shown on the right. 212
Figure 222: North Sequim Bay showing the sediment coarseness, bathymetry contours, and 100 x 100m box in which testbed targets have been arrayed. 213
Figure 223: Zoomed in portion of Fig. 222 showing the testbed site and the location of the calibration line south of the target field which runs east-west. The sediment within the testbed region was dominantly fine sand, but it became muddier toward the north end of the box. 214
Figure 224: All of the targets (excluding duplicates) deployed in the Sequim Bay testbed. 215
Figure 225: Target field (red square), calibration line (green points) and scan lines (black lines) used to survey the target field. The scan lines are along heading angles 280, 340, and 40 degrees. 216
Figure 226: Towfish locations for all four days of field testing superimposed on the intended scanlines shown in Fig. 225. 217
Figure 227: Alignment scatter plots showing measured *x* and *y* displacements vs. quality (top), and absolute displacements vs. quality (bottom). On left are the scatter plots for the raw navigation data, middle is after first-stage re-navigation and the figures on the right show the results after feature-based re-navigation. 218
Figure 228: Alignment scatter plots showing measured *x* and *y* displacements vs. quality (top), and absolute displacements vs. quality (bottom). On left are the scatter plots for the raw navigation data, middle is after first-stage re-navigation and the figures on the right show the results after feature-based re-navigation. 219
Figure 229: Gridded scan mosaics generated using the first re-navigation solution generated by re-processing raw navigation data (top) and generated using the feature-based navigation solution (bottom). The nominal locations of the targets in the calibration field are marked as yellow stars, and a 2 meter radius is drawn around their nominal locations. 221
Figure 230: 3-direction orthogonal MIP views for a SAS volume generated using the path-following volume geometry. The region shown here is east of the target field and bridges a very distinct layer transition several meters below the surface layer. Several clouds of fish can also be seen floating above the sediment. 222
Figure 231: Profile-MIP for sediment transition data with (top) and without (bottom) application of the de-glinting algorithm. 223
Figure 232: Z-MIP through the volume depicted in Fig.'s 231 and 232 with (top) and without (bottom) the de-glinting algorithm. 224
Figure 233: Various targets at nadir viewed with and without de-glinting. The target on the far right is the only buried target of the set, and is ~0.3m below the surface. It is completely obscured in the MIP image without the use of the de-glinting algorithm. 224
Figure 234: MIP profile comparing the default beamforming results with a 160 ping queue and a reduced 64-ping queue. 225
Figure 235: Depth MIP comparing the default beamforming results with a 160 ping queue and a reduced 64-ping queue. 226

Figure 236: A comparison of the beamforming results with the standard configuration on the top and a fast configuration with a small queue size, integration angle, and no de-glinting or SVA sidelobe reduction is shown on bottom. 227
Figure 237: the Depth-MIP version of Fig. 236, showing comparing the standard configuration beamforming results (top) with the fast configuration (bottom). 228
Figure 238: Example MIP layer data stored for a given target with layer centers ranging from almost 90cm above the surface (top) to 90 cm. below the surface (bottom). This figure shows every other layer, 15 layers are stored in the current beamformer configuration. 229
Figure 239: Sequim bay data mosaic, layer 6 (-0.25 m relative to sediment surface). The target area is outlined in the yellow box in the north-east corner of the image. The UTM coordinates of the targets in the calibration line are depicted with yellow stars and arrayed along the southern end of the mosaic. 231
Figure 240: Sequim bay data mosaic, layer 9 (+0.125 m relative to sediment surface). The target area is outlined in the yellow box in the north-east corner of the image. The UTM coordinates of the targets in the calibration line are depicted with yellow stars and arrayed along the southern end of the mosaic. 232
Figure 241: Sequim bay data mosaic, layer 11 (0.375 m). The target area is outlined in the yellow box in the north-east corner of the image. The UTM coordinates of the targets in the calibration line are depicted with yellow stars and arrayed along the southern end of the mosaic. 233
Figure 242: Sequim bay data mosaic generated using full Z-MIP images. The target area is outlined in the yellow box in the north-east corner of the image. The UTM coordinates of the targets in the calibration line are depicted with yellow stars and arrayed along the southern end of the mosaic. . 234
Figure 243: Detection set for a flush or mostly buried UXO in the target field showing different MIP's (Z, East & North) through the UTM gridded patch cube, a *phi*-MIP through the cylindrical patch, and the location of the detection in the target field (yellow star, upper right map). The set on bottom shows the 3D acoustic-color MIP's through *phi* and the accompanying spatial MIP's pre- and post-filtering. 235
Figure 244: Detection set for the 30cm x 60cm solid aluminum slug. 236
Figure 245: The detection set for a proud cylindrical object. The improvements from AC filtering are highly visible in this example. 236
Figure 246: Detection hits for a smaller buried object that could possibly be a UXO. Several of the specular glints in this set were not filtered out by the AC filtering algorithm because the specular glint detection and removal filter described in *Materials and Methods: Structural Acoustic Feature Extraction* was not developed when these data visualizations were created. 237
Figure 247: Detections for a UXO target. 237
Figure 248: Detections and acoustic-color data for a cinder-block, flat-side up. The phi-MIP second from right shows what appear to be two distinct chambers in the object. 238
Figure 249: The hollow 2-1 aluminum cylinder. Another object was placed very near this object, interfering with the clutter rejection algorithm used at the time of analysis. 238
Figure 250: The detection set for the object adjacent the aluminum shell shown in the previous figure. 239
Figure 251: The UXO detection score for the 2020 Sequim Bay blind testbed target list created by APL-UW. The scoring of the detection list was performed by The Institute for Defense Analyses (IDA). 240
Figure 252: MuST Towbody INS error statistics (Percent distance-travelled error vs. distance travelled) compiled during the 2021 INS integration tests. 241
Figure 253: A composite Z-MIP rendering of a sunken barge in Lake Washington, collected during the 2021 INS integration tests. 242

Figure 254: An Alpha-blended rendering (top left), and MIP rendering (bottom left), generated from a multi-scan dataset. A slice through the keel captured by one scan which travelled down the fishing boat is shown on the right along with a 6-foot tall polygon model for a man giving a size reference for the boat. 243
Figure 255: The navigation screen of MuSTLive showing the large numbers of objects detected after the first day of operation. Each of the green circles inside of the testbed region, demarcated by the orange circular line, is a detection noticed by the MuSTLive operator and tagged during data acquisition. 245
Figure 256: A snapshot of MuSTLive as the vehicle traverses the northern portion of the testbed region. Several prominent scatterers are visible in the SAS MIP images on the right, one of which (near the center) is clearly cylindrical and has a tapered end. This was the fourth day of testing and the detections created from the mosaics of the previous days' data are visible in the navigation map screen on the bottom left of the figure. 245
Figure 257: The day-1 mosaic for the target field and calibration line (left) and a zoomed inset showing the targets on the eastern end of the calibration line. 246
Figure 258: The detection list created from a cursory examination of the day 1 mosaic (Fig. 257). The target field detections are shown in the orange circle. The particular view tool, ("Waypoint_edit") was developed just prior to the experiment to help view marker data and view metadata associated with scan detections. 247
Figure 259: The mosaic generated from scans collected on Day 2 of the Sequim Bay 2021 survey. The zoomed plot on the right shows the targets on the eastern edge of the calibration line. 248
Figure 260: A combined sonar mosaic generated from data acquired in days 1 and 2 of the Sequim Bay 2021 testbed survey. Several prominent objects in the test bed and calibration lines are indicated with the zoomed snippets on the right and arrows pointing to their locations in the mosaic. 249
Figure 261: A view of the calibration line using the day 1 and 2 combination mosaic. The calibration line targets are circled and the calibration line key is shown on the bottom. Greyed targets in the calibration line key indicate that the targets were buried by divers. 250
Figure 262: The Sequim Bay day 4 detection list viewed in Waypoint_edit. 251
Figure 263: Sequim Bay 2021 test region mosaic set showing the max, mean, and contraharmonic (Lehmer, p = 2) mean for the lowest layer. Note the MAX mosaic is least impacted by scan beam pattern artifacts, however it gives you no indication of scan density. The mean is the most impacted, but is also useful for determining the regions of high coverage in surveyed area. 252
Figure 264: Sequim Bay 2021 test region mosaic set showing the max, mean, and contraharmonic (Lehmer, p = 2) mean for the second to lowest layer. 252
Figure 265: Sequim Bay 2021 test region mosaic set showing the max, mean, and contraharmonic (Lehmer, p = 2) mean for a layer that includes 30 centimeters below the sediment and 10 cm above it. 253
Figure 266: Sequim Bay 2021 test region mosaic set showing the max, mean, and contraharmonic (Lehmer, p = 2) mean for a very shallow layer of sediment. 253
Figure 267: Sequim Bay 2021 test region mosaic set showing the max, mean, and contraharmonic (Lehmer, p = 2) mean for a region that is nominally above the sediment. This area should only contain features above the sediment, however it clearly contains sediment features as well. This is due to the fact that the point-spread function of the system has a slant in the vertical direction for any points other than directly at nadir. The result is that there is a projection of intensity above the interface which shows up in this layer. 254
Figure 268: Sequim Bay 2021 test region mosaic set showing the max, mean, and contraharmonic (Lehmer, p = 2) mean for a layer that nominally has 30cm of clearance relative to the water /

sediment interface. Note that the prominence of proud scatterers relative to the sediment is significantly higher in this layer vs. the previous layer (layer 5). At this level of height above the sediment most of the scattering due to the tilt of the point-spread-functions and aliased energy is significantly lower. 254
Figure 269: Sequim Bay 2021 test region mosaic set showing the max, mean, and contraharmonic (Lehmer, p = 2) mean for the full sediment MIP. 255
Figure 270: Sequim Bay 2021 layer 1 mosaic, zoomed in on part of the calibration line, set showing the max (top), mean (middle), and contraharmonic (Lehmer, p = 2) mean (bottom). 256
Figure 271: Sequim Bay 2021 layer 2 mosaic, zoomed in on part of the calibration line, set showing the max (top), mean (middle), and contraharmonic (Lehmer, p = 2) mean (bottom). 257
Figure 272: Sequim Bay 2021 layer 3 mosaic, zoomed in on part of the calibration line, set showing the max (top), mean (middle), and contraharmonic (Lehmer, p = 2) mean (bottom). 258
Figure 273: Sequim Bay 2021 layer 4 mosaic, zoomed in on part of the calibration line, set showing the max (top), mean (middle), and contraharmonic (Lehmer, p = 2) mean (bottom). 259
Figure 274: Sequim Bay 2021 layer 5 mosaic, zoomed in on part of the calibration line, set showing the max (top), mean (middle), and contraharmonic (Lehmer, p = 2) mean (bottom). 260
Figure 275: Sequim Bay 2021 layer 6 mosaic, zoomed in on part of the calibration line, set showing the max (top), mean (middle), and contraharmonic (Lehmer, p = 2) mean (bottom). 261
Figure 276: Sequim Bay 2021 full depth MIP mosaic, zoomed in on part of the calibration line, set showing the max (top), mean (middle), and contraharmonic (Lehmer, p = 2) mean (bottom)..... 262
Figure 277: Color-by-depth MIP's (side view) through target ROI's for the Sequim Bay UXO testbed. A.1 – A.2: 155mm Howitzer shells, A.3 – B.1: 105mm HEAT rounds. B.2 – C.1: 105mm M60 shells. C.2 – D.2: 81mm M821 finned mortar rounds. D.3 – F.1: 81mm M889A1 mortar rounds. F.2 – F.3: 60mm M49 rounds..... 263
Figure 278: Color-by-depth MIP's (tilted) through target ROI's for the Sequim Bay UXO testbed. A.1 – A.2: 155mm Howitzer shells, A.3 – B.1: 105mm HEAT rounds. B.2 – C.1: 105mm M60 shells. C.2 – D.2: 81mm M821 finned mortar rounds. D.3 – F.1: 81mm M889A1 mortar rounds. F.2 – F.3: 60mm M49 rounds. 264
Figure 279: Color-by-depth MIP's (top) through target ROI's for the Sequim Bay UXO testbed. A.1 – A.2: 155mm Howitzer shells, A.3 – B.1: 105mm HEAT rounds. B.2 – C.1: 105mm M60 shells. C.2 – D.2: 81mm M821 finned mortar rounds. D.3 – F.1: 81mm M889A1 mortar rounds. F.2 – F.3: 60mm M49 rounds. 265
Figure 280: 155mm Howitzer shell facsimile. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum. 266
Figure 281: 155mm Howitzer shell. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum. 267
Figure 282: 105mm HEAT round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum. 268
Figure 283: 105mm HEAT round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect

composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum. 269
Figure 284: 105mm HEAT round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum. 270
Figure 285: 105mm M60 shell. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum. 271
Figure 286: 105mm M60 shell. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum. 272
Figure 287: 105mm M60 shell. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum. 273
Figure 288: 105mm M60 shell. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum. 274
Figure 289: 81mm M821 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum. 275
Figure 290: 81mm M821 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum. 276
Figure 291: 81mm M821 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum. 277
Figure 292: 81mm M821 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum. 278
Figure 293: 81mm M821 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum. 279
Figure 294: 81mm M889A1 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum. 280

Figure 295: 81mm M889A1 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum. 281
Figure 296: 81mm M889A1 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum. 282
Figure 297: 81mm M889A1 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum. 283
Figure 298: 81mm M889A1 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum. 284
Figure 299: 81mm M889A1 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum. 285
Figure 300: 81mm M889A1 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum. 286
Figure 301: 60mm M49 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum. 287
Figure 302: 60mm M49 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum. 288
Figure 303: 3D Mondrian detector results with color-by-depth encoding using the color-scheme on the right, and radially encoded prominence (i.e. a wider radius = more prominent target). The prominence threshold value is $\tau = 4.5$. 290
Figure 304: 3D Mondrian detector results with color-by-depth encoding using the color-scheme on the right, and radially encoded prominence (i.e. a wider radius = more prominent target). The prominence threshold value τ has been reduced to 4.0. 291
Figure 305: The entire set of 3D Mondrian detections for the Sequim Bay 2021 survey captured with τ = 3.5. 292
Figure 306: Example of an across-track MIP through beamformed data crossing through the testbed. Note that there doesn't seem to be a strong sediment interface ~0.5 meters below the interface, despite there being an increased number of prominent scatterers collecting at this depth as observed in Fig. 305. 292
Figure 307: A spatial histogram (25 x 25 cm bins) of the 3D Mondrian detections. Clusters of detections can be seen in regions containing prominent targets. The calibration line is visible in the southern end of the displayed region. 293

Figure 308: A spatial histogram (25 x 25 cm bins) of the 3D Mondrian detections *after a position update using feature-based navigation correction*. The detection clusters visible in Fig. 307 now collapse to very tight spreads. 294
Figure 309: 2021 Sequim Bay ROC curve for the APL-UW target list. 295

List of Acronyms

2D:	Two Dimensional
3D:	Three Dimensional
AM:	Ante-Meridian
APL-UW:	The Applied Physics Laboratory, University of Washington
ARL-PSU:	The Applied Research Laboratory, Pennsylvania State University
AUV:	Autonomous Underwater Vehicle
BP:	Back-Projection
BOSS:	Buried Object Scanning Sonar
CBA:	Color-by-aspect
CBD:	Color-by-depth
COTS:	Commercial off-the-shelf
CPU:	Central Processing Unit
CSAS:	Circular Synthetic Aperture Sonar
CW:	Continuous wave
dB:	Decibels
DoD:	Department of Defense
DVL:	Doppler Velocimeter
eBOSS:	EdgeTech BOSS
ENU:	East – North - Up
ES:	Executive Summary
ESTCP:	Environmental Security Technology Certificate Program
FFBP:	Fast Factorized Back-Projection
FFT:	Fast Fourier Transform
GB:	GigaByte
GPS:	Global Positioning Satellite

GPU:	Graphics Processing Unit
GUI:	Graphical User Interface
HF:	High Frequency
IDA:	Institute for Defense Analyses
IFFT:	Inverse Fast Fourier Transform
INS:	Inertial Navigation System
IRLS:	Iterative least squares
LAT:	Latitude
LBL:	Long baseline acoustic positioning
LF:	Low Frequency
LFM:	Linear frequency modulated
LON:	Longitude
MB:	MegaByte
MEMS	Micro-electromechanical systems
MIP:	Maximum Intensity Projection
MR	Munitions Response
MuST:	Multi-Sensor Towbody System
NED:	North-East-Down
NOAA:	National Oceanic and Atmospheric Administration
ONR:	Office of Naval Research
PM:	Post-Meridian
PMA:	Post Mission Analysis
PNNL:	Pacific Northwest National Laboratories
PoSSM:	Point-based Sonar Scattering Model
PSF:	Point-Spread-Function
PST:	Pacific Standard Time
RAS:	Real Aperture Sonar
ROC:	Receiver Operating Characteristic
ROI:	Region of Interest
RV:	Research Vessel

SAS:	Synthetic Aperture Sonar
SAR:	Synthetic Aperture Radar
SERDP:	Strategic Environmental Research and Development program
SNR:	Signal-to-Noise Ratio
TCP/IP:	Transmission Control Protocol / Internet Protocol
UI:	User Interface
USBL:	Ultra-short baseline acoustic positioning
UTM:	Universal Trans Mercator
UXO:	Unexploded Ordnance
VLA:	Very Large Array
WLS:	Weighted Least Squares

Keywords

Buried Munitions, Unexploded Ordnance, UXO, Buried Object Scanning Sonar, Synthetic Aperture Sonar, SAS, Beamforming

Acknowledgements

The work described in this document was funded by the Strategic Environmental Research and Development (SERDP) program under project number MR18-1051. The investigator would also like to thank the APL-UW Ocean Engineering group for their support of the development and deployment of the Multi-Sensor Towbody (MuST): Ben Brand, Nicholas Michel-Hart and Tim McGinnis. Additionally the investigator would like to thank Daniel Plotnick, co-PI through the early portion of this project, Kevin Williams, PI of ESTCP project MR18-5004, Pacific Northwest National Laboratories (PNNL) for collaborating to deploy a UXO test field in Sequim Bay, and Jason Sara, Rick Babicz and Kevin Rychert at EdgeTech for their excellent hardware and software support with the eBOSS.

Abstract

Many unexploded ordnance (UXO) in the aquatic environment are buried under seafloor sediment, making them difficult to detect and remediate. Though buried, these UXO can resurface after storms or be encountered during dredging. Low frequency, sediment penetrating acoustics has the potential to improve our ability to localize, detect and classify buried UXO, aiding in remediation. The aim of this project was to improve the UXO detection and classification capabilities of the Buried Object Scanning Sonar (BOSS), a low frequency, sediment penetrating sonar, by improving the beamforming, target response extraction, and data visualization methods used with this sonar.

Research initially focused on improving the beamforming methodology associated with data collected by the eBOSS. Back-projection of planar array data to a volumetric image can be a computationally burdensome task. A goal of this project was to enable the generation of complex three-dimensional imagery from which spectral response information of targets could be extracted and retain the generality and flexibility of time-domain beamforming while still remaining computationally efficient enough to be useful both to researchers and the end users of the system. The beamformer, which leveraged aperture factorization and GPU acceleration to achieve these goals, was initially developed using a basic point-scatterer simulator created as part of this project. Following testing on simulated data, the beamformer was tested with real data captured by two BOSS systems, one an AUV mounted system and the other a towed system which was actively being developed over the course of this project under ESTCP project MR18-5004 ("Multi-Sensor Towbody (MuST) for Detection, Classification, and Geolocation of Underwater Munitions").

The original project plan called for extensive research focusing on developing data-driven autofocusing procedures to improve the acoustic data products and augment the navigation. Due to multiple factors, however, it was found that the data-driven autofocusing was unnecessary, and research was diverted to feature-based mosaic alignment and multi-aspect data fusion.

Following successful beamformer development and testing, research branched into several related threads: 1) improving the signal-to-noise ratio of target spectral features by reducing the influence of volume scattering and specular glints from sediment interfaces, 2) interfacing with current SERDP-funded UXO classification efforts, and 3) user-oriented data processing, organization and data visualization methods for improving data interpretability, a large task that included everything from live data visualization for decision making, to post-processing tools for beamforming, filtering, viewing, combining and exporting data products. The final product resulting from these research threads is the "MuST toolbox," software library which greatly improves the UXO-remediation capabilities of the BOSS system.

Executive Summary

Introduction

Many unexploded ordnance (UXO) in the aquatic environment are buried under seafloor sediment, making them difficult to detect and remediate. Though buried, these UXO can resurface after storms or be encountered during dredging. Low frequency, sediment penetrating acoustics has the potential to improve our ability to localize, detect and classify buried UXO, aiding in remediation. Several sensors with different transmitter and receiver array designs have been developed for this purpose; the primary examples being the Buried Object Scanning Sonar ("BOSS," [1-5]), and the sediment volume search sonar (SERDP project MR-2545, PI: Daniel Brown).

The Buried Object Scanning Sonar combines an across-track real aperture and an along-track synthetic aperture with a down-looking transmitter geometry to create a ribbon-like aperture that can create volumetric images within the sediment near nadir (typ. +/- 60 degrees relative to nadir). A visual example of a BOSS array is shown in Figure ES.1.

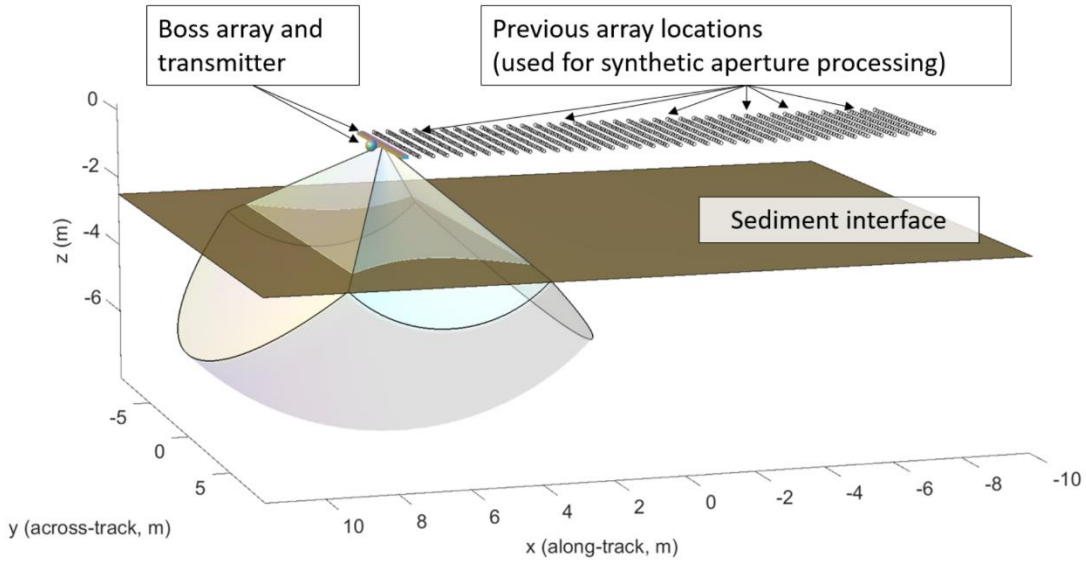

Fig. ES. 1: An illustration of the BOSS volumetric acoustic imaging. An across-track array is combined with an along-track synthetic aperture to create volumetric images of the sediment beneath the sonar. The figure shows the volumetric region imaged by a single real-aperture ping (with no along-track resolution). The synthetic aperture imaging result is an extended volumetric region tracking the location of the sonar and imaging the sediment beneath its trajectory.

Systems like the BOSS do not directly create images, but rather, images must be created in post-processing via beamforming. Beamforming is of significant importance to this project, and an example of its necessity is shown in Fig. ES. 2, which displays a rendering of raw, un-processed BOSS data corresponding to a scan of a prominent UXO, and a synthetic aperture image created by applying the necessary post-processing to create an image.

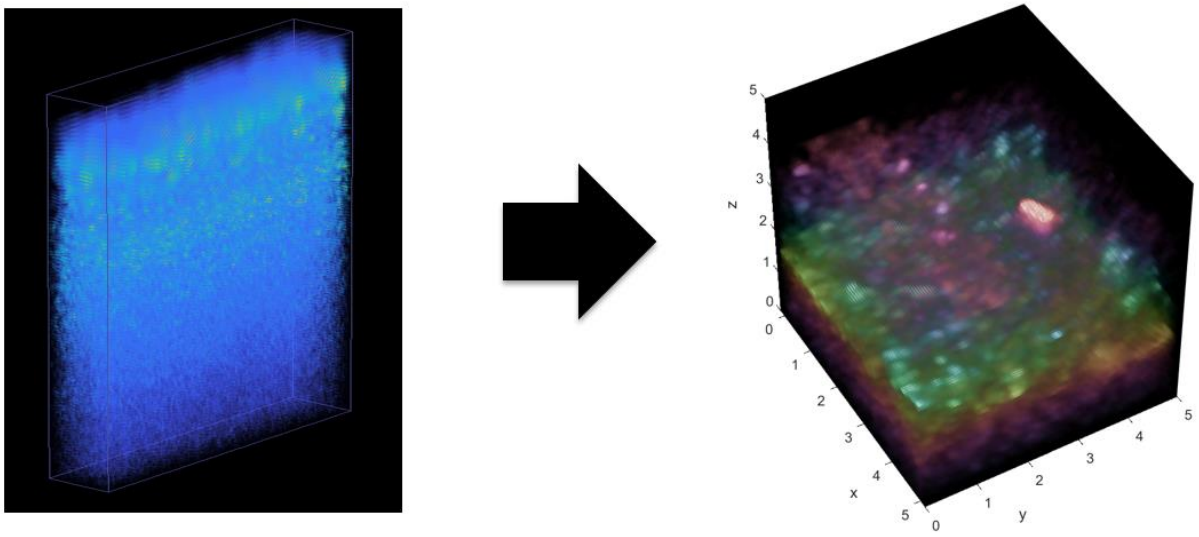

Fig. ES. 2: A volume rendering of raw data captured directly from the BOSS system (left) and a post-processed image (right). Note that though the raw data contains backscattering information from a very prominent UXO target it is essentially inaccessible in its raw form.

The beamformer plays a significant role in the quality of reconstructed imagery, the computational burden, and the imposition of limitations on further post-processing capabilities (e.g. how the data is beamformed can enable or prevent spatial-spectral analysis). In addition to creating volumetric imagery, the BOSS system has the potential to make measurements of the structural-acoustic frequency response of objects; this is especially the case if the system exploits broad-band transmitters. Unfortunately, both spatial domain and structural-acoustic frequency response features are often difficult to detect due to the high levels of volume scattering from the sediment that the scattering signatures are often competing against.

Another challenge faced by BOSS systems includes the very high computational burden associated with beamforming: the number of interpolation operations required to backproject data acquired by an N x N element sonar array to an N x N x N voxel grid is proportional to N^5, and small changes in voxel size or imaging volume can quickly scale the computational burden. Fast frequency domain beamforming methods could be exploited, however these have strict assumptions regarding array geometry. A beamformer with the speed of a Fourier approach but *generality* (i.e. flexibility with regard to imaging geometry) of a time-domain approach is a technology gap that needs to be filled.

Additionally, volumetric data – especially noisy volumetric data – can be challenging for end users to interpret. Unlike two dimensional data, which can be easily projected to standard computer screens, three dimensional data can't be easily visualized on screens, and effective information conveyance requires additional thought. These technology gaps were the primary motivation for the research performed in MR18-1051.

Objectives

The objective of this project is to enhance the ability of BOSS systems to detect and classify UXO by improving the acoustic data products created by the system and the methods by which

data, including spatial and structural acoustic features of detected targets, is extracted and interpreted.

More specifically, a primary objective of this project is to create an efficient and flexible time-domain beamforming algorithm that enables the large quantities of information collected by the BOSS system to be rapidly and accurately beamformed. This would enable researchers and system end-users to iterate beamforming, mosaicking, data representation, collection and post-processing tasks rapidly without having the beamformer be a productivity hindrance. From a research standpoint, this accelerates the development and testing of new algorithms and the creation of data libraries, and from a user standpoint, it improves the timeline between when data is collected and targets can be localized and identified.

Additionally, this project aims to improve on existing approaches for extracting the pose-dependent structural acoustic frequency response of detected objects by reducing the impact of volume scattering, nearby clutter, and interface specular contributions. Further goals include accurate multi-aspect data fusion, effective 3D data information portrayal, live data visualization for aiding decision making, and precision earth-coordinate localization for streamlining remediation.

Technical approach

Fig. ES. 3 is a flowchart showing the originally proposed technical approach for MR18-1051.

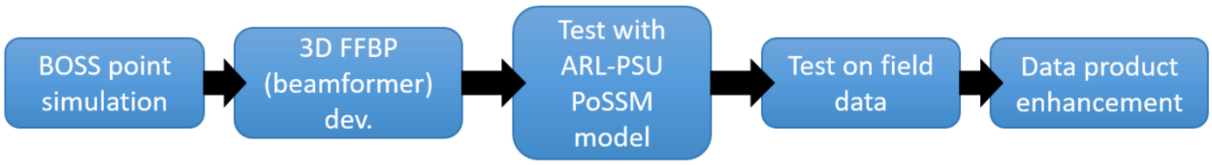

Fig. ES. 3: The proposed technical approach timeline for BOSS signal processing development in MR18-1051.

The project began with the creation of a Matlab-based BOSS simulator ("BOSSim") for the purpose of creating data with which prospective beamforming algorithms could be evaluated. The simulator allowed data to be generated using arbitrary BOSS array geometries, transmitter configurations, track trajectories, and point-scatterer distributions. Beamformers are primarily evaluated using point-spread-function (PSF) analysis and therefore the simulator was limited to beamforming sparse sets of points distributed around the simulated trajectory and are assumed to be suspended in an infinite medium with homogeneous propagation speed.

During the development of the simulator the mathematical framework for an extension of an accelerated time-domain imaging method called Fast Factorized Backprojection (Ulander et al., [18]), to the three dimensional BOSS imaging geometry was developed. Matlab implementations of the new beamformer (3D FFBP) and standard time-domain backprojection were then evaluated and compared using data from the simulator.

Standard FFBP leverages pair-wise sub-aperture image combinations of increasing size to achieve computational reductions. There were multiple routes toward extension of this idea to three dimensions that needed consideration. Eventually, due to the extremely fast manner in

which real-aperture imagery could be generated using standard approaches, factorization in only the synthetic-aperture dimension was implemented and a unique "asymmetry correction" for the bistatic BOSS transmitter and array geometries was developed that greatly improved the efficiency of the first several aperture factorization stages. In the end it was found that the BOSS imaging geometry was very amenable to factorized time-domain beamforming and the reduction in the number of computations required to create imagery was enormous; examples will be shown in the results section.

Though FFBP was found to be both efficient and general, an accurate knowledge of the platform trajectory is necessary to leverage its generality. In the original project proposal, research was planned towards correcting residual navigation error in an autofocusing framework; i.e. inverting track navigation corrections from data-driven observations of image focus aberrations. This was to be done via ARL-PSU's PoSSM simulator [15 - 17], which would be leveraged for algorithm development by providing realistic data with sediment refraction and many more scatterers. There was a significant delay in getting the PoSSM simulator, but this ended up being providential. Prior to receiving PoSSM, real data was acquired which had all of the necessary characteristics to evaluate beamformer performance and observe residual focus aberrations. Significantly, it was found that the position estimates provided by the onboard navigation system coupled to the BOSS were sufficient for creating well focused imagery without any additional navigation refinement. In fact, this was found to be true for *all* subsequent datasets acquired by BOSS systems integrating high quality commercial off-the-shelf inertial navigation systems: no additional data-driven processing was found to be necessary for creating well-focused imagery. A likely explanation for this lies in the combination of long wavelengths used by the BOSS and the short synthetic apertures used to create images. The ramifications were extensive: the effort planned toward developing autofocus algorithms, testing with PoSSM and then migrating to real datasets was instead put toward development of feature-based, data-driven renavigation routines that would allow multiple datasets from multiple scans to be combined in a manner that would reduce mosaic speckle and improve the signal-to-noise ratio of targets.

Subsequently, following successful testing of the beamformer on real data acquired by an AUV, the Multi-Sensor Towbody (ESTCP project MR18-5004, [11]) became available as a collection platform for BOSS data. Via collaboration with the MuST project various 3D FFBP beamformer implementations were tested, including one version which was integrated into a real-time data visualization software tool. The MuST system was deployed in Sequim Bay in 2019, 2020, and 2021 to survey UXO testbeds positioned by Pacific Northwest National Laboratories. These survey experiments were the optimal platform for development of live and post-processing tools for the data-product enhancement portion of this project, and via the Sequim Bay surveys the MuST Toolbox, which includes software tools for live data visualization and target tagging, batched SAS post-processing, mosaicking, feature-based "renavigation," and post-mission analysis, which includes structural acoustic feature enhancement tools, multi-aspect data fusion, data organization and other visualization tools, was developed.

Results and Discussion

Data simulation and beamformer testing

Fig. ES. 4 shows an example of a simulated BOSS array and platform trajectory, along with a set of points distributed around the array corresponding to the locations of simulated scatterers. Below the illustration of the simulated track, maximum-intensity-projections through the raw data generated by the simulator are also shown.

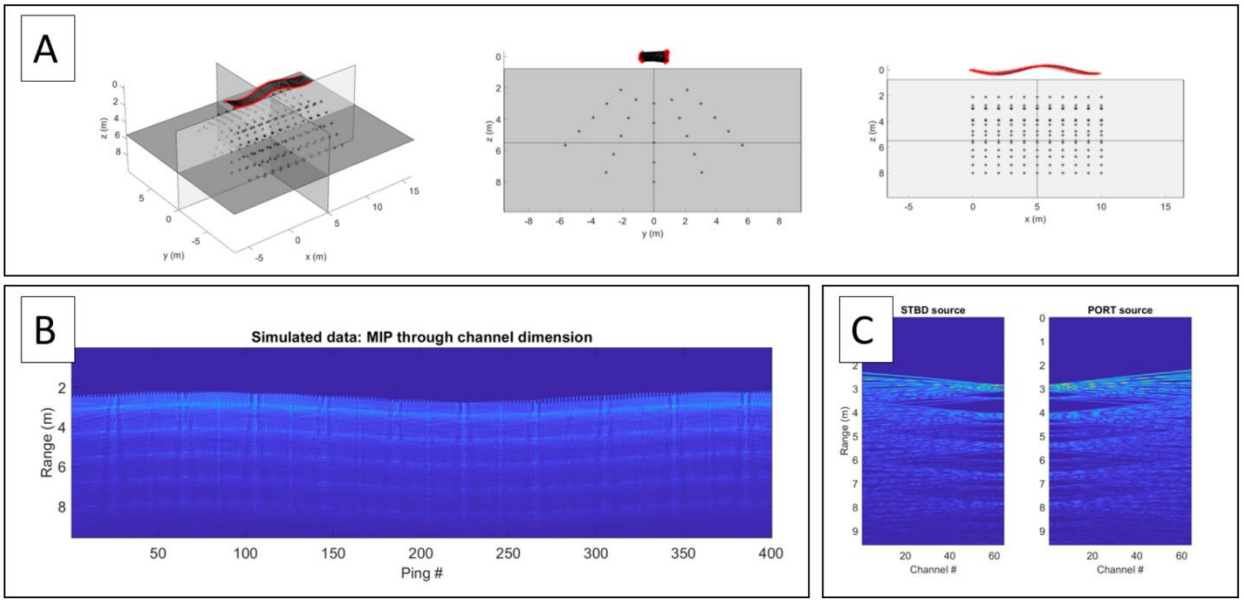

Fig. ES. 4: (A) shows the simulation array and trajectory geometry and scatterer distribution. (B) an (C) show maximum-intensity-projections (MIPs) through the array dimension (B) and along-track dimension (C) of the produced dataset. Note that the simulated system has an alternating port-starboard transmitter configuration and the port and starboard data channels have been separated in (C).

The accelerated 3D beamformer, 3D FFBP, was used to beamform the simulated data shown in Fig. ES. 4. For comparison, the data was also beamformed using a standard time-domain backprojection beamformer. A MIP rendering of the volume image created with FFBP from the previously depicted simulation is shown in Fig. ES. 5 alongside some quantitative comparisons with standard backprojection beamforming results. Standard backprojection was found to be impractically slow, however, so comparisons were performed by beamforming a single vertical slice through nadir using standard backprojection and comparing that with the common slice created using 3D FFBP.

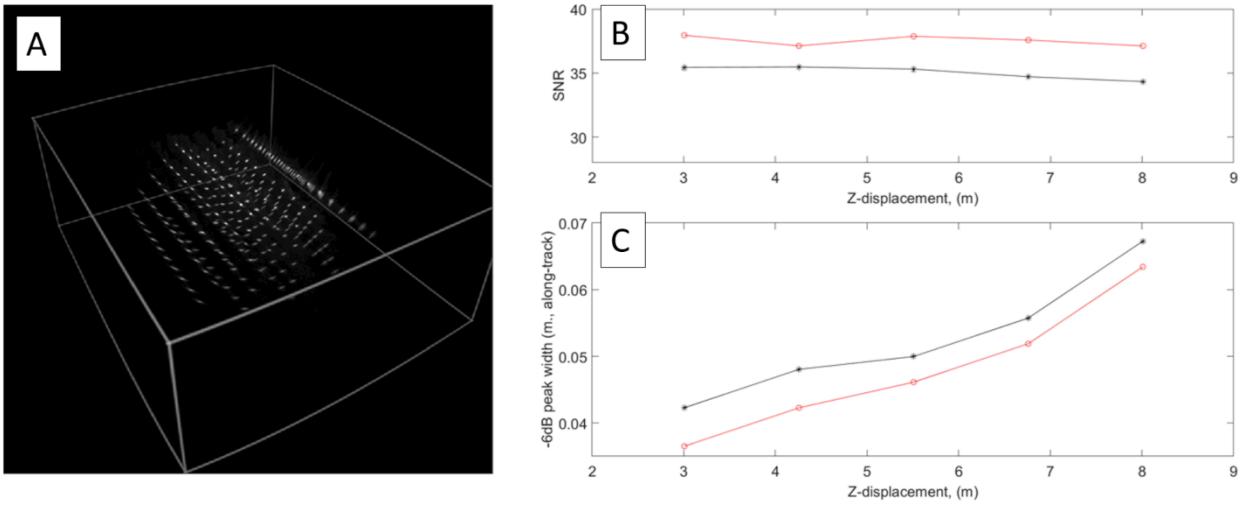

Fig. ES. 5: (A) A MIP rendering of the image created by beamforming the data shown in Fig. ES. 4 using the 3D FFBP algorithm. (B) is a quantitative comparison of the SNR (peak relative to interpolation noise background) for backprojection (red) vs. 3D FFBP (black) for points at different ranges at nadir, and (C) shows the resolutions of the point spread functions for backprojection (red) and 3D FFBP (black)

Qualitatively, 3D FFBP does not result in any major imaging artifacts, and the nonlinear track perturbations visible in Fig. ES. 4 do not produce focus aberrations in Fig. ES. 5A. Quantitatively, a couple dB of SNR is sacrificed relative to standard backprojection, however it is still sufficient high that the dynamic range in an actual field experiment will be set by reverb noise, ambient noise and volume scattering. The resolution of backprojection is also slightly higher, (Fig. ES. 4C). There are two likely sources for this discrepancy: the 3D FFBP implementation had integration angle limits of -60 to +60 degrees. In contrast, the backprojection implementation contained no integration angle limitations and therefore wider coverage in the wavenumber spectrum could be achieved. Wider wavenumber spectrum coverage results in higher resolution. Secondly, the multiple interpolation operations performed on each sub-aperture can cause minor resolution degradation depending on interpolator type; linear interpolation, for example (used in Fig. ES. 5), can attenuate high spatial frequencies.

	Completion time (s)	Standard dev. (s)
Back projection (linear)	859,476	1856
Back projection (spline)	924,690	14362
FFBP (linear)	344	26
FFBP (spline)	892	10

Table ES. 1: The completion times using different beamformers (backprojection and 3D FFBP) and interpolators (linear and spline) required to image the volumetric data shown in Fig. ES. 4.

Table ES. 1 compares the beamforming completion times of backprojection and 3D FFBP using different interpolators. Backprojection was excessively slow, therefore a limited number of pings was beamformed and an estimate for the duration of the entire beamforming process was made using these pings. Error statistics for the beamforming times in the form of standard deviations

were made by running the beamformer multiple times. Note that the backprojection operation was not memory limited by the machine; the bottlenecks of both implementations were CPU speed. Further hardware acceleration leveraging graphics processing units (GPU's) would be accomplished later in the project.

Field test results: AUV data

Figure ES. 6 shows examples of some of the later sub-aperture stages being combined in spherical coordinates to produce a single, full-resolution image which is then converted to Cartesian coordinates for summation with the rest of the imaged data. Typically, an array would be broken into 256 ping chunks, and each chunk would be beamformed using the progressive sub-aperture combination approach shown in Fig. ES. 6. The resulting Cartesian images could be combined to form a 3D volumetric rendering of the entire aperture. Figure ES. 6 contains well focused scattering formations from sand-ripples on the sediment. Importantly, none of these sand-ripple formations, or any of the discrete scatterers that appeared in other portions of the data, exhibited noticeable along-track blurring that would result from navigation error, despite the heavy track nonlinearity very visible in Fig. ES. 7.

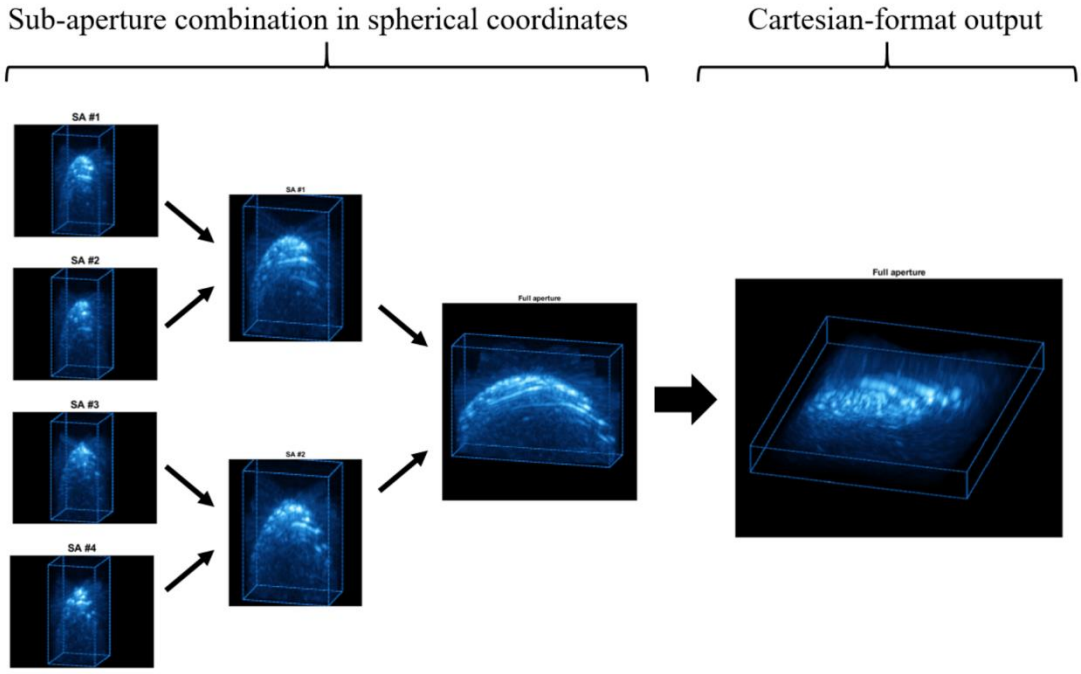

Fig. ES. 6: An illustration of the 3D FFBP beamforming process, using BOSS data from an AUV travelling over some sand ripples. The four images on the left are low resolution images corresponding to four non-overlapping 64 ping sub-apertures of a 256 ping data chunk. These sub-apertures are pairwise combined to form two higher resolution 128 ping sub-aperture images. This final pair is combined to form a full 256 ping spherical coordinate image, which is then interpolated into Cartesian coordinates (far right). The Cartesian format image can be combined with other chunks to form a volume rendering of the whole scan.

Figure ES. 7 shows a geo-rectified Z-MIP (i.e. projection through the height dimension) of AUV BOSS data mosaicked from two separate scans. Noticeable features of the mosaic include patches of sand ripples, a buried pipe, and two spherical markers at the end of the pipe. Due to the imaging geometry of the BOSS array the point-spread-function (PSF) is highly anisotropic,

and resolution is much higher in the SAS (i.e. along-track) dimension than the across-track dimension. No significant along-track blurring is visible in the mosaicked data despite the obviously strong non-linearity of both AUV scan trajectories.

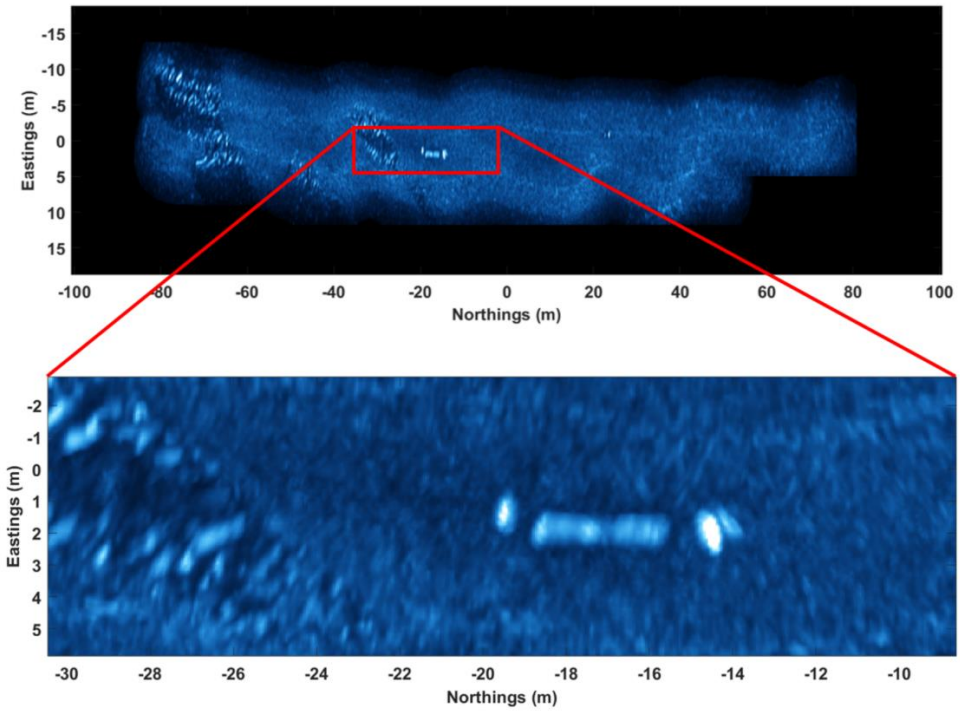

Fig. ES. 7: Mosaics of two datasets from the AUV BOSS (top) and a zoomed view of a target, a buried pipe (bottom). Two spherical end-markers were placed on the sediment surface at the ends of the pipe and are visible in the image. Note the large across-track oscillations of the AUV trajectory, which are visible in the mosaic. Despite the strong track nonlinearity, the image retains good focus.

Field test results: Multi-Sensor Towbody

The Multi-Sensor Towbody was deployed multiple times between 2019 and 2021, and a large amount of data was collected for analysis and signal processing development. Figures ES. 8-9 show different renderings of a scan conducted over a target field arrayed in a muddy region of Sequim Bay in 2019. It was soon found that, when imaging volumes with multiple sediment layers and large amounts volume scattering, gating projections in the height-dimension was necessary to image all but the brightest objects. This led to a standard practice of compressing and exporting geo-rectified volumetric datasets as a series of discrete projections through ~half meter layers. In Figure ES. 8, all of the UXO targets are visible in the deployed array of items, however all but the pipes at the ends of the array were hidden if a Z - projection through the entire sediment volume was used to render the data.

In contrast, the projection in the Eastings-dimension (approximately across-track) shown in Fig. ES. 9, easily reveals the proud and buried targets deployed in the target array, but their relative positions in northings and eastings are not visible in this view.

Note that in Figure ES. 8 the spherical floats at the north end of the array (outlined in the white box) have essentially perfect focus. This implies successful beamforming, but also negligible motion of the scatterers in the watercolumn.

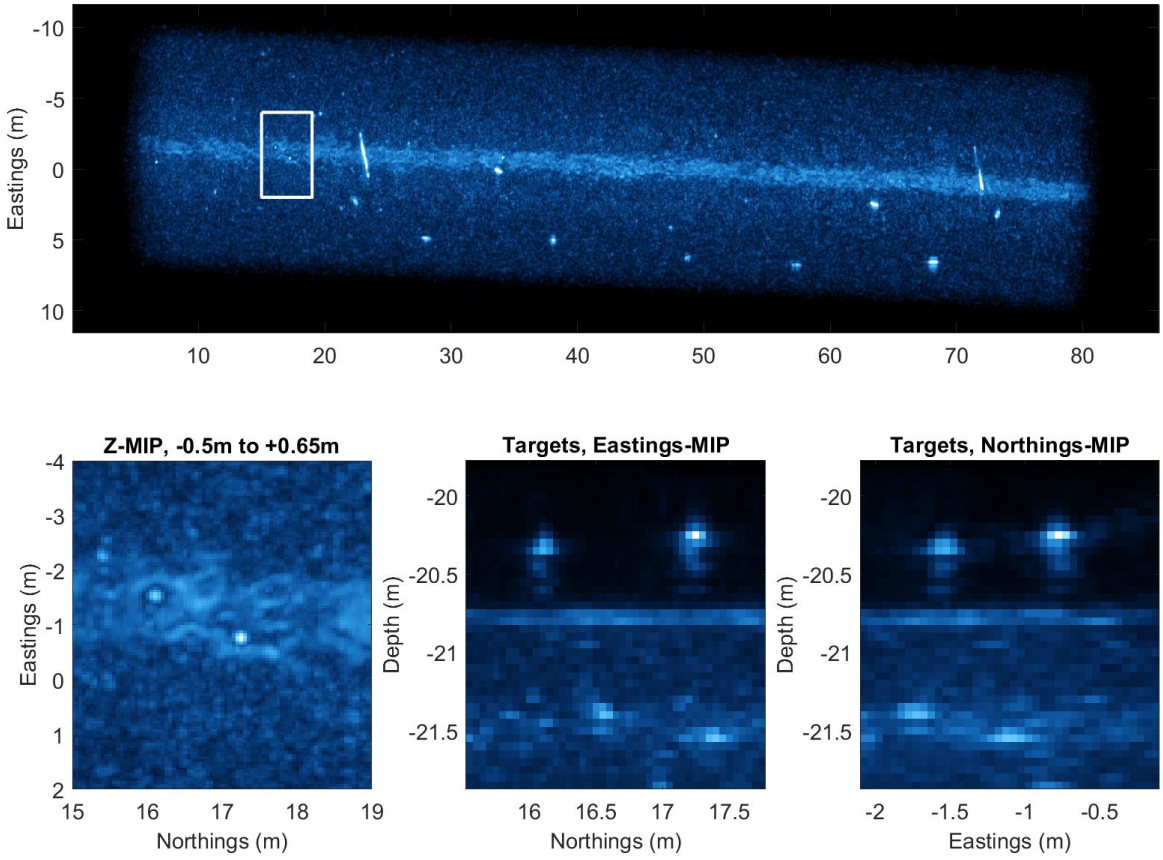

Fig. ES. 8: The top figure shows a near-surface Z-MIP of a target field arrayed in a muddy region in Sequim Bay. Targets are readily visible against the background. The three plots on the bottom show a zoomed in view of a region near the end of the scan (left) containing small spherical floats, and the other two plots are projections through northings and eastings of these floats. The focus of the imagery is superb, and the anchors tethering the floats are visible in the mud.

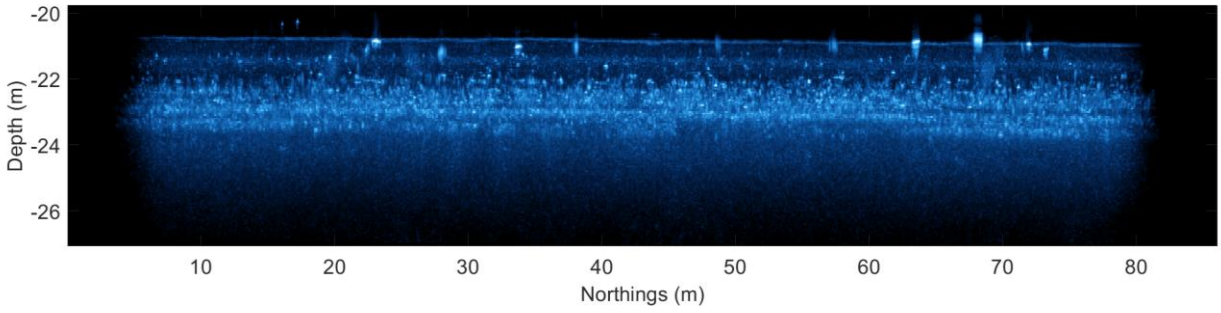

Fig. ES. 9: A projection through the Eastings dimension of the georectified volumetric SAS scan of the muddy target field. The targets, both proud and buried, are visible in the upper layers of the sediment. Different sediment layers are readily visible beneath the initial layer. The floats protruding above the sediment are also visible.

As more data was acquired, more post-processing tools were developed. Figure ES. 10 shows a set of multi-aspect data-products, created by fusing the spatial and enhanced (i.e. filtered to improve the SNR) structural acoustic spectral representations of a target captured by multiple scans performed during the 2021 Sequim Bay field test.

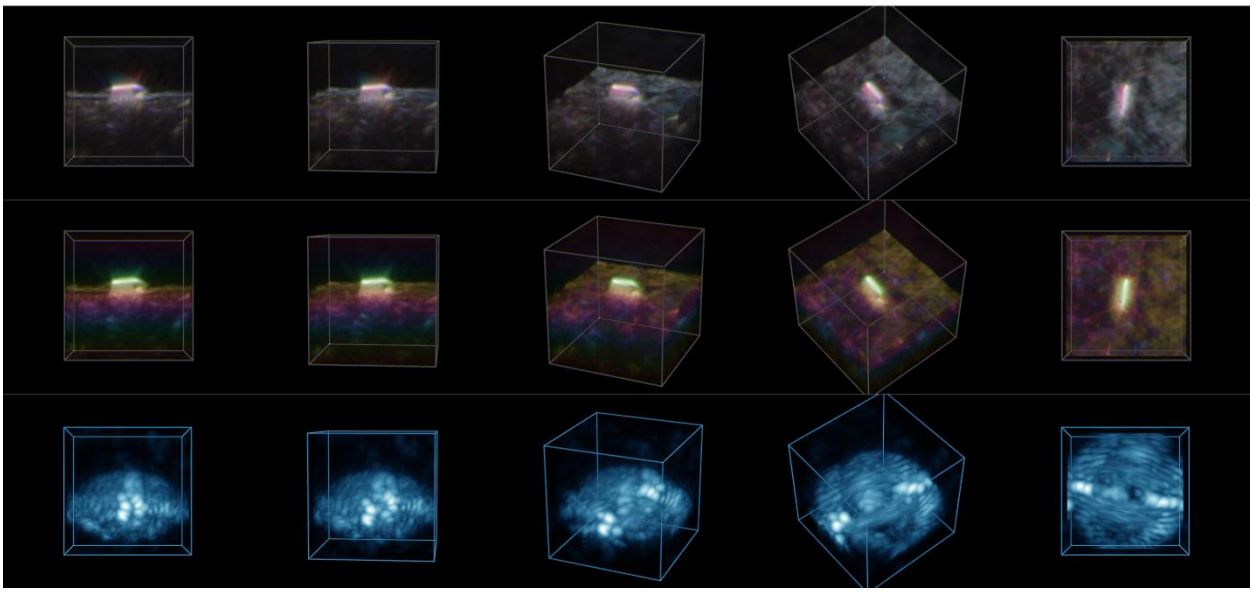

Fig. ES. 10: Color by aspect (top), color by depth (middle) and 3D spatial-spectral frequency response (bottom) fusions of multi-aspect data captured for a 155mm UXO facsimile placed in the Sequim Bay 2021 testbed. Note the strong linear spectral feature in the bottom row which corresponds to the broadside glint.

Figure ES. 10 is an example of some of the post-processing tools created to enable end-users and target recognition algorithms to access and evaluate data associated with detections made by the BOSS system. These tools and others were added to a post-processing software library, the "MuST Toolbox," designed to streamline the process of creating and evaluating the SAS data products with the BOSS system on the MuST.

Implications for Future Research and Benefits

The fast and generalized 3D FFBP beamformer will be useful for accelerating beamforming tasks associated with any volumetric sonar exploiting a similar imaging geometry to the BOSS systems evaluated in this project. Related to the MuST, it facilitated the development of both live and post-processing tools that enabled high resolution imaging of data, as well as structural acoustic analysis in post-processing. Related to UXO detection and classification, both from an operator standpoint and automatic classification standpoint, multi-aspect data fusion was found to provide huge benefits. These benefits were likely due to the more complete set of target-specific features (both in the spatial and spectral domains) observed by ensonifying the target from different orientations. Multi-aspect data fusion requires extremely precise relative localization of data captured by different scans, and is only enabled by data-driven, feature-based navigation, a process that is highly dependent on environment parameters. It is recommended that future research be conducted to make this a reliable post-processing technique.

Objective

According to the Statement of Need, US governed coastal regions potentially contaminated by UXO exceeds 10 million acres. These UXO pose a threat to humans and the environment. Many of these UXO are buried, though they can still potentially be encountered by dredging operations or by migration during storm events during which they may be uncovered or washed ashore. Remediating these UXO is important, however the associated cost is very high, making false alarm reduction a significant performance factor. Low frequency acoustics is a technology that can be used for detecting and classifying objects buried in the seabed, and one of the first low frequency systems designed explicitly to image and detect UXO-scale objects buried within the sub-sediment volume is the Buried Object Scanning Sonar. The objective of this project is to enhance the ability of the BOSS system to detect and classify UXO by improving the acoustic data-products created by the system and the methods by which data is interpreted. Specifically, an objective of this project is to make a very efficient and high fidelity beamforming algorithm that enables the large quantities of information collected by the BOSS system to be rapidly and accurately beamformed, enabling researchers to iterate beamforming, mosaicking, data representation, collection and post-processing ideas rapidly, without having the beamformer be a major bottleneck to further inquiry. Additionally, this project aims to improve methods by which structural acoustic information related to targets, for example, the pose-dependent frequency response, are extracted from detections within beamformed imagery, by reducing the impact of volume scattering, nearby clutter, and interface specular contributions. Further goals include accurate multi-aspect data fusion, effective 3D data information portrayal, live data visualization for aiding decision making, and precision earth-coordinate localization for streamlining remediation.

Background

A review of extant BOSS-related literature reveals that, while the system is capable of detecting buried objects on size scales relevant to UXO detection, shape features can be difficult to extract from imagery. To illustrate, the figure below shows images gathered from several publications related to BOSS-like systems, spanning from the early 2000's to more recent studies [1 – 5]

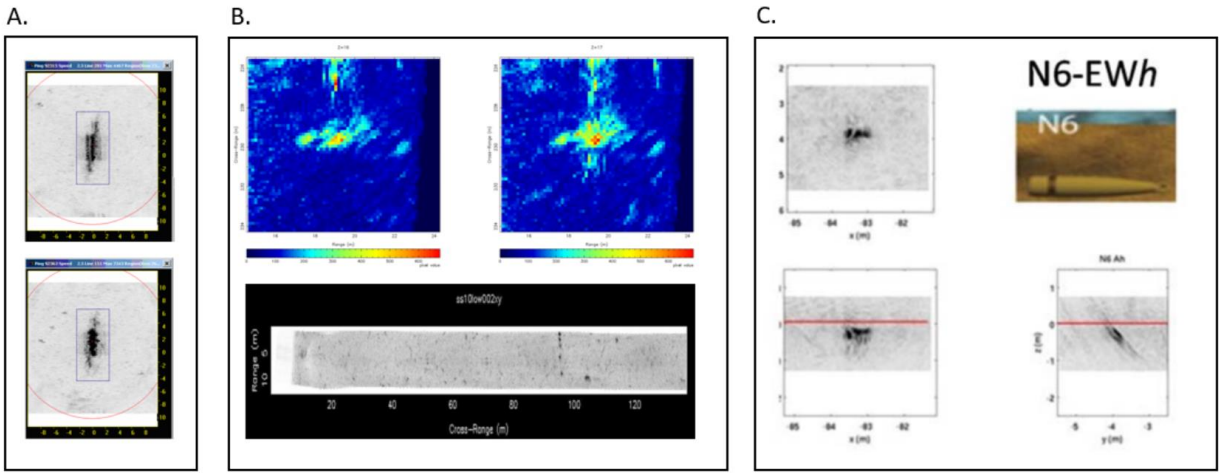

Figure 1: Example BOSS images gathered from the literature. Fig. A. is from one of the early papers discussing synthetic aperture processing on BOSS systems and shows two images generated using a 30-ping integration length [2]. The plots in Fig. B were gathered from SERDP project MM-1533 [3] and the plots in Fig. C. are from SERDP project MR-2103.

While low-frequency images are often unintuitive, it has been shown via multiple field and laboratory tests, as well as simulations, that many targets display useful shape features even at very low frequencies, especially when viewed from certain aspect angles. The examples below were taken from the final report of [6], a project in which methods for obtaining high resolution low-frequency acoustic imagery are detailed.

Figure 2: A 120-degree field of view low-frequency acoustics image of a solid 30 by 60 cm. aluminum cylinder, produced using simulated data [7], and viewed to catch the broadside and end-on glints (A). Figure B. shows narrow-aperture low-frequency "looks" at the same target but using real data from a field experiment, and Figure C. shows a three-dimensional low-frequency acoustic image of the same target as in (A) and (B). All figures are from [6].

The systems that generated the images in Fig.'s 1 and 2 are both low-frequency acoustic systems, however the image quality and features recovered for targets differ greatly. Some of these differences have to do with the imaging geometry (the system used to make Fig. 2 observes targets from low grazing angles, and the system used to make the plots in Fig. 1 observes targets from steep grazing angles), however both systems create imagery through the generation of a "synthetic aperture." A synthetic aperture is created by moving an array along a path and creating a virtual array in post-processing that is bigger than the physical array. Imagery created using synthetic apertures is often much higher resolution than imagery created using individual physical apertures and the information in the synthetic aperture imagery is also much richer,

containing multi-aspect spectral response information. While the theory for synthetic aperture processing is straight forward, in practice the quality of synthetic aperture imagery tends to be highly dependent on the accuracy of the beamforming assumptions, the key factors being navigation and propagation speed accuracy.

Acoustic data can be processed in different ways to accentuate different types of features. Figures 2.B and 2.C, for example, are generated from the same dataset, but processed to show different aspects of the target's acoustic behavior. Similarly, from a classification standpoint there can be advantages to representing the acoustic response of targets in other domains than the spatial domain. A well explored alternative to imaging in the spatial domain is the multi-aspect frequency response of targets, otherwise known as the "acoustic color" representation, which characterizes the frequency response of a target vs. pose. For example, the data used to create figures 2.B and 2.C can be processed to produce the three-dimensional acoustic color response shown in Figure 3:

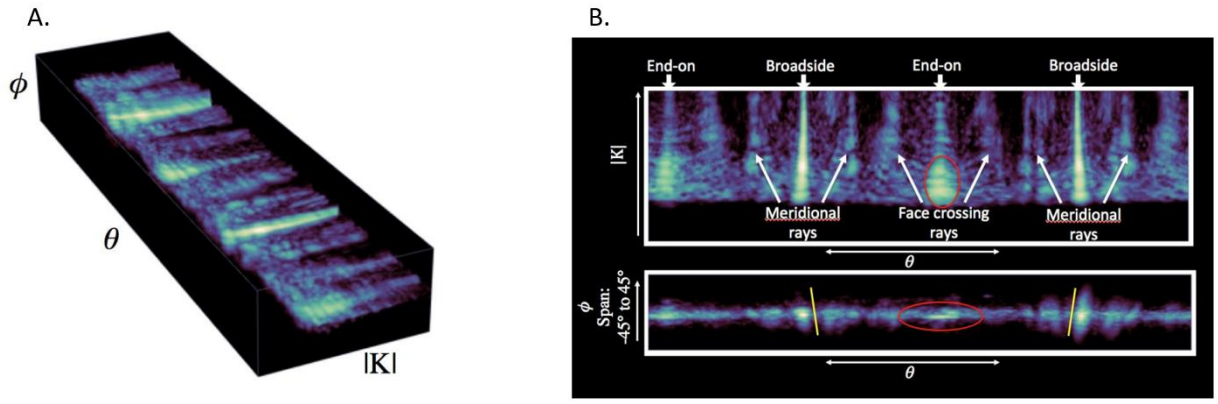

Figure 3: Plot (A) shows the three-dimensional acoustic color measurement extracted from a solid 2-1 aluminum cylinder using a low-frequency sonar mounted on an AUV. Dimensions are azimuth (theta), elevation (phi), and absolute spatial wavenumber |K|. Plot (B) shows the same acoustic color plot but viewed as two maximum intensity projections through the elevation dimension (top) and wavenumber dimension (bottom). Different physical rigid and elastic scattering mechanisms are labeled in (B). These figures were lifted from the final report of [6].

One advantage of these types of representations are that they are less affected by navigation error than SAS images. It was demonstrated in [8], for example, that the navigation and environment modelling errors that cause images of targets to defocus can be thought of as linear phase ramps passing through the origin in wavenumber space and extending in the directions of the sonar's varying look angles at the target. The slopes of these phases can be modified arbitrarily to defocus an image but they do not affect the magnitudes of the target's acoustic response vs. angle. In that sense, the acoustic color description of a target (target scattering strength vs. angle) can be fairly robust to synthetic aperture processing errors that spatial domain representations are sensitive to.

It has been demonstrated in numerous other places, however (see e.g. [9-11]), that achieving reasonable SNR for acoustic color estimates often requires pre-processing to reduce the influence of clutter and sediment backscattering. This isolation can be done in the image domain [9], or,

for particular scan geometries, using a set of transformations and wavenumber filters [6]. Without application of these filters the SNR of acoustic color information can be severely degraded. An example of an acoustic color plot generated with and without a clutter reduction filter is shown below in Fig. 4, again lifted from an interim report of [6]:

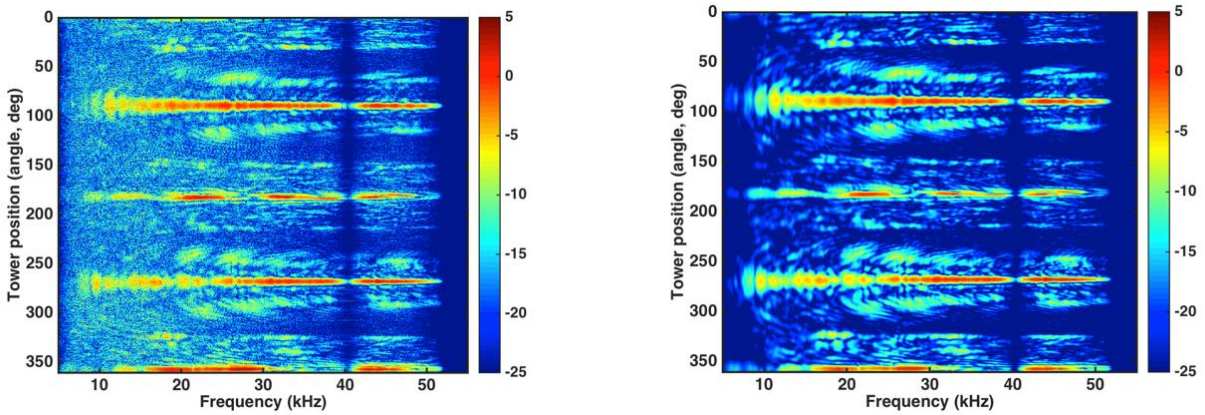

Figure 4: Acoustic color measurement for a target without (left) and with (right) a clutter reduction filter implemented in the wavenumber domain [6].

For acoustic systems scanning along a linear trajectory, an effective strategy for enhancing structural acoustic features is to image using a holographic technique [9, 11] which preserves the focus of ringing features that have long time-delays. Following application of this imaging method, spatial masks can be applied that isolate the acoustic response more effectively from surrounding clutter. This dependency on image domain processing for clutter reduction, or - even in the case of the wavenumber filter - a correct navigation solution (phase error disperses the energy in the domain of application for the wavenumber filter described in [6]), implies that there is great benefit to creating well focused imagery even when a desired data product is not an image: clutter reduction is enhanced by improving the definition of target borders and edges. Furthermore, the image must be sufficiently sampled in the spatial domain to prevent spectral aliasing, and the beamforming process must not corrupt the spectral response. This leads to the topic of beamforming.

The raw data collected by the BOSS are not images; rather they are the echoes that have been backscattered from the environment following the transmission of a nearly omnidirectional ping and are received by the hydrophones in the BOSS array. The data must be heavily post-processed in an operation called *beamforming* to determine the angle-of-arrival of the echoes and assemble a three-dimensional acoustic picture of the environment. The standard beamforming process is called “back-projection,” [12], “delay-and-sum” or the “exact method” [13], and is ubiquitous because it is simple and flexible: flexible in the sense that it doesn’t make strict demands about the shape of an array being used to create an image. The array could be a line, circle, curve, plane, or any other arbitrary shape. The quality of the image will depend on factors like array size and shape, accuracy of the array position estimates, accuracy of the propagation model, and how densely the array is sampled. In contrast, many other algorithms (e.g. the Range-Doppler, Chirp-Scaling, and Omega-K [13] algorithms) have strict requirements about the synthetic

aperture array geometry; for example, they may require that the array be perfectly straight and evenly sampled. These other algorithms are often employed despite their strict requirements because they are much faster than back-projection, which is a relatively slow algorithm. For a linear synthetic aperture array with N elements, back-projecting to an N x N grid requires a number of operations proportional to N^3, whereas the number of operations required for faster algorithms may be proportional to $N^2\log(N)$, which is a significant improvement. For a three-dimensional imaging system like the BOSS which creates a planar array the computational burden is even higher: backprojecting data from an N x N planar array to an N x N x N volume requires a number of operations proportional to N^5, (vs. N^3 for the linear 2D case). Despite the enormous ramifications of this difference, there are relatively few examples in extant literature describing extensions of efficient Fourier-based imaging to the volumetric synthetic imaging case represented by BOSS data. One potential reason for this may be that data captured by a system like the BOSS is inherently not as amenable to the pre-processing which enables Fourier imaging algorithms to be applied to standard 2D SAS and SAR data. One example that can be found in the literature is the application of the range-migration algorithm to BOSS data [14], but in general Fourier imaging is not preferred in the present case because the coupling of a nearly omni-directional source with significant array non-linearity results in an imaging problem poorly suited for beamformers with strict geometric requirements. In general, BOSS data have primarily been beamformed using time domain, delay-and-sum techniques and computational inefficiencies have primarily been dealt with by: 1) spatially under-sampling, 2) reducing the size of the imaged region, or 3) making overly simplistic assumptions about SAS array geometry, reducing beamforming to an interpolation problem using pre-computed look-up tables.

From a computational standpoint, it is easy to see why under-sampling is effective. Filling a $1m^3$ volume with voxels (3D pixels) 2 cm on edge requires 125,000 voxels. In contrast, if the voxel size is increased to 2.5 cm on edge, now only 64,000 voxels are required. Reducing the resolution[1] (i.e. *degrading* the resolution) by a small factor, only 25%, reduces the number of voxels by almost a factor of two. From a 3D beamforming standpoint, small resolution reductions are accompanied by large improvements in computational efficiency, and very large resolution reductions are accompanied by extreme improvements in computational efficiency (i.e. if a 10cm voxel size is used, computations can be reduced by a factor of 125!).

The problem with undersampling, however, is that spectral analysis (e.g. acoustic color analysis) imposes strict requirements on spatial sampling. A CPU might be able to beamform a scene at 10cm voxel resolution in 10 minutes, but if 2 cm voxels are required to prevent spectral aliasing, then beamforming with the required resolution would take 20 hours because 5^3 more interpolations are necessary.

[1] Terminology pertaining to resolution can easily be misleading or misinterpreted. The convention of the current document is that "reducing" resolution refers to reducing the ability to represent fine details or discriminate between adjacent scatterers. An example of how this resolution could be reduced is by *increasing* the size of the voxels in an image representation. Synonymously, reducing resolution means to degrade the resolution, or make it worse. Conversely, improving or increasing resolution refers to reducing pixel or voxel size and improving the ability to represent fine details or discriminate between adjacent scatterers.

The second shortcut, reducing the imaging region, is also straight-forward to understand. Rather than beamforming the entire ensonified region, it is much faster to only beamform perhaps the top 30cm of the sediment. The obvious downside of this approach is that a large amount of potentially relevant data is never utilized.

Lastly, the third shortcut: using a beamformer that is indifferent to platform motion, is unacceptable because a goal of this project is to create well focused imagery. Array motion within a synthetic aperture must be appropriately integrated into the beamforming operation to achieve well focused imagery.

A technology gap that this project aims to fill is the need for a beamformer that doesn't take any "shortcuts;" (i.e. it samples sufficiently densely to perform spatial-spectral analysis, it images the relevant acoustic swath out to full range, and it can beamform geometries that exhibit highly non-linear array motion), and yet is sufficiently efficient that a user of the system can beamform data acquired by a BOSS in a timely manner without specialized, use-case specific hardware.

Beyond the fundamental tasks of beamforming BOSS data and extracting structural acoustic information from targets, questions pertaining to how the 3D volumetric acoustic data will be accessed end users has not been addressed in much detail in prior literature. Data-interaction methods – by which is meant the methods and techniques leveraged by end users to obtain meaningful information from the data products created by the system – have mostly been limited to viewing orthographic sets of maximum intensity projections through volumes rendered by BOSS systems (see, e.g. Fig. 1). BOSS data – particularly multi-aspect datasets that contain the scattering responses of targets from multiple observation angles – are incredibly rich with information, and this project also aims to explore how the extra information in these datasets can be used to augment extant data interaction methods.

Materials and Methods

This project developed effective algorithms for beamforming data captured by the BOSS sonar both in post-processing and in a live-streaming manner, along with data visualization tools for interpreting results. Beamformer development was facilitated by the additional development of a simulation tool creating artificial BOSS data with fidelity suitable for point spread function analysis. Subsequent developments were enabled by the acquisition of data captured by various systems fielding BOSS arrays, and by the participation in UXO survey exercises conducted using the Multi-Sensor Towbody system (ESTCP project MR18-5004, [7, 11]). This section will describe, in the following order, the BOSS sonar hardware (including a description of the Multi-Sensor Towbody System), the point-simulator used to create data for beamformer development, a description of the 3D Fast Factorized Backprojection (3D FFBP) beamforming algorithm, a detailed description of the structural-acoustic feature enhancement process developed for improving acoustic-color data products, the components of the MuST Toolbox software library – the end-user toolbox for processing BOSS data – and end with some considerations for and descriptions of methods for visualization of 3D BOSS data.

I. The Buried Object Scanning Sonar (BOSS)

Design overview

The buried object scanning sonar is a down-looking sonar system that leverages either a planar or linear (depending on the design) array oriented parallel to the surface of the sediment, and a low frequency source, to generate three-dimensional volumetric images within the sediment volume. The original design, created by a Florida Atlantic University team headed by Steven G. Schock, contained a planar array of 32 hydrophones in four rows of eight and a linear array of projectors.

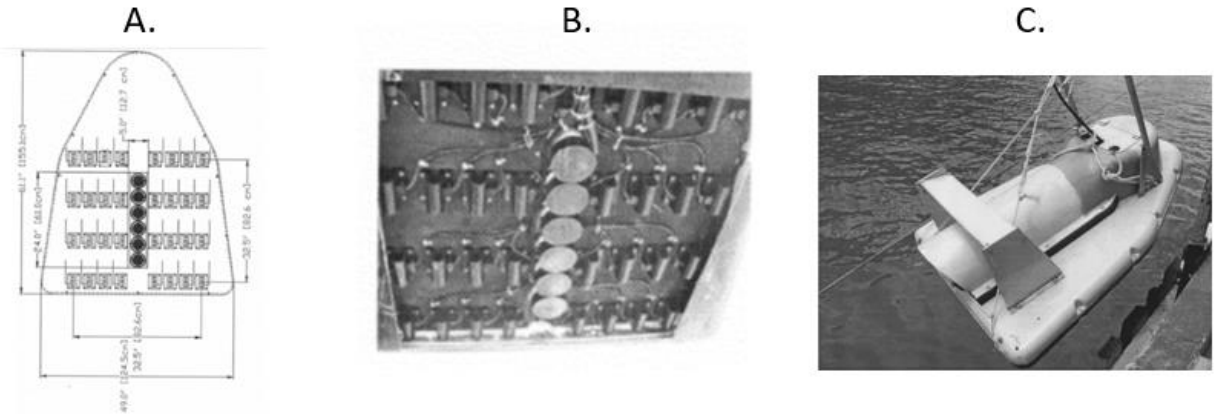

Figure 5: The original BOSS system designed by S. G. Schock et al at Florida Atlantic University. The figures, borrowed from [1], show A) the layout of the projectors (circular elements) and hydrophones, B) a photograph of the projector and hydrophone array on the underside of the BOSS platform, and C) a picture of the BOSS platform from above.

Following the design of the original BOSS system, subsequent design updates have generally utilized linear rather than planar arrays. An example is the winged design shown in Fig. 6, which has been used for a number of DoD funded research projects, including MR-2103, from which the figure was grabbed.

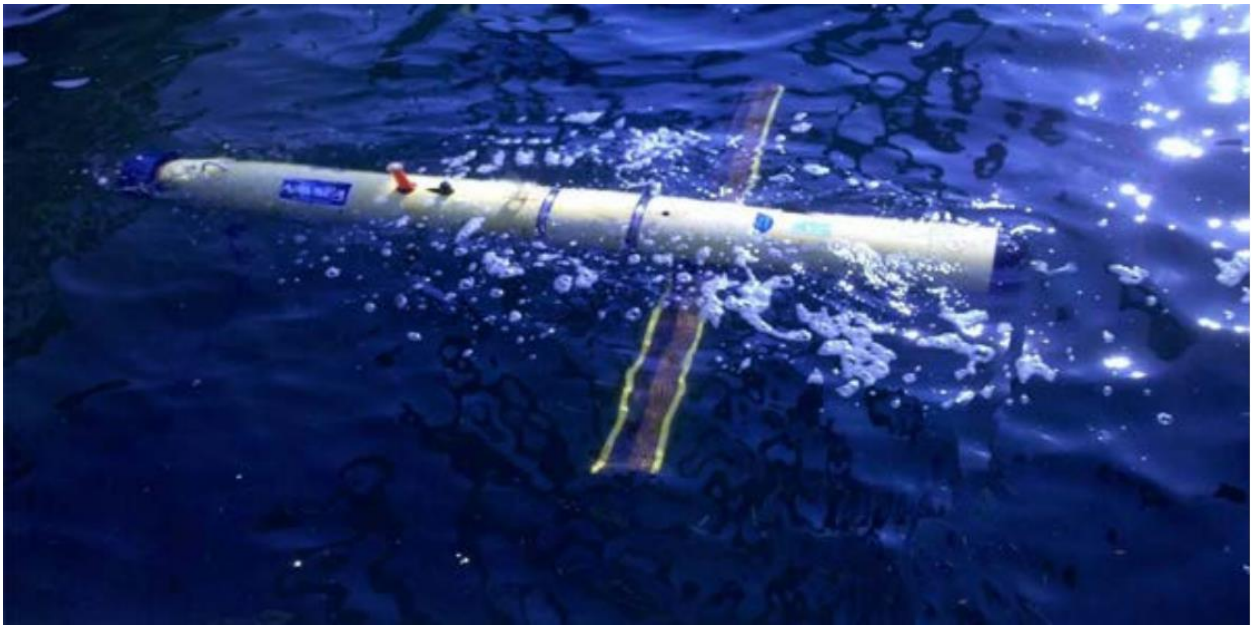

Figure 6: A subsequent BOSS design from FAU, which used a linear array of 40 receiver elements (20 elements on each wing) which, when combined with SAS processing, improved resolution over previous designs. Image taken from Appendix A of the final report of SERDP project MR-2103.

The general concept of operation leveraged by the BOSS design is to ensonify the sub-sediment volume using a low frequency source and sense the echoes with a ribbon-like two-dimensional planar array that combines the across-track extent of the physical array with an along-track synthetic aperture. Near-field volumetric imaging is accomplished by back projecting the acoustic data into the volume. This concept is illustrated in Fig. 7, which shows a linear array, the previous positions of its hydrophone elements as it scans through a linear trajectory, and a volumetric region being ensonified by the source.

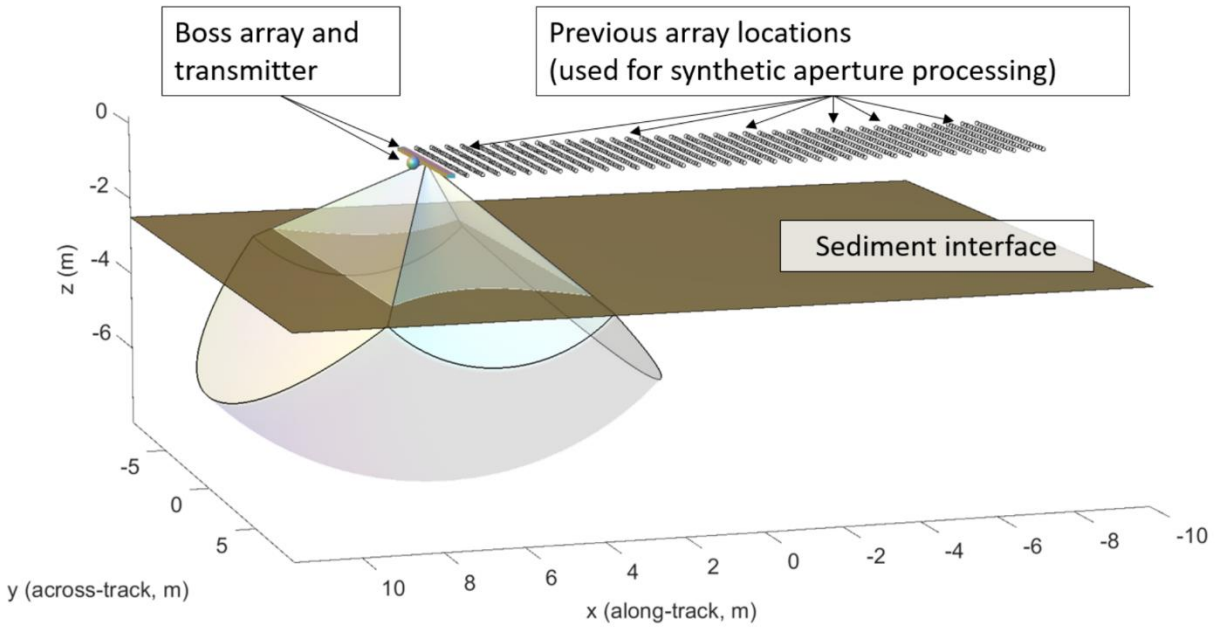

Figure 7: the BOSS concept of operation. A down-looking source with a very broad beam and low center frequency ensonifies the subsediment volume (sediment surface = flat brown rectangle), and a ribbon-like synthetic aperture array is used to backproject the data into the volume, creating a 3D image.

More recently, the company EdgeTech began collaborating with FAU to commercialize the BOSS sonar concept, resulting in the eBOSS system, a modular array system with 16 elements in each array which can be tiled according to user needs. Data from two eBOSS systems was acquired for development purposes in this project, one originating from an AUV platform and another originating from the Multi-Sensor Towbody, a UXO detection platform [11].

The previously described BOSS systems have a unified academic lineage that be traced back to the original FAU research group headed by S. G. Shock and hence have preserved the name "BOSS", however there are several very similar sonars that have independent family trees (though perhaps they may have been conceptually influenced by the original BOSS research). Examples include the Sediment Volume Search Sonar (SVSS, D. Brown, SERDP project MR 2545), the Pangeo Sub-bottom Imager (SBI, [8]), the Soacsy SeaChirp3D [9], and the Kongsberg GeoChirp 3D [10]. The SVSS is another SERDP funded down-looking volumetric sonar that is similar to the original BOSS design shown in Fig. 5 in that it uses a relatively small planar array, but it has the advantage of exploiting broader spatial diversity in its transmitter configuration. The SVSS is currently being evaluated for detecting and classifying ordnance in very shallow regions; recent results can be found in the final report of MR-2545, which is publically available on the SERDP website. The Pangeo SBI and Soacsy Seachirp 3D on the other hand are both very similar to the winged linear BOSS design, however the available literature related to the systems – which includes an OCEAN conference proceeding – omit any references to the BOSS system or any other prior art.

AUV BOSS

In January 2019, field data collected from a prototype BOSS system mounted on an AUV was given to APL-UW for analysis. The system is a NAVY / DoD owned system and specifics of the sonar and AUV platform are not available in this report; interested parties should contact Code 32 if more information about the system is desired. The data from this BOSS system provided the initial field data enabling testing of the accelerated beamforming algorithm developed as part of this project. The data was ideal because it contained all of the noise, calibration issues, and trajectory perturbations that accompany real-world systems. The trajectory, in particular, was highly unstable and non-linear, providing excellent material for evaluating beamformer generality.

The Multi-Sensor Towbody (MuST)

The multi-sensor towbody (MuST) is a system that is being developed by the Applied Physics Laboratory, University of Washington (APL-UW), for the purpose of detection, classification and geolocation of buried UXO [7, 11]. The primary sensing equipment currently installed on the towfish are a dual-frequency EdgeTech 2205 sidescan sonar and a four-panel EdgeTech BOSS ("eBOSS") low frequency buried object scanning sonar. The dual-frequency sidescan sonar is meant to provide shape-information over a wide area enabling detection of proud UXO. The eBOSS system is meant to provide detection and classification capability for buried objects using both shape analysis from beamforming and structural acoustic analysis. The eBOSS is located on the bottom of the tow-fish, and the array mounted in the across-track dimension as shown in the following figure. All four panels are placed in a row, forming an extended linear array with 64 elements.

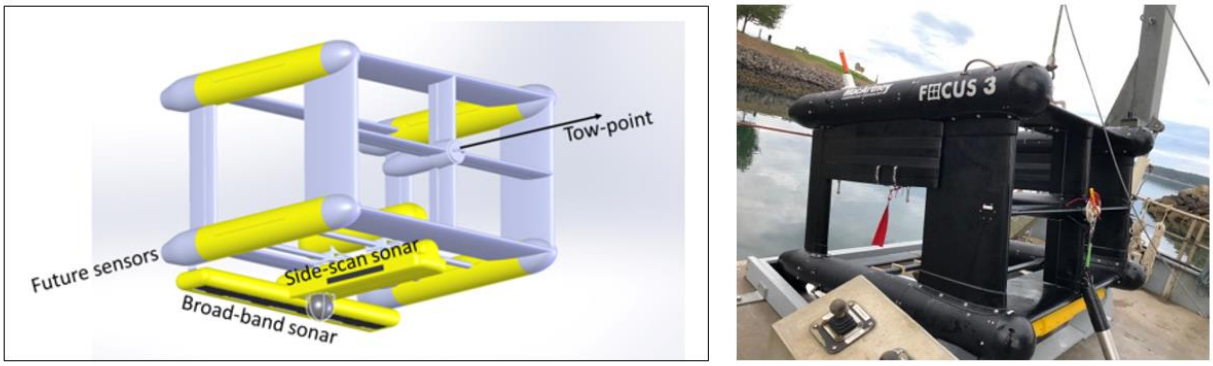

Figure 8: Concept art for the MuST system (left) and a photo of the actual system (right). The leading edge of the eBOSS array is the yellow strip visible on the bottom of the towfish on the figure on the right. Note that in the concept image the eBOSS array was located on the aft end of the towfish, whereas it is on the fore-side for the final product.

Figure 9 shows a photograph of the array as mounted for calibration in a calibration facility. Both the receiver array and unique transmitter configuration are visible.

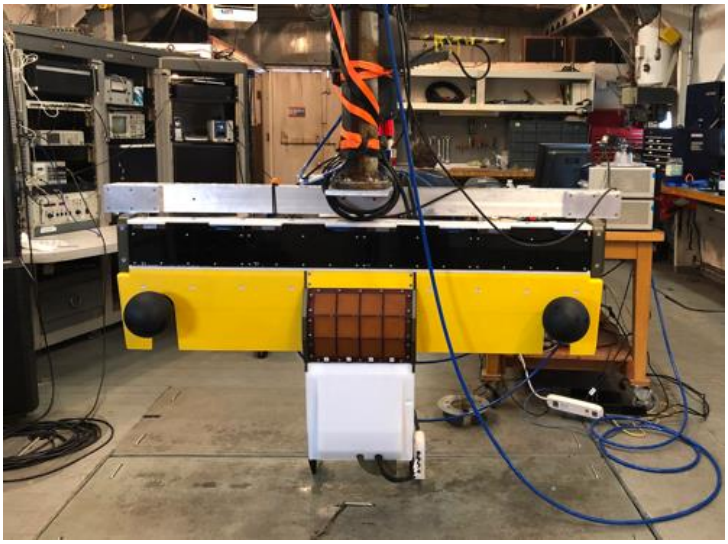

Figure 9: The eBOSS array dismounted from the towfish, viewed from the bottom. The three sources are visible in the image: the omnidirectional Neptune sources are the black spheres near the sides of the array. The VLA array is the brown grid located near the middle of the system.

The two black spheres on the right and left of the array are broadband omnidirectional low-frequency sources available from Neptune Sonar, a commercial transducer company [12]. During operation, transmission from these transducers alternates back and forth which, when combined with the appropriate post-processing, doubles, (i.e. improves by a factor of two – see footnote on pg. 16), the across-track resolution of the system. A third transducer, composed of an array of twelve sources, is visible near the center of the array. This transmitter, dubbed the Very Large Array or "VLA," was a custom broad-band transmitter that was used in several of the initial experiments with the MuST. It was determined, however, that there was no significant advantage gained by using the VLA vs. the Neptune sonars (data captured using both sources are shown in the results section) and its usage was eventually abandoned and the sonar removed from the system.

The eBOSS array – the black strip running across the top of the hardware shown in Fig. 9, is composed of four panels, each containing sixteen receiver elements that are spaced ~1 inch apart in the across-track dimension. Each hydrophone element is rectangular and approximately 1 inch (across-track) by 2 inches (along-track). The four arrays are organized adjacent to each other forming a linear array of 64 elements with a total length of ~1.6 meters. It is theoretically shown and experimentally verified that the dual source configuration approximately doubles the resolution of the system in [13].

Operational software written by EdgeTech is installed on a computer mounted on the towfish near the sonar. A user interface, *BOSScontroller.exe* allows users to input options such as the transmitted chirp signal, pulse repetition rate, file size and naming convention for the generated data files. The operational software has many other options and capabilities that are not used in the present study. Critically, the BOSScontroller.exe allows a TCP/IP connection to be set up with a top-side computer for streaming data during acquisition, enabling custom real-time analysis. It also has a "replay" mode that can be used to broadcast the data in saved sonar files

over TCP/IP mimicking live acquisition, which was critical for the development of MuSTLive, the realtime beamforming and analysis tool developed in this project as part of the MuST Toolbox.

The towfish component of the system, visible in Fig. 8, is a MacArtney FOCUS 3. The FOCUS 3 combines active control surfaces with a design meant to be hydrodynamically stable, providing an optimal deployment platform for a synthetic aperture system.

Four additional sub-systems are mounted on the FOCUS 3 to aid in navigation and environment characterization: a pressure sensor, an in-situ sound speed sensor, a Doppler velocimeter (DVL, the specific model was a Teledyne Pathfinder [14]), and an inertial navigation system (INS). The INS went through two manufacturer and model updates before arriving at the current solution: an IXBlue ROVINS NANO. The model initially incorporated into the MuST – the Inertial Labs INS-D – was found to have insufficiently reliable positioning accuracy for the needs of the project when the system was submerged, however the hardware included a dual-GPS positioning system that was retained on the MuST and was used by all subsequent INS systems as the absolute positioning reference when the towfish was surfaced.

MuST eBOSS Calibration

Calibration is an important sensor characterization process that enables the voltage signals measured by the digital acquisition system to be accurately mapped to acoustic pressures. Characterizing the transmitter and receiver sensitivities and beampatterns has the potential to improve classifier performance but additionally enables researchers to determine whether features in imagery or frequency space are characteristics of targets or simply a manifestation of the response of the system. Furthermore it can serve as the basis for correction filters that compensate for any irregularities in the system response. For these reasons the eBOSS transducers on the MuST system, including the Neptune sources and the eBOSS receiver array elements, were calibrated in a calibration facility at APL-UW during the spring of 2020. By this time in the project the VLA array had been replaced with another Neptune source, creating three omnidirectional sources that could be operated in two modes: a dual transducer mode that alternated between port and starboard sources and a single source mode that used the center transmitter. Figure 10 contains a photograph of the eBOSS array in the calibration facility. The eBOSS array has been removed from the towfish and mounted to a ram that can be raised and lowered into the water via a moon pool. The ram can rotate around its axis, enabling beam-patterns to be measured.

Figure 10: The MuST eBOSS array removed from the towfish and mounted on the ram used for suspending and rotating the array in the water column during calibration.

A calibrated F37 reciprocal transducer was used as a calibrated reference for evaluating the sensitivity of the receivers. The reference transducer was placed 3.17m from the array element being calibrated. Both the element being calibrated and the transducer were submerged at a depth of 2.13 meters. An illustration of the calibration geometry is shown in Fig. 11.

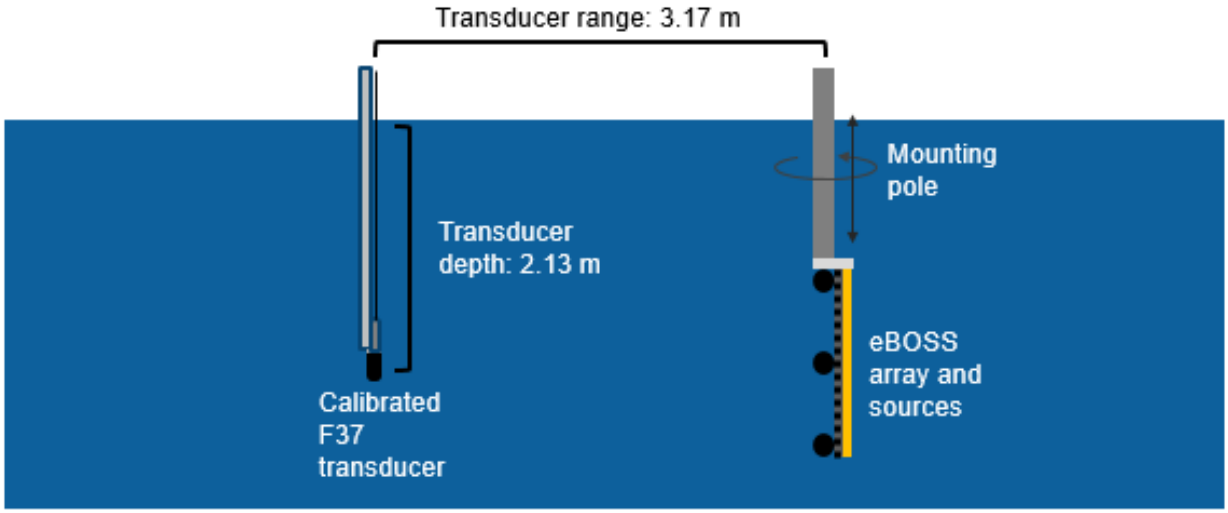

Figure 11: Illustration of the calibration setup used for calibrating the MuST eBOSS array showing the calibrated reference transducer on the left and the eBOSS array mounted in the configuration used to calibrate the center Neptune source on the right.

The following calibration measurements were performed on the array:

Measurement	Component	Rotation type	Signal
Broadside receiver response	Full receiver array	None (broadside)	Stepped CW
Receiver beam pattern	Single receiver panel	Pitch	Stepped CW
Receiver beam pattern	Full receiver array	Pitch	Chirp
Receiver beam pattern	Single receiver panel	Roll	Chirp
Neptune source beam pattern	Sources 1, 2 & 3	Pitch	Chirp
Neptune source beam pattern	Sources 1, 2 & 3	Roll	Chirp

Table 1: Table of calibration measurements performed on the MuST eBOSS array at the APL-UW calibration facility in the spring of 2020.

Multiple mounting configurations were necessary to be able to make the measurements listed in Table 1, shown above. During beam-pattern measurements the array was mounted to align the element being calibrated with the axis of rotation. The four array mounting configurations are shown in Figures 12 and 13. The ram was raised or lowered to place the element being calibrated on the same horizontal plane as the source transducer. All ram rotations were done counter-clockwise.

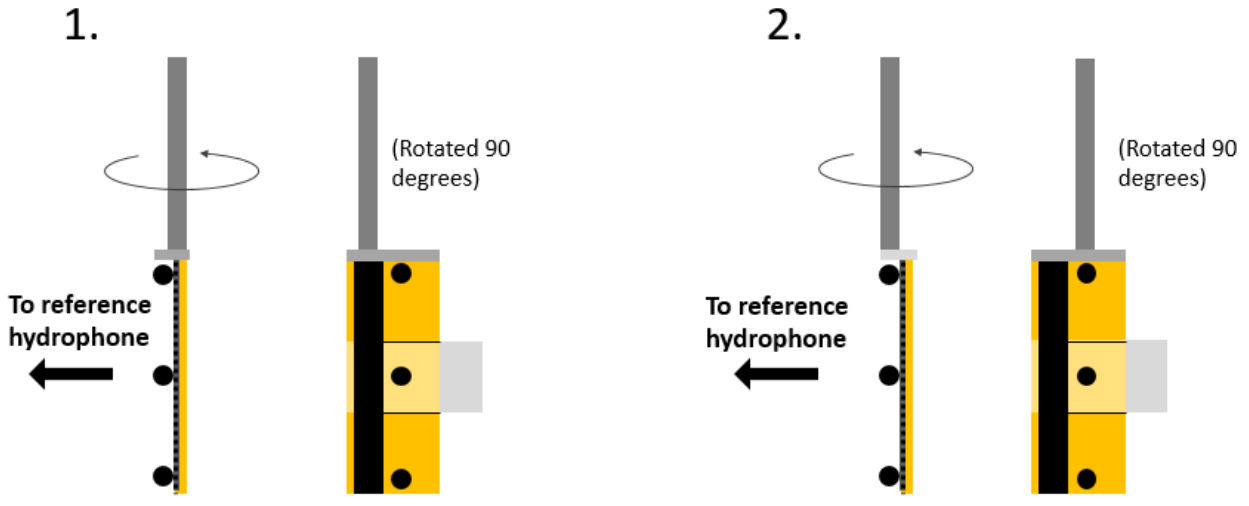

Figure 12: Mounting configurations 1 (left) and 2 (right) used to measure the aft-to-fore beam patterns of the receiver array (configuration 1) and transducers (configuration 2). The ram is attached to the port side of the

array in both configurations. The only difference in these configurations is that the panel is shifted more directly over the transducers in configuration 2.

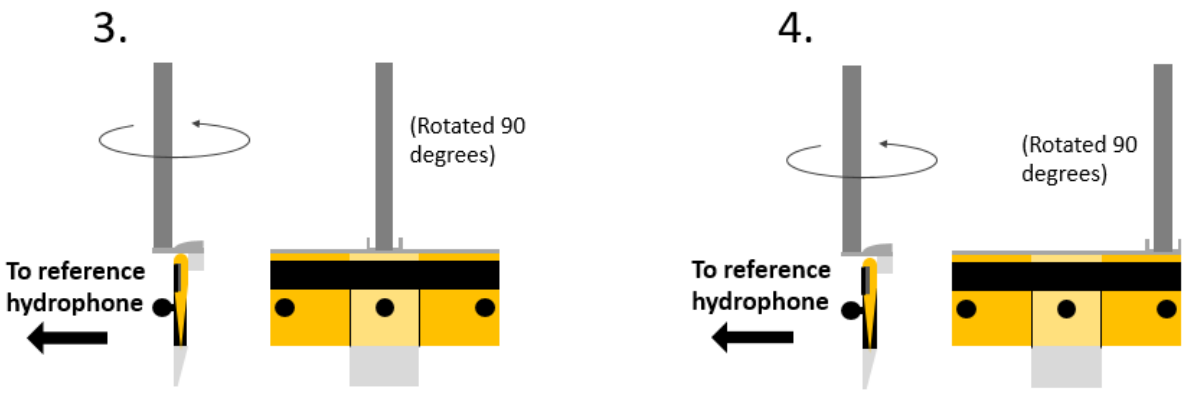

Figure 13: Mounting configurations 3 (left) and 4 (right) used to measure the port-to-starboard beam patterns of the central (left) and port-side (right) transducer. Configuration 3 was also used to measure the port-to-starboard sensitivity of the central receiver array panels.

Figure 14 and 15 are photos of the system mounted in configurations 3 and 4 both in and out of the water. A ram stabilizer is visible in Fig. 15. The stabilizer was used to prevent the ram from undergoing large oscillations during rotation.

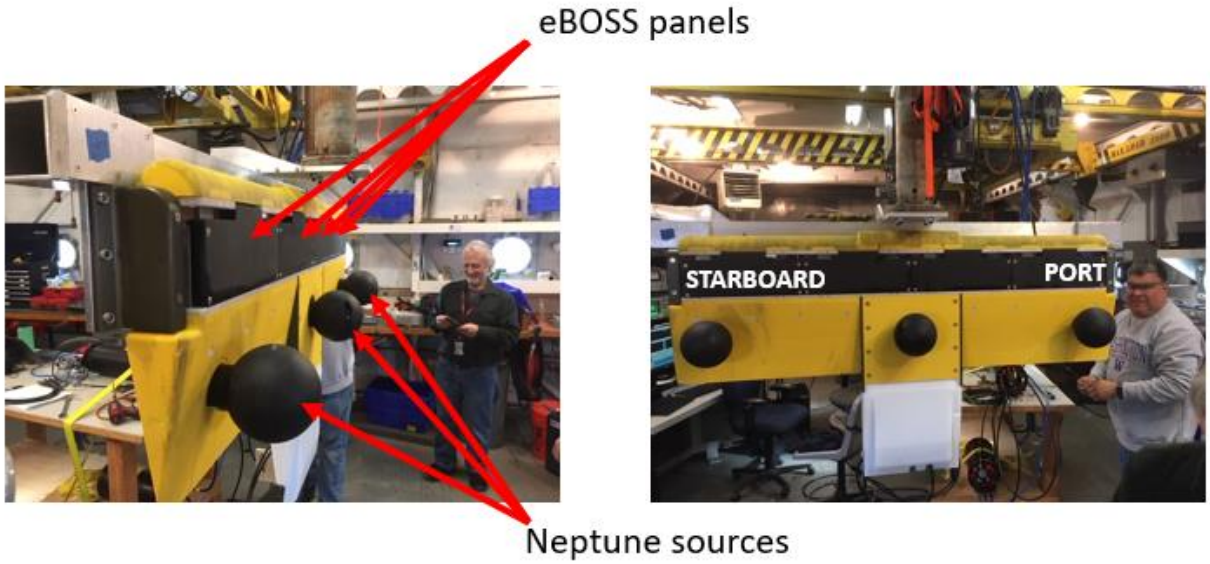

Figure 14: The MuST eBOSS array mounted in configuration 3 showing the eBOSS panels and Neptune sources (left), and the port and starboard sides of the array (right).

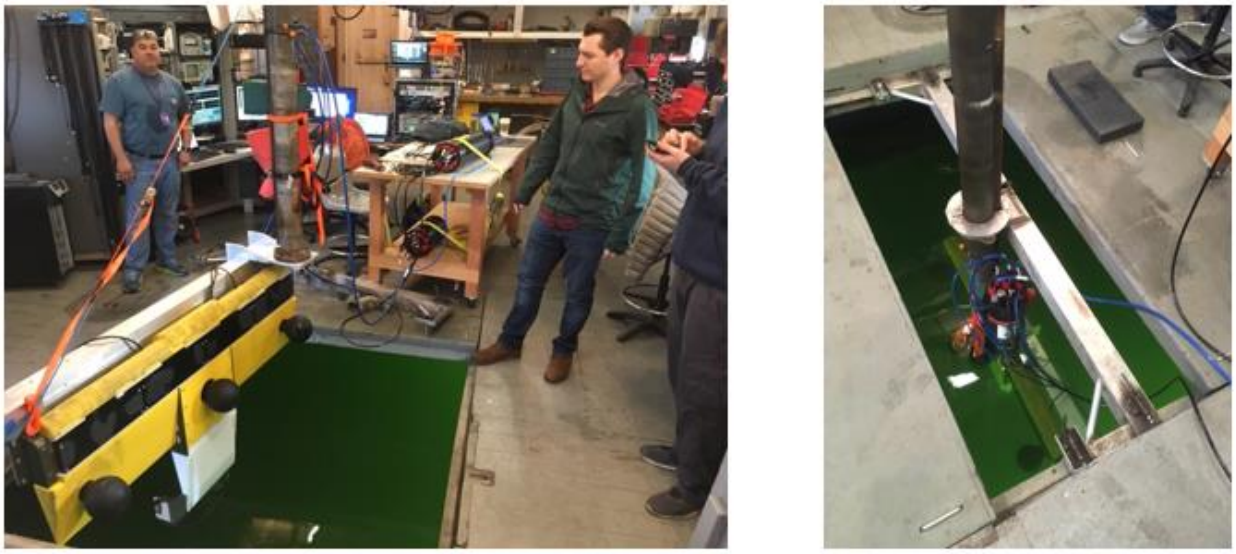

Figure 15: The array being lowered into the moon pole in configuration 4 (left), and the submerged array mounted in configuration 3 (right). The ram stabilizer, a metal collar which fits around the ram, is visible in the photo on the right.

The following sections provide more detail about the calibration measurements.

Broadside free-field receiver voltage sensitivity

Configuration 1 at broadside (0 degrees) orientation was used for this measurement, and the signals used were a set of continuous-wave (CW) pulses stepped in 0.25 kHz increments from 3 kHz to 35 kHz. The source was a calibrated F37 transducer located 3.17 meters from the array and driven at 88 Vrms. The mounting ram was stepped in increments of 4 inches to place elements in the center of the beam. The full series of CW pulses was measured by the array at each ram position. At the broadside configuration the receiver array panel was rotated such that the face of the hydrophones faced the reference transducer, i.e. the ray from the hydrophone to the array was normal to the surface of the array. The measured sensitivity in dB re. 1V/uPa is shown in Fig.16 for each element of the array and calibration frequency.

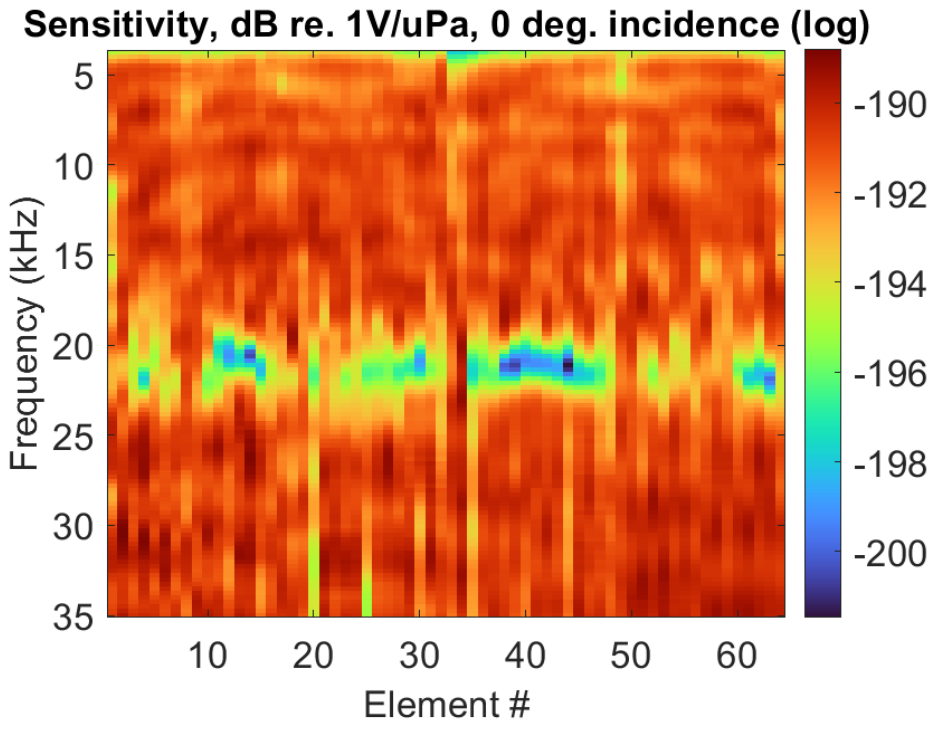

Figure 16: The broadside free-field voltage sensitivity of the eBOSS array elements 1 to 64 over the frequency range 3 to 35 kHz.

Significantly, over most of the band the response is fairly flat however there is a deep sensitivity null in most channels at ~20 kHz. In subsequent beam pattern measurements it was revealed that this null was a sample of a null that changes in frequency as a function of array orientation. Outside of the null the response is mostly uniform, only varying by a couple dB in the broadside configuration.

Receiver beam pattern: aft-to-fore (pitch) free-field voltage sensitivity (CW)

Configuration 1 was used for this measurement but the axis of the ram was rotated over 180 degrees rather than being fixed at 0 degrees. The beam pattern of the array was measured at five different frequencies: 7.5, 12.5, 17.5, 21, and 22.5 kHz. The signals transmitted from the reference transducer were CW pulses, and the array was rotated in steps of 1 degree. The array was lowered to a depth that placed the starboard-side panel in the main beam of the reference transducer.

Figure 17 shows a polar plot of the measured sensitivity in counts / decaPascals for the first 8 elements of the receiver array panel. The units counts / decaPascal were chosen to make the beam patterns interpretable in the polar plots shown in Fig. 17. Beam pattern sensitivities in dB re. 1V/uPa are shown in the next section showing calibrations using an LFM pulse instead of separate CW pulses.

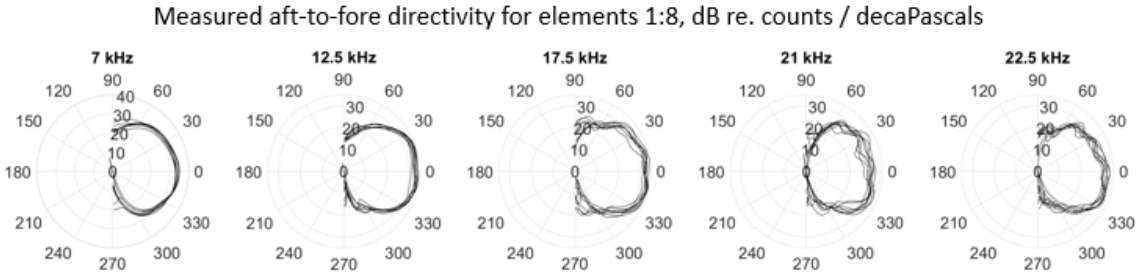

Figure 17: Starboard panel receiver array element directivity plots in counts / decaPascals at each of the tested CW pulse frequencies.

From Figure 17 it can be seen that the responses are quite uniform at 7 kHz with a slight bias in the aft dimension. The beam pattern becomes slightly more irregular at higher frequencies but is still fairly uniform.

Receiver beam pattern: aft-to-fore (pitch) free-field voltage sensitivity (LFM)

In this calibration measurement CW pulses were replaced with a single 5 – 25 kHz LFM (i.e. chirp) pulse that was 1 millisecond in length. The short duration of the pulse caused spectral leakage outside of the nominal transmit band, enabling an extended estimate of the beam-pattern response up to 35 kHz. Verification of LFM performance was performed by comparing sensitivities of the LFM measurements with the CW measurements. Figure 18 shows the measured sensitivities in dB referenced to volts per micro-Pascal for each of the 16 elements of the starboard side array. Figure 19 shows the average array response over the full starboard panel in polar and Cartesian coordinates. Note that a sensitivity drop is visible at 20 kHz at broadside. This frequency varies with the rotation angle of the array but maps to a straight vertical line in the Cartesian representation.

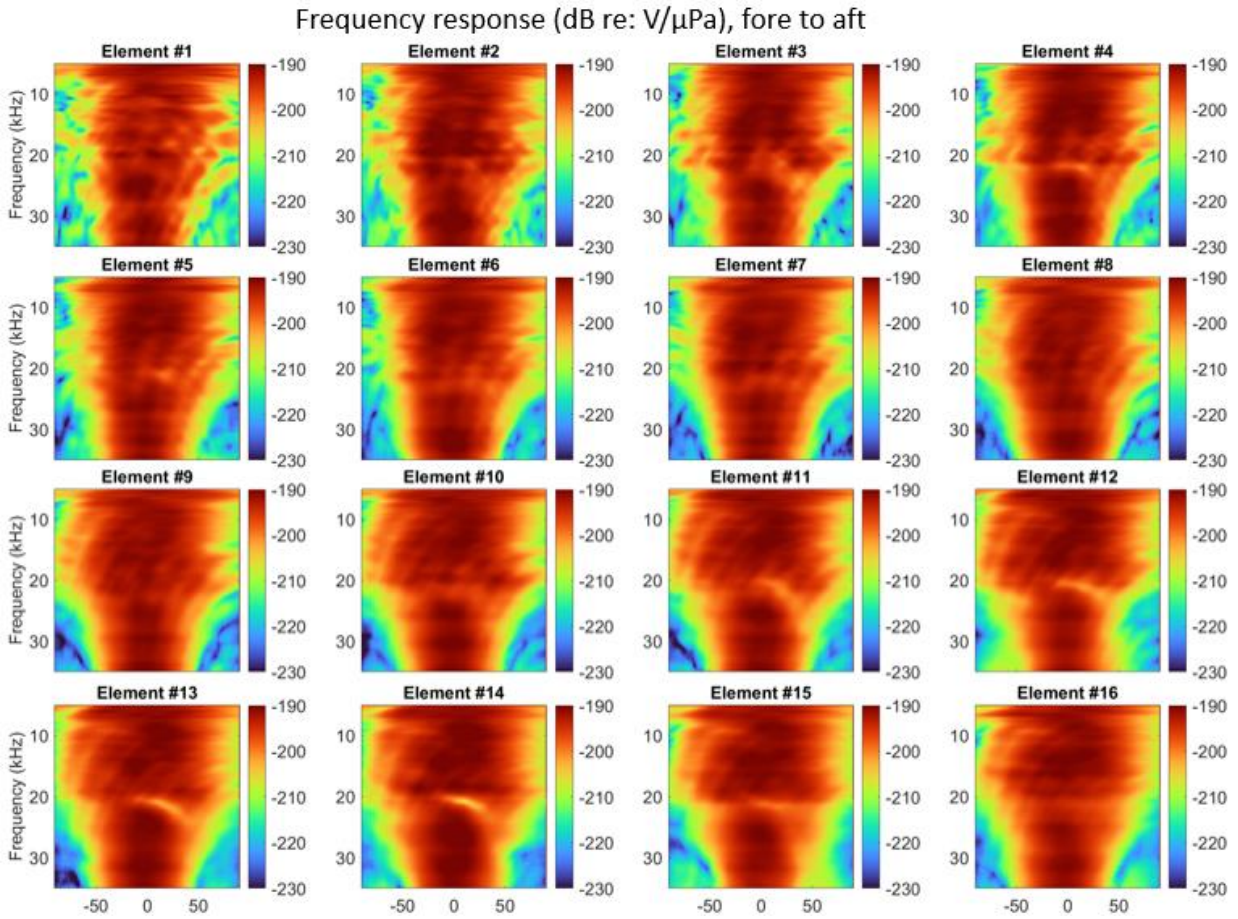

Figure 18: The individual beampattern sensitivities measured for elements 1 – 16 for the starboard side panel of the MuST eBOSS receiver array. Sensitivities are shown in dB referenced to V/uPa. The total dynamic range of the sensitivity plot is 40dB. In several receive elements a null near 20 degrees at broadside is visible.

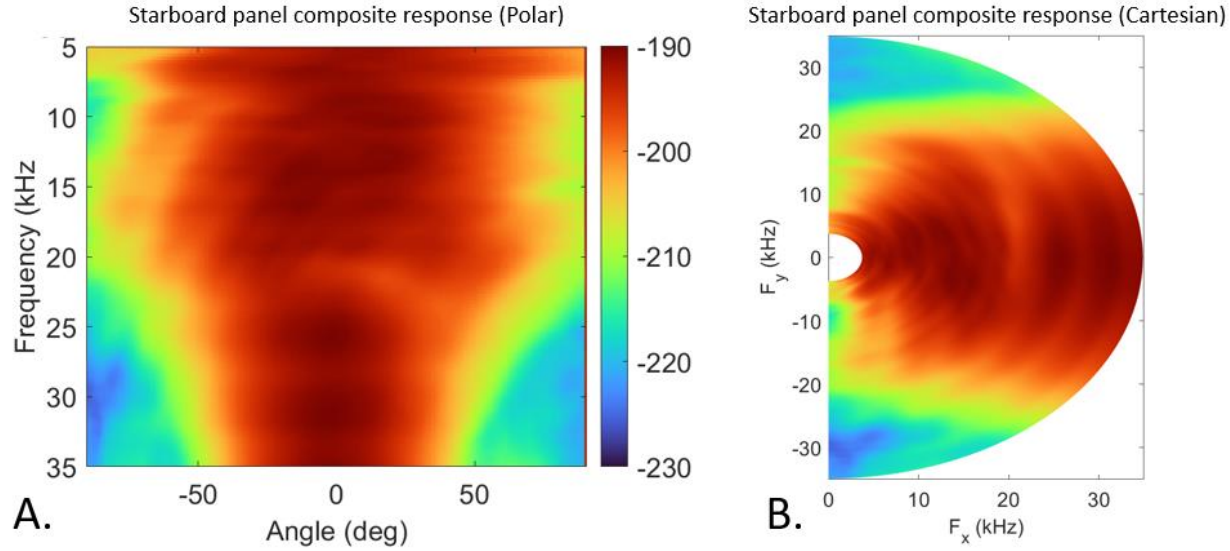

Figure 19: The composite starboard-side receiver panel beam pattern shown in polar coordinates (left) and Cartesian coordinates (right). Note the sensitivity drop near 20 kHz.

Receiver beam pattern: port-to-starboard (roll) free-field voltage sensitivity (LFM)

In this calibration measurement the MuST array was placed in configuration #3 and rotated from -90 to 90 degrees, counter-clockwise, measuring the receiver array sensitivity from port to starboard. The same 1ms 5 – 25 kHz LF pulse was transmitted from the reference transducer. Beam patterns were measured near the center of the array (i.e. near the axis of rotation) at elements 28 through 35. Elements seemed to have more uniformity in this test, however the total beam pattern seemed to be less uniform. Figure 20 shows polar plots in dB referencing ADC counts / decaPascal to illustrate the directivity shape at several discrete frequencies, and the fully populated sensitivities in dB referencing V/uPa and the aggregate response are shown in figures 21 and 22.

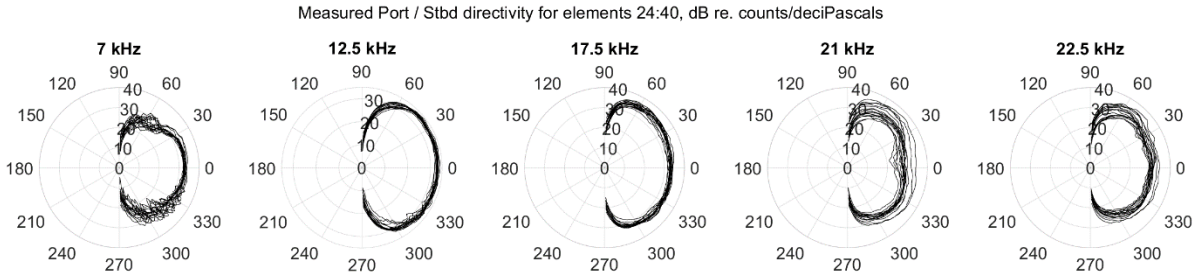

Figure 20: Port-to-starboard sensitivities at discrete frequencies for a set of elements near the center of the array. Note the dip in sensitivity near broadside at 21 kHz.

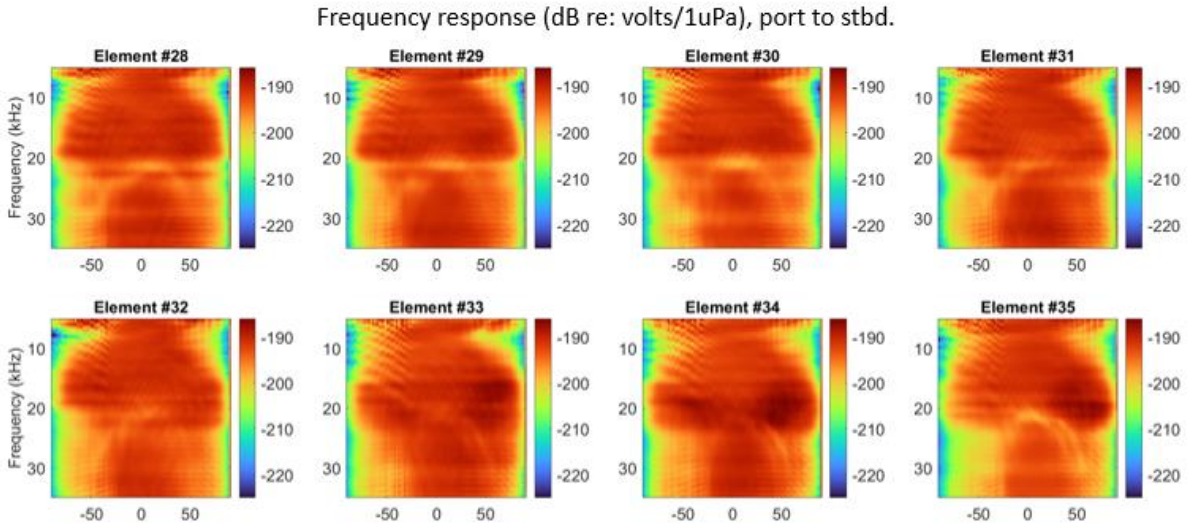

Figure 21: Port-to-starboard sensitivities in V/uPa for elements 28 – 35.

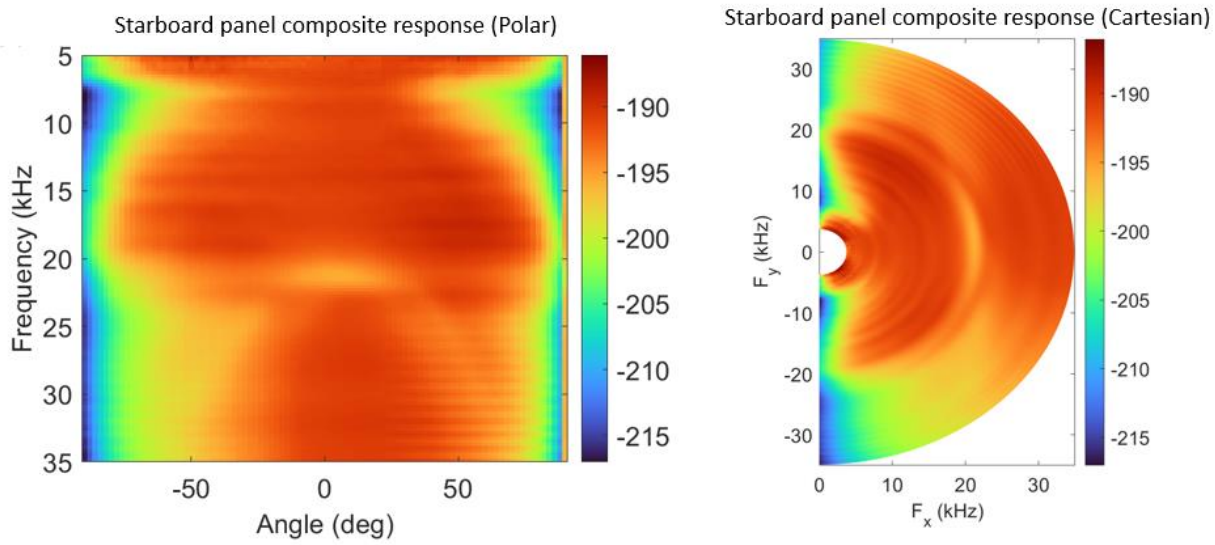

Figure 22: The composite port-to-starboard beam pattern for the two central panels of the array (i.e. elements 17 – 48). Note the very different directivity characteristics vs. the aft-to-fore sensitivity. In both cases, however, there is a dip in sensitivity near 20 kHz.

Neptune beam pattern and sensitivity: central transmitter

To characterize the beam pattern of the central transmitter a 1ms 5 – 25kHz LFM chirp was transmitted from the central Neptune and received on the calibrated reference microphone. The array was rotated from aft-to-fore in configuration #2 and re-mounted in configuration #3 to measure the port-to-starboard beam pattern. Results are shown in dB referenced to uPa/V at 1 meter.

Significantly, it was found that the transmitter beam pattern was extremely non uniform. Source sensitivities were roughly twice that shown in the transducer datasheet at some frequencies, but nulls occasionally over 30dB deep were present in the beam pattern. These are visible in both beam pattern plots shown in Fig.'s 23 and 24. It is almost certain that the cause of these nulls are constructive and destructive interference between the direct path signal and a reflection off of the transducer wing mount. As visible in Fig. 14 the centers of the Neptune spherical sources are ~10cm in front of the yellow eBOSS structural mount. Due to this offset a reflection with an approximately 0.13ms delay will be created, resulting in interference nulls spaced at approximately 7.5kHz apart at broadside. This is observed in Fig.'s 23 and 24. These nulls have a strong potential to impact the spectral responses detected in targets, and an advisable future update to the eBOSS system would be to either baffle the transmitters reducing the impact of the reflected signal or replace the omnidirectional spherical sources with hemispherical sources that have a more favorable transducer directivity pattern.

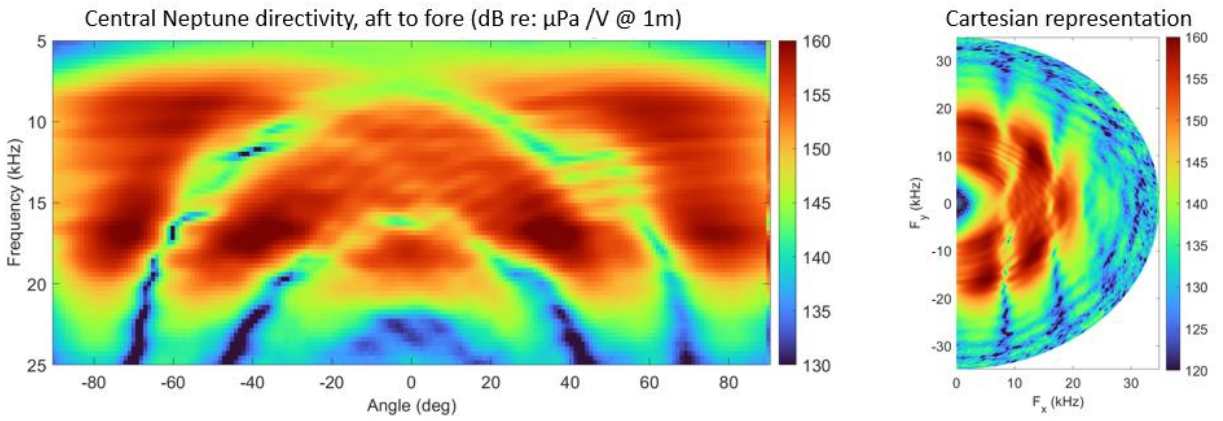

Figure 23: The aft-to-fore directivity measured for the central Neptune transducer. Note the strong interference fringes caused by interference with reflection off of the eBOSS array mount.

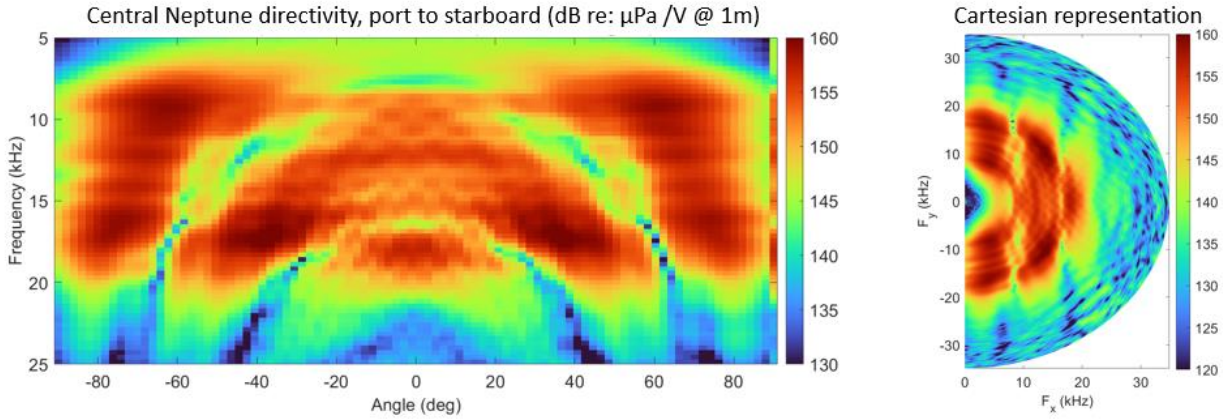

Figure 24: The port-to-starboard directivity measured for the central Neptune transducer. The interference pattern created by the array is almost identical to that shown in Fig. 23.

Neptune beam pattern and sensitivity: port side transmitter

The port-to-starboard transmitter beam pattern was also measured on the port-side of the array using the same approach described for the central transmitter. Only the port-to-starboard transmitter beam pattern was measured because the array is strongly asymmetric around the transmitters located at the extrema of the array and were anticipated to have a very different response. It was further assumed that the transmitter on the starboard side would have a similar (but reversed) beam pattern to the beam pattern measured on the port side, and only the port side was calibrated. The measured beam pattern is shown in Fig. 25. As anticipated, it is asymmetric, however the reflection nulls are still dominant features.

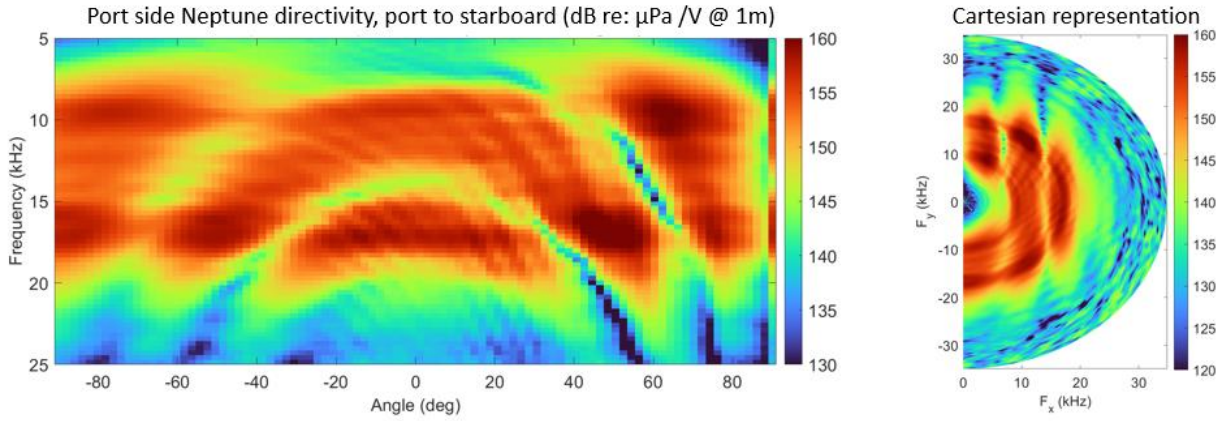

Figure 25: Port side transmitter port-to-starboard beam pattern in uPa /V at 1 meter. Unlike the central transmitter the port side transmitter shows a high degree of asymmetry in the response, however the reflection nulls are still a dominant feature.

Additional measurements

Following calibration the eBOSS array was mounted in configuration #3 and scanned vertically by lowering the ram to create a planar array simulating a scan. Targets were suspended in the water column in front of the sensor to determine detection potential by the eBOSS system. Targets included a 60 x 30 cm hollow aluminum cylinder, a 1.5 meter long section of nylon-coated ¼ inch steel cable, a 0.9 meter diameter circular loop of the same nylon coated cable, and a 0.9 meter diameter coil of nylon coated cable containing several loops. All targets were readily visible; volumetric images created from these scans are shown in Fig.'s 26 – 30, which show maximum intensity projection (MIP) renderings through the volumetric data. The data in each volume are displayed in dB referenced to the brightest point with a 30dB dynamic range. In contrast to the standard 3-view orthographic maximum intensity projection images commonly used for rendering volumetric eBOSS data, this rendering traces rays from a virtual camera location and a sense of depth is added by applying a 1/range artificial attenuation from the camera position. The visualizations were created by a custom renderer developed as part of this project. Rendering methods for volumetric acoustic data products are discussed later in the Materials and Methods section.

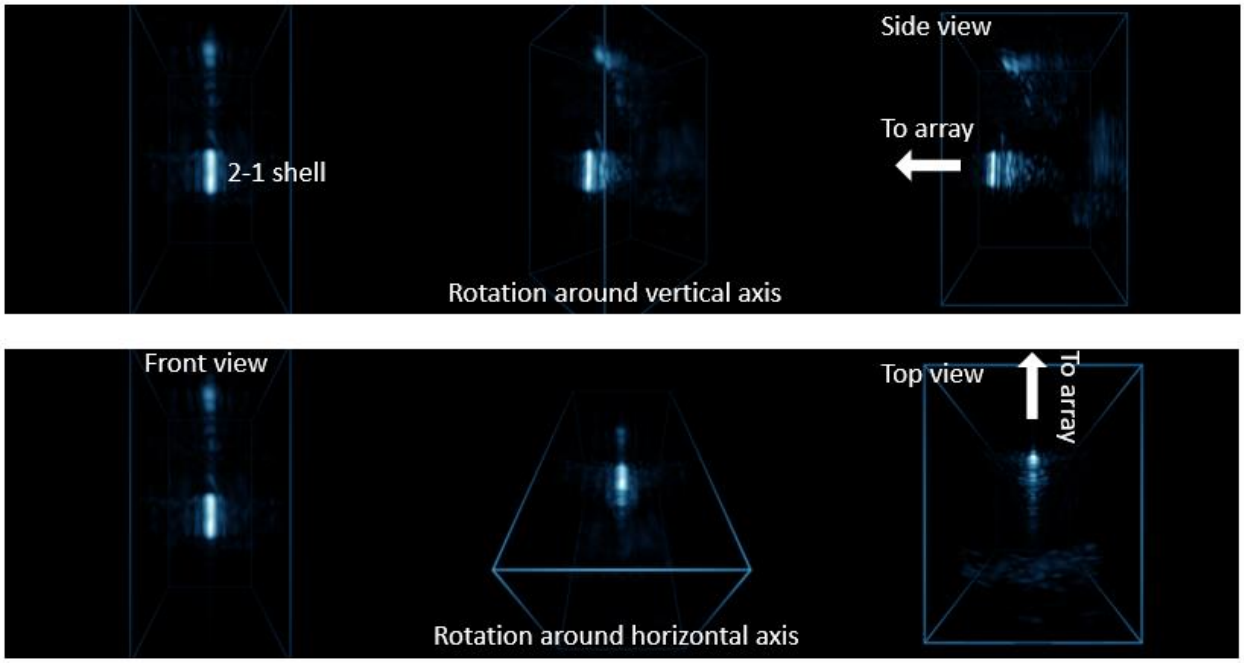

Figure 26: A volume rendering created in the calibration facility by moving the array vertically in configuration #3 to create a synthetic aperture. The target is a hollow 60 x 30 cm aluminum cylinder. The dynamic range of the image is 30dB referenced to the brightest point. The direction toward the array is illustrated in both figures which show the volume rotating around different axes.

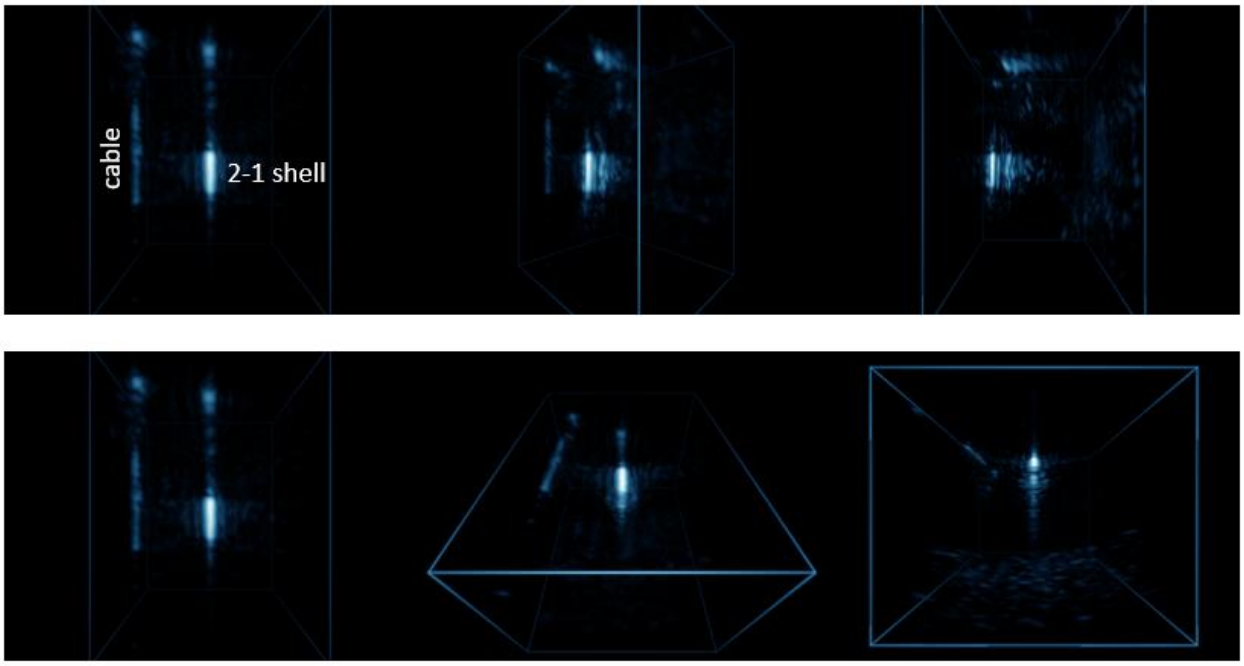

Figure 27: An image showing both the cylindrical shell and the ¼ nylon coated steel cable suspended next to it. The cable is not as strong of a scatterer as the cylinder, but it is still readily visible, with a positive SNR.

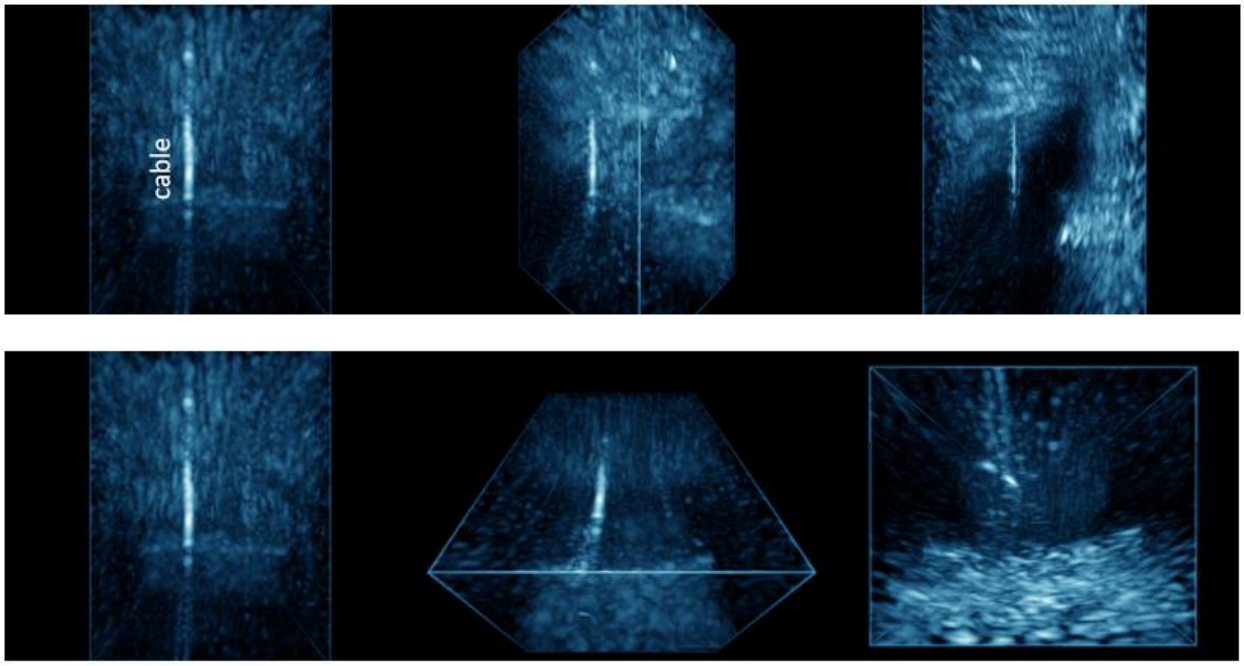

Figure 28: The 1.5 meter ¼ inch nylon coated steel cable suspended by itself. The 30dB dynamic range, now referenced to the brightest point on the cable, clearly shows reverberation noise within the volume.

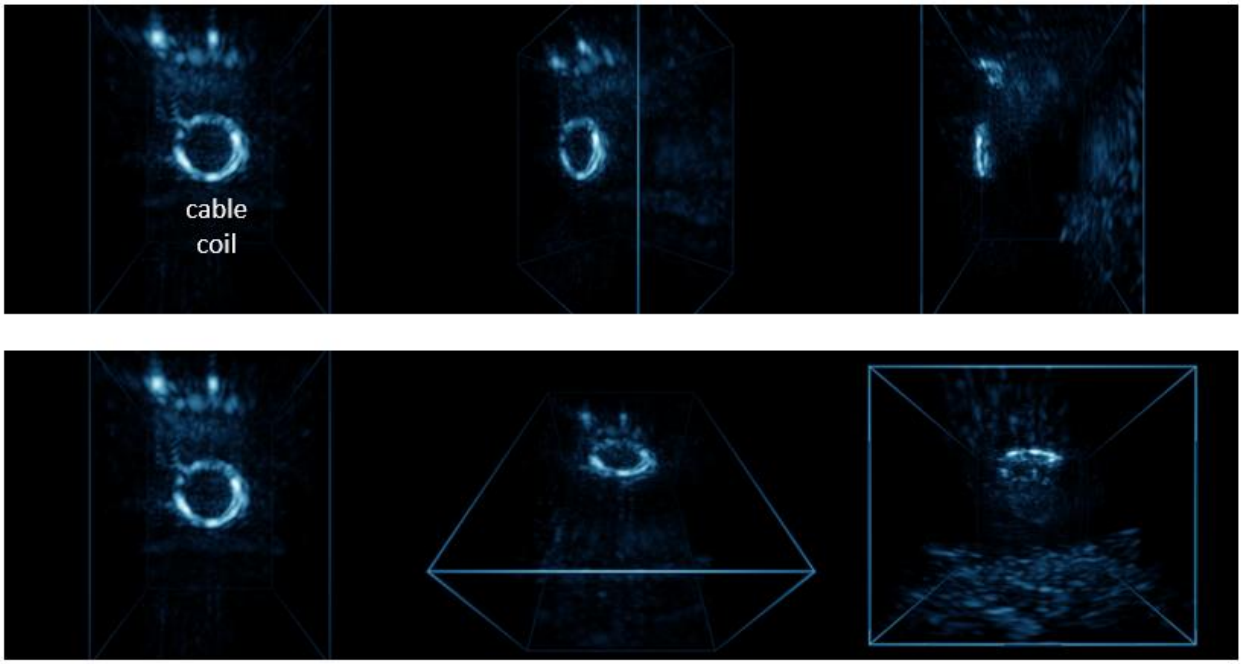

Figure 29: A suspended coil of ¼ inch nylon coated steel cable. The coil is 0.9 meters in diameter and contained several loops of the steel cable. The cable is noticeably brighter than the individual segment, as evidenced by the decreased relative level of the background noise.

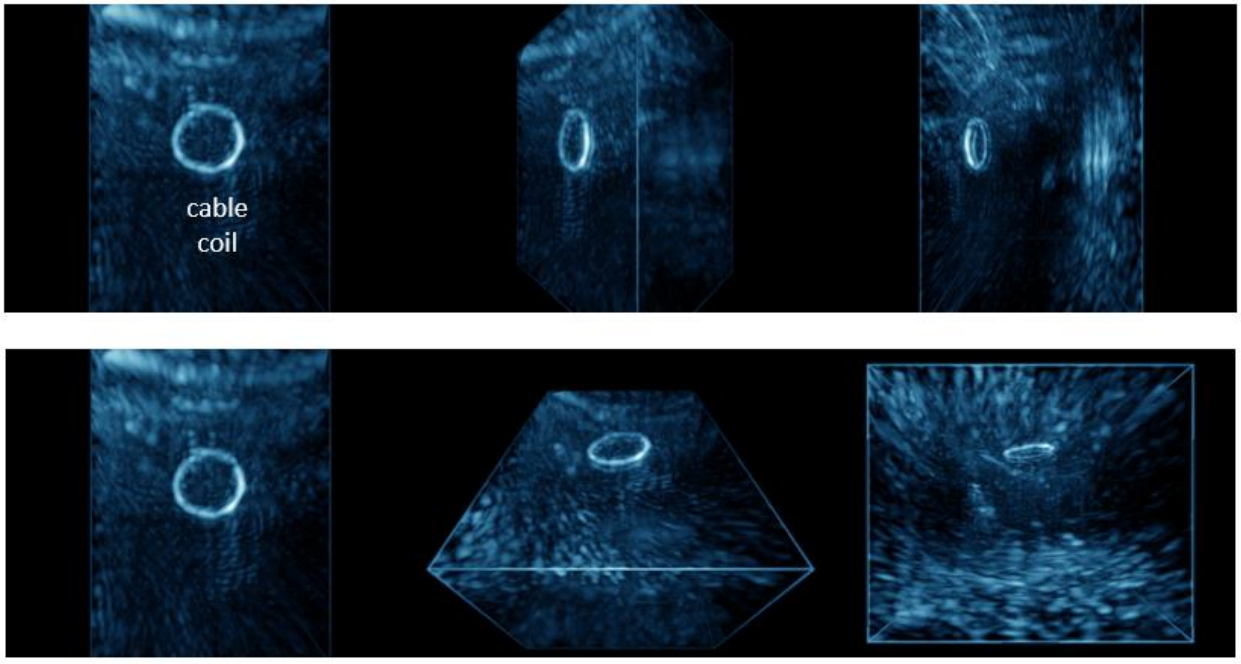

Figure 30: A single 0.9m diameter loop of cable. In several of the rendered orientations a faint ring is visible behind the direct return; it is possible that this faint ring may be the visible manifestation of the array mount reflection responsible for creating the nulls in the transmitter beam patterns (i.e. the plots shown in Fig.'s 23 – 25).

II. BOSS data simulation

Overview

A standard method used to evaluate the performance of a beamforming algorithm is to create a signature corresponding to an ideal point scatterer and use this signature to determine the beamformer's "point spread function" (PSF). Similar to determining a system's impulse response, the PSF is a measure of the beamformer's spatial response to a signal corresponding to an ideal point. Ideally the PSF is determined by the size of the aperture, but beamformers typically introduce some degree of distortion, (depending on any approximations made in the beamforming process, number of interpolation used, types of interpolations used, etc.), into the PSF, and figures of merit like resolution, signal-to-noise-ratio and sidelobe levels can be computed using these PSF calculations.

To test the functionality of beamforming algorithms by determining the sensitivity of focus to parameters like array shape, source directivity, source location, signal bandwidth, sample density, sensitivity to platform distortion (roll, pitch, yaw, surge, sway and heave), sound-speed mismatch and other system and environment parameters, a simple simulator called "BOSSim" was developed. Unlike many of the more sophisticated synthetic aperture simulation tools being developed, BOSSim is only meant for PSF evaluation and simulates the backscattering response of a user-defined set of sparse points in the free-field of a medium having a uniform propagation speed. This section discusses the scattering model and adjustable simulation parameters that can be defined in the simulator.

BOSSim, which is programmed in Matlab, was created in part to be a simulator very specifically tuned to the needs of this project (e.g. creating data suitable for evaluation of a BOSS-specific beamforming algorithm), however it was also created out of necessity because, at the start of this project, there was no other publically available simulation software readily adaptable to creating simulated BOSS data. As the project developed, this capability was integrated into ARL-PSU's PoSSM simulator [15-17], however by the time that capability became available a beamformer had already been successfully developed and was being applied to field data captured by a physical eBOSS system.

Scattering model

BOSSim is a free-field point-scatterer simulator that uses the following model to generate the acoustic response of a point target:

$$\boldsymbol{S(t) = \int_{-\infty}^{\infty} \frac{1}{d_1 d_2} \mathcal{T}(\varphi,\theta)\mathcal{R}(\varphi,\theta)\mathcal{H}(\omega)\alpha\left(\tilde{\varphi},\tilde{\theta},\omega\right)e^{j\omega(t-(d_1+d_2)/c)}d\omega} \quad \textbf{Eq. 1}$$

$$\boldsymbol{d_1 = \sqrt{\left(x_{transmit} - x_{point}\right)^2 + \left(y_{transmit} - y_{point}\right)^2 + \left(z_{transmit} - z_{point}\right)^2}} \quad \textbf{Eq. 2}$$

$$\boldsymbol{d_2 = \sqrt{\left(x_{r_element} - x_{point}\right)^2 + \left(y_{r_element} - y_{point}\right)^2 + \left(z_{r_element} - z_{point}\right)^2}} \quad \textbf{Eq. 3}$$

where t = time in seconds, ω = frequency (radians / second), c = medium propagation speed, $d_{1,2}$ are the Euclidean distances between the transmitter ($x_{transmit}$, $y_{transmit}$, $z_{transmit}$) and point (x_{point}, y_{point}, z_{point}), and receive element ($x_{r_element}$, $y_{r_element}$, $z_{r_element}$), $\mathcal{T}$ is the transmitter directivity and is a function of the vehicle-perspective azimuth (θ) and elevation (φ) to the point target, $\mathcal{R}$ is the receive array directivity and is likewise a function of θ and φ, $\mathcal{H}$ is the zero-phase amplitude of the transmit signal as a function of ω, and α is the target response, a function of ω and of azimuth and elevation angles $\tilde{\theta}$, $\tilde{\varphi}$, which describe the angular position of the array from the perspective of the target.

Several simplifications are made on top of (1). Because this is a sparse point target model, the values for α were set to unity for all scatterers computed in the scene. The transmitter is simulated as an omnidirectional source by setting $\mathcal{T}$ to unity. $\mathcal{H}$ is purely real and computed as a tukey window (e.g. cosine-tapered window) over the signal bandwidth, to imply that $S(t)$ is already pulse compressed. $\mathcal{R}$ is computed as a sinc-function in θ and φ using a rectangular aperture model determined by the size of the receive element panels in the BOSS array. The point-scatterers are assumed to be in the free-field of a homogeneous medium with propagation speed c and lastly, the distance values d_1 and d_2 are computed using both the location of the points distributed in the medium and the location of the transmitter and array elements as the vehicle path is simulated through the medium.

Simulation parameters

Though simplistic in terms of the scattering and environment model, (which is intentional: the point-spread function under ideal circumstances is what is needed for beamformer evaluation), *BOSSim* allows for the adjustment of a wide variety of simulation parameters related to the

transmitter, receiver and scan trajectory geometries, as well as environment and transmit signal parameters. These are described in the following sections.

BOSSim array configuration

The following transmit and receive array parameters are tunable in the simulator and affect directivity as well as the resolution by determining the real-aperture size:

Transmitter configuration

Number of transmitters (*integer*)

Position of transmitters (*vector with units meters, relative to centroid of receive array)*

Firing order (*array of integers, length equal to number of transmitters*)

Transmitter aperture size in vehicle coordinate *x* and *y* dimensions (*meters*)

Receiver configuration

Element size in vehicle coordinate *x* and *y* dimensions (*meters*)

Element spacing in vehicle coordinate *y* dimensions (*meters*)

Number of elements per panel (*integer)*

Number of panels in array (*integer*)

Panel kerf in vehicle coordinate *y* (across-track) dimension (*meters*)

BOSSim scan configuration

The following scan parameters are tunable in *BOSSim*:

Scan configuration

System speed (*m*/*s*, *x-dimension*)

Transmitter ping rate (*Hz*)

Total number of pings (*integer*)

Track perturbation: surge (*amplitude (m), # oscillations over scan & starting phase*)

Track perturbation: heave (*amplitude (m), # oscillations over scan & starting phase*

Track perturbation: sway (*amplitude (m), # oscillations over scan & starting phase*)

Array orientation perturbation: roll (*amplitude (rad.), # oscillations over scan & phase*)

Array orientation perturbation: pitch (*amplitude (rad.), # oscillations over scan & phase*)

Array orientation perturbation: yaw (*amplitude (rad.), # oscillations over scan & phase*)

BOSSim environment configuration
The following environment parameters are tunable in *BOSSim*:

Environment configuration:

Acoustic propagation speed (*m/s*)

Scatterer distribution coordinate system (*options: cylindrical, Cartesian, spherical*)

Scatterer distribution, dim. 1: (*boundary 1 & 2 (m), number of scatterers (integer))*

Scatterer distribution, dim. 2: (*boundary 1 & 2 (m), number of scatterers (integer))*

Scatterer distribution, dim. 3: (*boundary 1 & 2 (m), number of scatterers (integer))*

BOSSim signal configuration
The following transmit signal parameters are tunable in *BOSSim*:

Signal configuration:

Sampled frequency band (*lower bound (Hz), upper bound (Hz)*)

Signal frequency band (*lower bound (Hz), upper bound (Hz)*)

Tukey parameter (*real value* $0 < \beta < 1$)

Simulation example
BOSS systems come in different array configurations, and the platforms can experience different levels of motion and with different characteristics. The previous simulation parameters allow the beamformer to be tested under realistic array configurations and scan conditions. The following table contains an example scan configuration tuned to the parameters of the BOSS system mounted on the Multi Sensor Towbody (MuST) platform:

Simulation parameter	**Value**	**Notes**
Transmitter: *x*-position	[-0.8, 0.8]	Number of transmitters determined by vector length
Transmitter: *y*-position	[0.2, 0.2]	
Transmitter: z-position	[0, 0]	
Transmitter firing order:	[1, 2]	Alternating
Transmitter aperture: *x*-width	0.01	~omnidirectional in θ
Transmitter aperture: *y*-width	0.01	~omnidirectional in φ
Receiver element width, *x*	50.8e-3	
Receiver element width, *y*	19.1e-3	

Receiver element *y* spacing	25.4e-3	Elements slightly narrower than spacing which could cause undersampling at higher frequencies and wide angles
Panel kerf (*y*)	9.2e-3	
Number of panels	4	
Elements per panel	16	64 elements total
Along-track speed	2	m/s
Pulse repetition frequency	80	Pulses per second
Number of pings	400	10 meter track
x-dim. motion amplitude	0.01	SURGE, peak-peak = x2
x-dim. motion oscillations	6	# of oscillations over track
x-dim. motion phase	0	starts @ amplitude*sin(0)
y-dim. motion amplitude	0.1	SWAY
y-dim. motion oscillations	2	
y-dim. motion phase	pi/2	
z-dim. motion amplitude	0.3	HEAVE
z-dim. motion oscillations	1.3	
z-dim. motion phase	pi/2	
Yaw motion amplitude	0.04	Rot. about array z-axis (rad)
Yaw motion oscillations	3	
Yaw motion phase	0	
Pitch motion amplitude	0.02	Rot. about array y-axis
Pitch motion oscillations	2	
Pitch motion phase	0	
Roll motion amplitude	0.12	Rot. about array x-axis
Roll motion oscillations	1	
Roll motion phase	0	

Acoustic propagation speed c	1540	Sound speed, m/s
Scatterer distribution coord. sys.	'cyl'	Points defined in x, φ, r
Scat. distribution, dim. 1	[0, 10, 11]	Boundaries and number of points in x
Scat. distribution, dim. 2	[-45, 45, 5]	Boundaries and number of points in φ
Scat. distribution, dim. 3	[3, 8, 5]	Boundaries and number of points in r
Sampling rate band	$[0, 40*10^3]$	Sampled frequency band, $[f_{low}, f_{high}]$, fs = 40 kHz
Signal bandwidth	$[5*10^3, 35*10^3]$	Signal band
Tukey taper	1	Full cosine taper

Table 2: *BOSSim* parameters for a point-target data simulation demonstration.

As will be discussed in the *Results* section, the MuST system has an unusual source transmission configuration in which two transmitters located approximately at the port and starboard sides of the array alternate transmission between pings. The sample rate for either transmitter individually is therefore halved (40 Hz using the value in Table 2 divided by the number of transmitters), but due to an increased coverage in the spatial wavenumber domain (see [13]) the resolution in φ is doubled, i.e. improved by a factor of two. It should additionally be noted that the values for scan perturbation are vastly exaggerated beyond what are expected for the MuST system and are meant to test the *generality* of the beamformer, i.e. how well it can handle nonlinear array motion. Furthermore, the sample bandwidth is 40 kHz and signal bandwidth 30 kHz, whereas in reality the signal bandwidth is closer to 20 kHz and the sample rate will be ~25kHz following base-banding. The following figure shows the distribution of the points in Cartesian coordinates, along with the location of the array samples generated using the simulation parameters shown in Table 2.

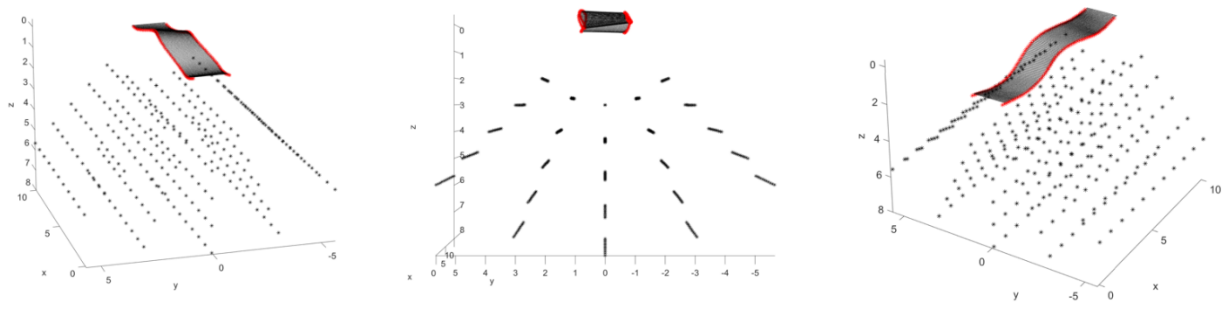

Figure 31: Three views of the distribution of simulated scatterers, array elements (dark strip down middle) and sources (red dots).

The following figures show a set of slices and MIP's through the dataset generated using the entry values listed in Table 2.

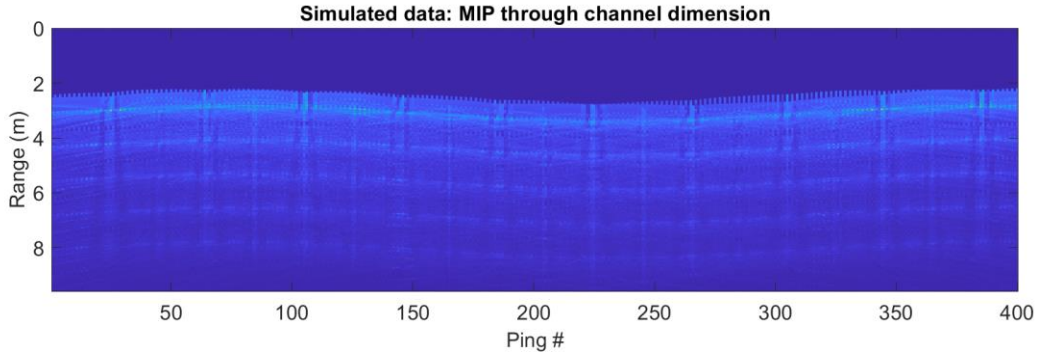

Figure 32: A MIP through the channel-dimension of the raw data matrix, showing the profile of the signal data in range and ping. Note that the undulating minimum range values match the vertical motion of the array shown in Fig. 31.

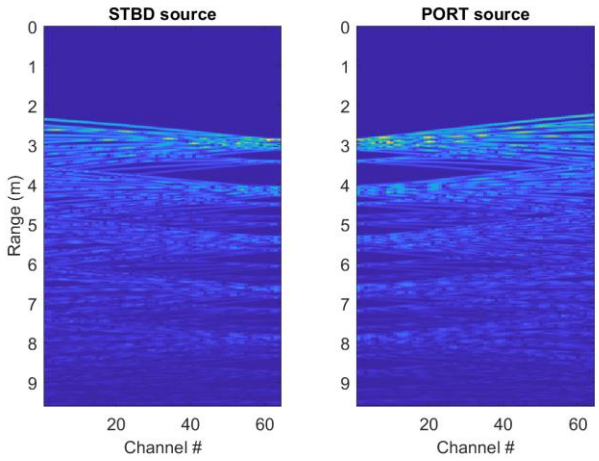

Figure 33: Slices of the volume corresponding to the raw data captured by the array on pings 51 (left) and 52 (right). Notice the convergence of the target delay arcs on different sides of the image. This is due to the shifting location of the source: on the figure on the left the source is on the starboard side of the array, and in the figure at right the source is on the port side of the array.

It is clear from Fig. 32 that the vertical profile has the expected delay structure to the closest scatterers one would intuitively imagine given the vertical displacement of the array visible in Fig. 31. Upon beamforming, the spatial wavenumber spectra of the images corresponding to these slices are non-overlapping but border each other, providing continuous wavenumber coverage with twice the area and resolution in φ [13].

Beamforming results for this data simulation will be demonstrated in the *Results* section of this document.

III. Beamformer development

Fast Factorized Back-projection (FFBP)

Time-domain algorithms are preferable for the 3D SAS image reconstruction problem because of their generality and ability to cope with non-ideal platform geometries. Even in the 2D beamforming case, however, standard time-domain beamforming is impractical without special hardware acceleration (e.g. GPU's). 3D volumetric beamforming adds an additional dimension both to the array and the volume, turning an $\mathbf{O}(N^3)$ process into an $\mathbf{O}(N^5)$, where N represents the length of the array (sides of the array, in the 3D case), and length of the sides of the image. The $\mathbf{O}(N^5)$ computational complexity is demonstrated via a simple Matlab script, shown on the following page, which implements three-dimensional backprojecting from an $N \times N$ planar array to an $N \times N \times N$ volume with customizable values for N.

```
% Unfiltered 3D Backprojection demonstration
% T. Marston, 4/7/2022

N   = 64;                          % Number of edge elements & voxels

% Array and volume locations
pos = linspace(-1, 1, N);                              % Position vector
[A_x, A_y, A_z]     = meshgrid(pos, pos, 1);           % Array locations
[V_x, V_y, V_z]     = meshgrid(pos, pos, pos-1);       % Voxel locations

% Vectorize:
A_xyz = [A_x(:), A_y(:), A_z(:)]';
V_xyz = [V_x(:), V_y(:), V_z(:)]';

% Signal range vector:
dx = pos(2) - pos(1);                  % Spatial sampling
r  = 0:dx:5;                           % Range vector

% Simulate a point in the center of the volume
t_dst   = sqrt(sum((A_xyz - mean(V_xyz, 2)).^2));
sig     = zeros(length(r), 1);
sig(2)  = 1;
isig    = interp1(r(:), sig, r(:)-t_dst, 'spline', 0);

% Image data
volume = zeros(N^3, 1);          % Reconstruction volume organized
                                 % as a vector
% Timer:
tic;

% Backprojection loop
for n = 1:N^2                     % Loop over all elements in the array
        % Distance from array element to volume
        dst = sqrt(sum((A_xyz(:, n) - V_xyz).^2))';

        % Backproject and add to volume
        volume = volume +...
            interp1(r, squeeze(isig(:, n)), dst, 'linear', 0);
end
T = toc;

% Reshape back to a cube
volume = reshape(volume, [N, N, N]);
```

Running the code snippet from the previous page and rendering the output using VOXview, a Matlab-based rendering tool developed as part of this project, the following image is generated.

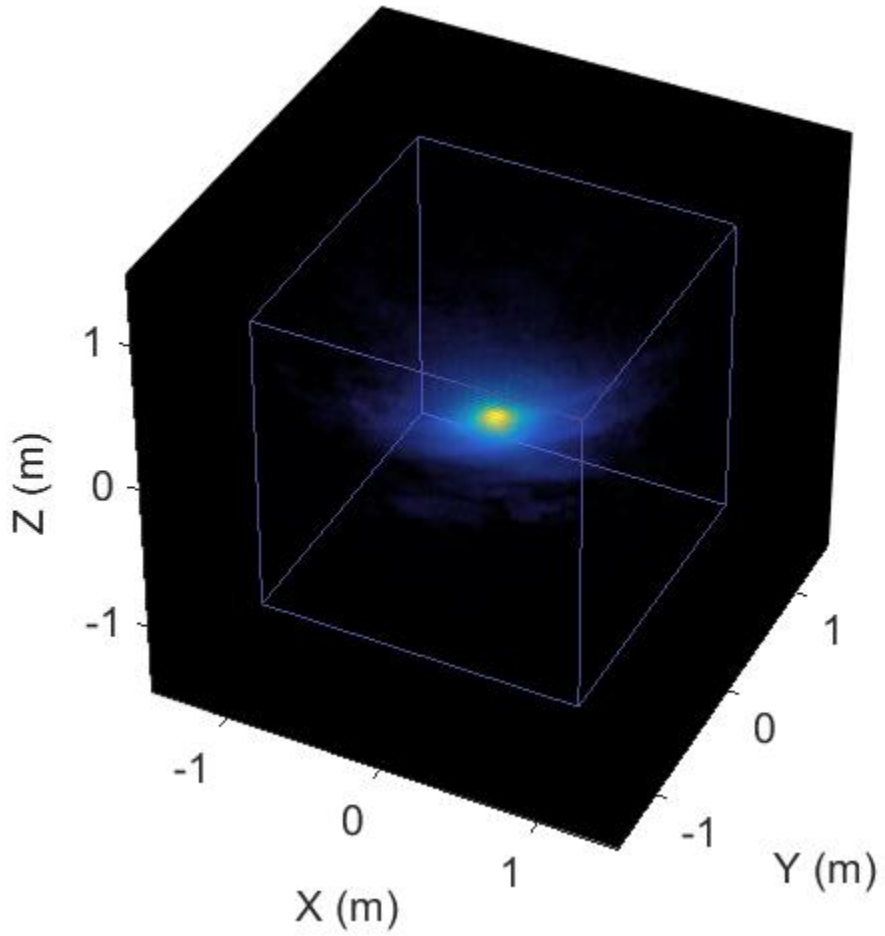

Figure 34: A volumetric rendering of the output of the 3D back-projection demonstration code from the previous page. The simulated scatterer located at the center of the volume is visible.

This demonstration code is not practical for real acoustic data which must account for basebanding phase and apply a ramp filter to preserve spectral calibration. These operations add additional complexity and overhead to the previous beamforming problem, but the demonstration code suitably shows the relative scaling of the computational burden for different sizes of beamforming problems. For example, in Fig. 35 the execution times are plotted for the following values of N: 8, 16, 32, 64, and 128.

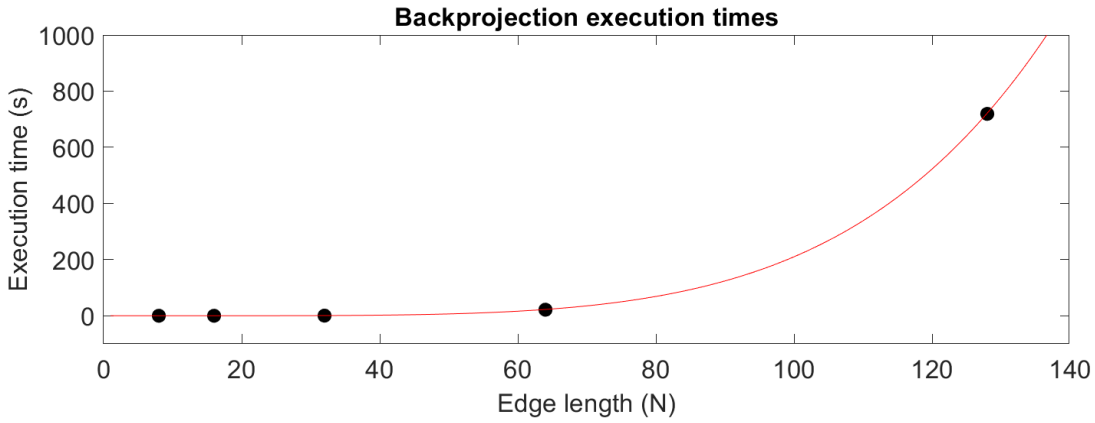

Figure 35: Beamforming execution times for the back-projection routine as computed on a RAZER Blade15 laptop using the following values of *N*: 8, 16, 32, 64, and 128. The red curve is a plot of the function $y = 2.1*10^{-8}x^5$.

The execution times shown in the previous graph were computed on the author's laptop, a RAZER BLADE 15 with a 2.20 GHz Intel® Core™ i7-8750H processor. Execution times may vary between systems, however the trend should be preserved and can clearly be modelled as a function proportional to N^5. This trend shows that increasing N by a factor of four results in an execution time *over 1000 times greater*. Related to BOSS processing, scaling the array (or adding transducers) to increase resolution, beamforming the full region ensonified by the sonar, and doing so at a resolution necessary for performing spectral analysis quickly becomes a computational bottleneck that, using the simple algorithm demonstrated on the previous page, becomes impractical very quickly.

To deal with the analogous two-dimensional back-projection problem from a linear synthetic aperture array, a technique called "Fast Factorized Back-Projection" (FFBP) was developed in the SAR community by Ulander et al [18]. Fast factorized back-projection is an algorithm specifically designed to provide the flexibility of generalized beamforming algorithms like back-projection but have the $N^2\log(N)$ efficiency of Fourier beamformers. The algorithm exploits the fact that for small sub-apertures (i.e. the length of the sub-aperture is significantly less than the distance to the closest range in the image), the fully sampled image can be more efficiently represented in polar coordinates than in Cartesian coordinates. The FFBP algorithm bypasses the high computational burden of directly back-projecting to a Cartesian coordinate image by dividing the backprojection operation into a set of smaller problems which are systematically computed and combined in stages. There are various implementations of the FFBP algorithm; one form divides the synthetic aperture into a set of very small sub-apertures which are independently beamformed to polar coordinate systems centered on the individual sub-apertures. Subsequently, pairs of polar coordinate images from adjacent sub-apertures are coherently combined to form polar-coordinate images with twice the resolution, and this process of sub-aperture combination is performed for $log_2(N)$ stages (assuming an array length of N). At the final stage, a polar-to-Cartesian transformation is applied to generate the final high resolution image in Cartesian coordinates. For strip-map images this is slightly modified in that the final stage off the FFBP polar-coordinate sub-aperture combination process may not represent the

entire aperture but only a small portion, and these polar-coordinate sub-apertures are converted to Cartesian coordinates before combination with the other final-stage polar-coordinate sub-apertures. In other implementations the scene is also factorized, see, e.g. section *F* of [18]. This prevents the need for a final polar-to-Cartesian transformation. The general FFBP process, however, is graphically illustrated in Fig. 36.

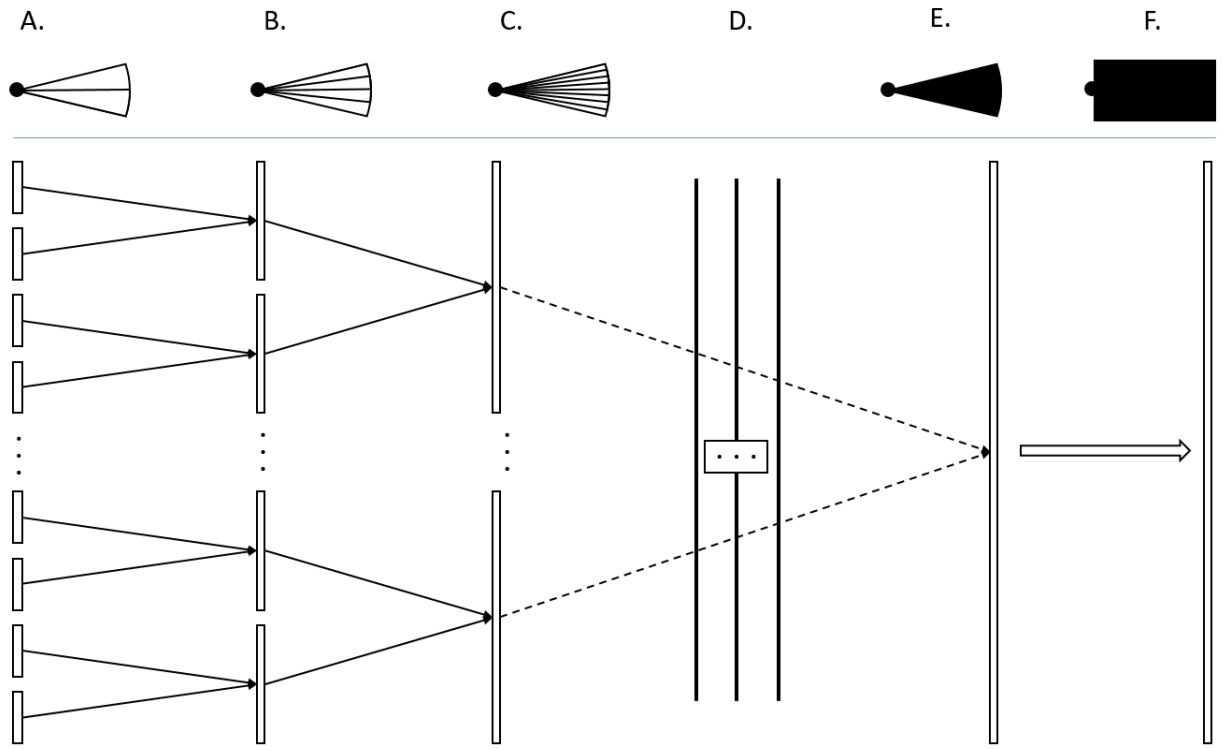

Figure 36: An illustration of the FFBP process for 2D SAS or SAR systems. The left-hand side, (column A.) shows an aperture that has been divided into small segments, each with two contributing channels. These are beamformed into polar aperture images that have only two "angles" (*A.*, top figure). In subsequent columns, *B.* and *C.*, the sub-aperture images are combined in adjacent pairs to halve the number of images but double the resolution of each image. In column *C*, for example, each sub-aperture image is comprised of 8 contributing pings and has 8 samples in the angle dimension. This process is repeated recursively through several un-illustrated stages (column *D.*) until the final stage in which all pings for the aperture or block of pings have been combined to form a single high resolution polar image (column *E.*). A polar-to-Cartesian transformation is performed to generate a final image in Cartesian coordinates and perhaps coherently combine with images from other blocks.

Typically, this beamforming algorithm is applied to base-banded complex data and each sub-aperture transformation is accompanied by an interpolation-conditioning phase transformation that centers the spectral content of the image on the origin. This is followed by the reversal of the operation in the transformed coordinates, to reposition the spectrum to the correct indices relative to the origin. The phase conditioners are described in equations (4) and (5):

$$\boldsymbol{\Phi}_{pre} = \boldsymbol{e}^{-i2k_0 r} \qquad \textbf{Eq. 4}$$

$$\boldsymbol{\Phi}_{post} = \boldsymbol{e}^{i2k_0 \tilde{r}} \qquad \textbf{Eq. 5}$$

where Φ_{pre} is the conditioner applied to the data prior to interpolation and Φ_{post} is the spectral re-positioner that must be applied if the sub-aperture images are to be combined successfully in the new coordinate system. In (4) and (5) $\boldsymbol{r}$ represents the ranges of the image samples relative to the source coordinate system, and $\tilde{\boldsymbol{r}}$ represents the ranges of samples interpolated to the new grid, but still relative to the source origin. Note that if (4) and (5) are not applied the image quality of the sub-apertures will degrade extremely rapidly as they are combined.

Though the FFBP algorithm has a theoretical computational complexity of $N^2\log(N)$, the number of stages and type of interpolation operations (e.g. nearest-neighbor, bi-linear, bi-cubic, etc.) have a large impact on both on the practical speed and resulting image quality of this beamforming algorithm. Even when tuned to produce high quality results, however, the speed gains can be considerable. For example, in Fig. 37, standard back-projection and a tuned FFBP algorithm produce imagery with nearly identical quality, yet the FFBP algorithm executed over forty times faster than the standard back-projection algorithm.

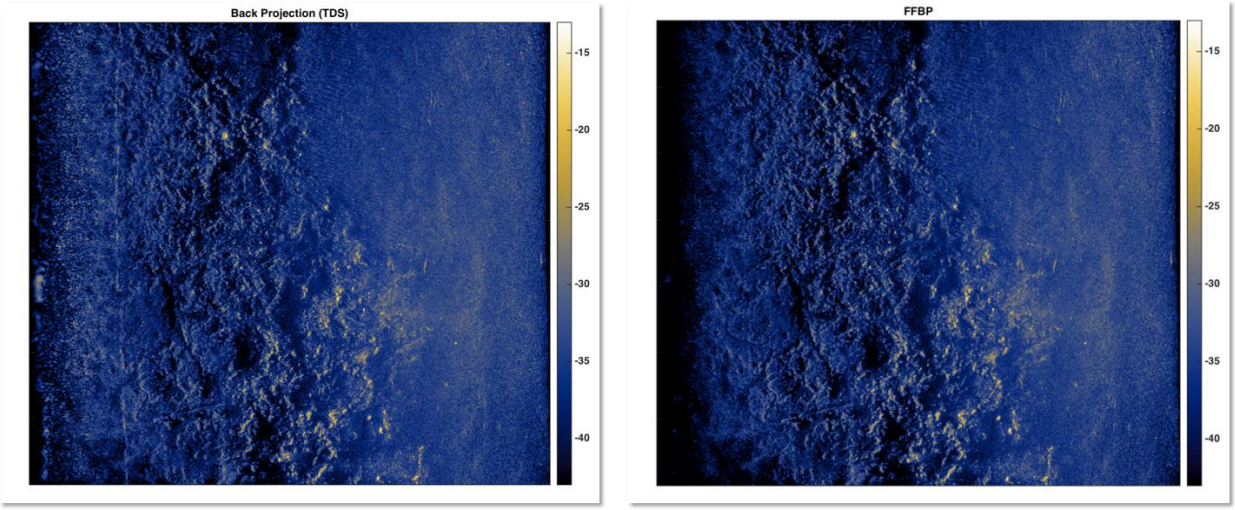

Figure 37: The figure on the left was generated using standard back-projection and the figure on the right was computed using a fast-factorized back-projection algorithm tuned to reduce interpolation artifacts. The FFBP algorithm executed over 40 times faster than the standard back-projection algorithm.

Extension of FFBP to 3D downward looking sonars

The scan an imaging geometry for BOSS systems is very different than the two dimensional, side-looking imaging geometry FFBP was originally designed for, however with the appropriate modifications it lends itself extremely well to the 3D downward looking case. Several facets of the problem are particularly convenient: the hydrophone array is perfectly linear in the across-track dimension (the dimension of the "real" aperture), scans are often approximately planar (note: the apertures are *not* approximated as a plane in the algorithm, but this approximately planar behavior prevents having to deal with some challenging sampling issues that can arise in the more general case), and the synthetic aperture extends in only one dimension. Similar to [19] the beamforming process can be broken into two stages, beginning with a real-aperture imaging stage that can be performed efficiently via pre-computed tables, and a synthetic-aperture beamforming stage that combines these real-aperture "frames." Note that pre-computed tables

can be used in the real-aperture case because, unlike in the synthetic aperture case, the imaging geometry is fixed mechanically by the design of the array.

To illustrate the 3D FFBP algorithm for BOSS imaging, consider a linear array with N evenly spaced elements in the across-track dimension located at $\boldsymbol{A}_{xyz} = (x_{a1,2,\ldots N}, y_{a1,2,\ldots N}, z_{a1,2,\ldots N})$, where $y_{a1,N} = -W/2, W/2$, $x_{a1\ldots N} = z_{a1\ldots N} = 0$, and W is the width of the array in the across-track dimension. Importantly, the array also has a transmitter off-axis to the receive array and located at $\boldsymbol{T}_{xyz} = (x_t, y_t, z_t)$. In the present case we are using the North-East-Down (NED) coordinate convention where x, y, and z, are north, east and down respectively when roll, pitch, and heading are all zero degrees, and the center of the array (i.e. the centroid of the array elements) is the origin of the body-centered coordinate system. At non-zero headings, x is in the along-track direction (direction of movement), and y is in the across-track direction, and is the dimension along which the real aperture receive array is distributed. This setup is illustrated in Figure 38.

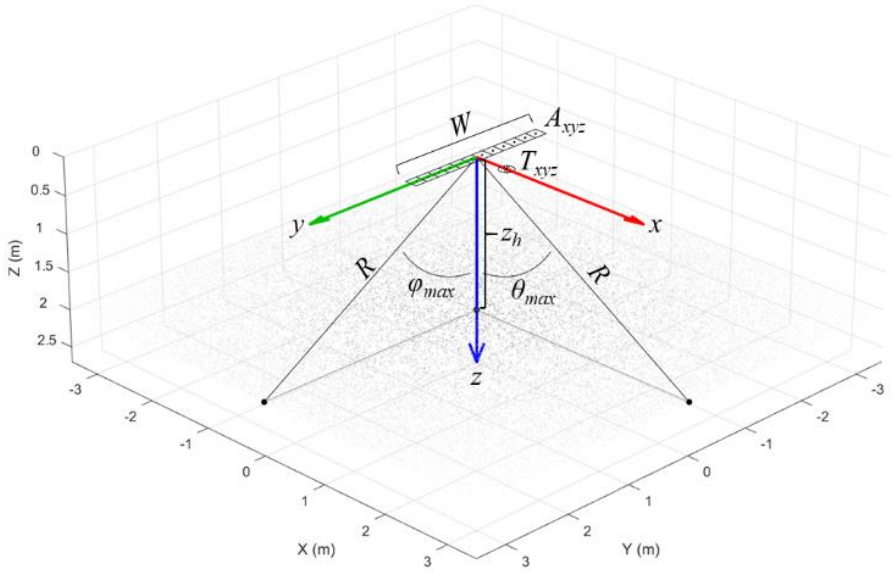

Figure 38: Illustration of the coordinate system of the down-looking synthetic aperture sonar, used for subsequent illustrations of the 3D FFBP beamforming process. Note that the coordinate system is in the North-East-Down (NED) convention, and that *x* is the direction of array motion.

For an omnidirectional source transmitting a signal with bandwidth f_{bw} and center-frequency f_0, a system altitude of z_h meters above the sediment, and a maximum ensonification range R determined by $R = c/(2f_p)$, (where c is sound velocity f_p is pulse repetition frequency), sampling intervals can be computed in each dimension to satisfy the constraint that the spatial spectrum of the 3D image should not be aliased. In the z-dimension, the region of wavenumber support will lie between the maximum value imposed by the sound speed and maximum frequency of the system: $k_{z_max} = 2\pi(f_0 + f_{bw}/2)/(c/2)$ and the minimum wavenumber, determined by the minimum frequency and maximum look angle supported by the altitude and system range: $k_{z_min} = 2\pi(z_h/R)(f_0 - f_{bw}/2)/(c/2)$. Using $\Delta z = 2\pi/(k_{z_max} - k_{z_min})$, the spatial sampling requirements in the z dimension can be expressed:

$$\Delta \boldsymbol{z} \leq \frac{2\pi}{k_{z_max} - k_{z_min}} = \frac{c}{2(f_{max} - \cos(\theta_{max}) f_{min})} \quad \textbf{Eq. 6}$$

where $f_{max/min} = f_0$ +/- $f_{bw/2}$ and $\theta_{max} = \cos^{-1}(z_h/R)$. Likewise, in the *y*-dimension the sample spacing is determined by the wavenumber span, but can also be directly found by recognizing that Δy must be less than or equal to half of the projection of the shortest in-band wavelength onto the horizontal plane:

$$\Delta \boldsymbol{y} \leq \frac{c}{2\left(f_0 + \frac{f_{bw}}{2}\right)\sqrt{1 - \left(\frac{z_h}{R}\right)^2}} = \frac{c}{2 \sin(\theta_{max}) f_{max}}. \quad \textbf{Eq. 7}$$

The sampling requirement for Δx is similar to Δy but increased by a factor of two due to the fact that the system is nearly monostatic in the *x* dimension, which effectively causes a doubling of the phase gradients:

$$\Delta \boldsymbol{x} \leq \frac{c}{4\left(f_0 + \frac{f_{bw}}{2}\right)\sqrt{1 - \left(\frac{z_h}{R}\right)^2}} = \frac{c}{4 \sin(\theta_{max}) f_{max}}. \quad \textbf{Eq. 8}$$

Note that vehicle velocity does not show up in (6 – 8), nor does the actual spacing between array points in the synthetic or real aperture dimension: this is because (6 – 8) represent the *image* sampling requirements assuming particular wavenumber coverages. Note, furthermore, that in (6 – 8) the maximum wavenumber coverage is being determined by the system geometry values z_h and *R*, which together specify a minimum grazing angle, and thus maximum horizontal wavenumber, the system is capable of observing. Raw data-collection is constrained to follow the same requirements to produce complex imagery free from grating-lobes caused by under-sampling, however even in the event that the data is under-sampled, eq.'s 6 – 8 must still be adhered to in the image domain to recover the impulse response over the populated region of the spatial spectrum.

Using equations (6 – 8), voxel size estimates can now be determined for images produced by systems with specified bandwidths, altitudes and pulse repetition frequencies. Consider, for example, a very reasonable BOSS system having parameters f_0 = 15 kHz, f_{bw} = 20 kHz, f_p = 20 Hz, z_h = 5 m, c = 1500 and R = 1500/(2*20) =37.5 m. For these parameters, Δz = ~0.031 m, Δy = ~0.030 meters and Δx = ~0.015 meters. Assuming we wish to back-project to the minimum-sized volume in Cartesian space encompassing only the sediment and sub-sediment regions ensonified by the sonar, we must back project to a volume $2(R^2 - z_h^2)$ $(R - z_h)$ m^3 in size. For the specified sampling parameters this equates to a three dimensional array containing approximately 6 billion voxels having edge lengths of $\Delta x \times \Delta y \times \Delta z$. Furthermore, this must be done *for each channel of data acquired by each element of the real-array*. This is why standard back-projection for 3D data is so slow.

The challenge of the beamforming directly in Cartesian coordinates now becomes clearer. In contrast, the number of voxels required to fully sample a back projected signal received by an element is greatly reduced if the space is represented in spherical coordinates (r, θ, φ) where r is range from the origin, θ is azimuth angle $\tan^{-1}(x/z)$, and φ is elevation angle $\pi/2 - \cos^{-1}(y/(x^2 + y^2 + z^2)^{-1/2}) = \tan^{-1}(y/(x^2 + z^2)^{-1/2})$. This expression of spherical coordinates varies slightly from the

standard version due to the axes conventions and the fact that φ is offset by $\pi/2$. The reason behind these unconventional spherical coordinate axes definitions is to account for the fact that array symmetry is along the *x*-axis (real-aperture), and we want elevation, which is ordinarily symmetric around the *z*-axis, to be symmetric around the physical array. Furthermore, by offsetting φ by $\pi/2$ then boresight θ, $\varphi = 0$ points directly down toward the sediment, as it ought.

The first task for back-projection in spherical coordinates is to determine the sampling constraints in each dimension. In the range dimension the sampling constraint is simply determined by the bandwidth of the system:

$$\Delta \boldsymbol{r} \leq \frac{\boldsymbol{c}}{\boldsymbol{2f_{bw}}}\,. \qquad \textbf{Eq. 9}$$

Assuming $r_s << z_h$ where r_s is the distance of the transmitter to the center of the array, sample spacing in φ can be determined by solving for $\Delta\varphi$ so that the maximum difference in range between any point on the array and two points at φ_n and φ_{n+1} is less than half the minimum in-band wavelength, i.e.:

$$\boldsymbol{\arg\max\left|r_{ap}(\varphi_p) - r_{ap}(\varphi_p + \Delta\varphi)\right| \leq \lambda_{min}/2,} \qquad \textbf{Eq. 10}$$

where r_{ap} is the distance the distance of a point on the array to some point $p = (r_p, \varphi_p)$, where φ_p lies between +/- φ_{max} = +/- θ_{max} = +/-$\cos^{-1}(z_h/R)$, and r_p lies between z_h and R. Using the standard spherical coordinate distance formula, the distance between the element at the edge of the array located at $A_{r\theta\varphi} = (-W/2 \leq x_a \leq -W/2, 0, \pi/2)$ to point p is:

$$\boldsymbol{r_{ap} = \sqrt{r_p{}^2 + x_a^2 - r_p x_a \sin(\varphi_p)}}\,. \qquad \textbf{Eq. 11}$$

Taking the partial derivative with respect to φ_p is useful for determining certain types of behavior. The partial derivative with respect to φ_p is:

$$\frac{\partial r_{ap}}{\partial \varphi_p} = -\frac{r_p x_a \cos(\varphi_p)}{\sqrt{r_p{}^2 + x_a^2 - 2 r_p x_a \sin(\varphi_p)}}\,. \qquad \textbf{Eq. 12}$$

When the distance to the scatterer is much greater than the dimension of the array, i.e. $r_p >> x_a$, the dependency on r_p drops out and equation (12) reduces to:

$$\frac{\partial r_{ap}}{\partial \varphi_p} \approx -\frac{r_p x_a \cos(\varphi_p)}{\sqrt{r_p{}^2}} = \boldsymbol{x_a \cos(\varphi_p)}\,, \text{ when } r_p >> x_a, \qquad \textbf{Eq. 13}$$

showing that the maximum rate of change is at $\varphi_p = 0$, (since the cosine function is at a maximum at 0), asymptotically converges to x_a, and in general has only a weak dependence on r_p for most of the imaged region since in general the distance to the sediment is greater than the width of the array. From (13) we can also see that $\frac{\partial r_{ap}}{\partial \varphi_p}$ is proportional to x_a and therefore the maximum occurs for elements at the edge of the array, i.e. x_a = +/- $W/2$. Thus, under the assumption that $r_p >> x_a$ we can determine the sampling constraint by combining eq.'s (11) and (13):

$$W/2 \leq \frac{\Delta r}{\Delta \varphi} = \frac{\lambda_{min}/_2}{\Delta \varphi}, \quad \textbf{Eq. 14}$$

$$\Delta \varphi \leq \frac{\lambda_{min}}{W}. \quad \textbf{Eq. 15}$$

For the previously specified system, this would equate to a sample spacing of $\Delta\varphi$ = ~1.7 degrees.

For a line-array and source distributed along the *y*-axis the array becomes symmetric in θ and therefore $\Delta\theta = 2\pi$ (its period), requiring no new interpolations to populate the volume. In reality the system is not axisymmetric, but a method for transforming imagery from a single back projection plane in θ to planes at other values will be demonstrated, eliminating the need for $\Delta\theta$ to be reduced from 2π and enabling direct projection to a 3D coordinate system with a different origin. The spherical coordinate representation for a single-ping from a system like the one described would therefore require interpolation to only $2\cos^{-1}(z_h/R)\,(R - z_h)/\,(\Delta r\,\Delta\varphi)$ = ~83 thousand voxels: a reduction of four orders of magnitude vs. fully sampled Cartesian coordinate back projection. This is an extreme reduction, but it is mildly deceptive: in the standard back-projection algorithm the Cartesian coordinate system to which the channel data is back projected is the final coordinate system in which the image will be expressed. In contrast, the spherical coordinate system to which the previous example was back-projected was only a temporary coordinate system and the data will be incorporated in several further interpolation stages. Nonetheless it is indicative of the relative efficiencies of the algorithms being compared and the new algorithm being described.

Asymmetry correction

In the previous section it was shown that, due to array axis symmetry, data from each ping could be back projected to a volume simply by beamforming the data in *r* and φ. Due to symmetry, there is no reason to populate θ with more than one value. Therefore, in the present algorithm, the first stage of 3D FFBP is simply to perform this beamforming operation in *r* and φ, creating a 2D matrix in sonar body coordinates (we will call these images "frames").

In reality, most real BOSS systems are not symmetric about the receive array if the transmitter is included; that is, generally the signal source is not located on the same axis as the receive array, breaking the symmetry. Consequentially, images back projected to planes of varying θ are not equivalent. This can be easily observed by considering the difference of the round-trip time of flight between the source located along the *x*-axis at $S_{r\varphi\theta} = (r_s, \theta_s = \pi/2, \varphi_s = 0)$, two points located on the image frames having different rotations in theta: $P_{1,2} = (r_p, \theta_{1,2}, \varphi_p)$, and the array: $A_{r\theta\varphi} = (r_a = -W/2 \ldots W/2, \theta_a = 0 \ldots 0, \varphi_a = \pi/2 \ldots \pi/2)$. Due to receive array element symmetry in θ the propagation distance from points $P_{1,2}$ to $A_{r\theta\varphi}$ are identical and cancel, leaving only the propagation distance from the source to the points:

$$r_{sp1,2} = \sqrt{r_s^2 + r_p^2 - 2r_s r_p\left(\cos(\varphi_p)\sin(\theta_{1,2})\right)} \quad \textbf{Eq. 16}$$

In eq. (16) a theta-dependency (i.e. frame rotation around the array axis) is introduced into the source-to-frame distance by the sine term. The implication of this theta dependency is that a single frame can no longer represent the data back-projected to spherical-coordinates, and

multiple frames spaced by some $\Delta\theta$ must be computed. The problem with this approach is that it greatly increase the number of interpolations: rather than requiring back-projection to only one frame, back-projection is now needed to M = ceiling($2\theta_{max}/\Delta\theta$) frames.

Fortunately an alternative exists, which will be called an *asymmetry correction.* The asymmetry correction, which is amounts to a range transformation with a (r, φ) dependency relating frames at varying θ, is possible to apply if a small amount of image distortion is allowed in the very near field of the array (i.e. similar to r_s, the distance between the source and receive array). In general, the distance to the sediment is much longer than this distance, (i.e. $r_p >> r_s$), and the impact on the image is negligible. The asymmetry correction is applied as follows: as in the truly axisymmetric case a single back-projection is performed, however due to asymmetry the projection plane θ_r must be defined: here, $\theta_r = 0$. The range stretch relating frames at different theta is defined as the distance to each point on the reference array plus the difference of the source-frame distances between a frame at arbitrary θ and the reference frame at θ_r:

$$\tilde{\boldsymbol{r}} = \boldsymbol{r} + \left(\sqrt{\boldsymbol{r_s}^2 + \boldsymbol{r}^2} - \sqrt{\boldsymbol{r_s}^2 + \boldsymbol{r}^2 - 2\boldsymbol{r_s}\boldsymbol{r}(\cos(\boldsymbol{\varphi})\sin(\boldsymbol{\theta}))}\right)/2. \qquad \textbf{Eq. 17}$$

In (17) the difference term reduces to zero when $\theta = \theta_r = 0$ and thus the stretch transformation $\tilde{r} = r$ on the reference plane, which makes intuitive sense (no transformation should be necessary at the reference plane). For non-zero values of θ however, $\tilde{r} \neq r$. One could now populate a volume in spherical coordinates by creating multiple frames spaced by $\Delta\theta$, though even this is not necessary: the second stage of the 3D FFBP is merely the projection of the frame onto a new 3D spherical sub-aperture coordinate system, and this can be done directly using eq. (17) without having to generate 3D sub-apertures or multiple frames per pings.

The very first stage, therefore, of the 3D FFBP algorithm with a source (or sources) not located along the axis of the receive array, is to back-project each ping p to a 2D frame $FR_p(r, \varphi)$ at $\theta_r = 0$. Each of these frames are back-projected in sonar body coordinates (i.e. fixed relative to the source and array), and therefore can be performed extremely rapidly using pre-computed interpolation tables and phase transformations. Examples of a set of these frames for a particular scan are shown in the following figure.

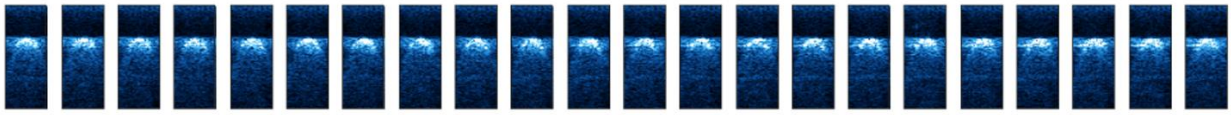

Figure 39: An example of a series of real-aperture frames generated in vehicle body coordinates.

Following the construction of these frames, the asymmetry correction in (17) can be used to map each frame into the 3D coordinates of the first sub-aperture. This frame-to-sub-aperture conversion will be discussed in the next section.

Stage 1 sub-aperture origin and orientation definitions

Following the creation of K two-dimensional real aperture frames $FR_{p\ldots(p+K-1)}(r, \varphi)$ where p is the ping index (i.e. $FR_p(r, \varphi)$ is the frame corresponding to the p^{th} ping) these frames are combined coherently in a new spherical coordinate system – the first-stage sub-aperture coordinate system

corresponding to the K frames – the origin of which is defined in Cartesian space as the average of the origins of the contributing frames, i.e.:

$$\boldsymbol{x_m^{SA_stage1} = \frac{1}{K}\sum_{k=p}^{p+K-1} x_k,} \quad \textbf{Eq. 18}$$

$$\boldsymbol{y_m^{SA_stage1} = \frac{1}{K}\sum_{k=p}^{p+K-1} y_k,} \quad \textbf{Eq. 19}$$

$$\boldsymbol{z_m^{SA_stage1} = \frac{1}{K}\sum_{k=p}^{p+K-1} z_k,} \quad \textbf{Eq. 20}$$

where the subscripts and superscripts on the left hand side denote that these values are for the origins of the m^{th} stage-1 sub-aperture, and on the right hand side (x_k, y_k, z_k) are the earth-coordinate locations of the real aperture array at the p^{th} through $(p + K – 1)^{th}$ pings.

In addition to the location of the origin, the relative orientation of the first-stage sub-aperture must be determined about its origin. The real and synthetic aperture angles φ and θ, for example, should conform as much as possible to the orientations of the contributing sub-apertures. The coordinates for each contributing frame are defined in vehicle body coordinates (i.e. referenced to the orientation of the array) and can be related by:

$$\begin{bmatrix} x^{Body} \\ y^{Body} \\ z^{Body} \end{bmatrix} = \boldsymbol{C^{E2B}} \begin{bmatrix} x^{Earth} - x_p \\ y^{Earth} - y_p \\ z^{Earth} - z_p \end{bmatrix}, \quad \textbf{Eq. 21}$$

or,

$$\begin{bmatrix} x^{Earth} \\ y^{Earth} \\ z^{Earth} \end{bmatrix} = \boldsymbol{C^{B2E}} \begin{bmatrix} x^{Body} \\ y^{Body} \\ z^{Body} \end{bmatrix} + \begin{bmatrix} x_p \\ y_p \\ z_p \end{bmatrix}, \quad \textbf{Eq. 22}$$

where x_p, y_p and z_p are the coordinates (in the earth-frame) of the origin of the array at the p^{th} ping and $\boldsymbol{C^{B2E}}$ and $\boldsymbol{C^{E2B}}$ are the orthogonal earth-to-body coordinate transformations:

$$\boldsymbol{C^{E2B} = \left[H * \Phi_{pitch} * \Theta_{roll}\right]^T,} \quad \textbf{Eq. 23}$$

$$\boldsymbol{C^{B2E} = \left[H * \Phi_{pitch} * \Theta_{roll}\right].} \quad \textbf{Eq. 24}$$

In (23) and (24), bold $\boldsymbol{H}$, $\boldsymbol{\Theta_{roll}}$, and $\boldsymbol{\Phi_{pitch}}$ denote rotation matrices around the z (yaw), y (pitch), and x (roll) axes respectively. Using the previously defined spherical coordinate conventions, frame spherical coordinates are always defined relative to the body as:

$$\boldsymbol{r_{frame} = \sqrt{x_{body}^2 + y_{body}^2 + z_{body}^2},} \quad \textbf{Eq. 25}$$

$$\boldsymbol{\theta_{frame} = \tan^{-1}\left(\frac{x_{body}}{y_{body}}\right),} \quad \textbf{Eq. 26}$$

$$\boldsymbol{\varphi_{frame} = \tan^{-1}\left(y_{body}/\sqrt{x_{body}^2 + z_{body}^2}\right).}$$ **Eq. 27**

Thus, via (21 - 27) the p^{th} frame's spherical coordinate grid points can be related the location of the origin and the orientation values (heading, pitch and roll) for the p^{th} ping.

As with the origin of the first sub-aperture, the orientation of the first sub-aperture can be found by averaging the orientation angles of the contributing frames. Because these values are angles, in the present algorithm the mean is found via complex phasors to prevent problems should the orientations contain a phase wrap:

$$\boldsymbol{\Theta_m^{SA\ stage1} = \angle\left(\sum_{k=p}^{p+K-1} e^{j\Theta_p}\right),}$$ **Eq. 28**

$$\boldsymbol{\Phi_m^{SA\ stage1} = \angle\left(\sum_{k=p}^{p+K-1} e^{j\Phi_p}\right),}$$ **Eq. 29**

$$\boldsymbol{H_m^{SA\ stage1} = \angle\left(\sum_{k=p}^{p+K-1} e^{jH_p}\right).}$$ **Eq. 30**

In fact, though wraps can be fairly common for heading it is highly unlikely that either roll Θ or pitch Φ will go through a phase-wrap, but eq.'s (28) and (29) are still used for roll and pitch for consistency.

Sampling constraints in θ

Now that the origin and orientations of the first 3D sub-aperture stages have been determined the sample spacing in θ must also be determined to ensure the spherical-coordinate representation is adequately sampled in the synthetic aperture dimension. Adequate sampling prevents artifacts from accumulating in subsequent interpolation operations and prevents aliasing in the spatial spectrum. Assuming that motion is dominated by motion in the *x*-dimension of each frame and that roll and translations in the *y*-dimension are small enough over *K* frames that the span and sample rate in φ don't need redefinition, the necessary sample rate $\Delta\theta$ can be found in a manner similar to $\Delta\varphi$ but recognizing that the extrema of the change in time delay occurs at the central element of the array (i.e. at $x = 0$), rather than at the end, and for $\varphi = 0$. Furthermore, the round trip distance must now also account for transmitter movement in *x*. Assuming $r_s << r_p$, an along-track sample spacing of Δx and a sub-aperture size of *K*-frames, the round-trip distance from the source to any point (r, θ, 0) and back to the central array element is:

$$\boldsymbol{r_{sp} \approx 2\sqrt{r^2 + (K\,\Delta x/2)^2 - 2r(K\,\Delta x/2)\sin(\theta)}\,.}$$ **Eq. 31**

Similar to the case of $\Delta\varphi$ but using $K\Delta x$ instead of W, $\Delta\theta$ can be found using:

$$\boldsymbol{\Delta\theta \leq \lambda_{min}}/(2K\Delta x),$$ **Eq. 32**

where the extra factor of two is a result of the moving source.

It can be illustrative of representational efficiency to determine the number of samples N_θ in θ necessary to represent a fully sampled image composed of *K* frames. The number of samples in θ that a sub-aperture contains will be equal to the total angle span divided by the sample spacing:

$$\boldsymbol{N_\theta = \left\lceil \frac{2\theta_{max}}{\Delta\theta} \right\rceil.}$$ **Eq. 33**

Substituting λ_{min} for $c/(f_{max})$ in eq. (8) gives:

$$\boldsymbol{\Delta x = \frac{\lambda_{min}}{4\sin(\theta_{max})}.}$$ **Eq. 34**

Inserting (34) into (32) we see:

$$\boldsymbol{\Delta\theta \leq \frac{\lambda_{min}}{2K\frac{\lambda_{min}}{4\sin(\theta_{max})}} = \frac{2\sin(\theta_{max})}{K}.}$$ **Eq. 35**

Assuming we are critically sampled and substituting (35) into (33) we can arrive at an equation relating the number of samples N_θ to K and the total integration angle in θ:

$$\boldsymbol{N_\theta = \lceil \alpha(\theta_{max})K \rceil,}$$ **Eq. 36**

where:

$$\boldsymbol{\alpha(\theta_{max}) = \frac{\theta_{max}}{\sin(\theta_{max})}\,.}$$ **Eq. 37**

In (37), α is a *representational efficiency* term that shows how the number of contributing frames relates to the number of samples required in a sub-aperture given a particular integration angle span. For the narrow and wide integration angle cases the limits can be found:

$$\boldsymbol{\lim_{\theta_{max}\to 0} \alpha(\theta_{max}) = 1}$$ **Eq. 38**

$$\boldsymbol{\lim_{\theta_{max}\to \pi/2} \alpha(\theta_{max}) = \pi/2}$$ **Eq. 39**

Eq.'s (38) and (39) illustrate why spherical coordinates are an efficient sub-aperture representation: for small integration angles the efficiency ratio is almost 1, i.e. the number of contributing frames and number of samples in θ are almost equal. For large integration angles, however, the ratio isn't that much worse, reaching, in the limiting case of a full 180 degrees of integration, only about 1.5x the number of contributing frames. An "optimal" representation would have a uniform efficiency vs. integration angle.

Frame to sub-aperture conversion

Following the definition of the sub-aperture coordinate system and image grid-points in r, θ, φ, the K contributing frames can be projected into the sub-aperture coordinate system and coherently combined to generate the first low-resolution 3D sub-aperture of the FFBP process. The particulars of this process are complicated by the fact that, generally speaking, interpolations operate fastest when the source sample points lie on a monotonically increasing, uniform grid. The grid points of the contributing frames, when described in sub-aperture coordinates, are *not* necessarily monotonically increasing or sampled on a uniform grid. However this is the case if the target coordinates, (the grid coordinates of the sub-aperture), are described in the coordinate system of the frame. Therefore, the projection of frame samples onto the sub-aperture grid is best

accomplished if the sub-aperture grid points are represented first in the coordinate system of each frame. The following process is used to perform the frame-to-sub-aperture conversion:

1) Map the sub-aperture grid points to Cartesian earth-coordinates using a spherical-to-Cartesian transformation.
2) Subtract the Cartesian earth-coordinate origin of the current frame from the sub-aperture grid coordinates (which now appear as a spherical grid that has been translated relative to the origin).
3) Map the sub-aperture grid points to the p^{th} frame spherical coordinate system using a Cartesian-to-spherical transformation.
4) Reducing the number of variables describing the sub-aperture coordinates from three to two by using the asymmetry correction in (14), i.e.: $SA(r_p, \theta_p, \varphi_p) \rightarrow SA(\tilde{r}_p, \varphi_p)$.
5) Perform the interpolation of the frame sample-points onto the grid coordinates (described in the frame's own coordinate system) using a two-dimensional interpolation.
6) Repeat for all frames, coherently summing the results for each interpolation

Note that step 6 can be accomplished because each frame is being back-projected onto the same set of spherical coordinate grid points – these points are just transformed into the native coordinates of each frame for each interpolation. Also note that, prior to and after the interpolation in step 5 the phase transformations in eq.'s (4) and (5) must be applied.

Hierarchical sub-aperture combinations

In the 3D FFBP algorithm defined in this paper the first stage three-dimensional sub-apertures are constructed sets of K frames, where $K \geq 2$. (In the present algorithm, $K = 8$). Origin locations and sub-aperture orientations are combined exactly as done for the frame-to-sub-aperture case, e.g.:

$$\boldsymbol{x_o^{SA\ stage\ n+1}} = \frac{1}{2}\sum_{k=m}^{m+1} \boldsymbol{x_k^{SA\ stage\ n}}, \qquad \textbf{Eq. 40}$$

$$\boldsymbol{y_o^{SA\ stage\ n+1}} = \frac{1}{2}\sum_{k=m}^{m+1} \boldsymbol{y_k^{SA\ stage\ n}}, \qquad \textbf{Eq. 41}$$

$$\boldsymbol{z_o^{SA\ stage\ n+1}} = \frac{1}{2}\sum_{k=m}^{m+1} \boldsymbol{z_k^{SA\ stage\ n}}, \qquad \textbf{Eq. 42}$$

$$\boldsymbol{\Theta_o^{SA\ stage\ n+1}} = \angle\left(\sum_{k=m}^{m+1} \boldsymbol{e^{j\Theta_k^{SA\ stage\ n}}}\right), \qquad \textbf{Eq. 43}$$

$$\boldsymbol{\Phi_o^{SA\ stage\ n+1}} = \angle\left(\sum_{k=m}^{m+1} \boldsymbol{e^{j\Phi_k^{SA\ stage\ n}}}\right), \qquad \textbf{Eq. 44}$$

$$\boldsymbol{H_o^{SA\ stage\ n+1}} = \angle\left(\sum_{k=m}^{m+1} \boldsymbol{e^{jH_k^{SA\ stage\ n}}}\right). \qquad \textbf{Eq. 45}$$

As in the frame-to-sub-aperture conversions, combination of sub-aperture pairs into the new, higher resolution coordinates must be accompanied by a transformation of the grid for the $(n+1)^{th}$ sub-aperture into the coordinates of each contributing sub-aperture to facilitate interpolation. The steps are as follows:

1) Origins and orientations for the $(n+1)^{th}$ sub-aperture stage are determined via eq.'s (35 – 40).
2) The grid points of the $(n+1)^{th}$ sub-aperture are converted to earth coordinates using a spherical-to-Cartesian transformation.
3) The coordinates are translated by the difference in location between the origins of the target and source contributing sub-apertures, to reference the target sub-aperture grid points relative to the origin of the contributing sub-aperture.
4) A Cartesian-to-spherical transformation to place the target sub-aperture grid coordinates in the same coordinate system as the contributing sub-aperture data.
5) A 3D interpolation is applied to project the contributing sub-aperture data into the new sub-aperture grid coordinates.
6) The process is repeated for the second contributing sub-aperture, summing into the projected data from the first contributing sub-aperture.

As with the frame-to-sub-aperture transformation, step (5) is accompanied by the complex phase shifts from eq. (4) and (5).

Cartesian image generation

Following N stages of sub-aperture combinations using the approach outlined in the previous section, the final sub-aperture stage, composed of $K*2^N$ individual frames from the same number of pings, is converted to Cartesian coordinates to be combined coherently or incoherently with imagery produced by the other portions of the synthetic aperture not included in this sub-aperture. While being relatively straight forward, care must be taken as to how the final image is defined in Cartesian coordinates to appropriately combine data from other full sub-apertures, and prevent interpolating with an excessively large target matrix in three-dimensions, which can result in very slow interpolations.

In the present algorithm, the imaging volume is defined as the minimum size box defined in Universal Transverse Mercator (UTM) coordinates and depth in meters that bounds the ensonified area of the scan, as defined by the maximum range allowed by the ping rate and sound speed, the integration angles of the SAS algorithm, and the maximum altitude of the sediment. An example of bounding box generation based on these parameters is shown for an example MuST system in Figure 40.

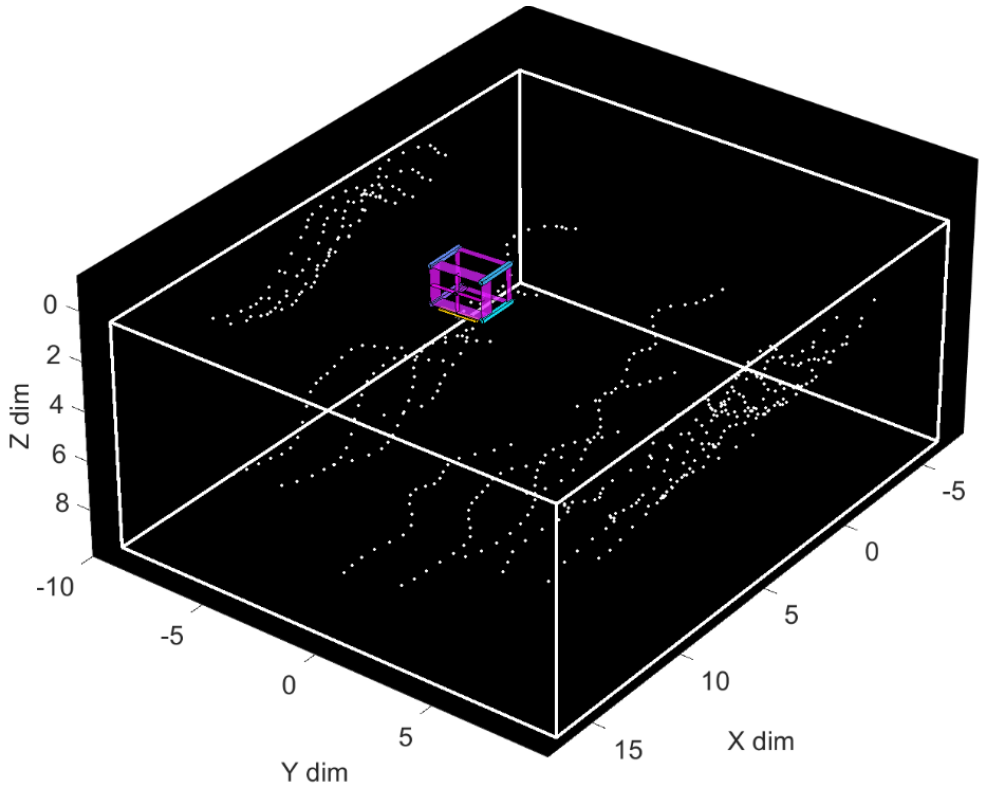

Figure 40: Example of a bounding-box constructed around the extrema of the visible sonar beam, used to determine the region of the 3D image that will be generated. This box should bound the region "visible" to the sonar. (Snipped is from an animation showing the vehicle, inset in the figure, traversing through the sonar path).

The voxels have equal size edges which are determined by the minimum value in *x*, *y*, *z* necessary to fully sample the spatial spectrum. It is easy to see that the data sizes can get very large very quickly using this sampling scheme. For example, a 100 x 100 x 10 meter volume with voxels 2cm on edge storing data as complex single values (i.e. 8 bytes per voxel) would require 100 Gigabytes of memory. (In contrast, 5cm voxels would only require 6.5 Gigabytes, hence the ubiquitous usage of under-sampling). To minimize the number of interpolations necessary for the final projection to the Cartesian-coordinate image domain, an indexing scheme is used in conjunction with pixel size Δx_{voxel}, Δy_{voxel}, $\Delta z_{voxel} = \Delta d_{voxel}$. The image origin is selected as the minimum value in *x*, *y*, *z* spanned by the bounding volume, and the origin of the target-coordinates for the spherical-to-Cartesian transformation of the final sub-aperture is determined via:

$$\boldsymbol{x^{SA\,final} = \Delta d_{voxel}\left\lceil (x^{SA\,stage\,N} - x^{origin})/\Delta d_{voxel}\right\rfloor,} \quad \textbf{Eq. 46}$$

$$\boldsymbol{y^{SA\,final} = \Delta d_{voxel}\left\lceil (y^{SA\,stage\,N} - y^{origin})/\Delta d_{voxel}\right\rfloor,} \quad \textbf{Eq. 47}$$

$$\boldsymbol{z^{SA\,final} = \Delta d_{voxel}\left\lceil (z^{SA\,stage\,N} - z^{origin})/\Delta d_{voxel}\right\rfloor,} \quad \textbf{Eq. 48}$$

where x^{origin}, y^{origin}, z^{origin} describe the origin of the image volume in earth-coordinates. As can be seen from (46 – 48), the origin of the target coordinate system for the final projection operation is simply the center of the voxel nearest to the origin of the final stage sub-aperture. A smaller Cartesian volume, encompassing only the grid coordinates of the final stage sub-aperture and with the origin defined in (46 – 48) and voxel edge size Δd_{voxel} can be used as the target grid for the final projection and merely indexed and summed into the total image to prevent excessive interpolations. The steps are as follows:

1) Determine size of Cartesian target image dimensions in voxels (perhaps fixed around some average value determined by beamwidth, max range and Δd_{voxel} or perhaps dynamically calculated for each final-stage sub-aperture)
2) Determine the relative locations of the points in the final image relative to the origin of the image. This origin may not be located within the volume; for example, if the boundary of the image is determined to be near the sediment.
3) Assign earth coordinate values to the origin using (46 – 48). Note that orientations do not need to be defined because the volume is always fixed in earth coordinates.
4) Translate the origin of the Cartesian earth-coordinate target grid by the difference its origin and the origin of the final sub-aperture.
5) Perform a Cartesian-to-spherical transformation on the translated target grid to place the target grid points into the coordinate system of the final sub-aperture.
6) Perform a 3D interpolation of the final sub-aperture onto the target grid points.
7) Via indexing, sum the target grid points directly into the final 3D image cube.

As in the previous cases, the interpolation step must be accompanied by the phase multiplications from eq. (4) and (5), *provided a coherent output is desired*. A de-speckled image with reduced along-track resolution but lower speckle noise can be produced by computing the envelope via an absolute value operation applied to the data at step (6) prior to the interpolation. In this case the along-track resolution is limited by the aperture length spanned by $K*2^N$ pings. The final process, from frame generation to final 3D SAS image generation is shown in flow-chart form in Fig.41.

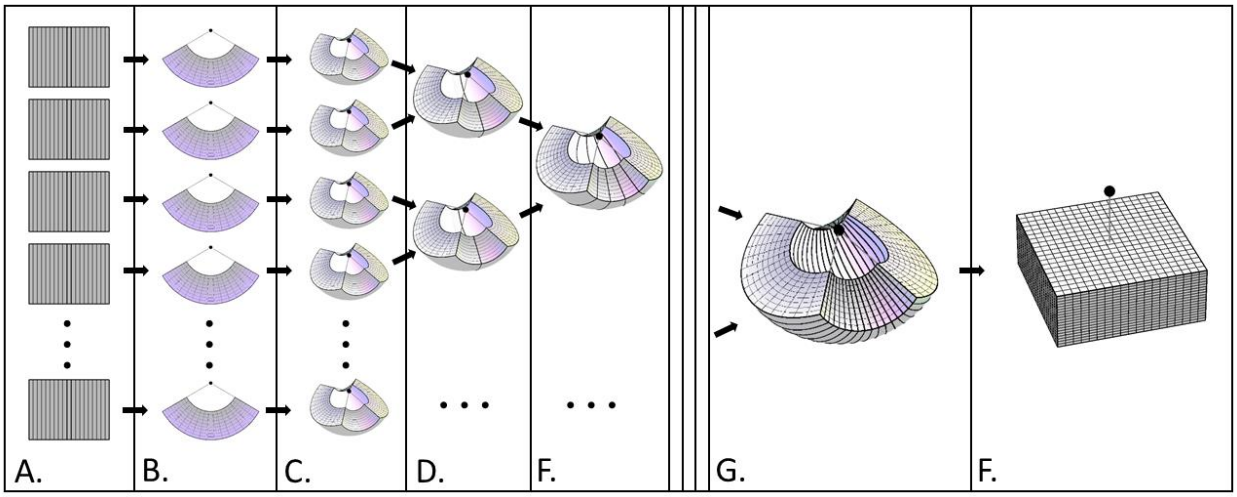

Figure 41: A flowchart illustrating the 3D FFBP algorithm. Two dimensional raw data matrices captured every ping in subplot A) are beamformed to real aperture in polar coordinates *B*). The axi-symmetry correction $r \rightarrow \tilde{r}$ allows these frames to be represented across θ without multiple back projection operations C). Adjacent frames are combined in a common spherical coordinate frame to form the first sub-aperture stage D), equivalent to column *A* in Fig. 36. As in Fig. 36, this process is iterated until the specified set of pings (e.g. 256) forming a full chunk are combined (G). A three-dimensional spherical-to-Cartesian interpolation is performed to bring the chunk into the Cartesian image domain, where it can be combined with other chunks.

Acoustic camera images and additional data products

In addition to contributing to the full synthetic aperture image, the final high resolution spherical coordinate image (last stage of the sub-aperture combination process) is a useful data product in its own right. The volume visualized is much longer in the along-track dimension than the length of aperture synthesized, and thus when replayed as a movie many objects on the seafloor and in the sediment volume are visible multiple times, and from multiple different angles in these final-stage spherical coordinate images. Often, certain targets have a much higher target-to-clutter ratio at certain angles than others, making them very bright against the background in particular spherical-coordinate sub-apertures. Related to this, one of the chief difficulties accompanying interaction with 3D data is visualization, and in conjunction with some relatively simple volume visualization techniques, a set of data-products can be generated from the high-resolution spherical coordinate representation of SAS data that are helpful for detecting objects and understanding the local environment.

Regarding three-dimensional data visualization: this is not a straight forward task because standard data visualization hardware (e.g. monitors, laptop screens, projectors) project only two dimensional images. Three-dimensional data must therefore be projected to a two-dimensional representation prior to viewing, a process that involves information loss. A variety of volume rendering methods have been developed in the image-processing community, some of which are accompanied by significant computational burden (e.g. ray-casting). By far the most common methods for visualizing three-dimensional acoustic data, (verifiable via a literature review, e.g. [1-5]), are “Maximum Intensity Projection” (MIP) and data slicing, the latter of which is simply visualizing a 2D slice of the 3D volume. In MIP rendering the maximum intensity value is computed along one axis of the three-dimensional volume forming a 2D array containing the

brightest pixels along the MIP axis. This removes a large amount of spatial information and therefore MIP images are often presented in sets of three: one MIP through coordinate axis, with the goal being to prevent bright voxels from obscuring other objects of significance.

The last stage of the 3D FFBP sub-aperture combination process – the high resolution spherical coordinate representation – lends itself to MIP representations readily, and in a way that is helpful to operators, data analysts, and potentially automated target recognition algorithms as well. The axes of the final stage sub-apertures are range r (meters), θ (along-track / SAS angle), and φ (across-track / real aperture angle). Taking maximum intensity projections along each of these dimensions creates three separate intensity projections, each of which are useful for certain types of analysis.

The maximum intensity projection through *range* creates a 2D image with vertical and horizontal axes being the angles θ and φ. This is very similar to images created from an optical camera, which have azimuth and elevation angles as axes but lack range information. Because of this similarity, these images are called "acoustic camera" images. Examples of these images are shown below for the Multi-Sensor Towbody platform.

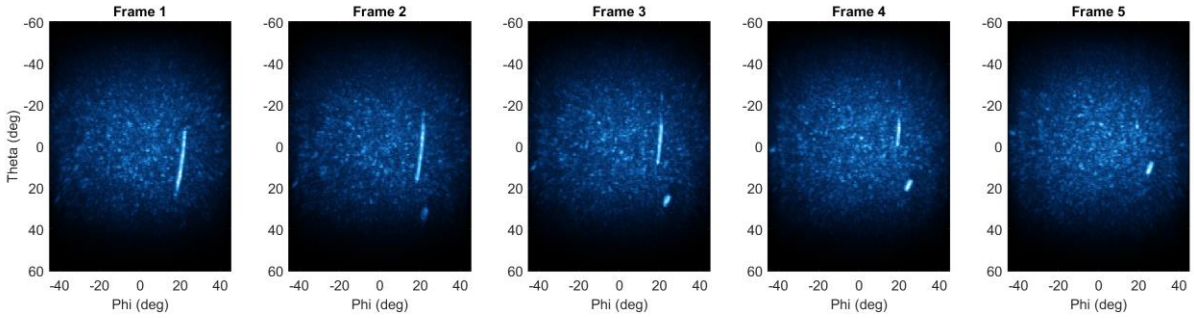

Figure 42: A series of *acoustic camera* frames from the MuST system as it travels over a pipe and solid aluminum cylinder. The vertical axis *Theta* is angle in the synthetic aperture dimension, and the horizontal axis *Phi* is angle in the real-aperture dimension.

There are several advantages that acoustic-camera images have vs. Z-dimension MIPs through images in the Cartesian domain: they are faster to compute, the point-spread function is more uniform throughout the image, and shape features (vs elastic) features are going to be the dominant features visible in targets, since range is effectively "in" to the screen. Furthermore, there is no ambiguity about the synthetic and real-aperture dimensions. Lastly, it is simpler to assign back-scattering strength and signal-to-clutter estimates to targets in the acoustic-camera domain than in Cartesian coordinates, because the angle to the target can be directly determined from the pixel bin. The acoustic camera view is essentially what the sonar sees over the number of pings used to generate the image.

Taking a MIP through the θ dimension produces a 2D image with axes r (range) and φ (real-aperture angle). This image produces a projection of the data that has low sidelobe pollution (due having the real-aperture angle as one of the image dimensions), low image distortion from any SAS processing defects, and targets that are passing through the image tend to stay in roughly the same spot (with a shift primarily in range) throughout apertures over which they are visible. Due to range being an axis, elastic features of targets can be visible in this domain. Because it is

the physical aperture, this projection is the least useful for performing shape analysis but is the most reliable aperture from a beamforming standpoint. It is also potentially the most useful for determining the relationship between angle of incidence and scattering strength from the sediment. An example is shown in the following figure:

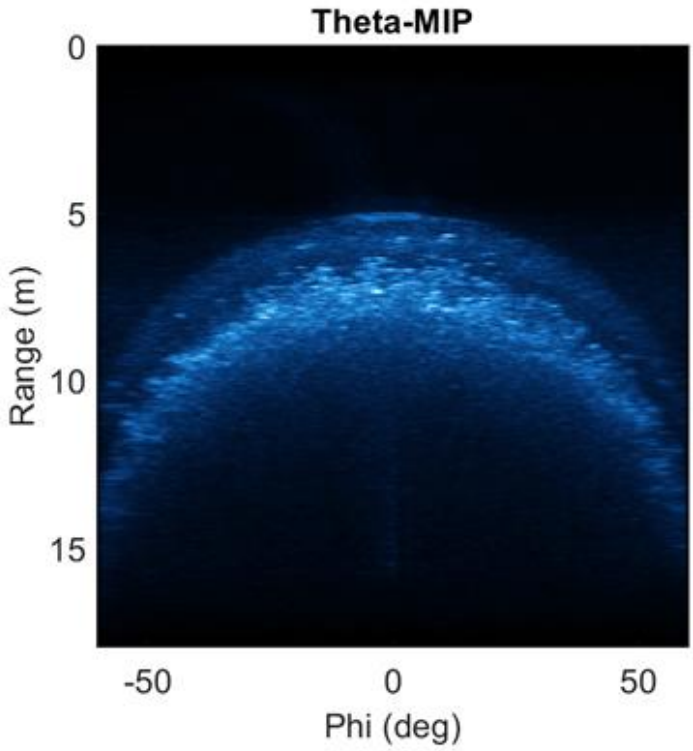

Figure 43: An example of a MIP through the θ dimension of a final-stage spherical-coordinate sub-aperture. This looks along the synthetic aperture and visualizes the scattering strength vs. angle of incidence for the different sediment layers.

Lastly, taking a MIP through the real aperture angle φ creates an image with axes range r and synthetic aperture angle θ. This is the SAS sediment profiling view. It has the best combination of structural features and elastic features of the various spherical coordinate MIPS because it combines range with the high focus of the SAS processing. If the data is under-sampled then sidelobes also tend to be most visible in these images. When imaging in regions with a mud layer, this projection is the most useful for detecting targets because targets that may be obscured by the bright scattering from a layer below a target are no longer obscuring the target in this view. Examples are shown below:

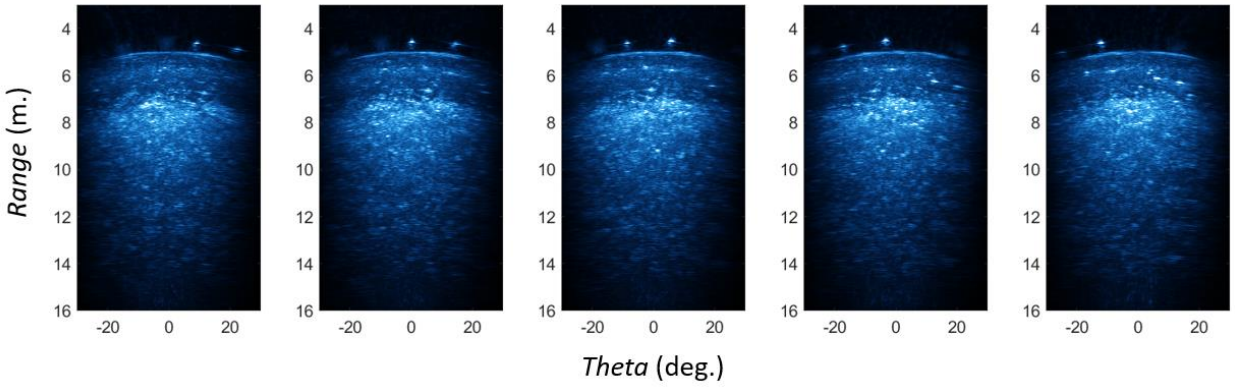

Figure 44: An example of a set of MIP's through the φ (real-aperture) dimension, captured frame-by-frame showing a set of floats moving through the field of view of the sonar. Some details of the sediment layering can also be observed.

Since these intermediary data-products are a byproduct of 3D FFBP programing they have very little overhead associated with their creation provided no overlap between frames is desired. In fact, overlap is very desirable for interpreting the shifting positions of objects moving through the acoustic-camera image, especially since some features are strongly anisotropic and may completely disappear in an adjacent frame. The next section outlines an efficient frame updating method that reduces the computations associated beamforming a partial shift of the synthetic aperture sub-aperture, allowing for smoother acoustic camera image updates.

Efficient sub-aperture shifting

Partial overlap of sub-apertures allows anisotropic scattering features to be visible over multiple apertures, aiding visual tracking of objects passing through acoustic-camera frames. If each shifted sub-aperture is processed independently, however, this could be accompanied by a significant computational penalty: e.g. a 50% overlap results in a 2X computational burden increase for any given scan. Alternatively, by recognizing that each sub-aperture image is generated through an iterative combination of lower-resolution sub-aperture pairs, overlapping sub-apertures can be generated more efficiently by recombining pre-beamformed low-resolution images generated during previous stages of the FFBP algorithm. The shifting penalty becomes the computational time associated with beamforming only the added data and combining it with the other pre-computed aperture stages. The efficiency gain can be computed as follows. Assume for a set of 2^N aperture points in the along-track dimension that the computation time can be expressed:

$$\boldsymbol{T_{full} = CN2^N,} \qquad \textbf{Eq. 49}$$

where C is a constant scaling with the number of pixels in each 2D frame composing the along-track set of real-apertures. For an along-track shift of 2^M aperture points, an entirely new subaperture must be generated from the added components, requiring:

$$\boldsymbol{T_{sub} = CM2^M.} \qquad \textbf{Eq. 50}$$

When the new partial aperture has been beamformed, (*N*-*M*) combinations must be performed with previously beamformed partial sub-apertures. The combined apertures double in size each time resulting in the total:

$$\boldsymbol{T_{shift} = C\left(M2^M + 2(2^N - 2^M)\right).} \qquad \textbf{Eq. 51}$$

The ratio, $Q = T_{full}/T_{shift}$ is the relative speed gain attained via combining the new sub-aperture with the pre-computed sub-apertures vs. re-computing the shifted aperture entirely from scratch:

$$\boldsymbol{Q = \frac{N2^{N-1}}{M2^{M-1}+(2^N-2^M)}.} \qquad \textbf{Eq. 52}$$

Figure 45 plots the speed gain factor *Q* vs. shift length in samples for an aperture size of 256 samples, along with a "fractional efficiency" equivalent to the product of the speed-up factor vs. the faction of the aperture shifted:

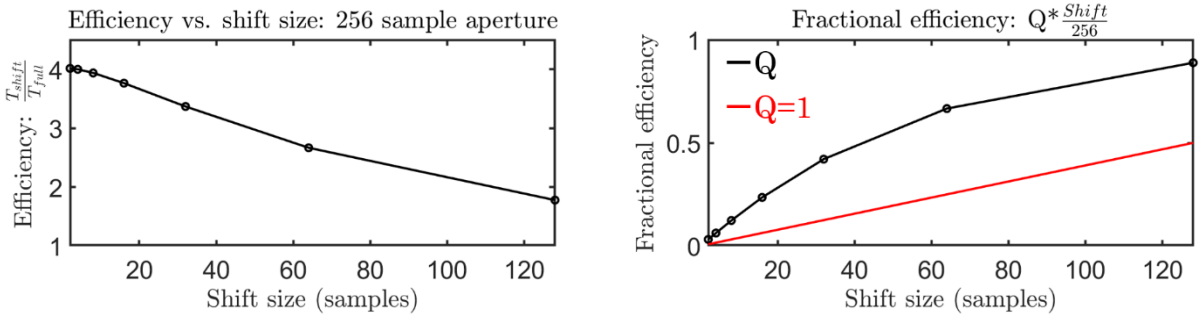

Figure 45: Efficiency vs. shift size for a 256 sample sub-aperture showing the relative speedup vs. complete re-computation using the stage combination method described previously (left), and the fractional efficiency, showing product of the speed gain and the fractional aperture shift (right). The red line is *Q*=1, the full-aperture reconstruction case.

It is seen that the efficiency of the sub-aperture recombination method improves as the percentage overlap increases, however the fractional efficiency decreases. A fractional efficiency equivalent to unity would imply no computational overhead for a fractional shift; e.g. a half aperture shift can be computed at twice the speed of a full aperture. This is clearly not the case, however it is close to the case for a half aperture shift, and though it decreases it is still considerably better than the case in which the full aperture is recomputed every shift. By inspection it can also be seen that (52) approaches a fractional efficiency of 1 for the case $M = N-1$ as N grows. This method of aperture shifting, with an overlap of 75%, was used to create the frames in Fig.'s 42 and 44. If partial overlap was not used, only the figures at the end of the series would be displayed in the acoustic-camera frame update.

Beamforming summary

The 3D FFBP algorithm has been presented in this section, along with some secondary data products that are potentially useful byproducts of this beamforming process. The algorithm proved extremely successful for real data, and results from various field tests and varying systems will be shown in the *Results* section of this document.

IV. The MuST Toolbox

Overview

The MuST Toolbox is a software library containing implementations of the beamforming, post-processing and data-visualization tools that have been developed as part of this project. Components of the toolbox span functionalities ranging from real-time SAS beamforming and data visualization for use during field operations, to data-driven mosaicking for multi-aspect data synthesis and structural acoustic feature enhancement for detected regions-of-interest. The toolbox makes heavy use of graphical user interfaces ("GUI's") to simplify the usage of the various tools associated with the eBOSS. Over the course of the project and deployment of the eBOSS on many field operations a functional and efficient data workflow evolved and a flowchart describing this workflow is shown in Fig. 46.

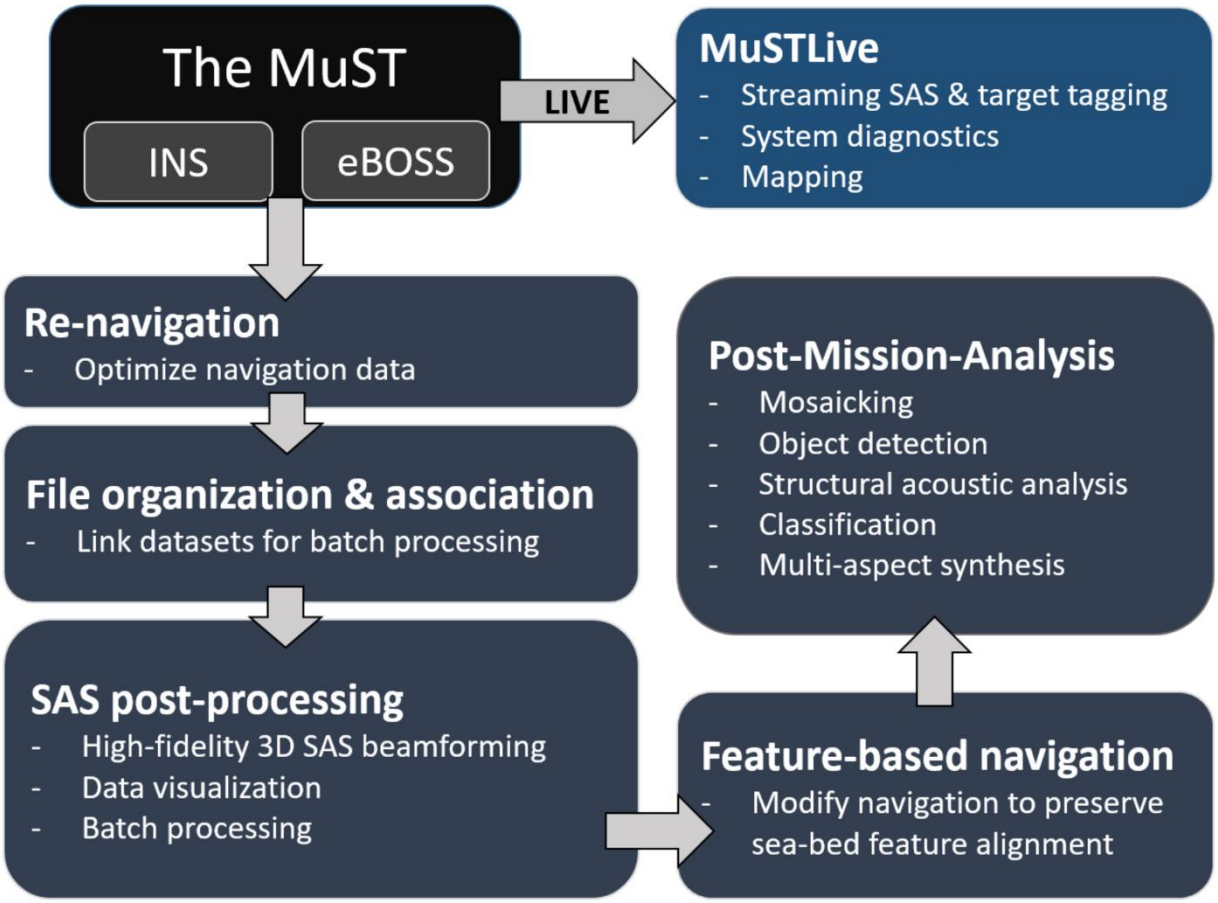

Figure 46: A flowchart showing the MuST eBOSS data workflow used for live and post-mission analysis.

Each of the blocks in the flowchart correspond to a specific tool in the MuST software library. MuSTLive, the real-time data processing software, enables the data being captured by the system to be viewed as it is being captured as well as providing system diagnostics, the ability to determine the location of the tow-body relative to previously encountered targets, and to tag targets observed during operation. This tool has proven indispensable during field operations for performance monitoring and as a decision-making aid.

Following data-collection, typically the first thing performed is a “re-navigation” of inertial navigation system data. SAS processing and target localization are both heavily dependent on accurate navigation data. The navigation filters incorporated into the inertial navigation system are tuned to deliver an optimal real-time estimate for the towbody location, but an improved solution both for SAS processing and target localization can be determined in post-processing by allowing navigation uncertainties to propagate both forward and backward in time. eBOSS and navigation files acquired during the field operation are linked in a software organization tool provided by the MuST software library, and in combination with the renavigation solution, are batch-processed to generate SAS data suitable for feature detection and mosaicking. A detector (Dave Williams, SERDP MR21-3543) has also been integrated into the navigation stage.

Following batch-processing, common features between scans are auto-detected, and a new “feature-driven” navigation solution is generated that minimizes the displacements between observations of common features, while being constrained to follow certain behaviors consistent with vehicle dynamics.

Finally, a combined mosaic is generated from the relevant datasets contributing to a region and displayed in a post-mission analysis tool which provides the end user with detection, classification, region-of-interest processing, data visualization, and data exportation capabilities. The next several sections will describe in more detail the capabilities and methods of operation behind each of these different tools.

MuSTLive

MuSTLive is the realtime SAS processing, data visualization, target-tagging and diagnostics tool developed to work with BOSS systems. It was designed in conjunction with the development of the MuST system but it is not limited to working with systems that have the array geometry used on the MuST. It provides the user with a number of SAS beamforming, data visualization and mapping options. Figures 47 – 49 show the settings that can be defined by the user in the MuSTLive graphical-user interface.

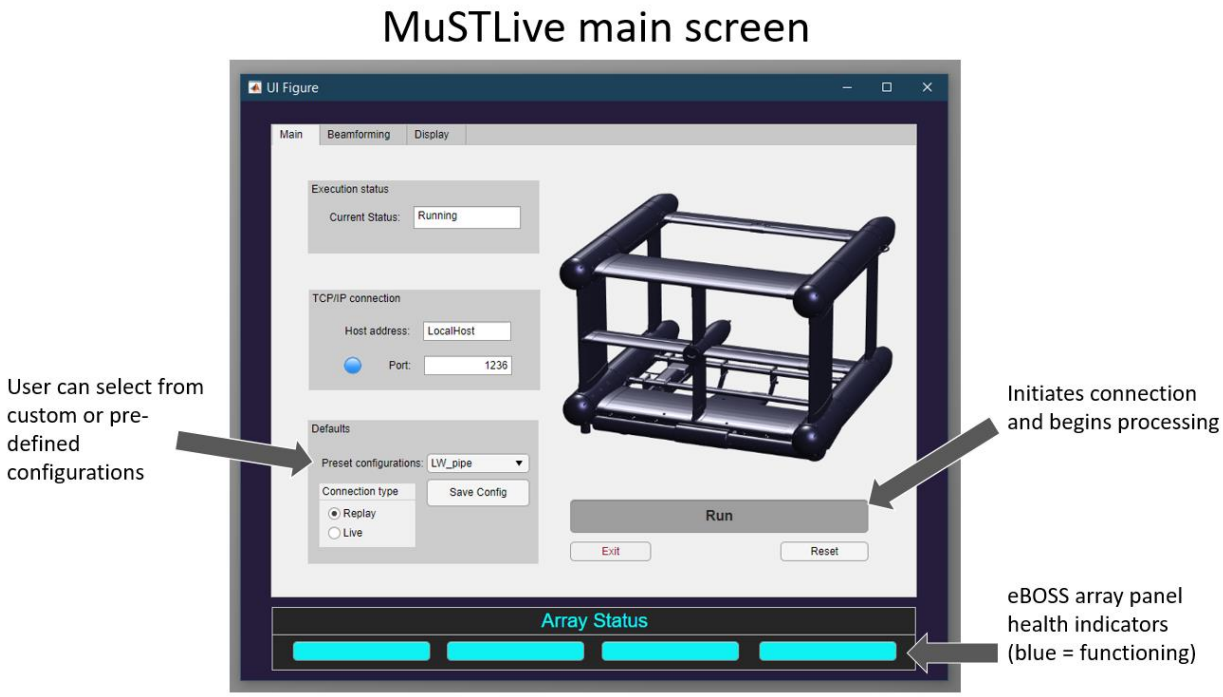

Figure 47: The main screen of MuSTLive.

The MuSTLive main screen, shown in Fig. 47, displays the system status, allows the user to quickly select settings, determine connection type, and initialize communications with the towbody and begin data visualization. Via the annotated dropdown menu ("Preset configurations"), users can select from a number of pre-defined configuration options or create their own. These configuration files populate the beamforming parameters and visualization options that are configurable on the other tabs of the GUI.

On the main screen users can also determine the source of the data stream: data can be replayed from a saved eBOSS file by selecting "replay," or piped from an active connection with a towbody by selecting "Live." In both cases, the tool is working alongside EdgeTech's *BOSScontroller.exe* program; in the replay case BOSScontroller is broadcasting a user-selected file over a localhost connection, and in the live case BOSScontroller is used to set the actual operation values of the BOSS system.

Hitting "Run" on this main panel establishes the connection and opens up the data visualization windows. Hitting "reset" closes all active data visualization windows and connections and re-initiates a connection with the new user-defined settings.

Lastly, *execution status* is a verification that a connection has been successfully made with the data stream, and the four blue lights along the bottom turn red if an array doesn't get properly initialized or is malfunctioning. These lights are visible regardless of which screen / panel is selected so that the device health can be monitored regardless of which panel is being displayed.

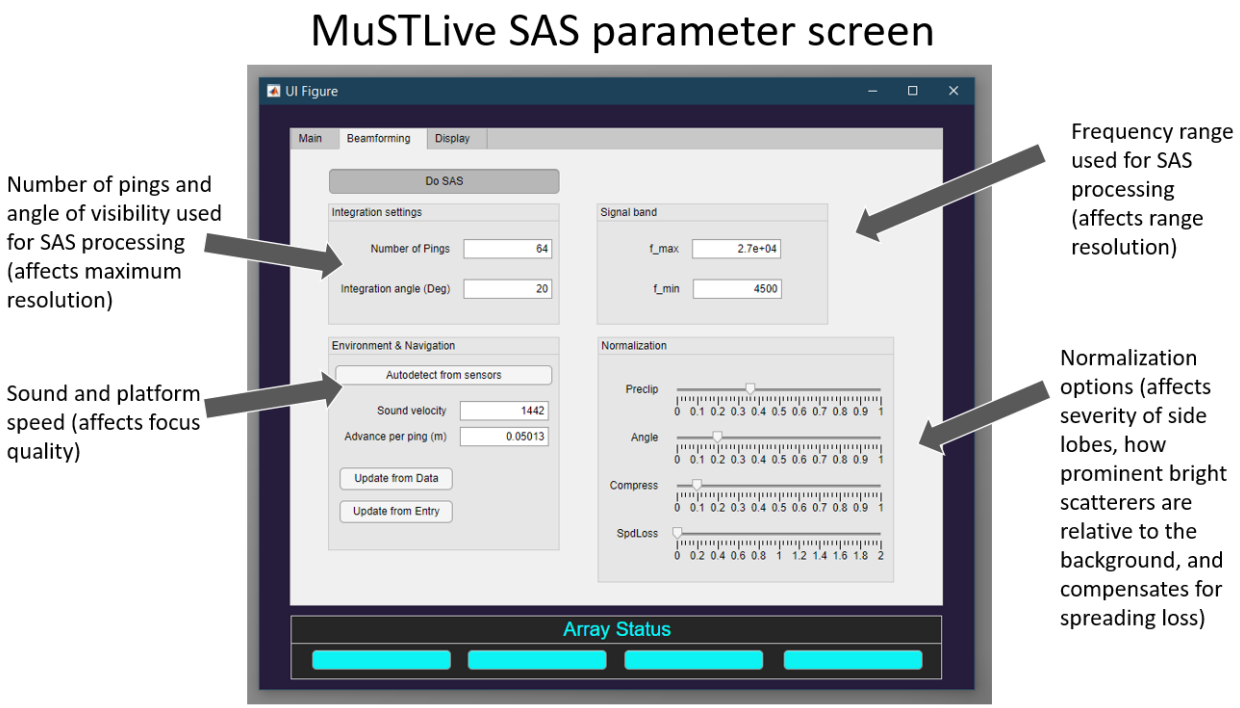

Figure 48: SAS parameters screen. This screen enables the user to define the length of the synthetic aperture used for beamforming (in number of pings), along with various options that can affect focus quality and image contrast.

The SAS parameters screen enables the user to define how the data is processed to create SAS images. Of primary significance is the number of pings and integration angle which, in the example, is set to 64 pings and 20 degrees (the latter is the angle from nadir of the SAS integration beam; the total beamwidth is 2X the value displayed, or 40 degrees in the present case) in the current configuration. The beamforming is GPU accelerated, and using a laptop with an integrated RTX 2070 Max-Q NVIDIA graphics card, array lengths of up to 192 pings can be supported. The integration angle setting determines the extent of a shading window applied to the beamforming matrix in the SAS beamforming engine and does not affect processing time. Importantly, beyond certain ranges (determined by the physical length of the synthetic aperture) the SAS integration angle is limited by the length of the aperture and not by the shading function.

The MuST platform, a MacArtney FOCUS III, has proven to be a very stable and, especially for lengths on the order of 128 pings or less, frequently delivers nearly diffraction limited quality imagery provided the platform is not being actively maneuvered and the sound velocity and advance-per-ping (i.e. the array sample spacing) are accurately specified in the beamformer. The user can input these values manually, or auto-detect these values from the onboard sensor. Pressing “Update from Data” will re-define the beamformer parameters based on the current values detected by auxiliary sensors onboard the MuST.

Users can also select the minimum and maximum values of the frequency bandwidth used for SAS processing and can modify normalization settings that affect beamforming and data visualization. A common difficulty with SAS systems is that coherent features can have signal levels many decibels above neighboring features that may potentially be of interest. Sometimes

the user may care to only see these bright features; other times the user may want to be able to see many smaller details of the sediment. These sliders give the user a great amount of control over how the data is compressed to accentuate features of interest. The sliders have the following functionality:

- High pre-clip values reduce the impact of sidelobes from bright features on imagery and prevent sidelobes from overshadowing features of less intensity - at the cost of total signal dynamic range.
- The “angle” slider increases the intensity of off-nadir scatterers
- “Compress” reduces the dynamic range of an image by applying the compression function $x./(x + \sigma)$, where the slider controls the value of σ. Low values of σ result in strong dynamic range compression (without modifying the relative levels of sidelobes as “pre-clip” does), and high values of σ reduce compression.

Lastly, the user can turn off SAS processing altogether by unselecting the “Do SAS” button on the top left; in this case all streamed imagery will be real-aperture data.

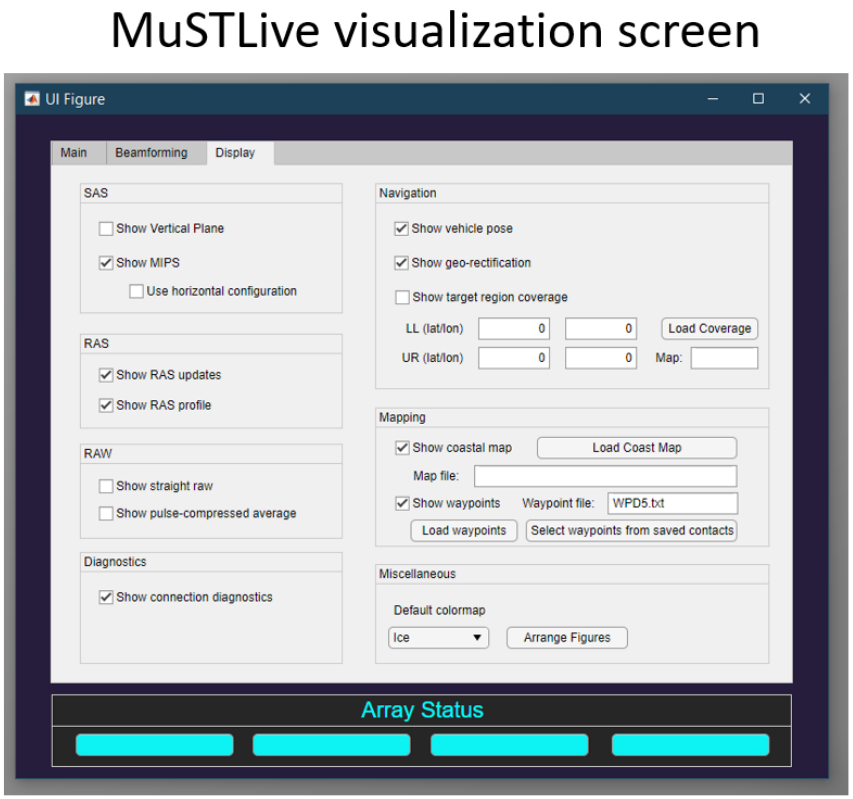

Figure 49: MuSTLive’s data visualization options screen. A large variety of options are available to the user.

The third panel of the MuSTLive GUI, shown in Figure 49, allows the user to select what sources of information are displayed and select user-defined features of the data visualizations. Beginning with the *SAS* sub-panel on the top left, the user has the option of specifying which projections are displayed through the beamformed data. Figure 50 shows the two projections through the 3D volume that are visualized if the “Show MIPS” box is checked.

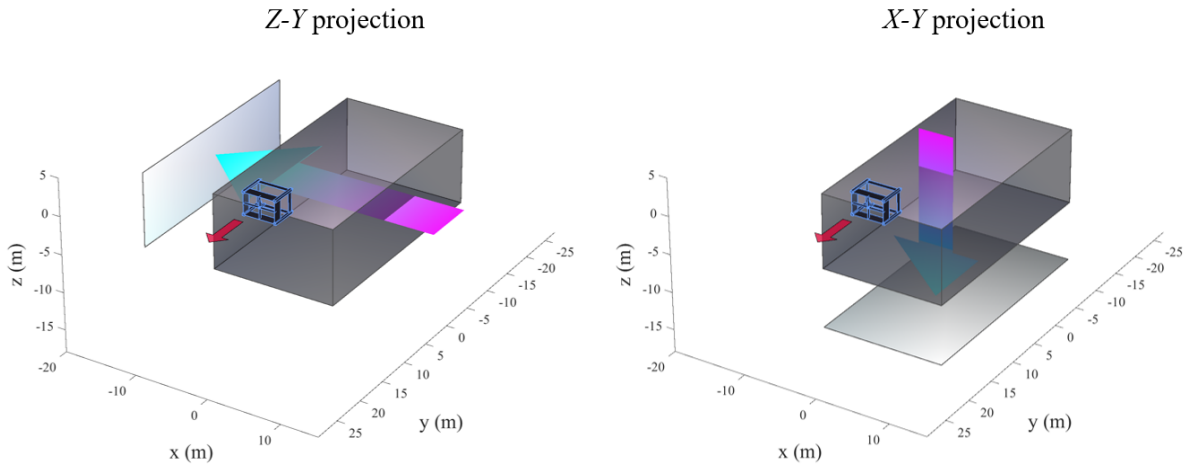

Figure 50: The two MIP projections that can be visualized in MuST Live: a projection of the SAS 3D image (represented by the translucent gray cuboid) along the X-axis onto the Z-Y plane, and a projection along the Z-axis onto the X-Y plane. The Z-X projection was not found to be as operationally useful as these two. Note the items are not to scale (the actual dimensions of the MuST are smaller, and actual dimensions of the SAS volume are larger).

Figure 51 is an example of the data visualization window that is created if the "show MIPS" option is selected. The top window is the Z-Y projection and the bottom window is the X-Y projection.

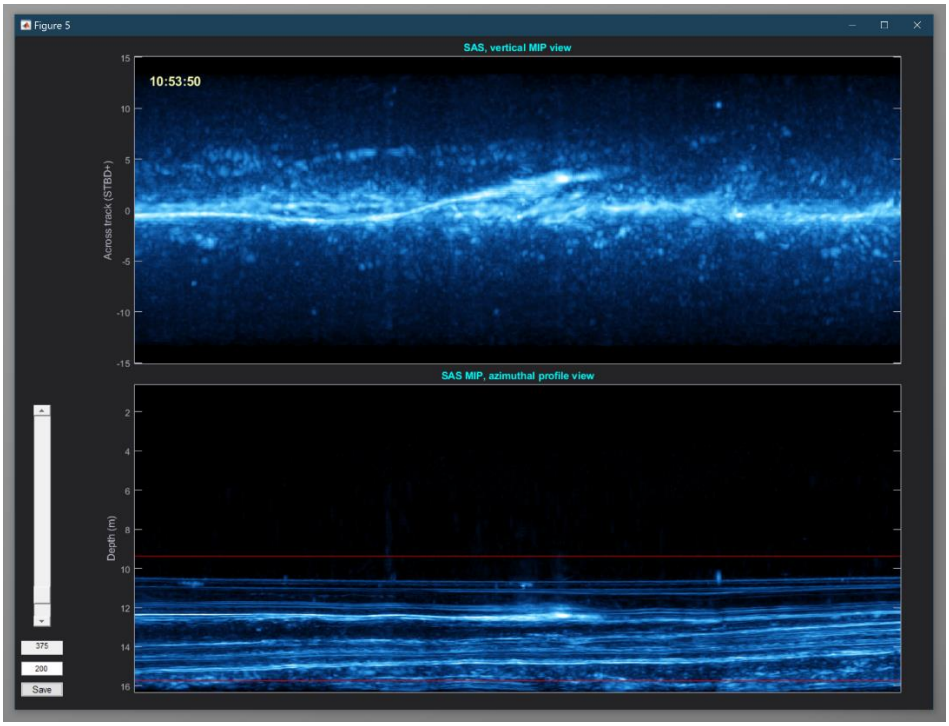

Figure 51: The SAS MIPs data visualization window. The top window represents the Z-projection of the 3D volume onto the X-Y plane, *limited to the region bounded by the red lines shown in the Z-Y projection on the bottom*. This region shows a random segment of sediment and sub-sediment layering from Lake Washington.

(Convention note: if two coordinates are listed in reference to a MIP, these specify the directions of the MIP image axes. If *one* coordinate is listed, that represents the axis through which the MIP was projected, e.g. a "Z-MIP" creates an X-Y MIP projection).

A significant feature of the data visualizations shown in Fig. 51 are the two horizontal lines shown in the bottom MIP (the X-Y projection). These red-lines represent the Z-span, or depth span, over which the Z-Y projection in the top plot is calculated. This limiting of the span over which the MIP is calculated is very important due to the layered nature of the sediment. It was frequently found that features can be masked by other features in brighter layers, and this gating tool allows the user to specify exactly where to look for features. It can be widened to span the whole volume or narrowed to show contents of a specific depth range. The user can specify the span of the region (the spacing of the red lines, in pixels) using the numerical entry box on the bottom left, and the slider positions the vertical position of the center-point between the lines.

The beamformed SAS data scrolls from left to right through these images as data is acquired and beamformed. If a feature of interest appears in the screen the user can click the "save" button (lower left), and the figure will be exported as a .png file, along with an accompanying matlab file containing the current navigation and time-stamp metadata associated with the aperture at the time of capture.

Lastly, the default relative orientation of the two MIP's is vertical stacking, as shown in Fig. 51, however checking the "use horizontal configuration" box in the MuSTLive visualization selection panel underneath "Show MIPs" will place the two figures adjacent to each other horizontally, rather than vertically.

The other SAS data visualization option visible in Fig. 49 is "show vertical plane." This is a slice of the image volume (vs. MIP projection), in the X-Z plane. In general, it was found that X-Z projections of any sort were not very useful for finding targets of interest and didn't add significantly to the information already shown in the MIP projections. This option, therefore, was rarely used during field operations, but due to the possibility of special cases arising in which this projection becomes more useful its function was retained.

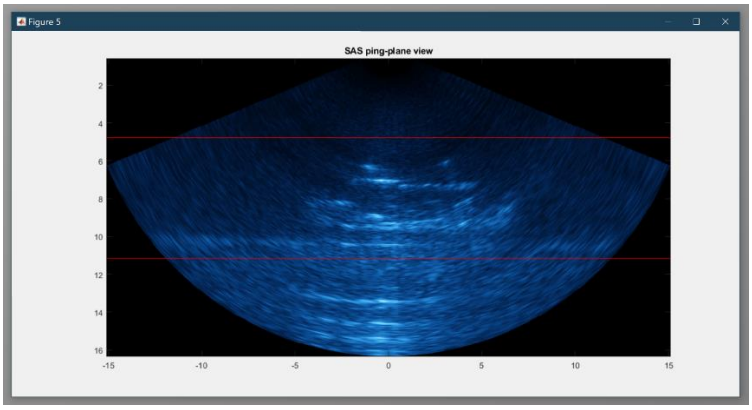

Figure 52: An X-Z slice through a fishing boat scanned by the MuST. The field-of-view of the sonar (+/- 60 degrees across-track) and range (16 meters, limited by the 40 Hz pulse repetition frequency) is clearly visible in this window. The parallel red lines represent the vertical boundaries of the range for which a MIP is computed and projected in the X-Y plane in the MIPs view.

The next visualization options provided by MuSTLive exploit real-aperture sonar (RAS) processing. Checking the “Show RAS updates” and “Show RAS profile” options bring up the data displays shown in Fig. 53.

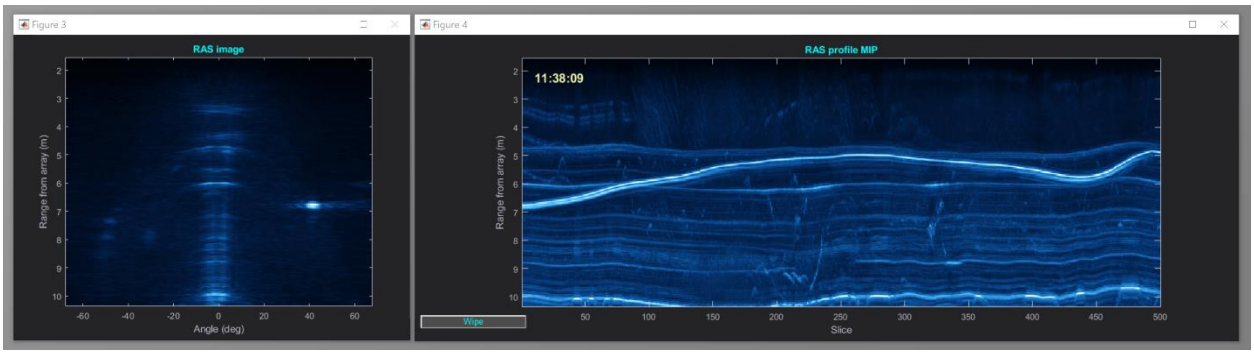

Figure 53: RAS data visualizations: the figure on the left shows the current “real aperture image,” a polar-format single-ping image displayed in range vs. angle, and the figure on the right shows a projection through angle of a ~3 minute span of RAS images, resulting in a range-vs-ping plot that is very similar to the output of a standard sub-bottom profiling sonar.

The RAS image shown on the left in Fig. 53 is a very similar data-product to the one shown in Fig. 52, but because RAS data, unlike SAS data, is unfocused in the Y- (along track) dimension, information from different along-track positions contributes to the RAS image, making it more useful than the SAS vertical slice, because discrete objects do not suddenly appear or disappear in the image.

The RAS profile, which is similar to a standard sub-bottom profiling sonar data product; i.e. it shows the level of return vs. range (or time), shows data over longer time-spans than the SAS MIP views. The example in Fig. 53 shows the unfocused profiles over a span of approximately 3 minutes.

The next set of visualization options available to the user deal with displaying raw, un-beamformed data with minimal processing. This type of data is most useful for array health checks. Data from individual sonar channels can be compared to the other channels to see if they are malfunctioning or noisy, if an array has switched off, or to observe any other abnormalities. Checking “Show straight raw” shows the positive values of the unmodified data corresponding to analog-to-digital voltage counts being output by the sonar and streamed over the connection. Checking “Show pulse-compressed average” shows the average of the base-banded, pulse-compressed envelopes with a decay constant. This allows the behavior of the channels over a limited period of time to be monitored. Examples of these windows are shown in Fig. 54.

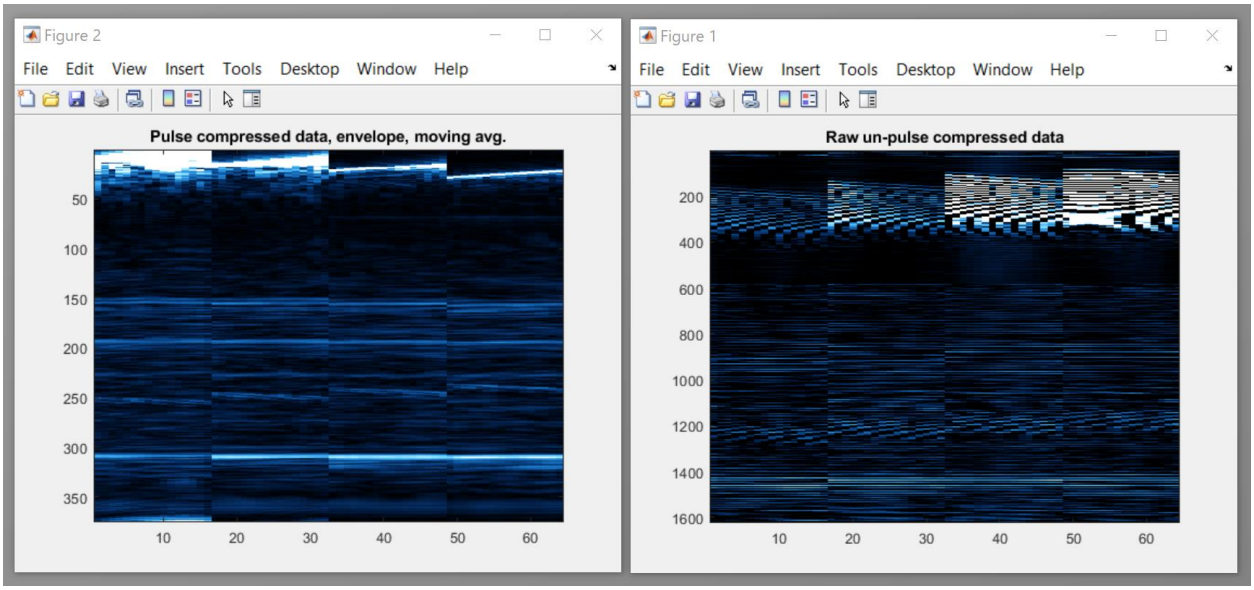

Figure 54: Moving average of pulse-compressed envelope data, shown on the left, and raw data on the right which has not been pulse-compressed. No dynamic range compression is being applied in the image so the color-scale has been pre-set to enable visualization of lower amplitude signals at the cost of clipping the direct-path signals, which occur in the first 400 samples. The vertical axes of both figures correspond to range bins, and the horizontal samples correspond to elements of the eBOSS array, which contains 64 channels total.

In addition to providing diagnostics tools for checking on array health and visualizing raw data outputs, MuSTLive also allows the processing to be monitored for connection lags which can be induced by overburdening the GPU processor. If, for example, too many pings are added to the ping-queue for SAS beamforming and the processing cannot keep up, then the TCP/IP buffer will not be read prior to a refresh containing new data from the sonar. The connection diagnostics tool monitors the difference in ping number, (each ping has a unique counter value associated with it), and if that difference rises above 1 then overburdening is indicated. The ping FPS window in the diagnostics tool monitors the frames-per-second being received from the sonar. This can be used to monitor the rate at which ping data is being acquired from the sonar. The diagnostics window is shown in Fig. 55.

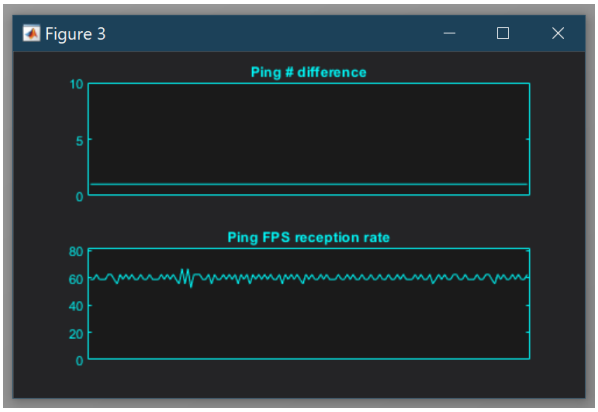

Figure 55: Connection diagnostics window showing ping difference values and frame-per-second (FPS) reception rate.

The previous user-selectable data visualization tools dealt with sonar data products and system monitoring. MuSTLive also provides mapping tools enabling the position of the towbody relative to user-defined waypoints or a coast-map to be determined, as well as vehicle orientation relative to Earth-coordinates to be monitored. Options for visualizing these types of data-products are found on the right hand side of the data-visualization panel (Fig. 49). In the "Navigation" sub-panel, checking "Show vehicle pose" enables a display of the MuST Towbody in earth coordinates. This window shows the real-time vehicle orientation in earth coordinates along with the cardinal directions and a pointer showing heading. Figure 56 shows an example of this data visualization window.

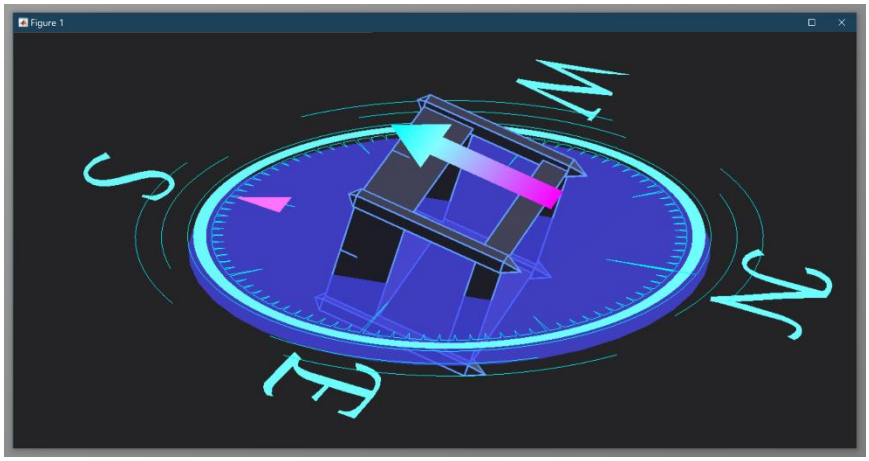

Figure 56: Vehicle pose window. The MuST towbody and cardinal directions are visible. The orientation of the towbody updates in realtime based on the feedback from the inertial navigation sensor. When the vehicle is surfaced, as in the example, the position of the waterline is shown as a blue disk.

Checking the "Show georectification" box in the MuSTLive display screen opens up the map screen. This is one of the most useful features of MuSTLive leveraged by field operations deploying the MuST. This window shows the position of the MuST vehicle relative to a coastline map (loaded by the user in a manner that will be described in shortly), historical position data showing the previous locations of the vehicle for a set span of time (~10 minutes of navigation history), and positional markers specified by an editable text-file list containing feature locations, labels, and types. The map window contains a number of useful features that help interpret data in other visualization screens:

- A map bounding box shows the size of the region visualized by the X-Y MIP in the SAS MIP visualization window. If a marker lies inside the box it should be contained within the SAS window and visible to the user.
- A vehicle location marker that shows the current position of the vehicle. The marker changes color if the beamforming matrix does not match the speed and sound-speed values indicating the imagery will be out of focus and the beamforming matrix needs to be updated.
- 10 and 15 meter range-rings around the target. The outer range-ring can be expanded using the (+) and (-) buttons on the left hand side of the window, or directly specified

in the numerical entry field. The range rings can be used to determine when the towbody needs to dive to acquire an image of a target.

- A zoom feature that allows the user to expand or zoom the map.
- An arrow indicating the current heading of the vehicle.
- A column of navigation data that shows:
    - Speed (m/s)
    - Displacement per ping (meters)
    - Sound speed (m/s)
    - Altitude above the sediment
    - Depth below the water surface
    - Heading, pitch and roll (degrees)
    - Distance submerged since last GPS acquisition (meters)
    - Offset statistics between predicted location using INS velocity and orientation data vs. current INS location estimate, which is useful for evaluating inertial navigation drift and performance.
    - A time stamp in local time.

An annotated example of this map window is shown in Fig. 57, and a zoomed out window showing coastline is shown in Fig. 58.

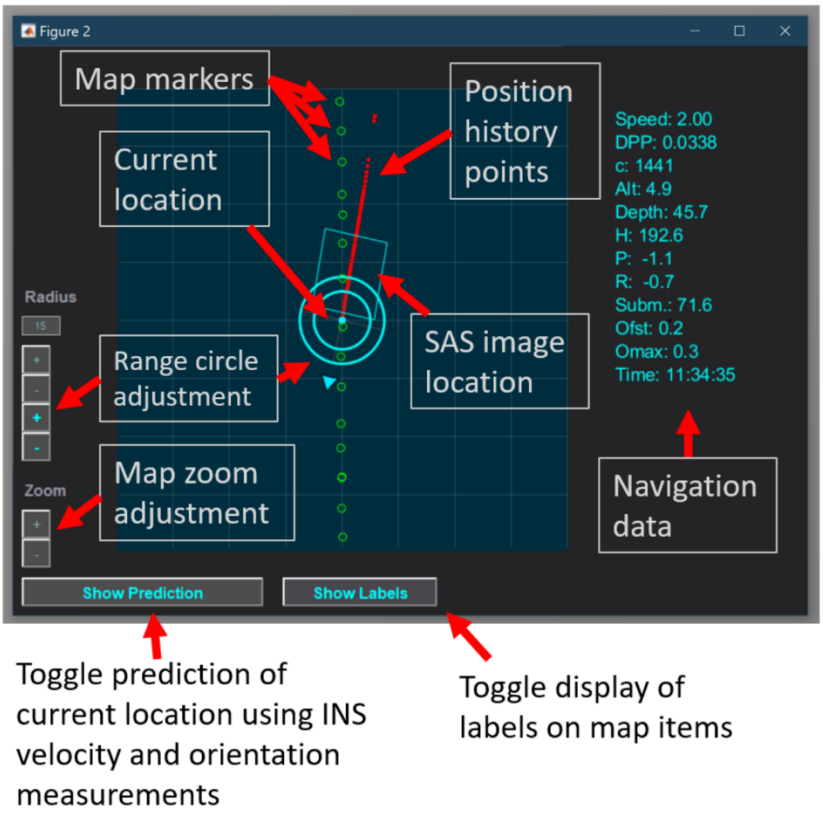

Figure 57: The navigation map window, annotated to show features.

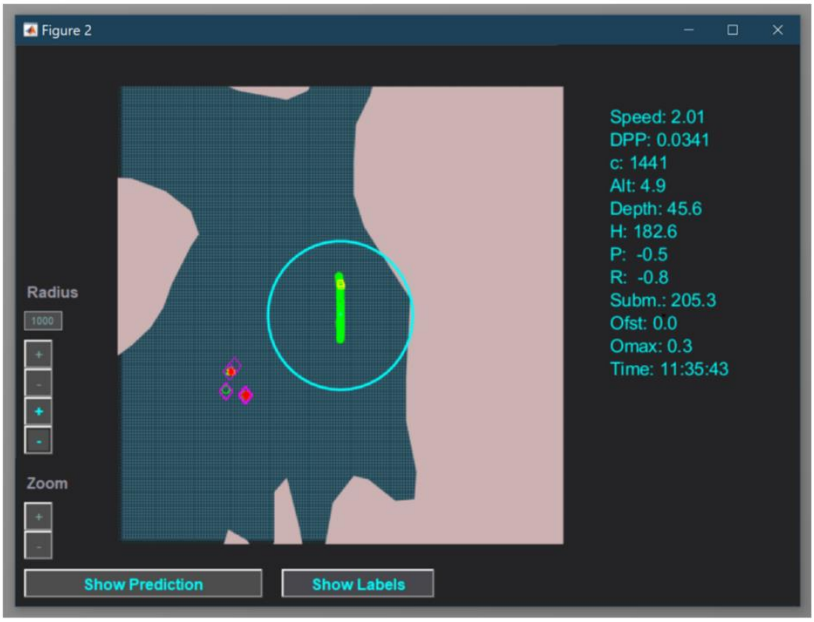

Figure 58: A zoomed out navigation map window, with the outer ring radius set to 1 km. This is the zoomed-out version of Fig. 57. The map shows the location of the Lake Washington shoreline. Map markers shown the location of shipwrecks and a cable (green line) are visible.

The markers and coastline map file, specified as a polynomial, are user-definable and are specified by the user in the "mapping" panel beneath the navigation panel in the data visualization screen of MuSTLive (Fig. 49). If the user checks the "Show Coastline" box a coastline polynomial file can be loaded by clicking "Load Coast Map." Currently maps are only available for the local Puget Sound region and the vicinity of Sequim Bay. A future expansion would be to include ESRI costal maps or NOAA charts as backgrounds, expanding and simplifying the process of incorporating publically available coastline information.

Additionally, if the user checks "Show Waypoints" in the same sub-panel and selects a waypoint text file, then a list of labelled waypoints will be added to the map. This is extremely useful for determining whether the vehicle is positioned correctly for acquiring known targets during field operations. The waypoint file is text file with comma-delineated variables containing data for:

- Latitude (decimal degrees)
- Longitude (decimal degrees)
- Marker class
- Marker label
- Time stamp of acquisition
- Time stamp of marker creation (used for identification)

A marker can currently be assigned to one of six classes by populating the marker column with one of the following letters. The class indicated by the letter and the marker drawn to the map are shown adjacent to the class letter.

- 'u': UXO (blue triangle)
- 'n': Unknown (green circle)

- ‘c’: Clutter (yellow square)
- ‘a’ Archeology (purple diamond)
- ‘b’ Boundary (orange line)
- ‘h’ Hazard (red star).

Figure 59 shows a screenshot of one example waypoint text file.

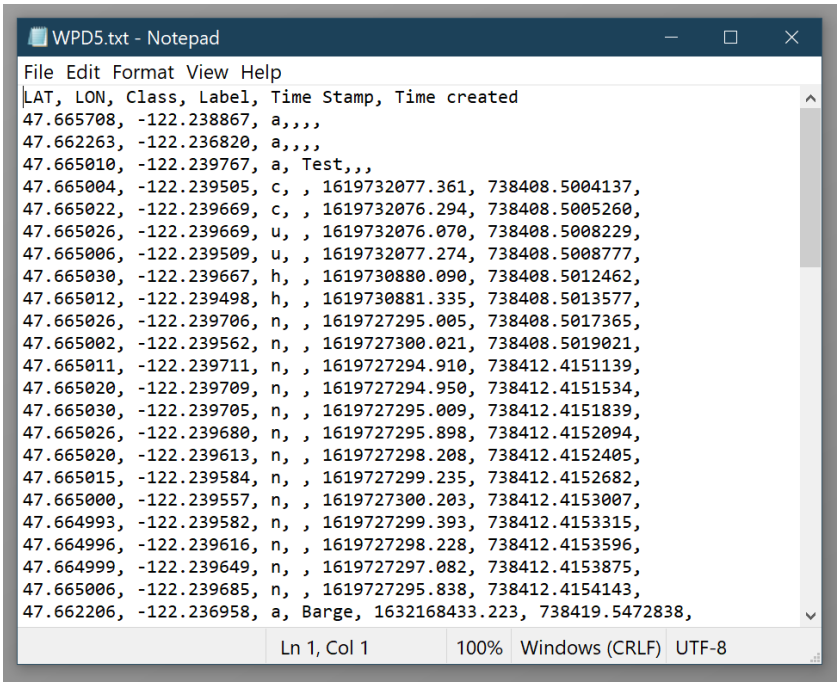

Figure 59: A waypoint text file imported into MuSTLive. This waypoint text file corresponds to the detections of objects found during a system check in Lake Washington.

The waypoint files can be modified during operation to include objects that the user has identified in the real time SAS imagery. Previously it was shown that in the SAS data visualization screen a “save” button exists that exports current imagery and corresponding navigation data. If the user opens up the data visualization panel in MuSTLive and clicks on the “Select waypoints from saved contacts” button, then a selection page will open and the user can select any of the data files that have been saved during live streaming (this selection process can be done while MuSTLive is running). Upon selection of a contact file a window like that shown in Fig. 60 will be opened.

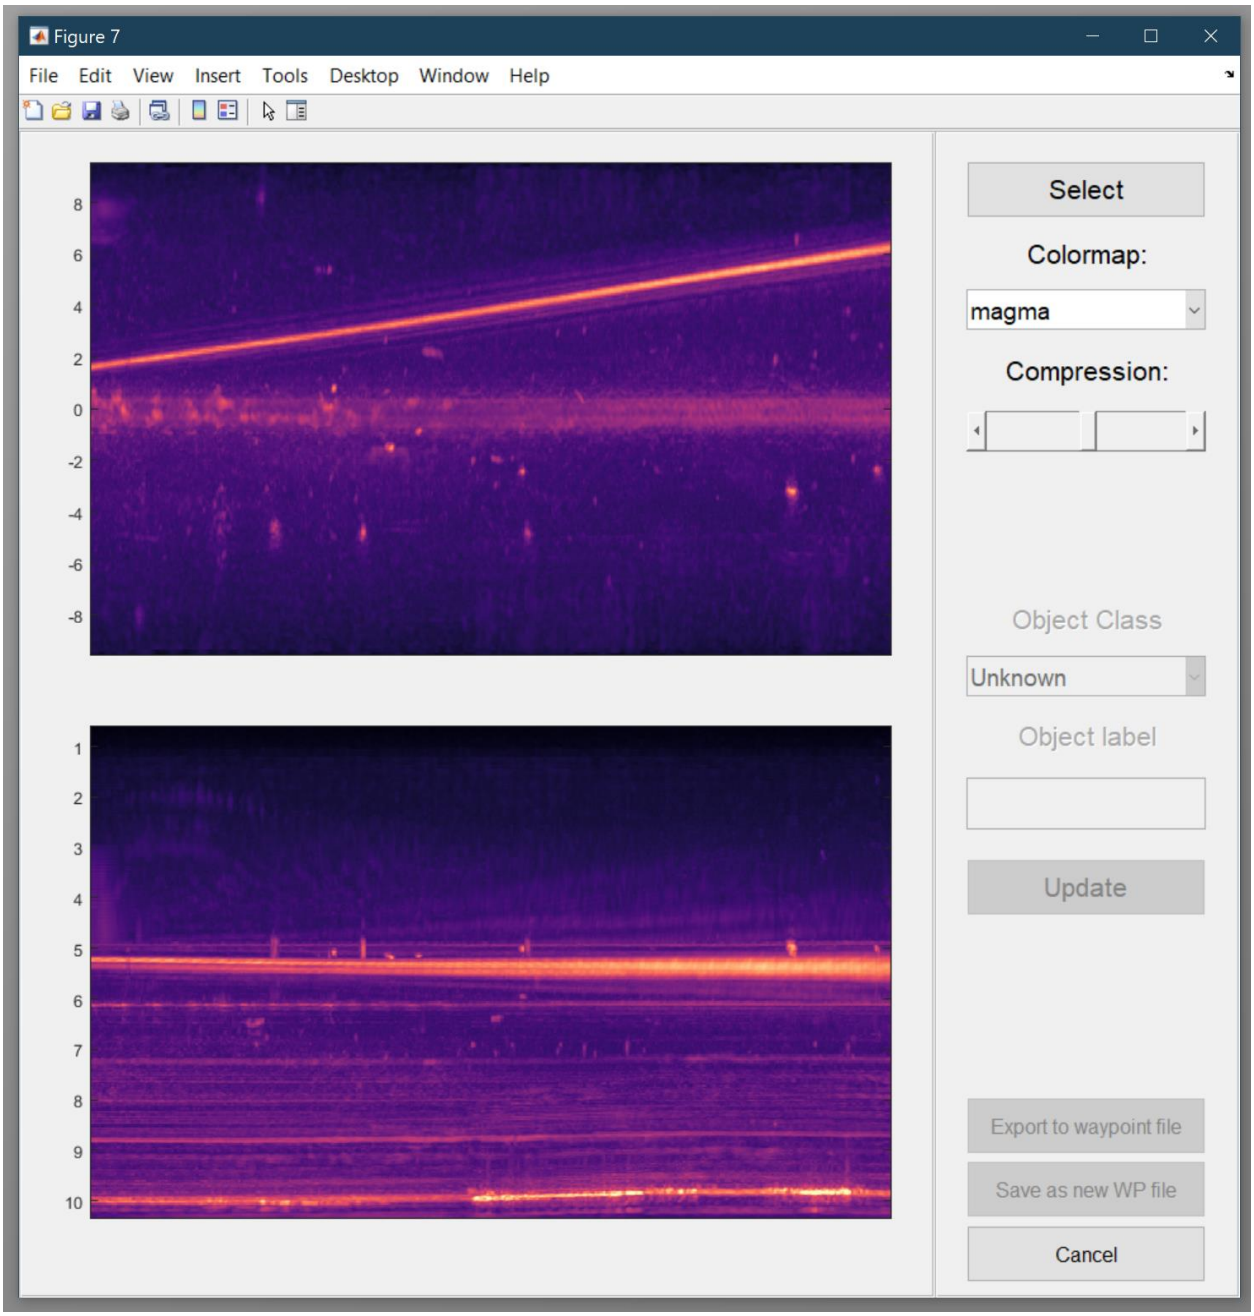

Figure 60: Waypoint update window: by clicking "Select waypoints from saved contacts" in the MuSTLive data visualization panel and selecting from a contact, this window can be opened. New detections can be added to the contact list from this page. Metadata for the contact in the form of location and local imagery is also saved as an external file that is reloadable in a visualization tool that is part of the MuST software library.

In the waypoint selection tool, shown in Fig. 60, a number of options are available for the user. Various colormaps and dynamic range compression are available to make targets of interest more visible, and hitting the "Select" button will enable the user to select a location in the top (X-Y projection) figure.

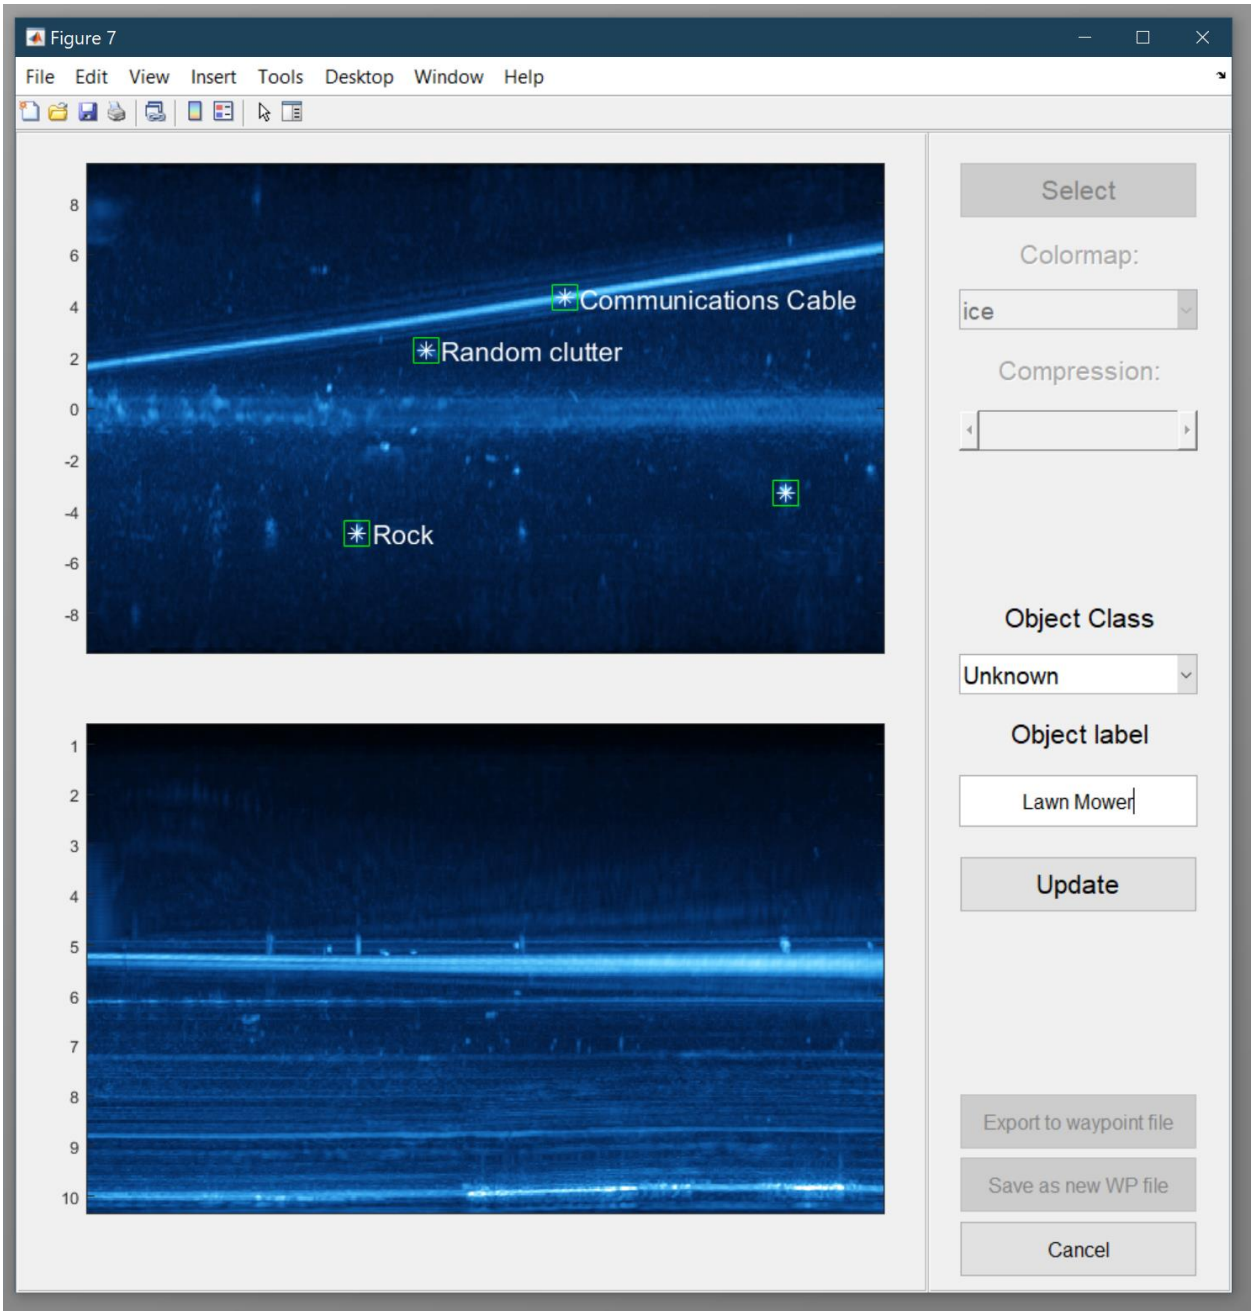

Figure 61: The waypoint selection window after the colormap has been modified and several objects have been tagged.

Once several objects have been tagged and their metadata modified, the user can select “Export to waypoint file” or “Save as new WP file” on the bottom right. If the former is selected, then the new contacts will be added to the current text file; if “Save as new WP file” is selected, then the user is prompted for a new file name, and a contact list duplicating the current contacts plus the new contacts is created. In both cases, the points are added to the current list of contacts in the navigation menu and become visible in the map screen.

MuSTLive can also display the current amount of survey coverage that has been performed over a region of interest. In the Navigation sub-panel, the user may enter the lower left and upper right corners of a survey region of interest, and by checking the “Show target region coverage” box MuSTLive will keep track of how much of the region has been scanned by the system. This is useful for ensuring that no “holidays” have prevented the region of interest from being completely scanned. Figure 62 shows the coverage map next to the navigation map screen. The path of the vehicle is visible in the map screen, and the corresponding population of the coverage map is also visible.

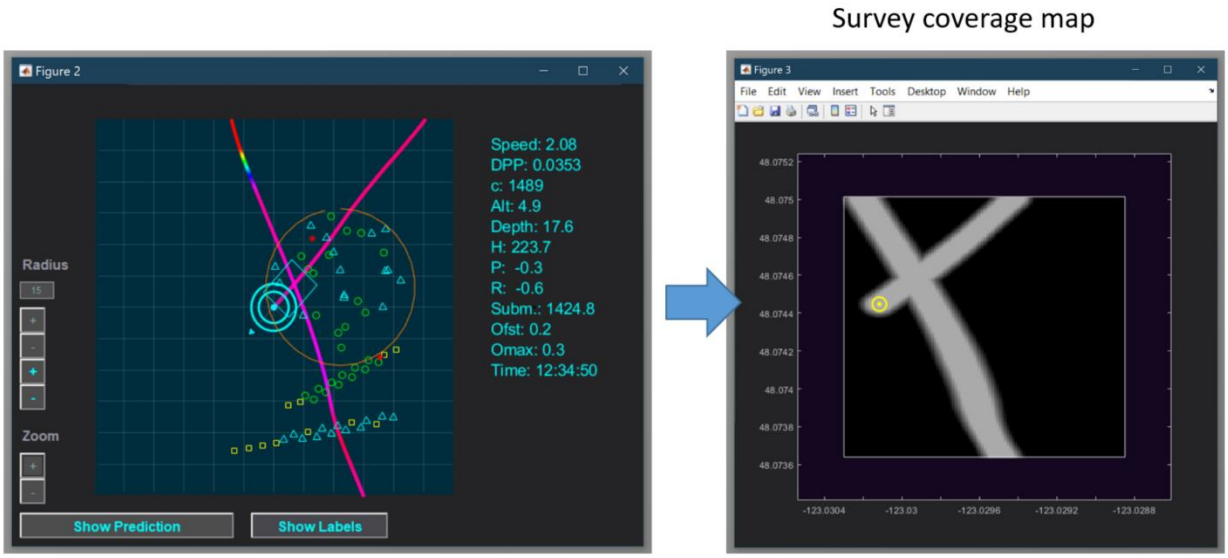

Figure 62: The navigation map for Sequim Bay (left) showing two passes through a UXO testbed, and the corresponding coverage map produced by the scans (right).

In the coverage map, two regions are visible: a dark gray box which depicts the region of interest and shows coverage, and a purple buffer region. The buffer region does not show coverage but is present to give an indication of when the tow-fish is about to enter the region of interest. The location of the towfish is marked by a yellow dot. The yellow circle surrounding the dot indicates that valid SAS data is being collected. When the sediment is out of sonar range and no SAS data is present, the circle will be orange and the coverage map will not be populated, even if the MuST is inside the region. The width of the swath is dependent on the altitude of the vessel. The swath of the MuST is determined by its beam width, which is being limited to +/-65 degrees from nadir. If the vehicle is at higher altitude, the swath is wider; however if it is out of range than no coverage is added to the coverage map.

A few additional convenience items have been added to the “Miscellaneous” sub-panel of the MuSTLive visualization screen. The user can select different colormaps from the popup menu; this set the colormaps used for all of the SAS data displays.

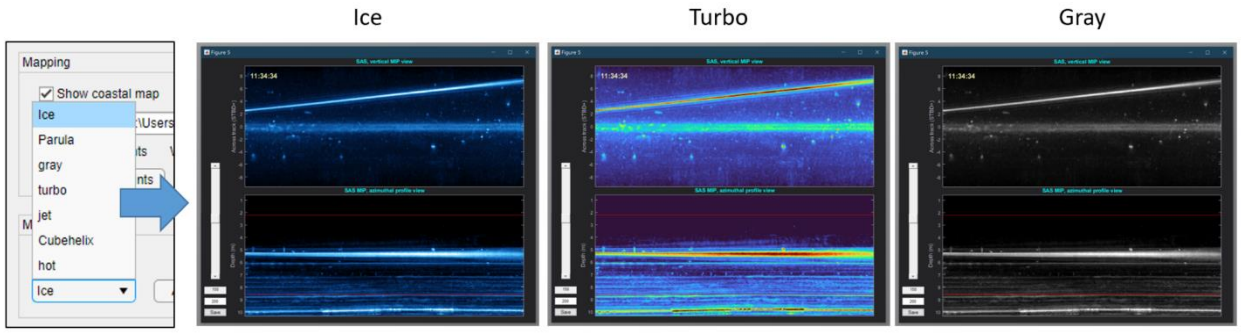

Figure 63: Examples of different colormaps that can be selected from the colormap popup menu under "Miscellaneous" in the MuSTLive data visualization screen.

Lastly, an "arrange" function is also present, which auto-arranges the most commonly leveraged data visualization windows in an effective pattern on the user's screen (this defaults to the second monitor, if available). A future expansion of this capability would be to have multiple layout options available in a popup menu and let the user customize their own window layout. The default layout is shown in Fig. 64.

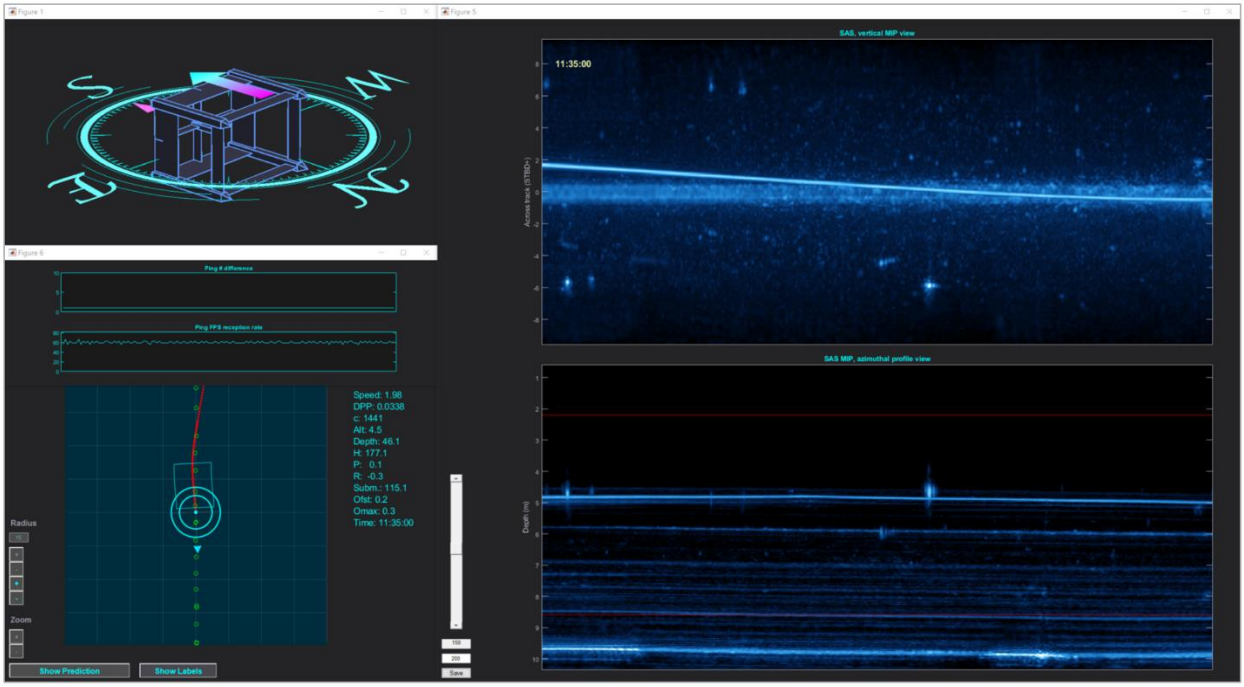

Figure 64: Default data visualization layout created by clicking the "Arrange" button. Visual preference is given to the scrolling SAS display; the map and orientation features are present and the processor can be monitored for overburdening using the connection diagnostics display.

This has been an overview of the functionality of MuSTLive. The *Results* section of this report will contain some examples of the ramifications of different settings and differences between MuSTLive beamforming results vs. post processing results. MuSTLive has proven invaluable for field work, yet would greatly benefit from being able to incorporate information from additional

sensors; future follow on work would emphasize integration of data from additional sensors into the live data processing framework established by MuSTLive.

Renavigation

Precision geo-rectification underwater is very difficult due to the fact that GPS signals do not penetrate underwater to any useful depth. At the expense of mechanical and operational complexity and range some systems exploit acoustic transponders linked top-side to GPS positioning, (e.g. LBL, USBL acoustic positioning) but often, as is the case with the MuST Towbody, surveying systems acquire GPS fixes on the surface to determine an accurate position before diving and then rely on dead reckoning (i.e. estimate position via integrating gyroscope-oriented velocities) to estimate underwater position. The navigation systems, often known as *Inertial Navigation Systems* (INS), contain extremely precise gyroscopes. Data from these gyroscopes along with accelerometers and velocity data from Doppler velocimeters are combined in a Kalman filter to estimate position. Because no absolute position estimate is being measured underwater the uncertainty in position always grows as a function of time. Furthermore, almost inevitably, small sensor biases exist that increase position error as a function of distance travelled. As a result, when a system resurfaces and re-acquire GPS, a large positional correction is discovered which causes the position estimate produced by the navigation filter to "snap" to the GPS augmented estimate. An example using MuST navigation data is shown in Figure 65.

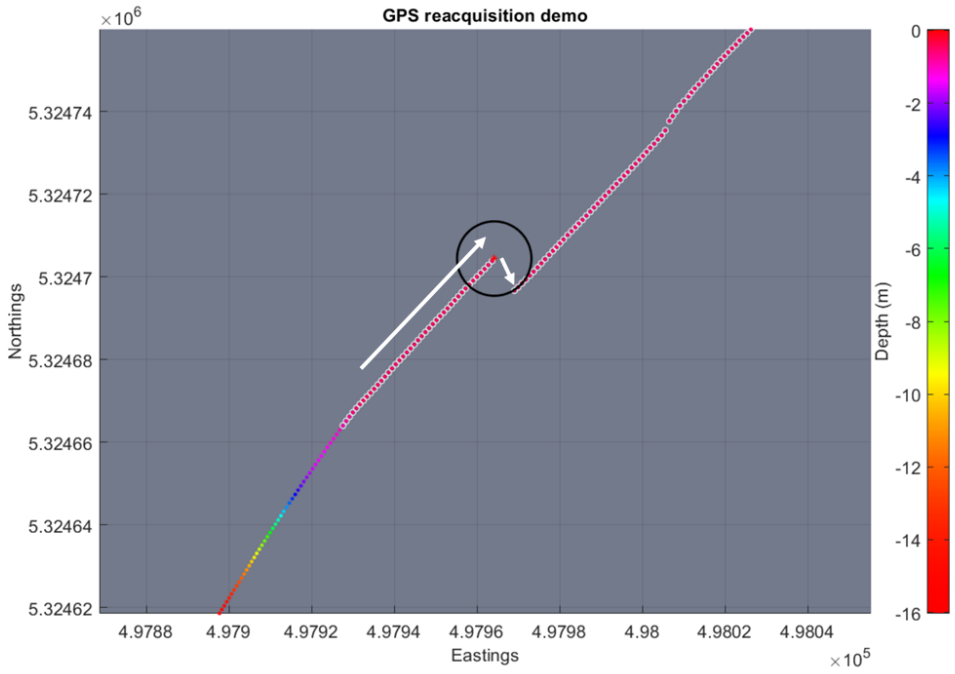

Figure 65: A GPS re-acquisition jump. The path of the towfish is plotted in this figure with depth color coded using the scale on the right. When the vehicle is less than one meter deep a white ring is placed around the marker, indicating the antennae are exposed. The black ring is a 9-meter radius circle placed just prior to the GPS update being incorporated into the position estimated by the INS.

Figure 65 makes it clear that the position estimate produced by the navigation filter of the INS system is meant to give an optimal estimate of location without any regard to the physical realism of the navigation solution; i.e. the path in Fig. 65 is not meant to be a best estimate of the vehicle trajectory, just a best estimate of the current towfish location. It can readily be acknowledged that a better estimate of the vehicle path history could be reconstructed from the data by incorporating the acceleration, velocity, and other navigation information and propagating navigation information both forward and backward in time. This would eliminate the discontinuity and improve the position uncertainty prior to GPS reacquisition.

In fact, there is a standard method for solving this problem: a fixed-interval smoother such as a Rauch-Tung-Striebel filter [20], also known as a *Forward-Backward* Kalman filter. These types of filters often come in optional post-processing software packages available from high-end INS vendors (e.g. IXblue Delph INS [21] or Sonardyne Janus [22]), and essentially apply a standard Kalman filter forward and backward in time, where the backward computation, when combined with the forward, eliminates the problems imposed by causality for the standard forward Kalman filter, eliminating the discontinuities and generally making the regions of largest position uncertainty the points of greatest distance from GPS acquisition (i.e. the central point of a scan).

The MuST Towbody has undergone several INS system changes, and most of the systems did not have off-the-shelf re-navigation solutions available, which prompted the development of a renavigation solution for the MuST and its incorporation into the software library used to post-process eBOSS data. The current INS Model on the MUST towbody, the IXBlue Nano, is compatible with commercially available renavigation software, however compatibility issues have thus far prevented the COTS renavigation solution from being exploited. The renavigation solution integrated into the MuST software toolbox operates by re-solving for improved position estimates using the stored data corresponding to the INS system filter states. The INS navigation filter state outputs for earth-coordinate velocities are combined with absolute position estimates values in a weighted least-squares solver where the weights are determined by estimates of the GPS sensor noise and uncertainty of the velocity states (weights go to zero when no GPS signal is present). The position estimates are found by solving the equation below.

$$\boldsymbol{P_{x,y} = (A'WA + \lambda_{accel})^{-1} A'WB} \qquad \textbf{Eq. 53}$$

In Eq. 53, *A* is a large and sparse transformation matrix encapsulating the linear equations relating the measurement vector *B* (INS position state and velocity estimates) to position, *W* is a sparse diagonal matrix of weights corresponding to each measurement, and λ_{accel} is an acceleration regularization term. INS systems are often optimized for minimum bias, not variance. If, as is the case with the MuST system, it is known *a priori* that a particular system is massive and stable, the position estimates can be improved still further by adding a regularizer to the solver that penalizes acceleration. This penalizer is represented in Eq. 53 by λ_{accel}, and it increases the local smoothness of the estimate, which greatly improves SAS processing by reducing artificial jitter introduced by measurement noise.

Matlab contains some very efficient solvers for large, sparse problems like that shown in Eq. 53, and the computation time required for performing re-navigation on multiple days' worth of data

sampled at 40Hz (i.e. B and $P_{x,y}$ are column vectors millions of entries long) is trivial: on the order of several seconds or less. As a demonstration of the improvement gained by renavigation, the re-navigation solution for the segment shown in Fig. 65 is plotted on top of the INS filter output in Fig. 66.

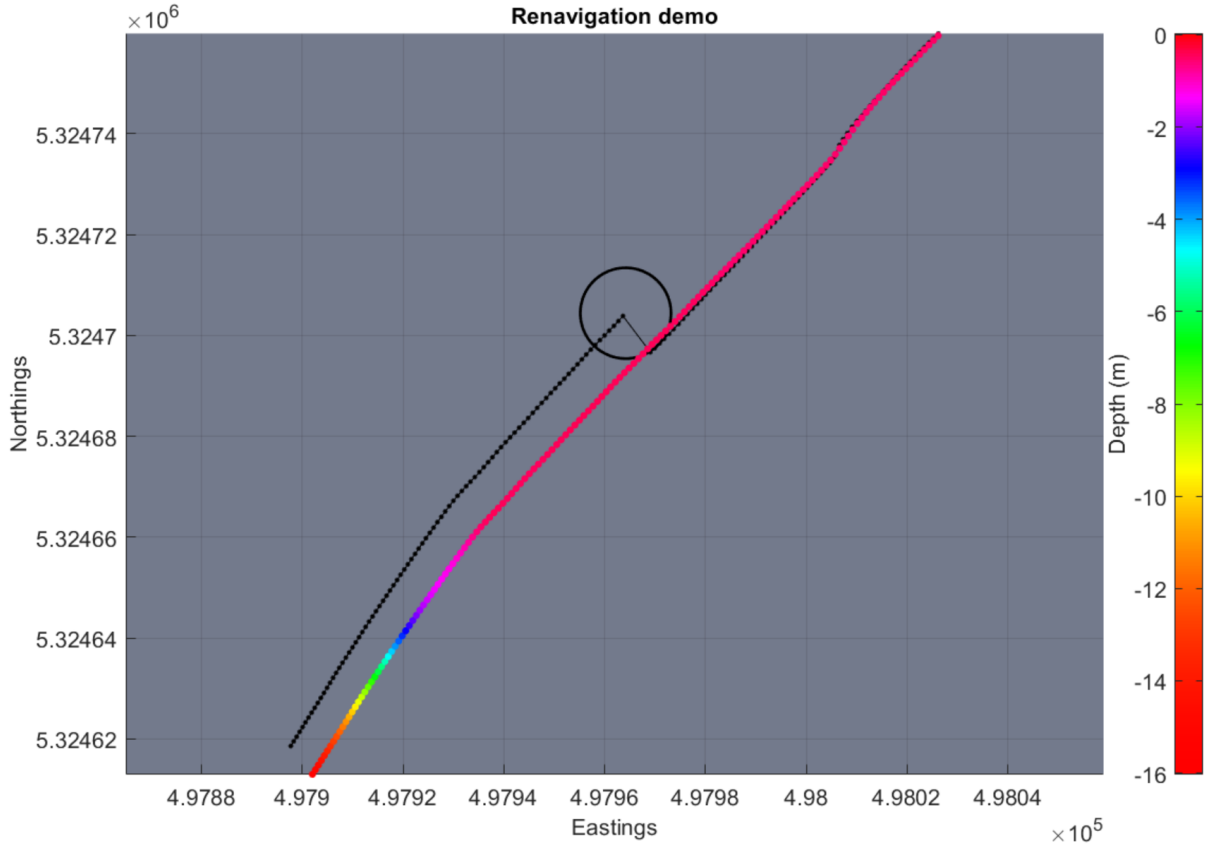

Figure 66: The renavigation solution (colored) vs. the original INS navigation position estimate (black). The reacquisition discontinuity has been completely removed.

The renavigation filter is a simple and fast method attaining immediate improvements in positioning accuracy following a survey operation and it was exploited on every MuST field deployment with INS data available. A quantitative evaluation of the improvement provided by renavigation is given in the results section of this document.

To create a renavigation solution and store the results in a renavigation file that is accessible by the rest of the programs in the MuST toolbox, the user begins by running the RENAV program. The user interface that appears is shown in Fig. 67.

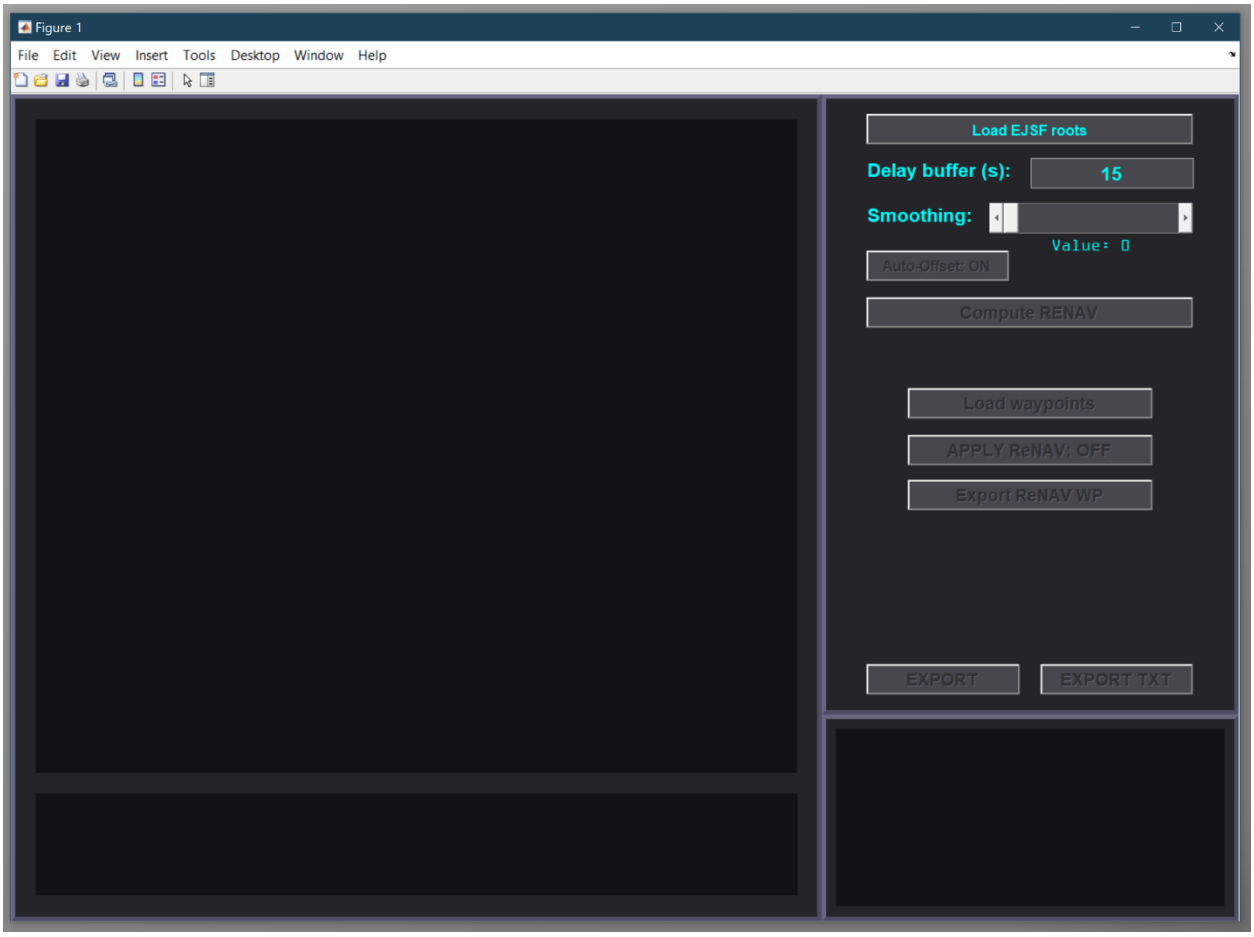

Figure 67: The user interface that appears when the RENAV program is run. RENAV is used to construct a renavigation solution from MuST INS data.

During field operations, MuST INS state data is continually being monitored and stored in *.ejsf* files which package the navigation information in the same format as the navigation data stored in eBOSS *.jsf* files but with some fields added for additional information provided by the INS for which there are no placeholders in standard *.jsf* files. The *.ejsf* files are broken into 3 megabyte chunks as the data is being recorded, and a day of field operations may result in 60 or more files that need to be loaded into memory prior to performing renavigation. Clicking on "Load EJSF roots" will open up a selection window that the user can browse. "EJSF roots" refers to the root name for the solutions corresponding to a particular operation. Occasionally there will be more than one root, so the user selects an EJSF file with the appropriate root for all the datasets that are to be combined in a renavigation solution and the program imports all the associated EJSF's with the same root. An example is shown in Fig. 68.

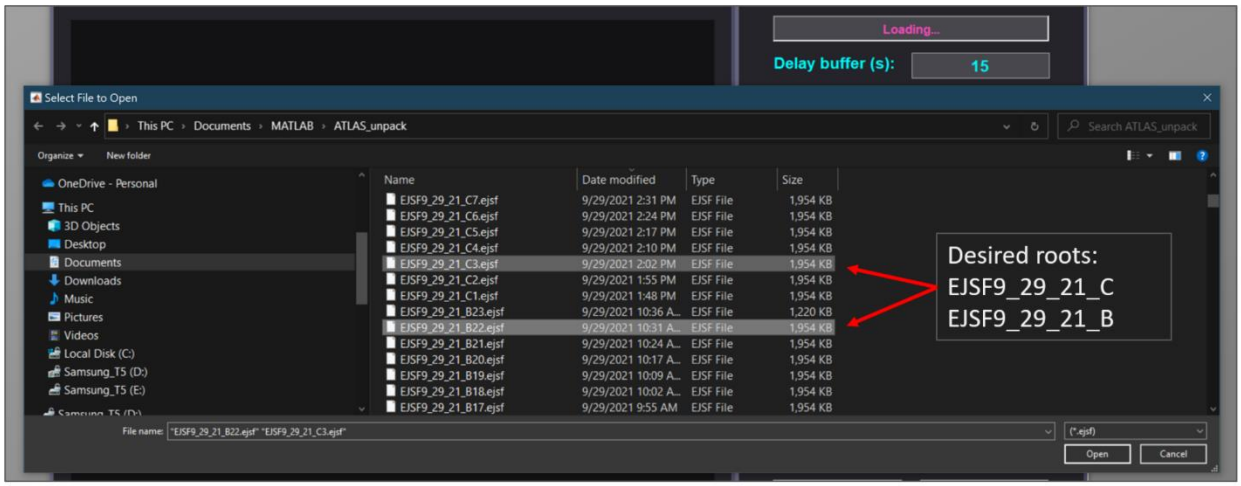

Figure 68: Specifying the roots for importing navigation data from *.ejsf* files into memory

Once the files are selected all files having the same roots are parsed, navigation data is concatenated in memory and the current position estimates are loaded into the position plots as shown in Fig. 69.

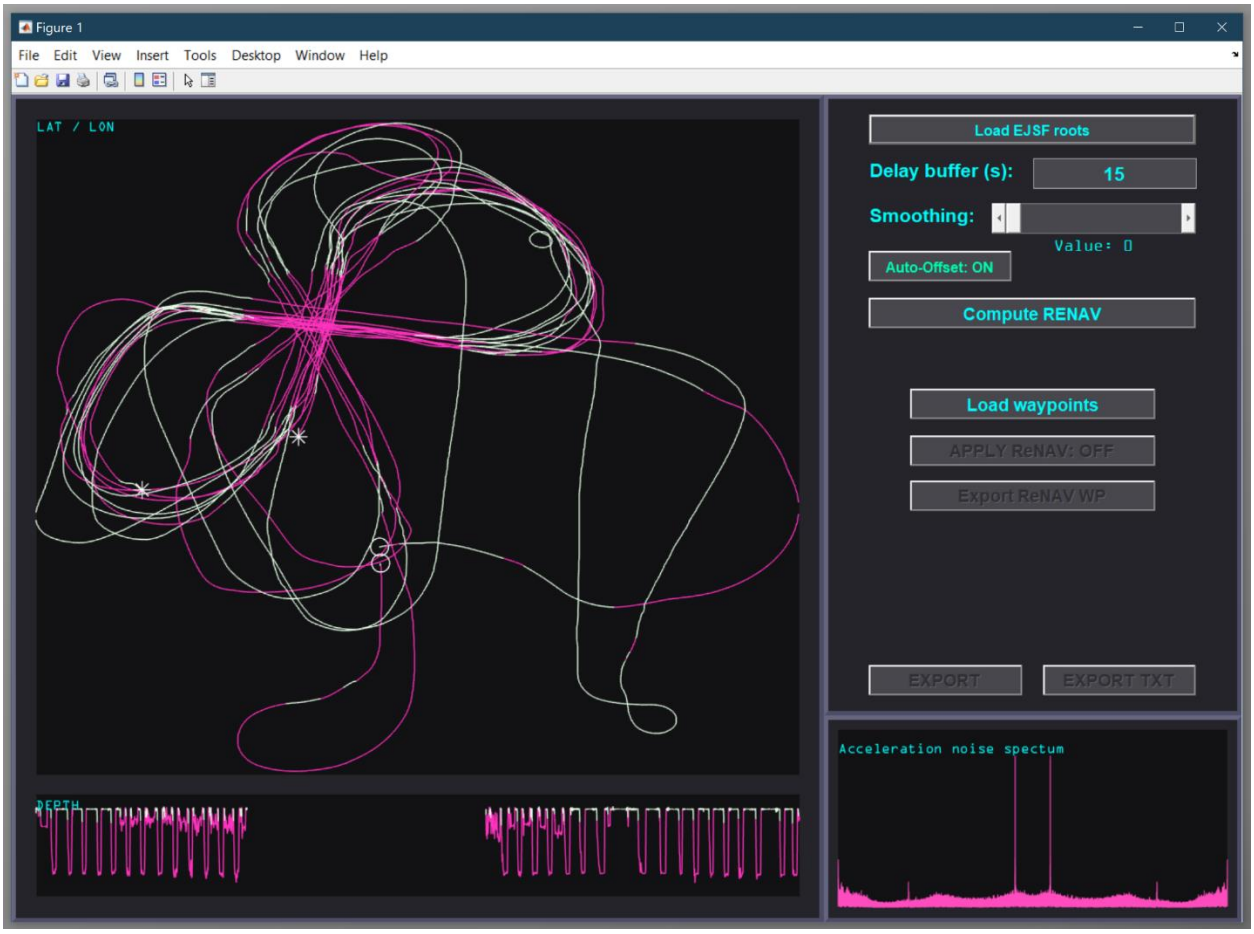

Figure 69: INS navigation data displayed in the ReNAV program following loading data from *.ejsf* files.

In Fig. 69 the top left plot shows the latitude and longitude of the towfish position as estimated by the INS system. The bottom left plot shows the depth vs. time. A large gap is present indicating that no data exists for an extended duration of time. In the present case this was due to inclement weather preventing the MuST system from being deployed. The color-coding refers to when the INS system has flagged that a valid GPS solution is acquired and being incorporated in to the navigation solution. White positions indicate a GPS solution is acquired, and ping solutions indicate that no valid GPS signal is being acquired.

The bottom right plot shows a noise spectrum associated with sensor acceleration estimates. Near the Nyquist rate there is a prominent narrowband spike indicating some sort of oscillation in the navigation solution is occurring that is almost certainly non-physical.

In INS systems and for some filter settings the GPS position estimate is not integrated into the navigation system immediately after signals are acquired, but after several seconds of statistics regarding the GPS signal have been gathered. The "Delay buffer (s)" numerical entry allows the user to specify how many seconds to wait after GPS acquisition has been acquired before the signal is used as an absolute position reference in the renavigation solution.

The "smoothing" function determines the value of λ_{accel} in equation 53. The results of the smoothing function are visible in the acceleration noise spectrum; pushing the slider to the right reduces the higher frequency content in the spectrum plotted on the bottom right.

Toggling the "Auto-Offset" function on prompts the renavigation routine to search for any biases in the velocity estimates by looking for long-term trend differences between purely navigation-driven position estimates and the position estimate determined by the INS navigation solution. Bias compensation terms are added to the velocity values that eliminate any long-term trend differences.

Pressing "Compute RENAV" solves for the improved navigation solution using the user-specified delay buffer, smoothing value, and incorporating the bias estimates if the automatic offset compensation is enabled. The resulting renavigation solution is plotted on top of the INS solution in cyan as shown in Fig.'s 70 and 71.

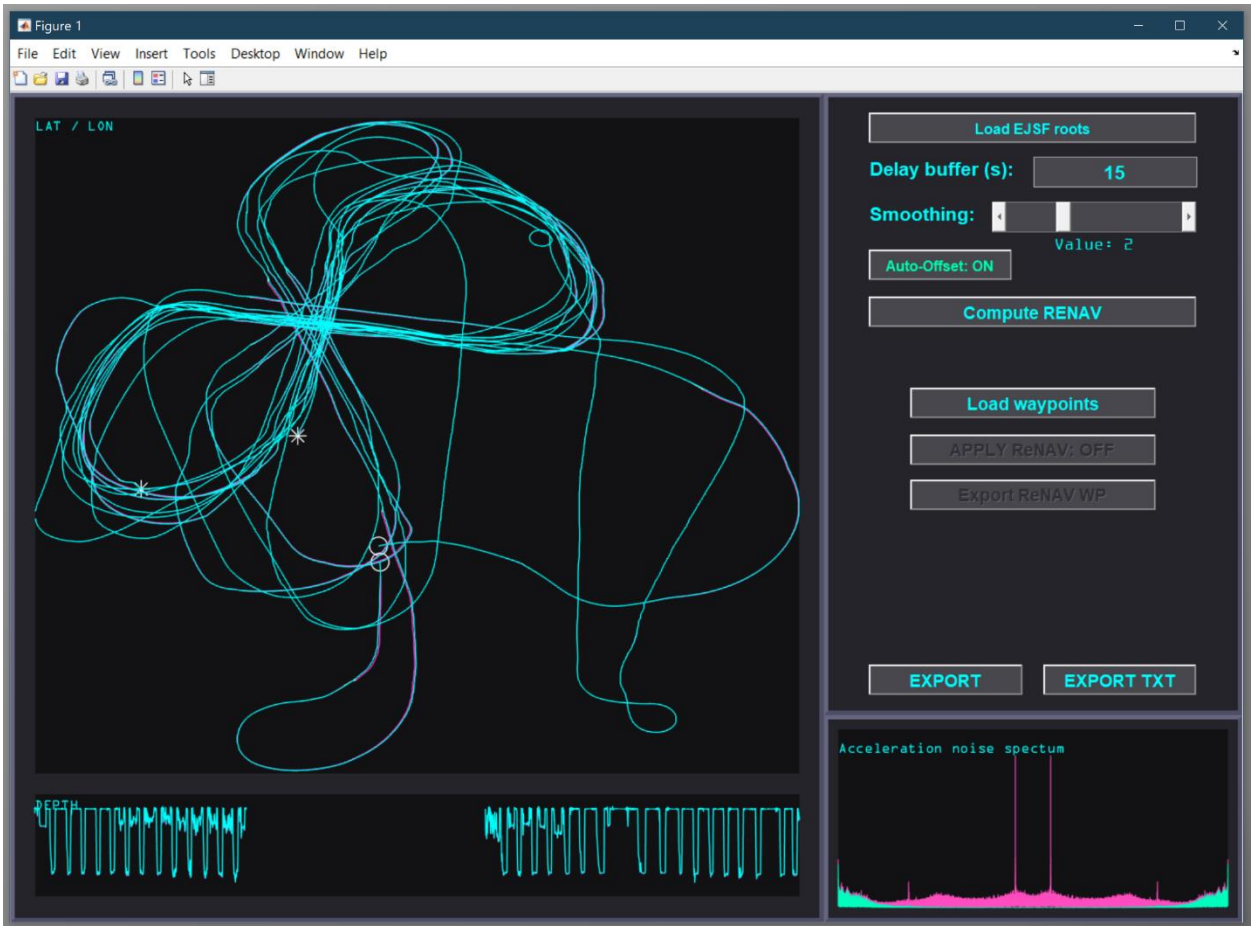

Figure 70: The renavigation solution (cyan) overlaid on the INS navigation solution (pink and white).

At the size scales used in Fig. 70 it is difficult to see any of the corrections resulting from renavigation, however there are some obvious differences in the acceleration spectrum shown in the bottom right plot. Most of the high frequency acceleration components have been nearly eliminated in the smoothing process. Figure 70 zooms in on the plot to show the resolution of some reacquisition discontinuities.

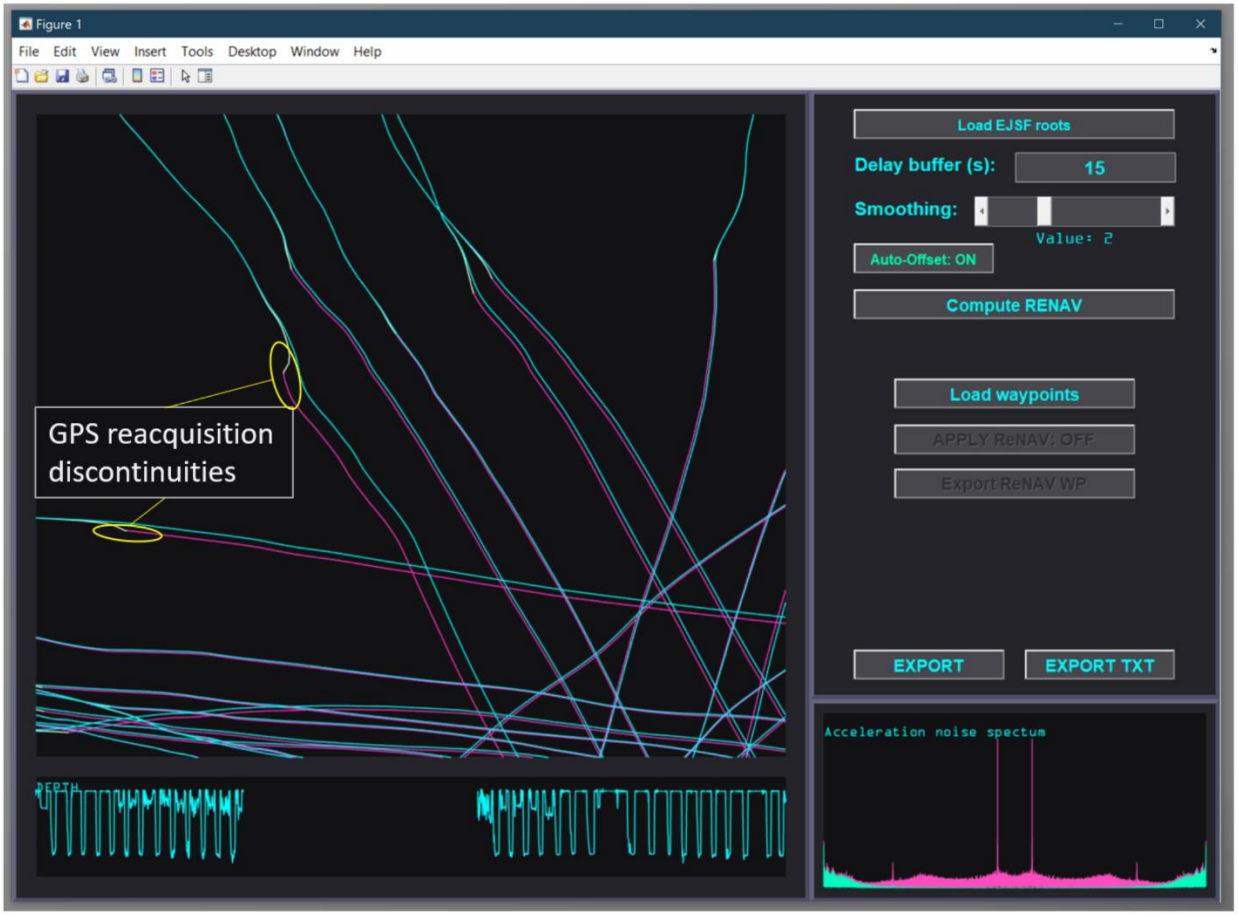

Figure 71: A zoomed in view of the track data shown in Fig. 70. Locations in which GPS reacquisition has caused discontinuities are highlighted in the yellow ovals. The renavigation solution smoothly transitions between the submerged and surfaced states with no position lurch.

RENAV also provides the ability to import waypoints by clicking on the "Load waypoints" button. The user can then select any of the textfiles used by MusTLive to display contacts in the navigation map display, and the same contacts will be displayed in the upper left track map window in RENAV. This is demonstrated in Fig. 72.

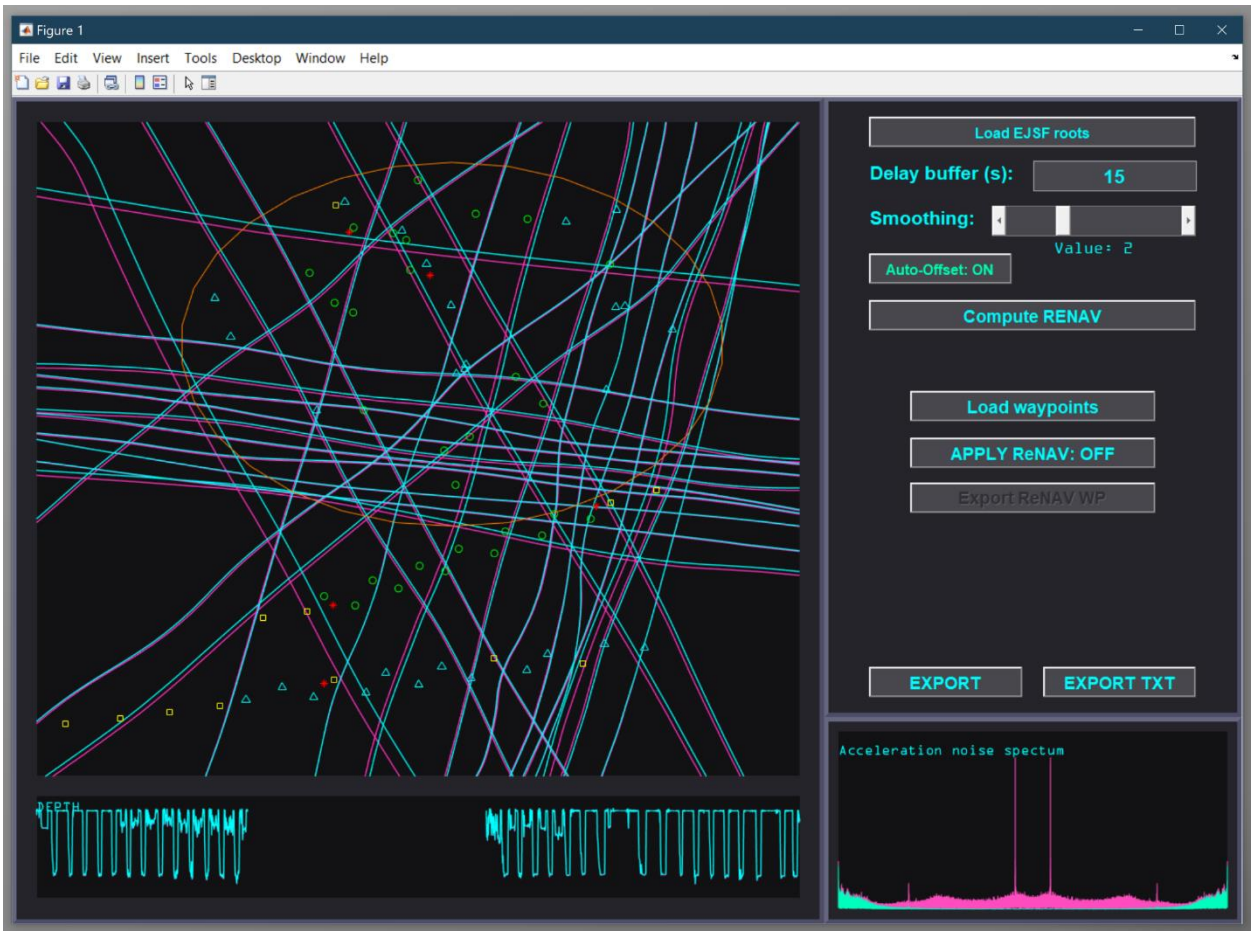

Figure 72: Map contact data loaded into RENAV via the "Load waypoints" option. Note the oval shape of the circular testbed region. The image is being displayed on an equal-axis scale which, for latitude and longitude, will cause a stretching based on the latitude of the map.

When map contacts have been loaded into RENAV the user has the option, via the "APPLY ReNAV" toggle switch, to apply the offsets estimated from renavigation to the positions of the contacts for the purpose of improving geospatial localization accuracy. Recalling from a description of the waypoint file data (summarized in Fig. 59) that each contact has the option of including a sonar acquisition time-stamp (which is auto-populated if the contact was created via detection in MuSTLive), RENAV will search through the list of contacts for any markers that were captured during the time-stamps associated with the renavigation solution. The renavigation offsets associated with the marker time-stamp are applied to the marker position as a navigation bias correction. Markers with no time stamp or non-overlapping time stamps are ignored in this correction process. The updated list of waypoints can be exported by the user by clicking the "Export ReNAV WP" button, which becomes enabled if the renavigation solution has been applied.

Lastly, as demonstrated in Fig.'s 70 – 73, when the user generates a renavigation solution the export buttons on the bottom right of the RENAV user interface become enabled. Clicking the EXPORT button will prompt the user to create a file name and will save the renavigation data as a *.mat* file that is accessible by other MuST Toolbox functions (e.g. beamforming and

mosaicking). Alternatively, clicking “Export TXT” file will save all of the navigation data in an ASCII text file with nine columns: 1) UNIX time stamp, 2) Latitude (decimal degrees), 3) Longitude (decimal degrees), 4) UTM northings (meters, zone specified in column heading), 5) UTM eastings (meters, zone specified in column heading), 6) Depth (m), 7) Altitude (m), 8) Heading (deg), and 9), sound speed. This alternative storage format makes the renavigation data available to COTS sonar software (e.g. Hypack, SonarWIZ, etc.) and was used by collaborators during MuST field experiments.

File association and linking

The eBOSS data files created by the MuST are typically on the order of ~1 GigaByte of data per file and new files are created approximately every minute. Loading a file into memory can take a non-trivial amount of time and fast sorting of sonar files based on location, altitude, pulse repetition frequency or other variables can take a long time if the original sonar files must be accessed to ascertain relevant information about the file. Data organization tasks can be greatly accelerated by generating a list linking relevant metadata to the sonar files; the list can be queried very rapidly to determine, for example, which data files captured data over a specific region of the testbed. The MuST software library provides a tool (“JSF_table”) for file - metadata linkage and list generation. Figure 73 shows the very simple *JSF_table* user interface.

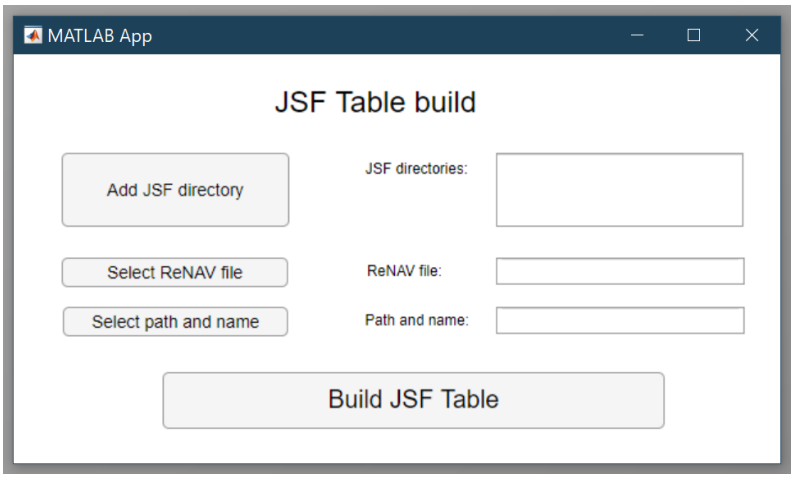

Figure 73: The *JSF_table* user interface. Users input the relevant sonar data directories, link a renavigation file and specify a path and name for the new table file.

In *JSF_table* the user is required to select directories containing the eBOSS *.jsf* files that will be linked in the table. The user also has the option of linking a renavigation file to the list which enables the renavigation solution to be utilized later during SAS batch-processing. Finally, the user can specify a save path and table name by selecting “Select path and name.” Selecting “Build JSF Table” will prompt the program to access each of the JSF files in the selected directories and parse the navigation and time stamp information to tabulate the following information:

- File name

- File path
- Time stamp: first ping
- Time stamp: last ping
- Average heading
- Heading standard-deviation
- Median sonar altitude
- Minimum sonar altitude
- Ping rate (pulse repetition frequency)
- Median speed of sound
- Maximum sonar range of visibility
- Starting and ending latitude
- Starting and ending longitude
- Starting and ending northings (UTM)
- Starting and ending eastings (UTM)
- Renavigation success flag (i.e. file time-stamp overlaps with associated renavigation file)
- Renavigation file name and path
- Renavigation feature-based flag (Boolean)

This information is organized in a Matlab *struct* variable and saved in the specified location as a *.jsftable* file with the user-determined file name. JSF table files are used extensively for SAS batch processing and region-of-interest post-processing when it is desirable to determine all files that contribute to observations of a particular region of interest.

SAS post-processing

User Interface and beamforming options

Recognizing that the data products required for performing different tasks can be very different, the MuST Toolbox provides the user with flexibility regarding how eBOSS SAS data is post-processed. Data can be processed individually and the volume data interrogated using the *SASview* program, or data can be batch-processed for large-scale target detection operations and creating data products for mosaicking. Both functionalities are accessed by the *MuST SAS post-processing tool* shown in Fig. 74.

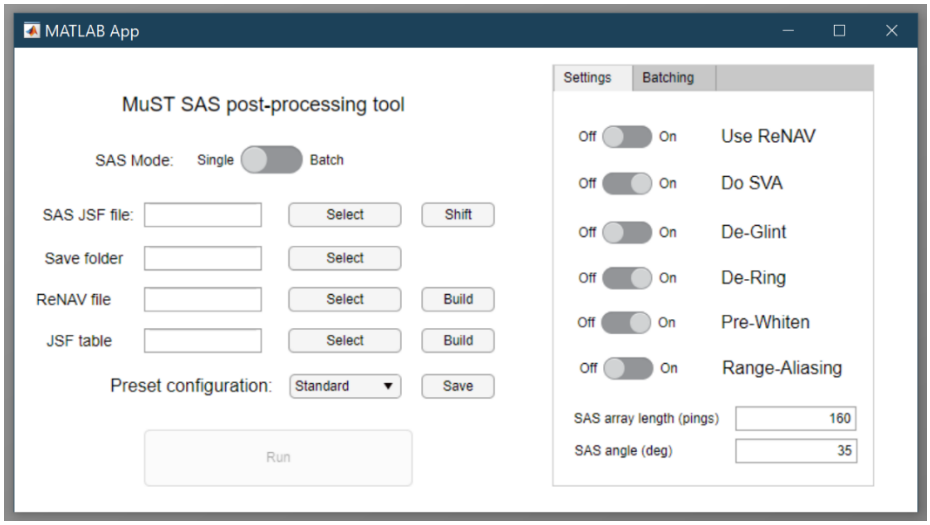

Figure 74: The MuST SAS post-processing tool user interface.

The "SAS Mode" option in Fig. 74 allows the user to toggle between single-file beamforming mode and a batch mode. If "Batch" is selected, then the "SAS JSF file" field becomes "SAS JSF *directory*," and by clicking "Select," the user is prompted to select a folder containing JSF files that will all be processed during the batch processing operation. In contrast, if "SAS Mode" is set to "Single" then the user is prompted to select a single *.jsf* file. The "Shift" button that, which is only enabled in "Single" mode, combines data from the selected file and the file immediately following it. It truncates the concatenated data to be the portion spanning files and processes this shifted dataset. This is useful when, for example, acoustic data from a feature of interest spans the end of one *.jsf* file into the beginning of the next.

The user has the option of entering a "save" directory to which data from the SAS processing will be saved. This is only necessary if batch mode is enabled, in which case the Run button will not be enabled unless this field is populated with a valid directory. Additionally, fields for renavigation and JSF table files exist. The JSF table file must be specified to be able to perform batch processing. In the event that the user needs to build a renavigation or JSF table, the user can select the "build" operations adjacent to each of these fields and the ReNAV user interface or JSF_table user interface described previously will open.

The tabs on the right of the user interface allow specific SAS processing parameters to be modified by the user. The user may opt to use renavigation data or simply use the data packaged with the *.jsf* file (or files, in batch mode) depending on whether the Use ReNAV toggle is set to *on* or *off*. The user also has the option of applying a non-linear sidelobe reduction algorithm called *Spatially Variant Apodization* [23], and applying a de-glinting operation. Enhanced backscattering from coherent near-nadir sediment surface reflections often dominate eBOSS imagery, hiding targets located near nadir. It also tends to reduce the interpretability of multi-scan mosaics and can bias feature-based renavigation algorithms by creating large, artificial features. For these reasons, de-glinting was found to be an extremely important post-processing step to apply to SAS data during UXO test-bed survey operations, and an algorithm description is provided in a subsequent section.

De-ringing, the next option available to the user, was intended to provide the user with the ability to subtract certain repeatable transmitter artifacts from the SAS data. The artifacts were not particularly strong and didn't cause noticeable corruption to most of the dataproducts created by the system and the function of this toggle hasn't yet been implemented into the beamformer.

The *pre-whiten* option tells the beamformer to access a calibration file, the data from which is used to compensate for the response of the transmitter and receiver array and normalize acoustic power across the frequency bandwidth. Currently the calibration file uses clutter calibration data, i.e. the spectral curves created by averaging the backscattered signals over a large amount of pings; a future update would be to incorporate the calibration curves collected for the MuST array and described earlier in the methods portion of this report. The typical result of the pre-whitening operation is to improve the resolution of the resulting imagery.

Lastly, toggling the *range-aliasing* option on will prompt the SAS beamformer to compensate the beamforming operation for the special case when the dominant scatterers lie beyond the maximum range of the sonar. In this case the ordinary range migration curves used to perform spatial matched-filtering are not applicable and the beamforming process needs to be modified to produce any useful imagery. This capability opens up some interesting opportunities for the MuST system, and a subsection will be devoted to detailing how this process works.

Just like in MuSTLive, the settings panel also allows the user to specify how many pings are used in the SAS beamforming queue and also to determine what the SAS integration half-angle is. Currently the system is set to contain 160 elements in the ping queue and integrate over a 35 degree angle. It is often the case, especially when operating at higher altitudes and lower pulse repetition frequencies, that the data from the MuST is undersampled, which has the result of producing distracting along-track aliasing artifacts. Sometimes these artifacts can be mitigated by narrowing the integration angle. Examples of this will be shown in the results section of this document. To achieve highest system resolution, which is important during UXO remediation, it is necessary to increase the integration angle of the beamformer and use a ping-queue size that doesn't limit the integration angle. The number of points in the ping-queue can be calculated using the following equation:

$$\boldsymbol{N_{queue}} = \left\lceil \frac{\boldsymbol{2PRF} * \tan(\boldsymbol{\theta_{half}}) \boldsymbol{R_{max}}}{\boldsymbol{velocity}} \right\rceil \qquad \textbf{Eq. 54}$$

Where *PRF* stands for the pulse repetition frequency, θ_{half} is the user-specified SAS integration half-angle, $R_{\max}$ is the maximum range at which the integration angle must be valid, ⌈ ⌉ represents the *ceiling* operator, and velocity is the vehicle velocity in m/s. Plugging standard survey values (PRF = 60 Hz, θ_{half} = 35, R_{max}= 6 meters, *velocity* = 2 m/s) it can be seen that the queue length must be no less than 253 elements long to prevent the integration angle from being queue-limited.

Figure 75 shows the additional options that become accessible if the user selects the *Batching* tab.

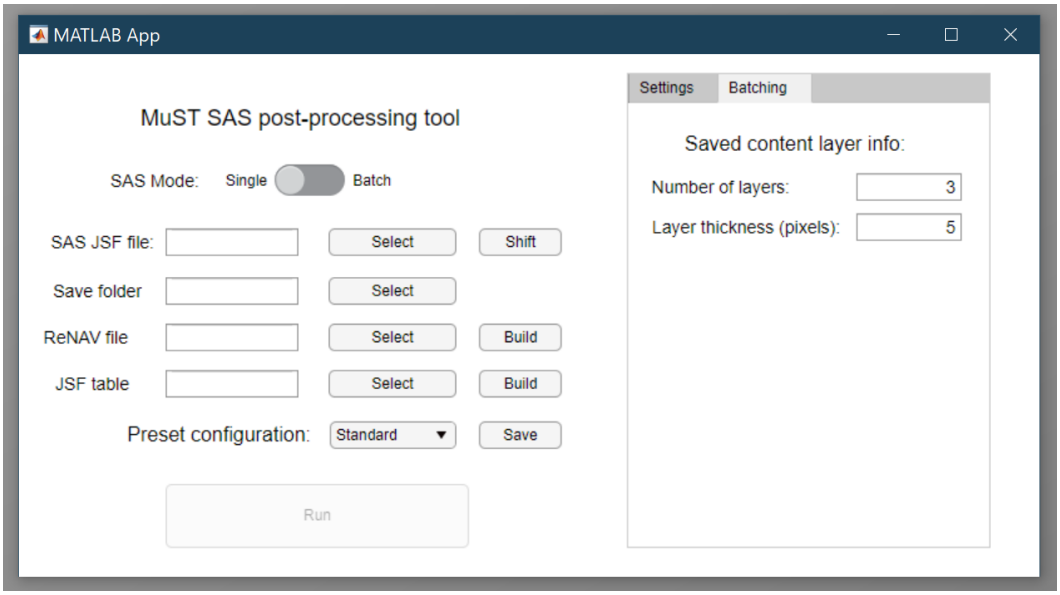

Figure 75: The SAS processing toolbox showing the options available when the Batching tab is selected (right side).

In the Batching panel the user can select the number of X-Y plane MIP layers that are generated, and how thick each layer is in pixels. The layers are arranged to overlap adjacent layers by half of the layer thickness, and to be positioned around the water/sediment interface which is tracked using the MuST altimeter or, if de-glinting is enabled, by tracking the backscattering directly from the sediment. Figure 76 shows a profile overlaid with the vertical position of four different non-overlapping layers to illustrate how they are juxtaposed relative to the sediment/water interface.

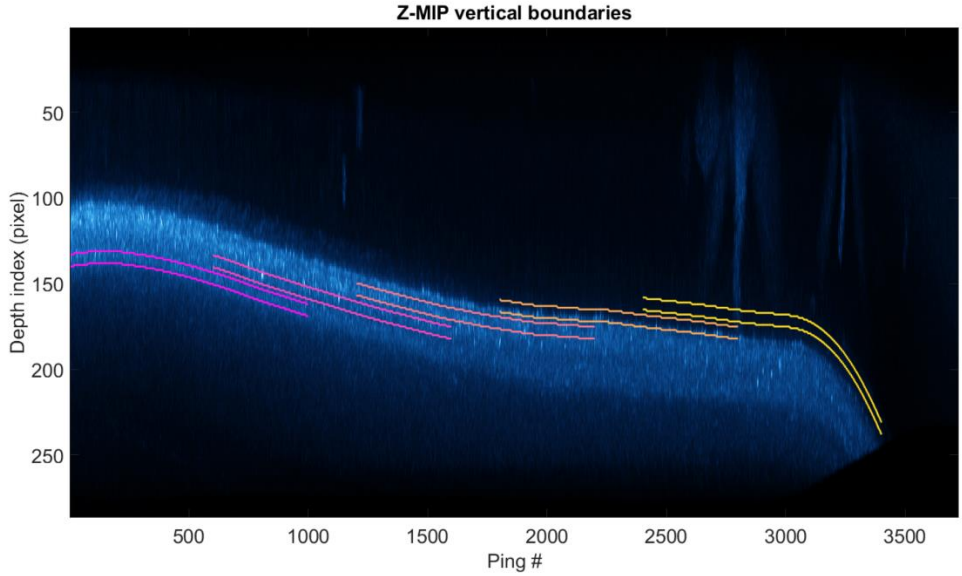

Figure 76: An illustration of the positioning of layers relative to the sediment interface. Rather than representing constant range swaths relative to the sonar, the layers represent depth slices relative to the

water/sediment interface. The layers above are truncated to improve interpretability, but they span the entire length of the imaged region.

Note that the full 3D SAS volume is *not* saved by the batch processing routine. The volumetric datasets created by each file are very large, especially when represented in Cartesian coordinates, potentially requiring tens of Gigabytes per file for storage. Instead, layer data corresponding to the preferences selected by the user in the panel above are stored along with navigation metadata. This vastly reduces size. The metadata stored with the layer images also allow geo-rectification of each image pixel, a process that is used for mosaicking the batch-processing results. Additionally, full range X-Y and Z-Y MIPS are stored in the layer data as well.

SAS coordinate system

The eBOSS SAS data is beamformed using synthetic aperture processing directly into an eastings-northings-up (ENU) coordinate system using the in-situ sound speed estimates measured onboard the MuST, the re-navigation coordinates for the vessel, orientation estimates from the gyroscope (including the appropriate Mercator convergence angle correction for heading), and lever-arms describing the relative position of the array elements vs. the INS.

For convenience, voxels are beamformed in a path-following grid: each ping frame is projected on to a vertical plane determined by the location and heading of the vehicle averaged over the size of the ping-queue used for SAS processing. The volumetric SAS image seamlessly follows the path of the sonar and, generally speaking, this method of volume representation – though not corresponding to a regular sampling in geo-rectified coordinates – results in a much smaller volumetric representation in terms of numbers of voxels.

Rather than storing the specific UTM eastings, northings and height values for each pixel, the path of the vehicle and heading are stored. Geo-rectified coordinates for pixels can be rapidly re-assembled from the path using the vehicle location, heading and pixel spacing values. This reduces file size and also creates the possibility for very simple re-rasterization of imagery with different navigation solutions. This feature is heavily exploited during data-driven mosaicking. Figures 77 and 78 show Maximum Intensity Projections (MIP's) through a data volume before and after SAS processing, showing the path-centered geometry of the volumetric imagery.

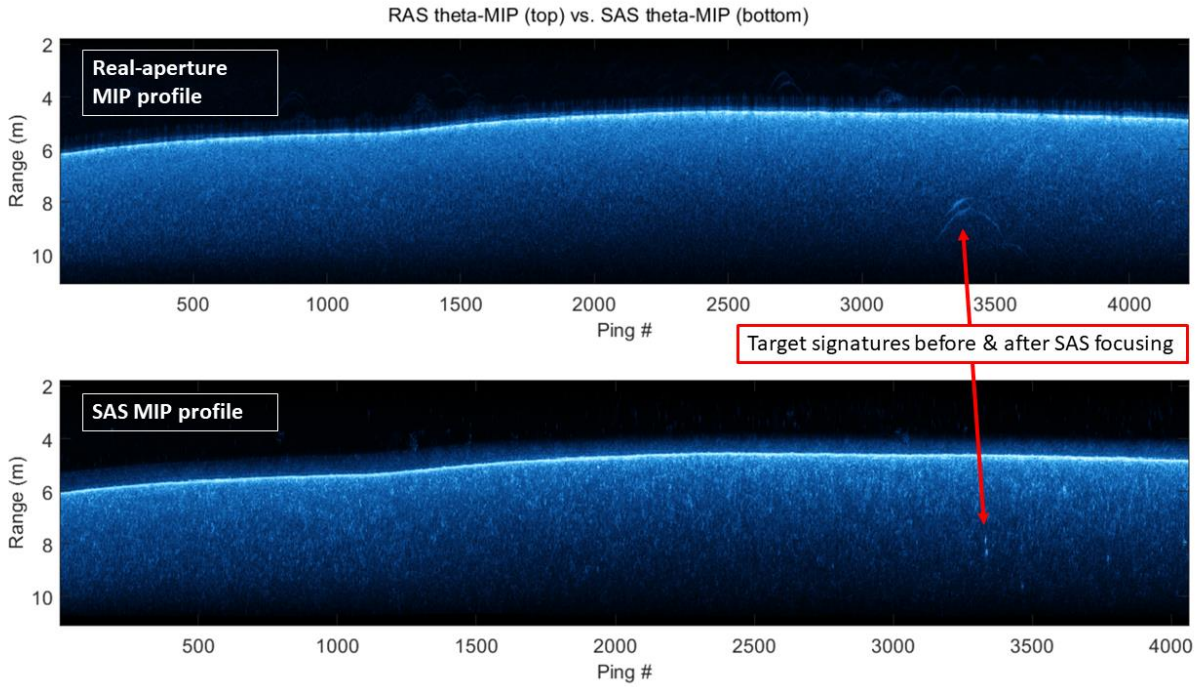

Figure 77: Example real aperture and SAS maximum intensity projection (MIP) profiles, projected through real-aperture angle phi (across-track). A scatterer before and after SAS processing is highlighted by the red arrow: upon SAS beamforming the range-migration arcs collapse into a well-focused scatterer.

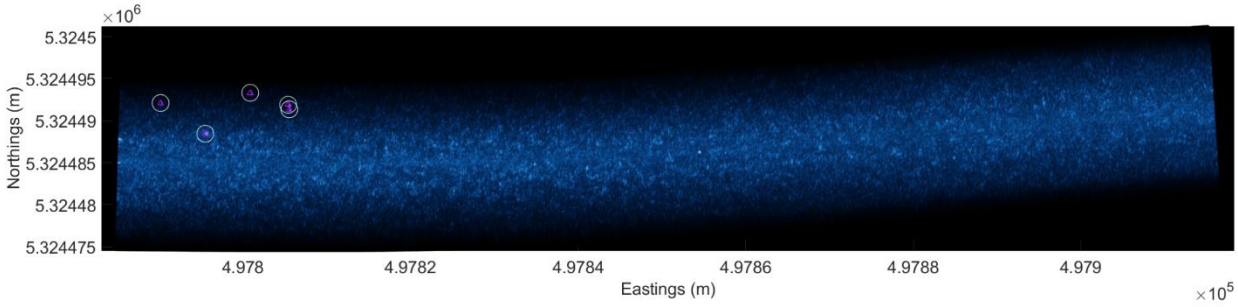

Figure 78: An X-Y MIP layer rasterized into UTM Northings and Eastings coordinates. Note that sediment texture is visible in the imagery and that the nadir glint, which ordinarily manifests as a bright stripe down the middle of the swath, is completely absent due to the application of the de-glinting algorithm.

De-glinting

The specular glint from sediment direct-path reflections often obscures targets at nadir. Specular glints tend to be a significant problem in areas where the impedance discontinuity between water and sediment is large. Normalizing the image to remove the effects of the specular glint can improve the visibility of buried targets however normalization also has the potential to attenuate targets near the sediment boundary. Preferably, the de-glinting process should filter out the specular glint but preserve features associated with targets. Figure 79 shows the filtering and normalization process that was developed to accomplish this.

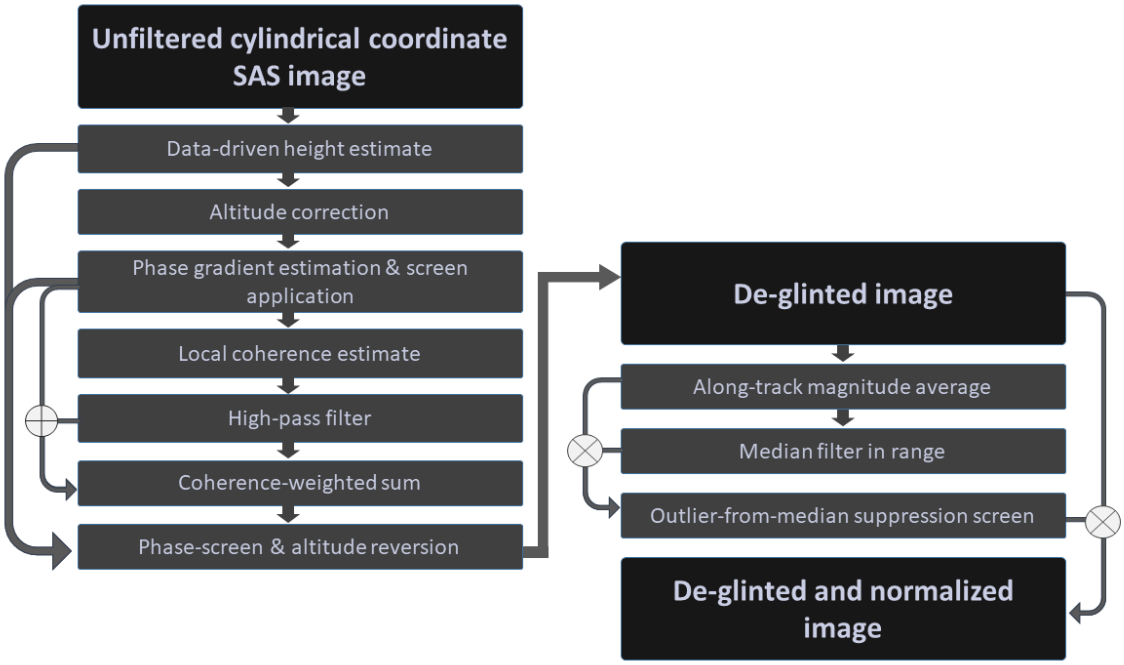

Figure 79: A flowchart of the de-glinting and normalization process. Some relationships (such as the application of the coherence-based weighting functions to the altitude-compensated and high-pass filtered data) have been omitted to preserve flowchart clarity.

The process begins by improving the imprecise altimeter-based estimate of array altitude relative to the sediment. The new altitude estimate is found by applying a matched-filter bank and a series of non-linear operations designed to accentuate features associated with the first return from the sediment. A strip of filtered near-nadir data spanning +/- 0.5 meters surrounding the altimeter estimate for the interface is extracted and the range associated with the brightest return is selected. Outliers from the set of range measurements across the sonar path are rejected by applying iterative-reweighted-least-squares filtering with a penalty on the second-order difference of the solution. A graphical illustration of the first-bounce isolation, detection and curve-fitting processes is shown in Fig. 80. Note that several significant assumptions are made by this algorithm:

- The sediment interface is relatively smooth. This process will not work over a shipwreck or other large interface discontinuity.
- The first bounce is, within +/- half a meter of the interface, the brightest interface.

With regard to the second assumption, environments which violate this assumption have already been scanned by the MuST system (for example, Lake Washington), however in these environments the soft top layer often doesn't conceal targets at nadir as strongly as a sand interface and de-glinting does not seem as important. A de-glinting approach that can be flexibly applied to a multilayer environment exhibiting discontinuities would be an interesting and potentially fruitful avenue of future research, however.

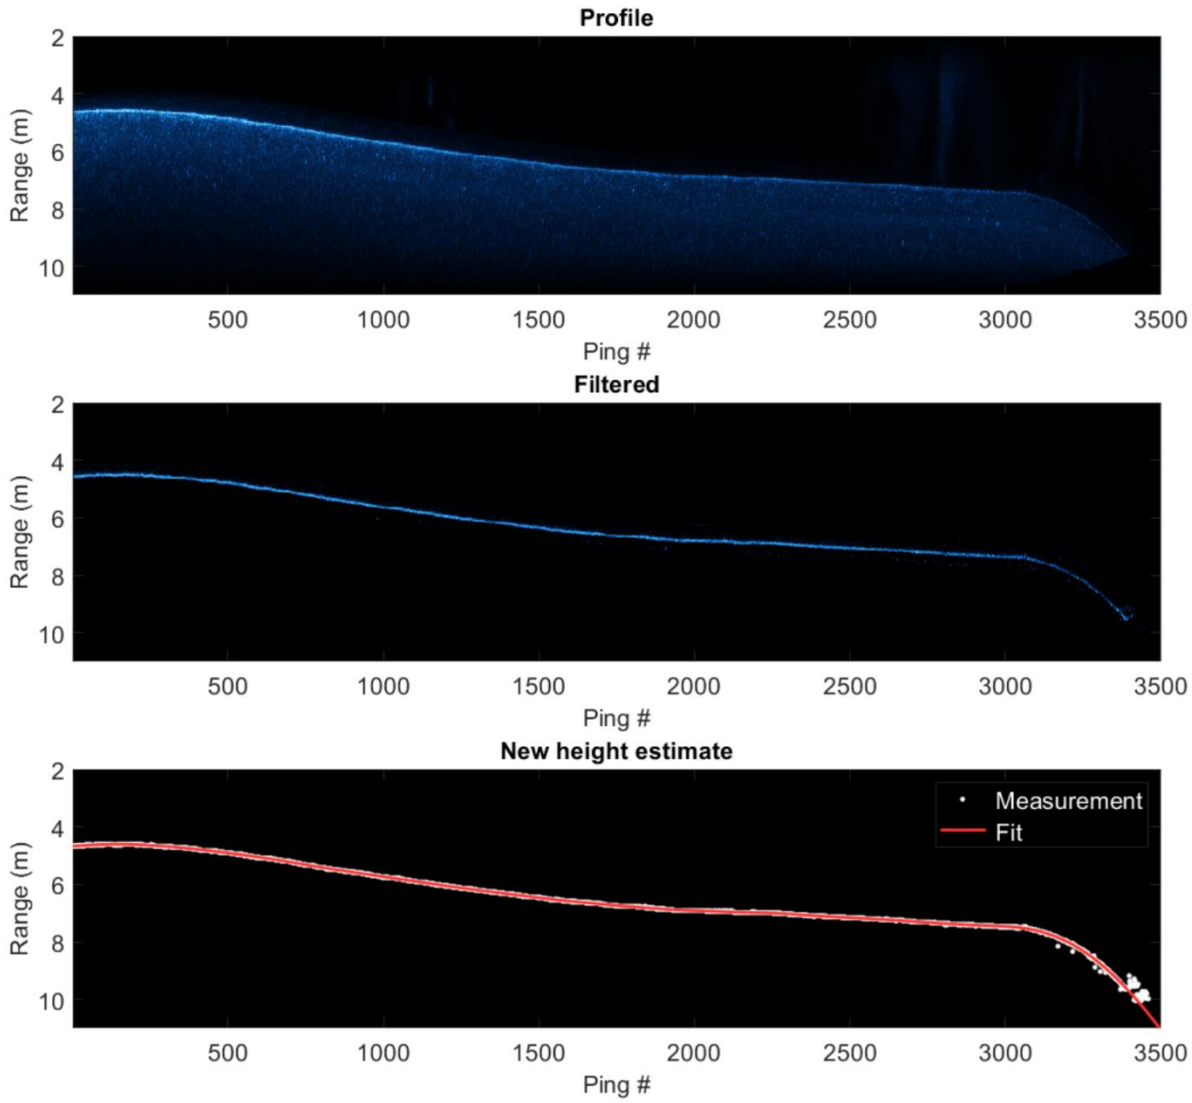

Figure 80: The data driven altitude estimation process - the profile (top) is filtered and masked to accentuate edges near the sediment transition using the DVL altitude estimate (center). The maxima of this masked region are fitted with a robust curve fitting solution that extrapolates outside of the measurement region using the derivative of the DVL altitude estimate (bottom).

The next step of the de-glinting algorithm is to propagate all of the data to an artificially levelled plain in which the specular reflections lie roughly at the same range bin. This is followed by cross-correlation with adjacent data and an extraction of the phase of the correlation peaks. This phase gradient estimate is integrated to produce a "phase screen" unifying the phase for coherent scatterers near the surface. This phase-screen is applied to the altitude-compensated data to generate the "phase-screen compensated data."

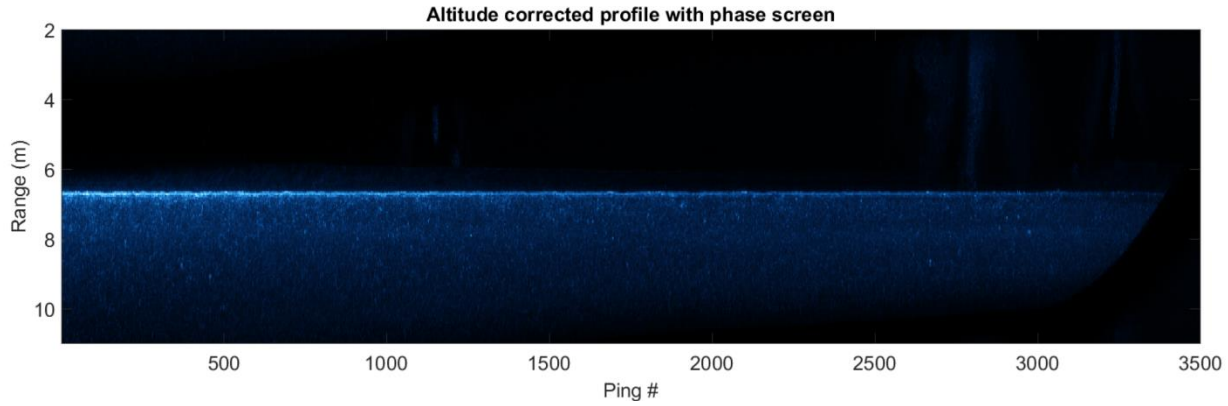

Figure 81: Data (MIP through theta) after altitude and phase screen corrections: the specular glint from the direct path to the sediment lines up across all pings.

A local coherence estimate is computed by applying a boxcar moving average filter in the along-track dimension to both the complex and magnitude values of the phase-screen compensated data, where the coherence can be determined by dividing the magnitude of the filtered complex data by the filtered magnitude data. A weighting function that eliminates all coherence values below 0.5 is constructed from this coherence estimate. A de-glinted dataset is then formed by subtracting the coherence-weighted low-pass-filtered data from the phase-screen compensated data, i.e.:

$$\boldsymbol{S_{psc_deglint} = S_{psc} - W_{coh}\mathrm{F}_{lowpass}\{S_{psc}\}} \qquad \textbf{Eq. 55}$$

In Eq. 55 $\mathrm{F}_{lowpass}$ represents the boxcar window filtering operation in the along-track dimension being applied to the complex phase-screen compensated data S_{psc} and W_{coh} is the thresholded coherence weighting function.

At this point the image is de-glinted. Despite the application of de-glinting there still tends to be a strongly biased degree of back-scattering originating from near-nadir and some level of normalization is also desirable. The next steps normalize bright transitions in the range that have continuity in the along-track dimension so that they don't mask lower-level scatterers in the maximum intensity projection images. The normalization process begins by filtering the magnitude volume image in the along-track dimension by a very long box-window (e.g. ~250 along-track samples). This smooths over potential target features and makes long-term trends in the sediment visible.

Following the smoothing, a median filter is passed over the data in the range dimension. Median filters are non-linear and have very important properties for the present application: 1) they are edge preserving, and 2) they aggressively filter outliers. Thus, general volume scattering level differences between layers are preserved, but residual spikes due to high reflections at impedance boundaries are filtered out. The normalization function is found by computing the ratio of the data pre- and post- median filtering. The original data is normalized by multiplication with this ratio.

$$S_{psc_deglint_normalized} = \frac{F_{median_range}\left\{F_{lowpass_alongTrack}\{|S_{psc}|\}\right\}}{F_{lowpass_alongTrack}\{|S_{psc}|\}} S_{psc_deglint} \quad \textbf{Eq. 56}$$

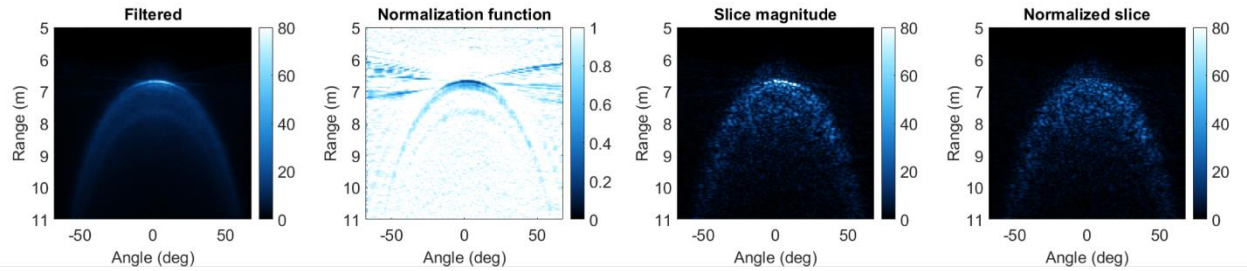

Figure 82: Moving from left to right: 1) a slice of the filtered volume used to determine general trends and discontinuities in range; 2) the normalization function which ranges from 0 (full attenuation) to 1 (no attenuation), 3) the magnitude of the de-glinted slice, still showing elevated levels of backscattering at the surface near nadir, and 4) the slice after applying the normalization function.

Fig. 82 shows the pieces of the normalization process and the effect on an individual slice. Note that the normalization function targets locations where rapid change in range are stable over long distances in the along-track dimension. Fig. 83 shows a vertical MIP through the SAS volume after 1) the coherent de-glinting process, and 2) the normalization process.

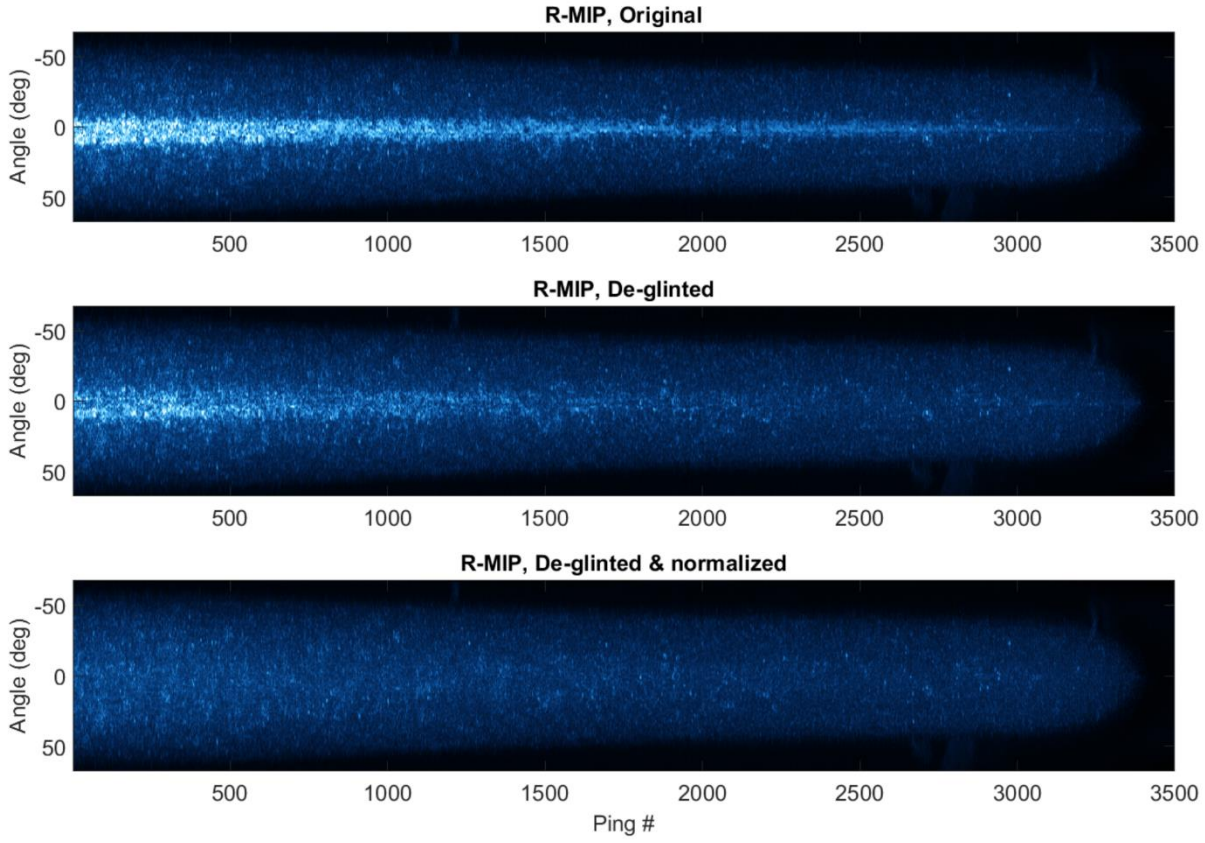

Figure 83: The three stages of de-glinting: a range-MIP through the original volume (Top), a range-MIP through the de-glinted volume (Center), and a range-MIP through the normalized volume (Bottom).

As a final step, the phase screen that was used to increase coherence for specular features is removed by multiplication with its complex conjugate, and the data is shifted back to the proper range by undoing the altitude-leveling propagation operation. In the sandy Sequim-Bay sight this de-glinting process was found to greatly improve visibility of near-nadir targets. Importantly, in the multi-layered transition region bridging muddy and sandy sediments it also seemed to work well, indicating this is process is not limited to working in a single environment, though it still needs to be tested over a broader range of environments. It also greatly improved the clarity of mosaics generated from the data. Figure 84 shows side-by-side examples of an early iteration of a mosaic-algorithm that did not include a de-glinter, and a subsequent algorithm that did. (Further improvements were added to the mosaicking process, however for the sake of illustrating the effects of de-glinting this illustration is included).

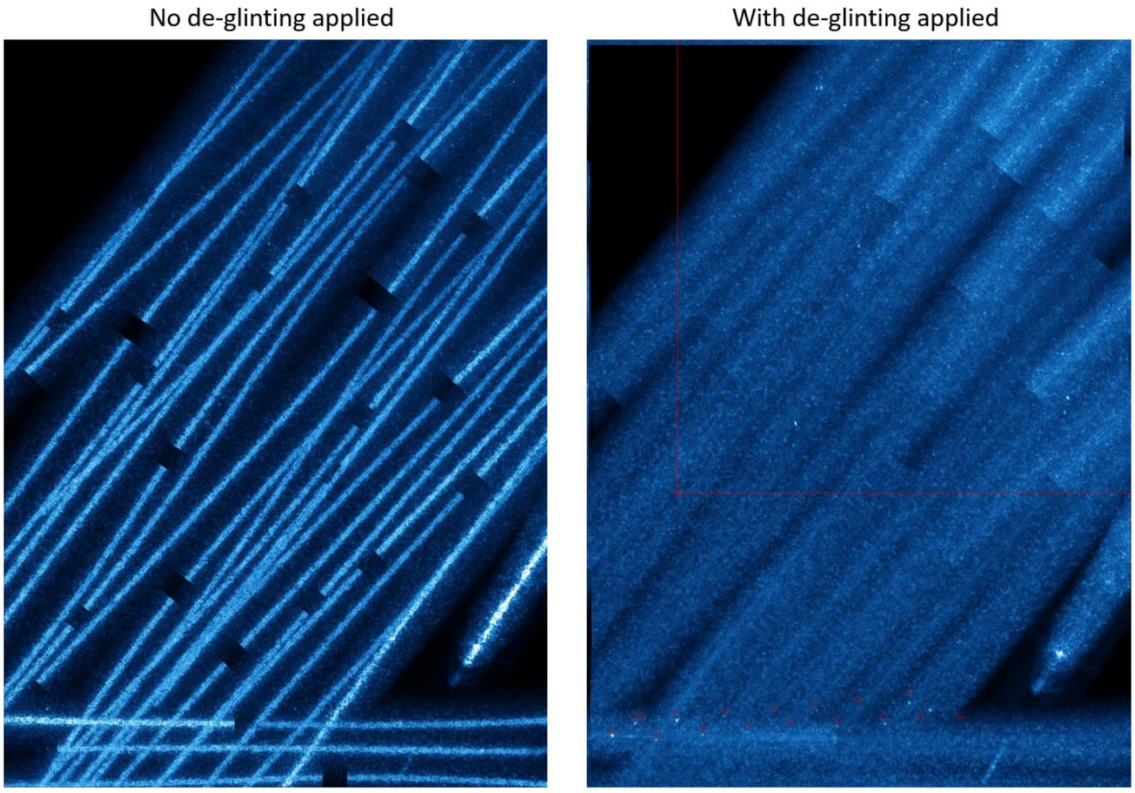

Figure 84: An example of a mosaic generated without the de-glinting algorithm (left) and with the de-glinting algorithm (right) applied. Note that though there are other differences (the figure on the right has additional measures for set continuity and registration applied), the effect of the specular glints is overwhelming.

BOSS data sampling and swath tradeoffs

The speed of sound in water limits monostatic SAS systems from achieving both long range and fast transit speeds; this handicap has been recognized for a long time (see, e.g. the comparison with radar in the introduction of [24], and equation 4 and point #1 in the "summary of results" of [25]). To circumvent this difficulty most deployed SAS systems use an along-track real-aperture array to boost range potential and mapping rate. The BOSS, on the other hand, does not, which means that the edge between being well-sampled (from a synthetic aperture processing

standpoint) and being under-sampled is often passed in the effort to boost range and meet the minimum speed requirements for towfish stability.

Consider the pulse-repetition rate required for adequately sampling a monostatic SAS aperture in which the sonar is travelling at a speed of 2 m/s: for an omnidirectional real aperture system with a maximum frequency of 25kHz a spatial sampling of $\lambda_{min}/2 = 0.03$ meters would be required, however due to an effective wavenumber doubling effect in synthetic aperture systems this value is *halved* again, requiring a spatial sampling of $\lambda_{min}/4 = 0.015$ meters. This equates to a 134 pings-per-second pulse repetition frequency. At this pulse repetition frequency the sonar is limited to a maximum range of only c/134/2, or *5.6 meters.* That is impractical for a surveying system. The difficulty is abated to some degree by the along-track directivity of the receive elements in the array, and thus the standard SAS sampling requirements are often expressed in terms of the length of the receiver elements (see e.g. equations 1 and 2 in [26], most SAS systems are sampled at *D/*2 where *D* = along-track element length). In the present case, the eBOSS along-track element length is ~0.05cm, therefore the pulse-repetition frequency required for adequate sampling could potentially be reduced to 2/0.025 = 80 Hz, extending the maximum range to 9.3 meters. Despite the improvement this is still impractical, and this hasn't even accounted for the fact that the MuST system uses an alternating transmitter design that doubles the across-track resolution at the cost of doubling the PRF rate required for adequate along-track sampling. Clearly, there is a tradeoff between sampling for good SAS reconstruction and range and swath-width, and the end result is that practical operation of the eBOSS accepts some degree of synthetic aperture undersampling – especially at the upper ranges of the frequency band – for the sake of extending range and swath. In many cases, the effects of undersampling can be mitigated by limiting the integration angle; e.g. for UXO survey work the pulse repetition frequency is often set to 60 Hz and the integration angle is limited to ~+/- 35 degrees. At 40 Hz, a pulse repetition frequency that was often used for test surveys of wreckage, aliasing artifacts are noticeable but frequently low enough in level due to sidelobe dispersion (resulting from the wide bandwidth of the system) that they are not extremely noticeable. Figure 85 gives an example of undersampling sidelobes visible in an image of a plane wreck.

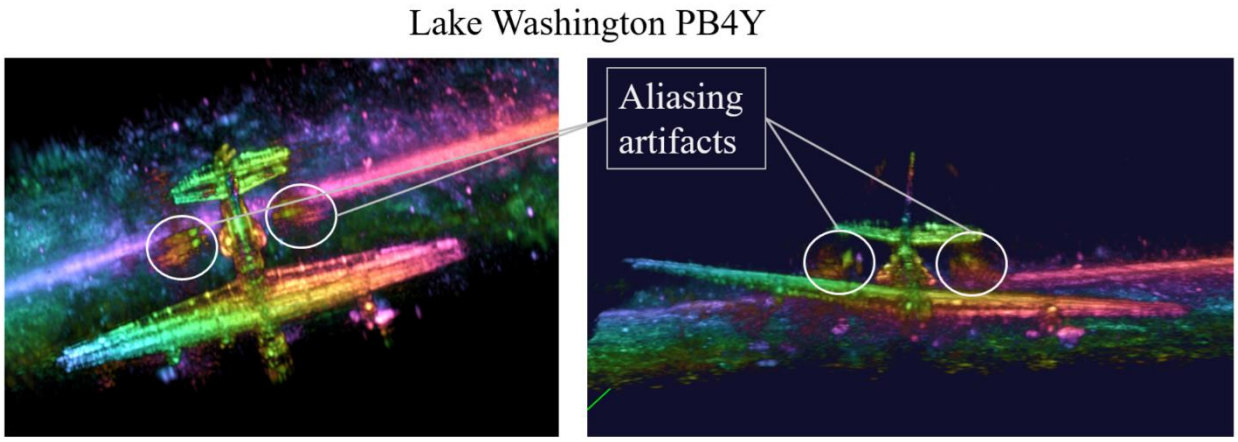

Figure 85: A MuST eBOSS image of the Lake Washington PB4Y captured at a 40 Hz pulse repetition frequency while moving at 2 m/s. The image is clear and well-focused, but some aliasing artifacts are present, resulting from undersampling the synthetic aperture array.

Range-alias imaging

In the previous section the swath-length limitations imposed by pulse-repetition frequency were discussed. During operation of the MuST system a frequently observed phenomena was the presence of high-SNR signals being received from pings prior to the most recently transmitted ping. This often occurred when the MuST system was travelling mid watercolumn when the sediment was out of range of the sonar based on its current ping repetition frequency. In such cases, the received signal seen in real-aperture processing appeared to have a very high signal-to-noise ratio but the SAS image was completely defocused due to an application of the wrong range-migration curvature corrections: the SAS beamformer assumed the signals map to the range [0, R_{max}) (where R_{max} = c/(pulse repetition frequency)/2), but in reality the signals mapped to [R_{max}, $2R_{max}$) or [$2R_{max}$, $3R_{max}$). Due to the high SNR of the backscattered signals the possibility arose that SAS images could be generated in these range-bands, or "zones," despite them nominally violating the "maximum range" limitations imposed by the pulse-repetition frequency. Figure 86 is a graphical illustration of these extra "range-aliased" imaging zones.

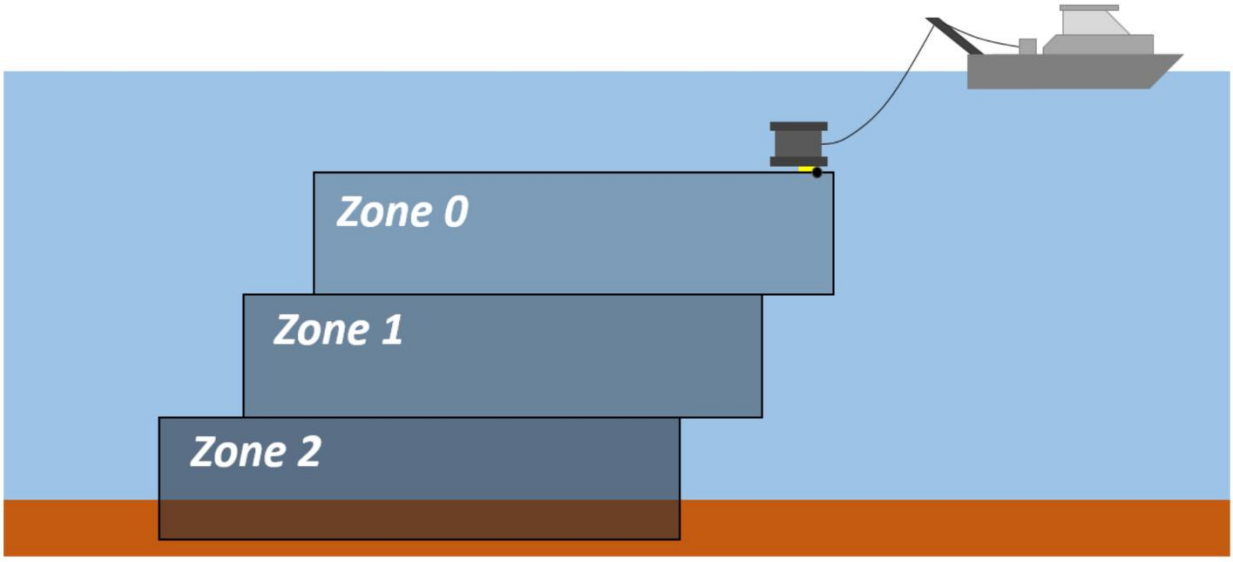

Figure 86: Range zones for alias imaging. Standard beamforming occurs in zone 0, and the maximum range is determined by the round-trip distance travelled by sound as limited by the pulse-repetition frequency. Zone 1 corresponds to the swath of ranges ensonified by the previous ping, and zone 2 from two-pings previous.

The beamformer was modified to detect which zone the sediment interface lies in (based on the MuST altimeter readings), and adjust the beamformer to focus data with range-migration properties corresponding to the appropriate range zone. This option was added to the SAS post-processing code and can be enabled by toggling the "Range Aliasing" switch in the beamforming user interface shown in Fig. 74. Potential applications include hazard detection at safer altitudes and high resolution imaging at standoff distances that reduce the potential for towfish collision. Tradeoffs include reduced across-track resolution (due to the increase range), the abnormal, cylindrical-shell shape of the swath carved zones beyond zone 0, and phase noise from transmitter timing jitter. Despite these tradeoffs, this extension of standard beamforming works

surprisingly well and an example of using this form of beamforming to image a sunken shipwreck is given in the results section.

SAS_view

In many cases the user may desire to examine the volumetric data created by a single eBOSS data file rather than perform batch processing. To facilitate this, an effective visualization tool called SAS_view was developed. SAS_view is automatically opened if the user is performing SAS processing in “Single” mode vs. “Batch” mode (i.e. the toggle on the top of the BOSS SAS post-processing tool). Figure 87 shows the SASview interface presented to the user upon synthesis of a 3D eBOSS image. The GPU-accelerated visualization tool provides rapid volume projection, slicing, scaling, and volumetric rendering features to the user that allow both small and large scale features of scanned objects to be examined.

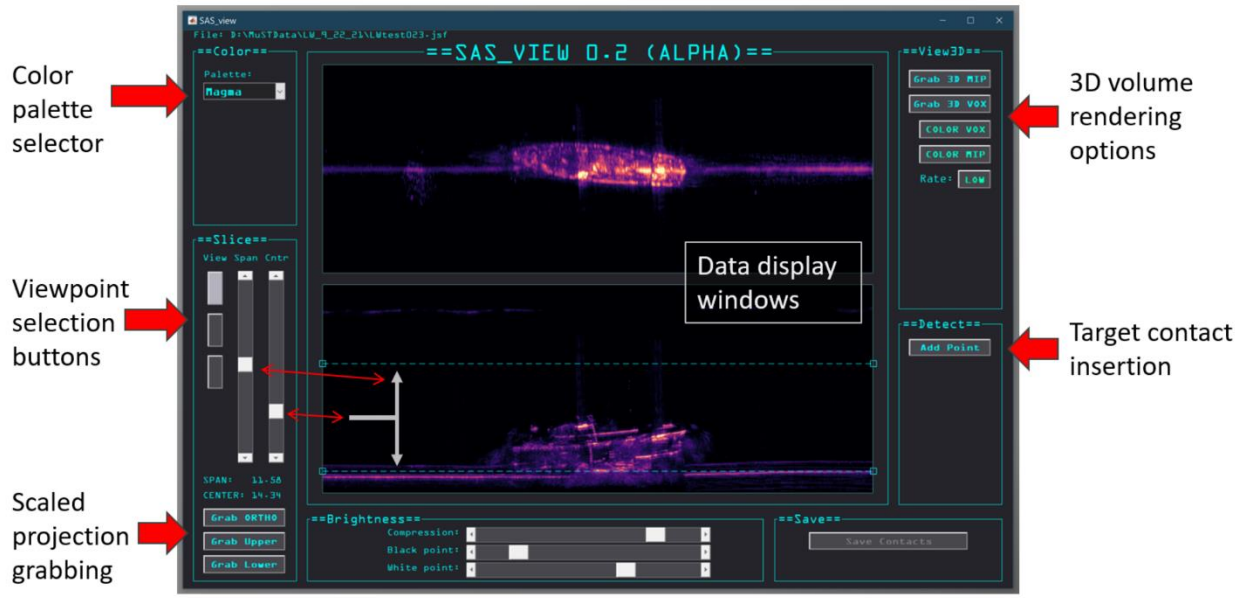

Figure 87: The SAS_VIEW program main screen with some of the features annotated.

The main screen shown in Fig. 87 displays many visualization customization features. The most important feature of the visualization tool is the pair of linked projections in the center of the screen. The top figure always represents the projection through the vertical dimension of the projection shown in the bottom screen, *gated by the dashed blue lines*. As indicated by the red arrows, the span and the centroid of the dashed lines are linked to the position of the sliders on the left-hand column. Figure 88 shows examples of different slider settings that 1) show a projection through the entire exposed region of the ship, or 2) show a projection through the sediment beneath the ship revealing an acoustic shadow. The white-point has been lowered to increase the perceptual brightness of the ship.

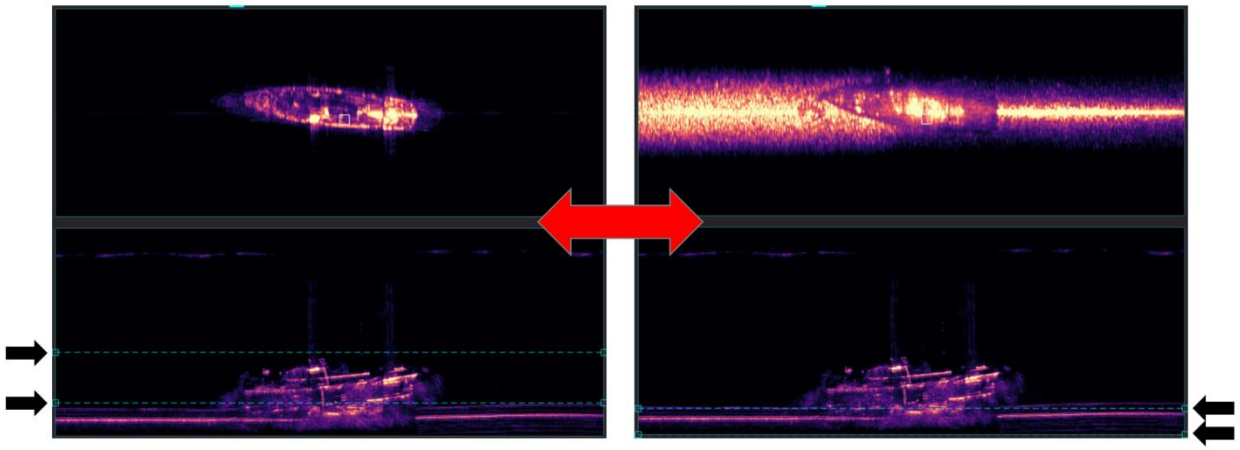

Figure 88: Demonstration of the linking between the projection in the top window and the spacing and location of the projection limit lines in the bottom window. In the figure on the left the projection lines (annotated with the two black arrows on the left) bound the exposed portion of the shipwreck. On the right, the projection lines bound the bottom of the keel and the sediment below the ship. The acoustic shadow produced by the ship is visible in the projection.

The projection on the top of the display is updated continuously as the slider values are shifted, giving the user active feedback about the projection being cast through limit lines. The three gray buttons on the left side of the panel modify the axes of projection for the images in the display panels. The default axes of projection are through the across-track axis (bottom) and through the vertical axis (top). By selecting the middle button the projection axis of the top screen is changed to the *along-track* axis, and the sliders on the bottom window are now vertical rather than horizontal. The third (lowest) button modifies the upper projection dimension to be in the across-track dimension. To accommodate adjustment of the projection limit boundaries the projection axis of the lower image is modified to be along the height (z) dimension. The user is therefore given the ability to perform customized projections through each of the three orthogonal image coordinate dimensions. Examples are shown in Fig.'s 89 and 90.

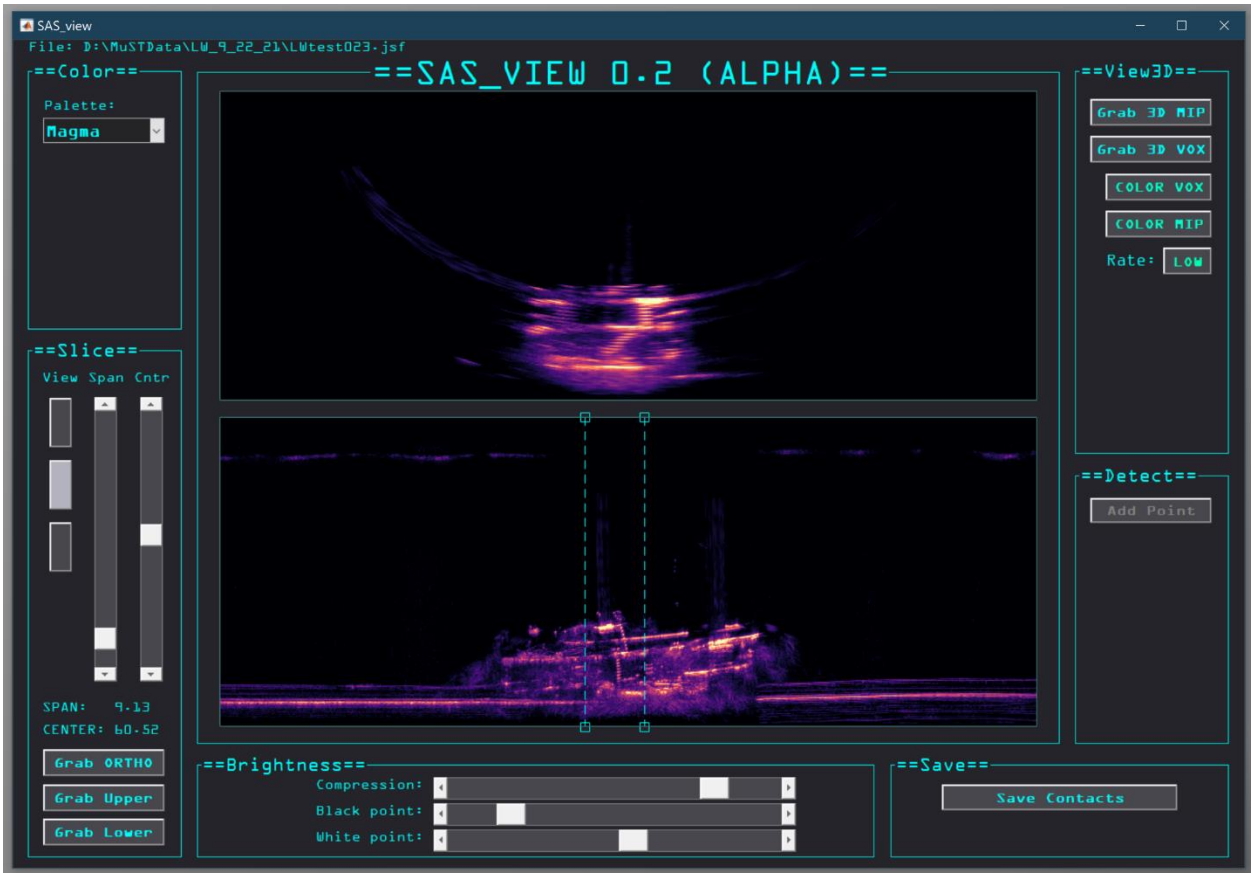

Figure 89: Modifying the axis of projection to be in the along-track dimension by selecting the central view-point button in the left-hand column. The projection-limit lines now become vertical dashed lines and are visible mid-ship. In the upper projection window different decks of the ship and ladder-like features become visible.

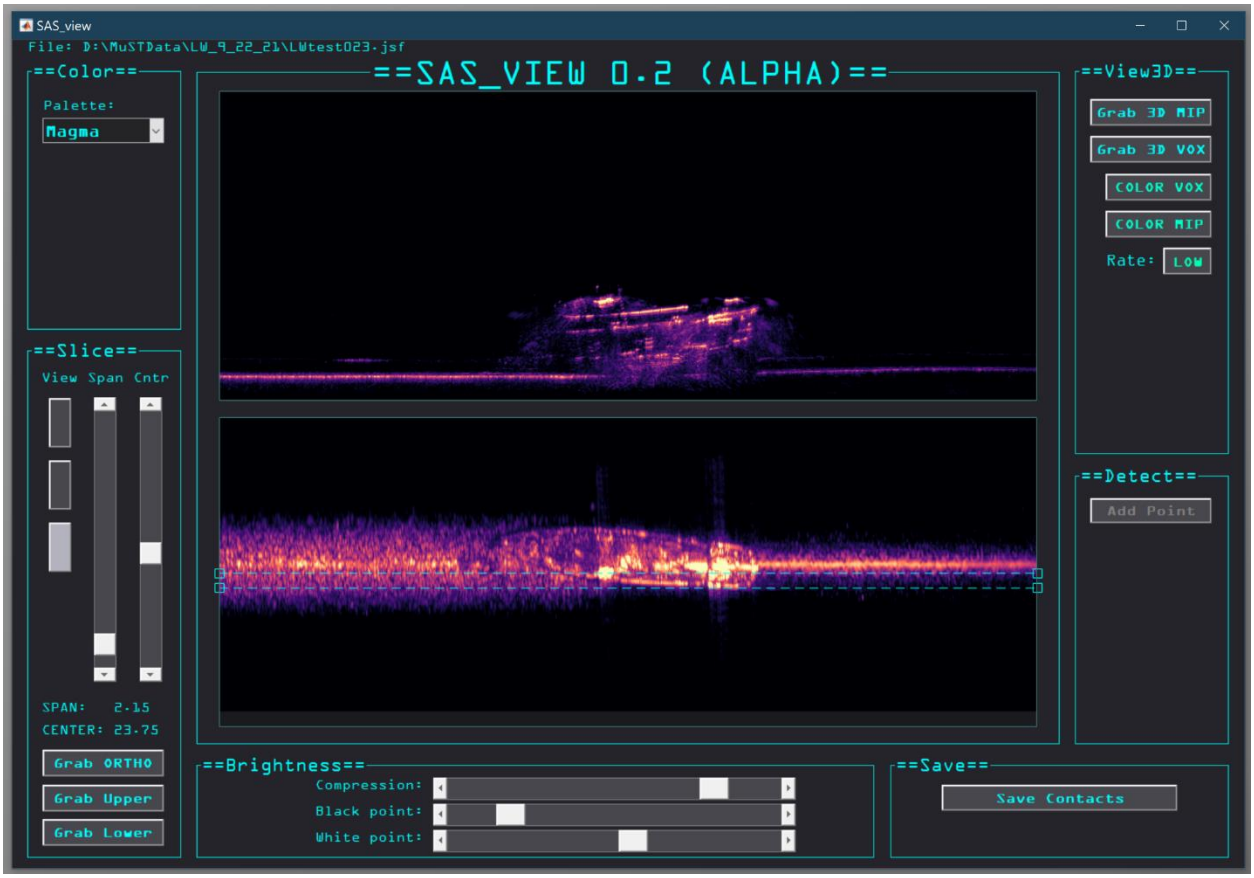

Figure 90: The third (lowest) viewpoint button changes the upper projection axis to be in the across-track dimension. To accommodate adjustment of the projection limits the bottom projection axis is altered to the z (height) dimension. The projection limit lines are visible as horizontal dashed lines in this viewpoint. The limit lines are currently set to isolate the port-aft section of the wreck highlighting a deck rail and some lower-deck features.

SAS_view also provides the ability to perform arbitrary orientation 3D volume renderings of user-defined regions of interest using the options shown in the top-right column of the main screen. Selecting “Grab 3D MIP” will highlight the upper projection. The user can drag a box in the upper projection to determine the limits of the region-of-interest; the limit along the third axis is determined by the projection limit set by the user and visible in the lower window. Figure 91 shows an example of the amidships region of the wreck being selected for volume rendering. The volume rendering appears in a new window.

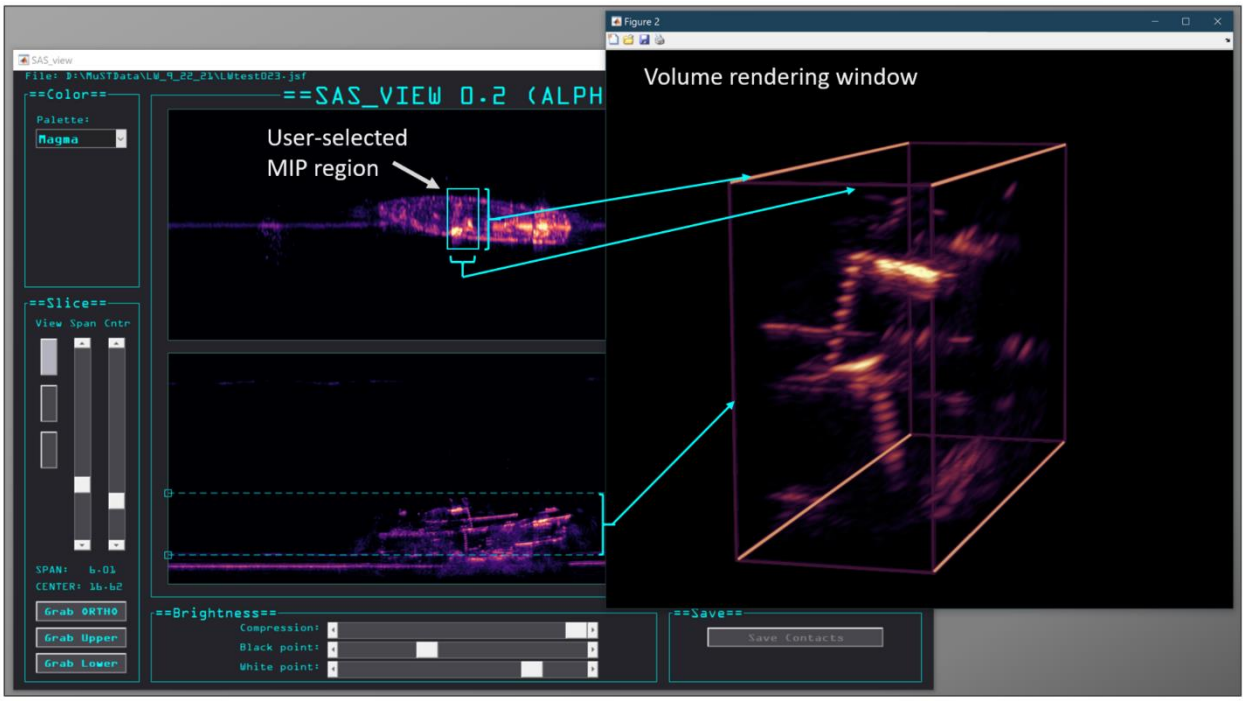

Figure 91: A 3D maximum intensity projection rendering created from bounds specified by the user after selecting the "Grab 3D MIP" option in the upper right corner of SAS_view. The blue annotation arrows show the relationship between the user selections and the dimensions of the rendered region.

The "Grab 3D MIP" option leverages Matlab's internal 3D rendering engine which is powerful but limited in some key ways: it doesn't allow dimensional units to be shown in the image and it doesn't allow a decoupling between voxel coloration and intensity. This is problematic if the user would like to vary voxel color with altitude, for example, to simplify feature differentiation. Another problem with maximum intensity projections is that feature depth becomes very difficult to ascertain without any sort of dynamic figure rotation occurring (even in these instances it can still be difficult). Alpha blending, which renders voxels after assigning each a transparency value, helps overcome this depth issue. Clicking "Grab 3D VOX" generates an alpha-blended volume rendering of the user's specified region of interest. This selection calls "VOXview," a Matlab-based volume renderer written as part of this project that overcomes some of the limitations in Matlab's internal volume rendering engine. Most importantly, it allows voxel intensity, transparency, and color values to be defined completely independently from each other. It also allows units to be included in the rendering axes. Figure 92 shows an example of a VOXview rendering of a UXO object captured during the 2021 Sequim Bay testbed survey, selected via the "Grab 3D VOX" button (left) and the "COLOR VOX" button (right), which rotates the color-scheme used by the renderer based on the height plane of the voxel.

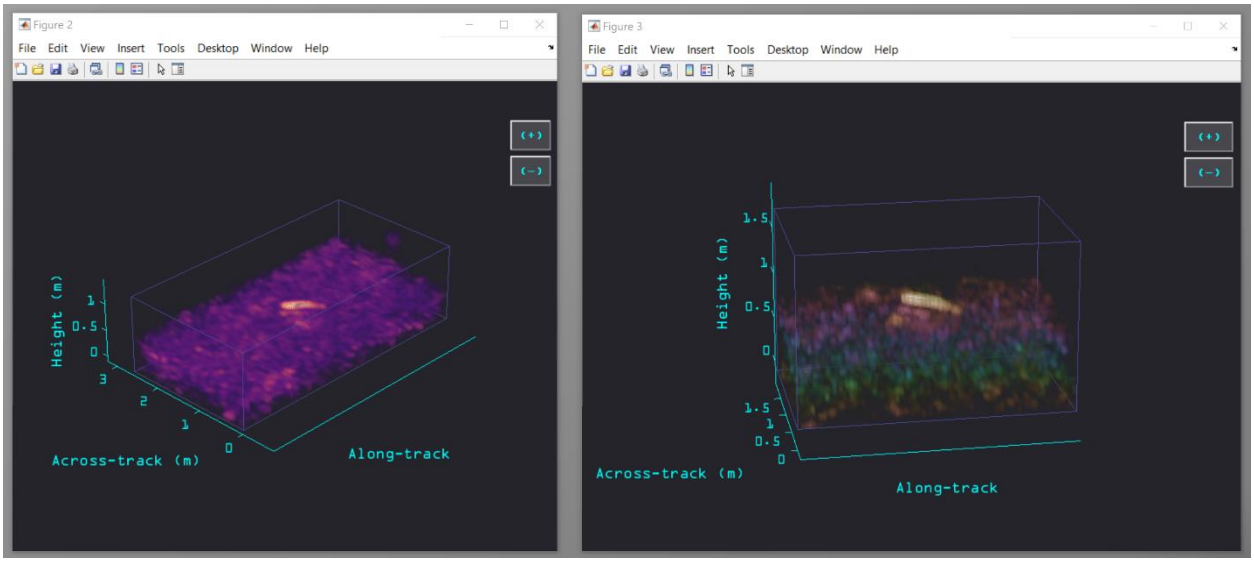

Figure 92: Two volume renderings of a UXO generated in SAS_view using the "Grab 3D VOX" button (left) and the "COLOR VOX" button (right). The colorized option rotates color vs. height.

In Fig. 92 the height and across-track axes show length scales. In contrast the along-track axis does not. This is due to the fact that the 3D SAS volume is gridded uniformly in these dimensions but is not gridded uniformly in the along-track dimension: the sample spacing is defined by the vehicle speed divided by the pulse repetition frequency. Estimated length scales in the along-track dimension can be determined using the "Grab ORTHO" button that will be discussed shortly. Two buttons are visible in the upper right corners of the rendering screens shown in Fig. 92. These buttons zoom in or out on the rendered image.

Making colorized maximum-intensity projection images is more challenging because of Matlab's lack of support for independently defining voxel color and intensity in its internal volume renderer. To overcome this a color-MIP renderer was written that generates a colorized volume using the same colorizing technique used in "COLOR VOX" (Fig. 92, right plot), but the RGB color channels are independently rendered using Matlab's internal rendering engine and the projections are combined to form a 3D color MIP. This kind of image can be generated using the "COLOR MIP" button and an example for the same target is shown in Fig. 93.

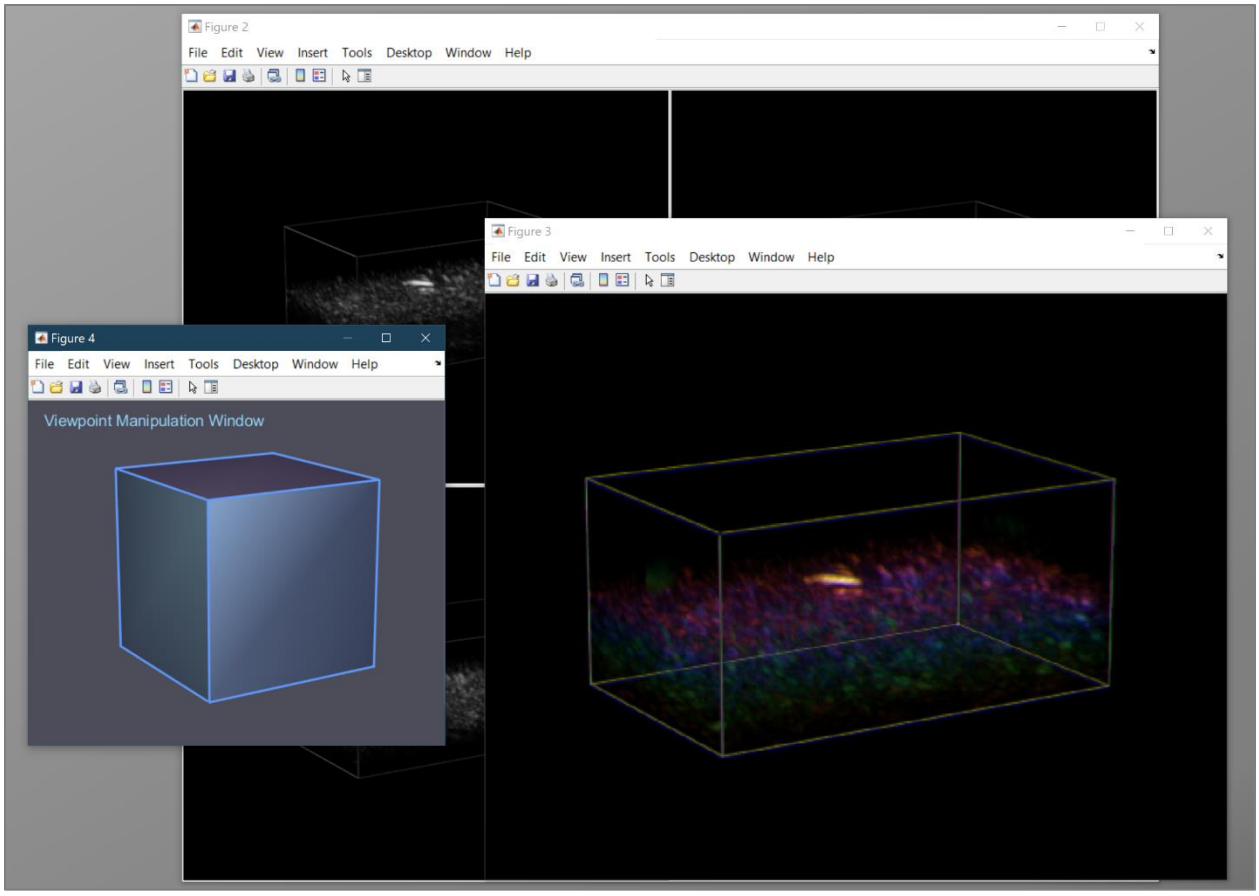

Figure 93: The COLOR MIP volume rendering tools opened by clicking the "COLOR MIP" button in SAS_view and selecting a region of interest. The box in the left-most window, which can be dragged by the user's cursor, modifies the viewpoint perspective. The three color-channel images used to create the RGB composite are visible in the background.

The "rate" button underneath the volume rendering options determines the rotation rate of the color scheme as a function of height through the region of interest: the default value is "low" which creates a single cycle through hue. Toggling the button to "high" results in two full color rotations being cycled across the height dimension of the volume.

For dimensionally accurate renderings the user can select from the buttons on the bottom right of SAS_view: "Grab ORTHO," "Grab Upper," and "Grab Lower." If the first is selected then the user is prompted to select a box in the upper window (as with the previous volume rendering options). In this case the volume renderer is not called but the 3 orthogonal views through the selected volume are shown and the along-track axis is scaled based on the average sample spacing of the SAS array in the along-track dimension. An example is shown in Fig. 94. Selecting "Grab UPPER" opens the upper plot in a new window (example shown in Fig. 95) that allows the user to zoom, pan, etc., and gives the user the option of rendering the projection with equal axis scaling by setting the toggle button to "Axis mode: EQUAL". "Grab Lower" does the same thing but with the lower, rather than upper plot.

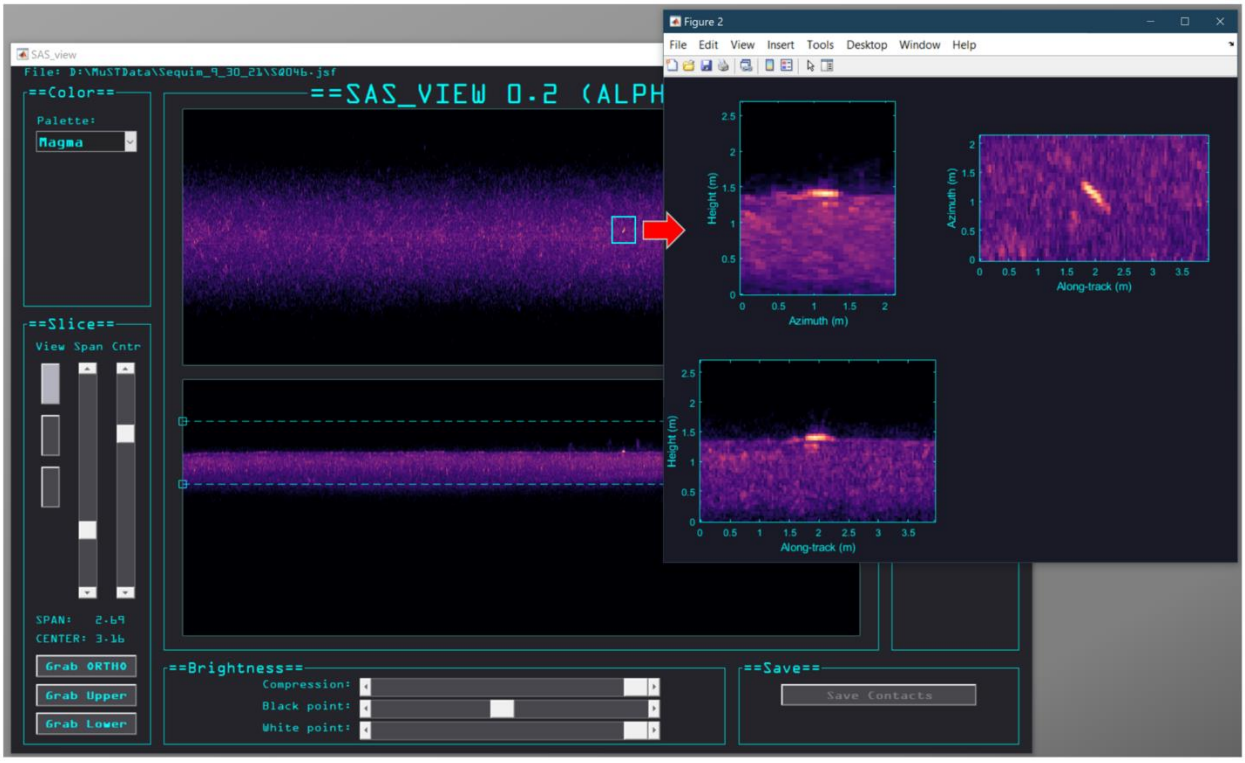

Figure 94: An example of the visualization produced by selecting "Grab ORTHO" from the bottom left tool panel and selecting the region outlined by the cyan box. Three orthographic projections are plotted with units on uniformly scaled images.

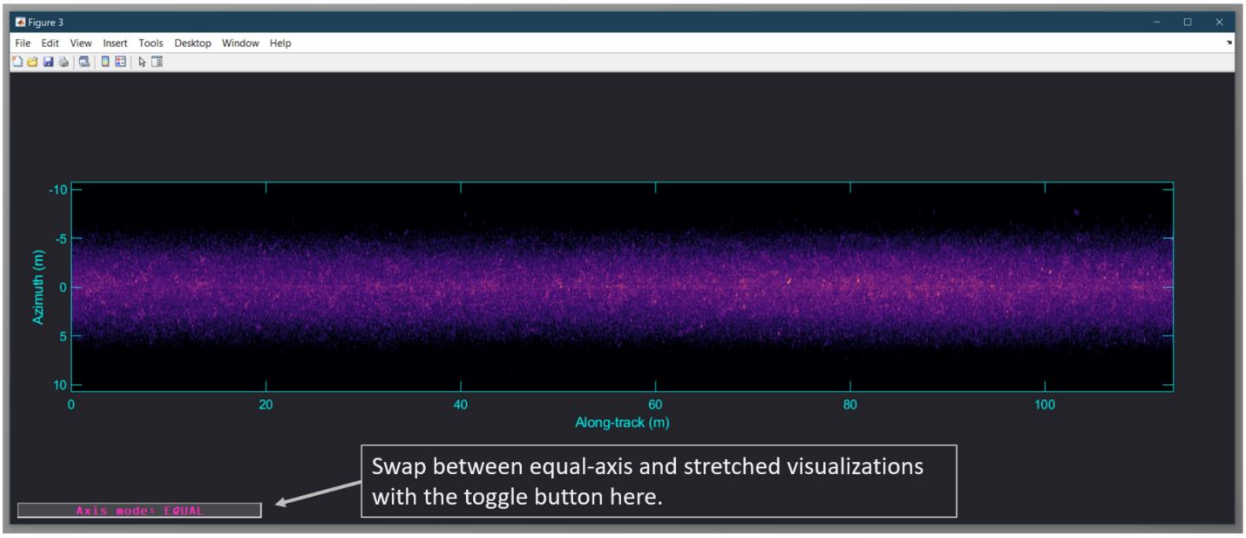

Figure 95: An example of the visualization produced by "Grab Upper" with the Axis mode set to equivalent scaling. Note that de-glinting was applied to this dataset and the distracting Nadir-glint has been removed.

Finally, SAS view also gives the user the option to select detection points in a similar manner to the waypoint selection tool included with MuSTLive. This can be done by clicking the "Add Point" button. This button is only enabled when the first viewpoint is selected so that a latitude and longitude can be defined for the selected point in the upper window. As with the MuSTLive

waypoint selection tool, the user can define the object's class and add a label, however SAS_view has not been developed to the extent that the marker can be exported to a list yet and the "Save Contacts" button remains disabled. Adding this functionality, along with additional data export options, is planned as a part of a follow-on effort.

The Williams Detector

In a collaborative effort with the Penn State Applied Research Laboratory, (ARL-PSU, SERDP project MR21-3543, PI: David Williams), a 3D detector was developed for locating UXO-scale objects in the volumetric 3D SAS datasets. This detector was integrated into the batch processing SAS imaging routine described in previous sections and results in a useful detection list of points. Figure 96 shows an example of a geo-rectified MIP computed along the height axis (i.e. a Z-MIP) accompanied by the detections produced by the Williams detector.

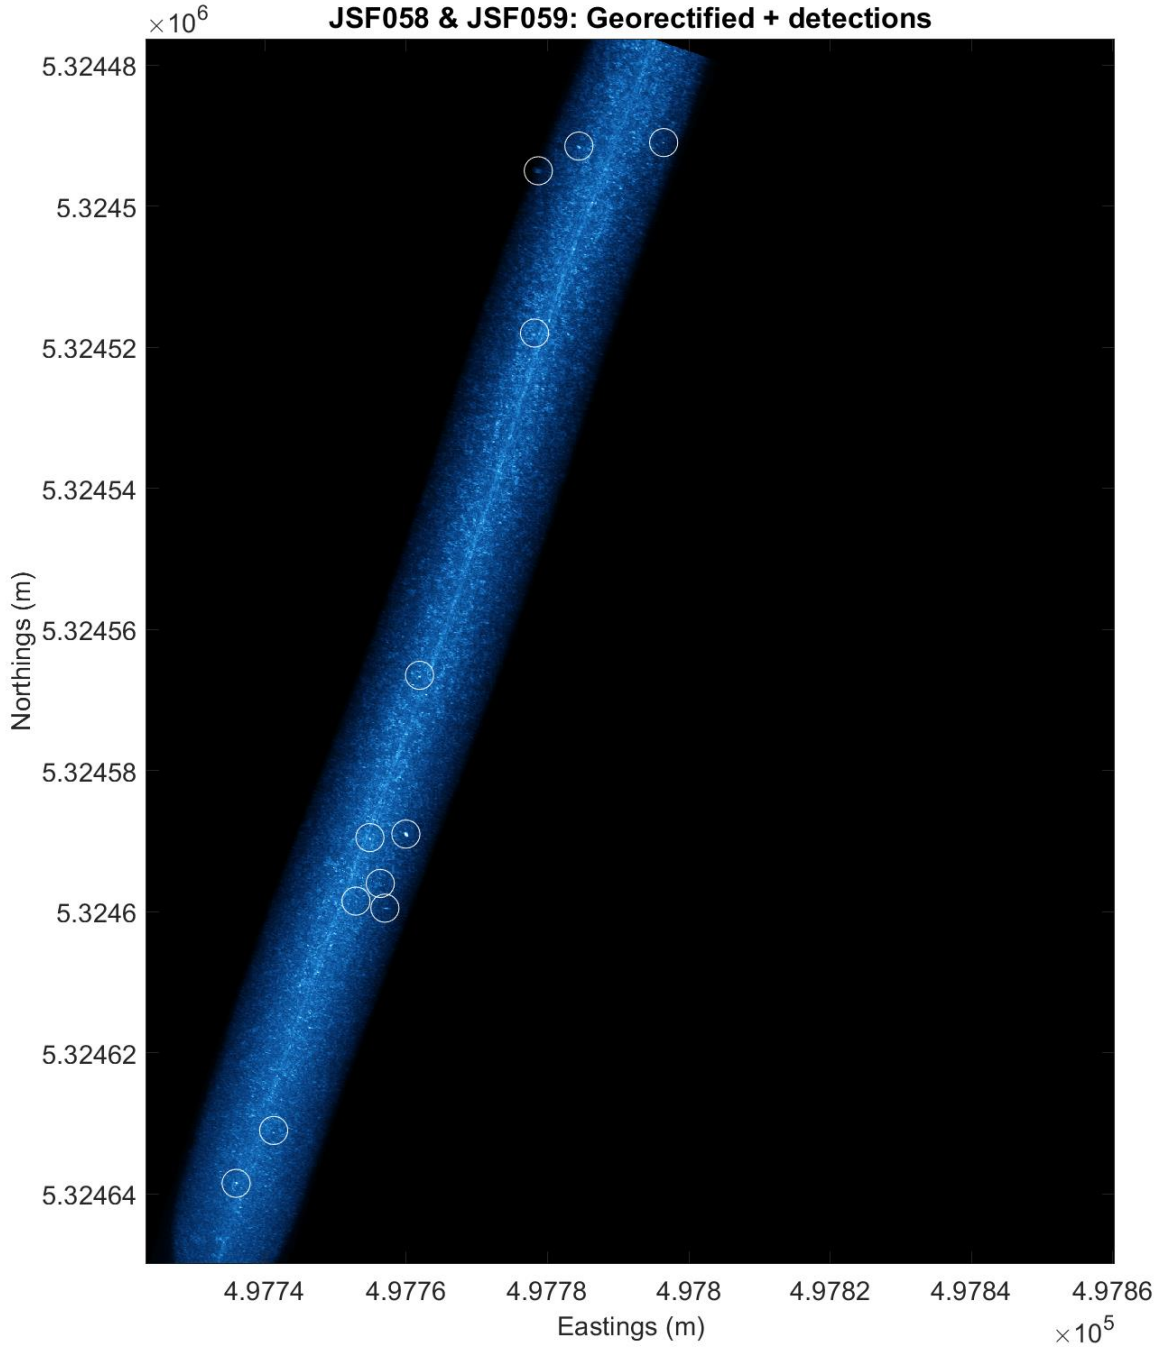

Figure 96: An example of the Williams detector results overlaid on the geo-rectified scan from which the detections were made. The detector operates on the volumetric data rather than the projected data and can be used to detect objects that are not easily detectable in a MIP projection that spans the entire scanned sediment volume.

Results from using this detector are very interesting, especially if the threshold for detection is set low, and will be described in the results section of this document.

Feature-based renavigation

Motivation

The Buried Object Scanning Sonar is a low frequency synthetic aperture sonar and the objects of interest – unexploded ordnance – typically lie in the "low *ka*" regime, (a = target radius, $k = 2\pi/\lambda$ where λ is the acoustic wavelength), intuitively meaning that the size of the objects being imaged are similar in size to the acoustic wavelengths being used for imaging. For example, the center frequency of the eBOSS sonar used on the MuST is approximately 15 kHz, which has a wavelength in water of about 10 cm. The cross-sectional radii of many UXO are on the order of 4 to 8 cm. In the low-*ka* regime meaningful shape information cannot be gathered from rough-surface scattering, like it typically is in very high frequency applications: the surface appears smooth to the signals at these wavelengths. As a result, rigid structural information is found in the form of specular components which are frequently only visible from narrowly defined observation angles. This was previously discussed in the background section of this report and is demonstrated in figures 2 – 4 of that section. Pertaining to BOSS post-processing, at low frequencies the data acquired by individual target scans becomes more meaningful in the context of other observations, and the ability to combine observations from different acquisitions is extremely powerful for object identification. Figure 97 shows an example of a multi-aspect data product captured for an object of interest in the 2021 Sequim Bay UXO textbed. Individual looks at the target may be useful, but the combined product, which combines intensity data in a color-by-aspect scheme [27], clearly shows more shape information, less speckle, and less volume noise.

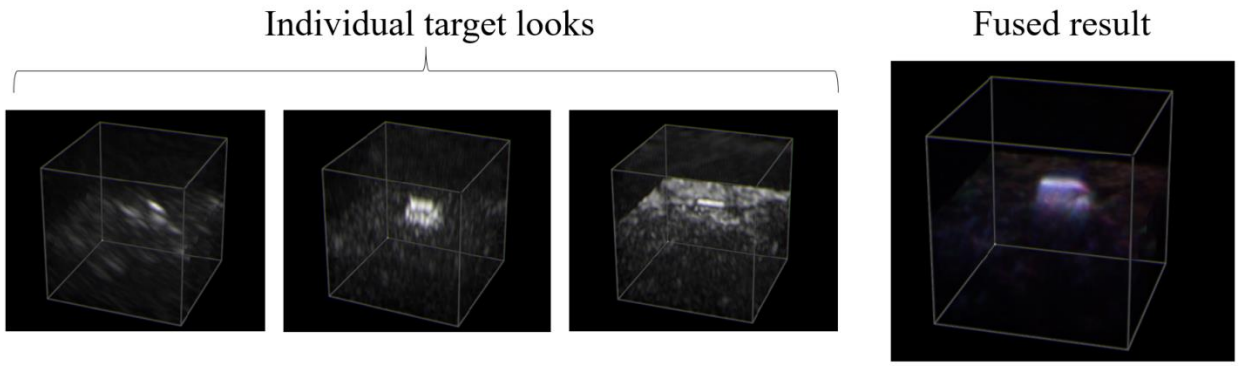

Figure 97: A sampling of some looks at a target captured in the 2021 Sequim Bay UXO test bed (left) and the result of fusing the set in a color-by-aspect imaging scheme (right). The taper at the nose of the UXO is visible in the fused data-product shown on the right.

Even at ordinary imaging frequencies of several hundred kilohertz, the improvements to classification gained by fusing data from multiple looks has been recognized for quite some time (see e.g. [28 – 32]), and this improvement is especially noticeable at low frequencies due to the limited angular range over which particular features may be excited.

A well-known difficulty with multi-aspect information fusion is the associated *registration* problem: the advantages of combining multiple datasets to piece together different features of an object are hindered by the navigation precision of the sonar acquisition system [32]. Unless the navigation precision is better than the size scale of the features being re-assembled then the geometric relationships between features in the different looks are lost due to the random translations imposed during image combination. An example of this problem is shown in Fig. 98, in which multiple looks of a target acquired by the MuST system have been randomly translated to simulate decreasing levels of navigation error.

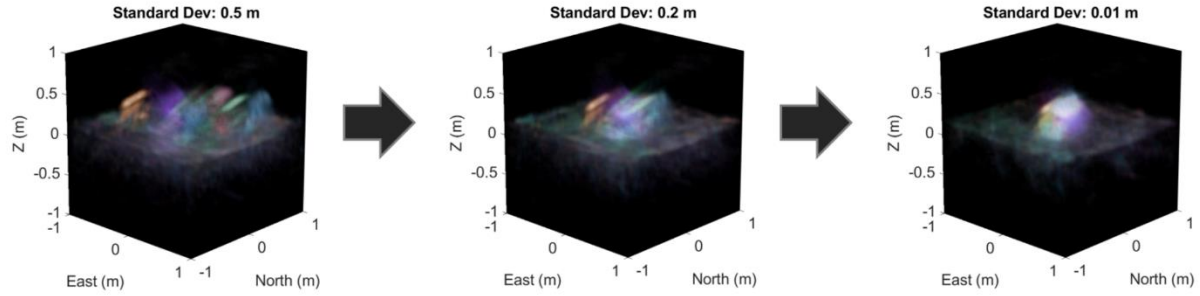

Figure 98: A multi-look fusion example showing the results of combining looks with progressively better navigation accuracy.

In Fig. 98, individual looks may contain shape information but with a navigation error of only 0.5 meters (standard deviation) the resulting combination adds no additional value for object identification. Even with a standard deviation of 0.2 meters the combination adds very little value. Only when the error is much smaller than the size of the target (a 30 x 60 cm solid aluminum cylinder) does the shape of the object – a solid cylinder – become readily apparent after the fusion process.

Even for spectral (e.g. acoustic color) analysis, which is less affected by translational alignment issues in the spatial domain, alignment has the advantage of being able to consistently apply spatial masks to reduce volume noise, so navigation refinement to sub-feature scale is almost universally beneficial to the classification process.

Unfortunately, the precision of the navigation systems used by the MuST system, or even generally commercially available for underwater applications, are not precise enough for this type of combination: the standard deviation of the alignment errors between images should be on the order of a couple centimeters. In the contexts in which the MuST system is operated, this type of precision positioning must be accomplished using the observed features of the sediment themselves. This section describes the methods by which this "feature-based" alignment is performed. Before elaborating on the algorithm further several points must be clarified:

- The algorithm was developed solely from data captured at Sequim Bay. Data-driven algorithms are notoriously environment-specific and no efforts have been put forward to rigorously test the algorithm across multiple environments and expand robustness

towards different environmental variations. This type of expansion could benefit from multi-sensor fusion and would be recommendable for follow-on efforts.

- From a software standpoint this is the least developed portion of the MuST Toolbox. Individual Matlab scripts have been written to perform the operations described in this section but there is not yet a user-interface that has been written to simplify and streamline this process for an end-user.

Despite the previous two caveats, the data-driven algorithm developed in this project worked consistently and reliably during field operations at Sequim Bay and on a nightly basis, following batch SAS data processing, data-driven mosaics from the previous days' efforts were provided each morning to inform subsequent survey operations.

Algorithm overview

Image registration is a common problem and many solutions combine standard feature detectors such as SIFT [33] or SURF [34] with a robust parameter estimation technique like RANSAC [35]. It is worth noting that SURF and SIFT features are intended to be scale invariant and at least somewhat rotation invariant. In the present context this adds two degrees of freedom that BOSS data mosaicking doesn't require: scaling is consistent and rotation is known. Even standard issues that accompany most sonar registration problems, such as determining the correct bathymetric plane for projection, are avoided: projections accrue no lateral shifts from height-plane projection errors because they represent a direct projection in the height dimension through a volume. MuST data registration is therefore a much simpler data-driven task than many, and reduces to a rigid-translation optimization problem. In this special case the very simple two-dimensional cross-correlation operation becomes a robust option for making data-driven region associations and alignment estimations; especially combined with appropriate pre-processing to normalize spatial-spectral power distribution. The feature-based re-navigation algorithm developed for BOSS MIP layer alignment follows the procedure outlined in the following flowchart.

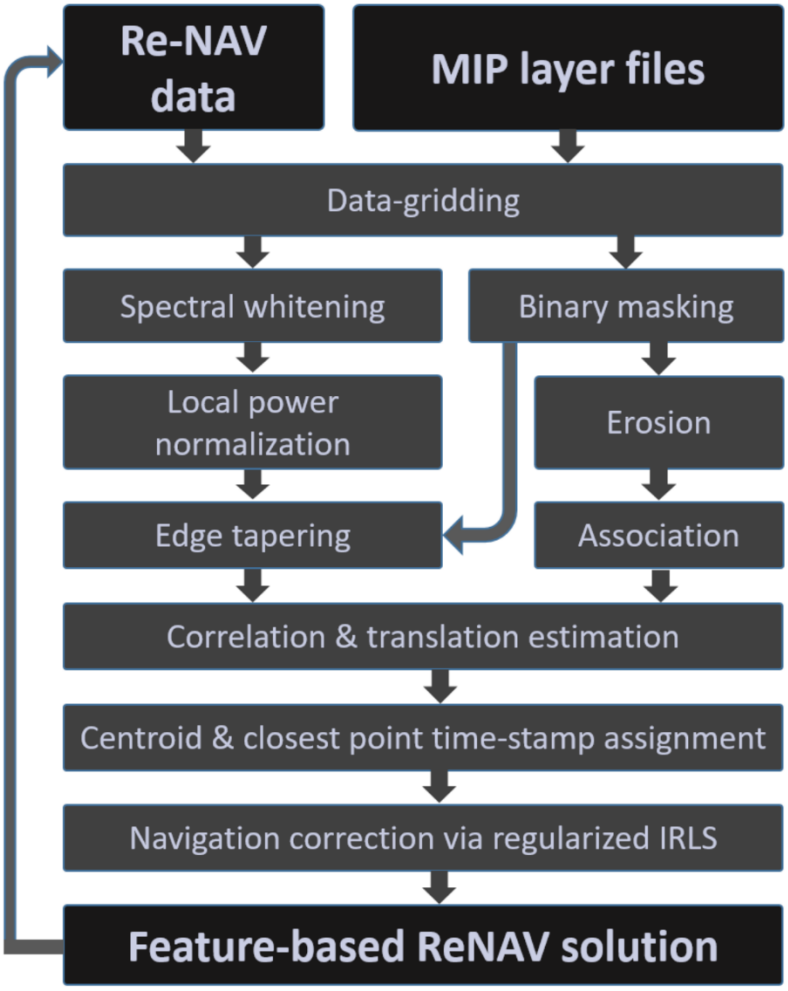

Figure 99: An outline of the renavigation algorithm developed for working with BOSS data.

Data gridding

The geo-rectification process used in the present approach generates a grid of points in UTM coordinates over the target field with a specified resolution d_{xy} in northings and eastings. This pixel size is intentionally specified to be slightly larger in size than the along and across track sample densities of the MIP layer data. Currently, d_{xy} is set to 0.05m.

Layer data is unpacked from the hard-drive and the location of MIP layer pixels in northings and eastings are calculated using translations based on the re-navigation data and heading information for the appropriate timestamps. The user specifies which layer to use for feature alignment and this layer is projected onto the grid via spatial binning. Figure 100 is an example of a scan gridded to geo-rectified coordinates using this method.

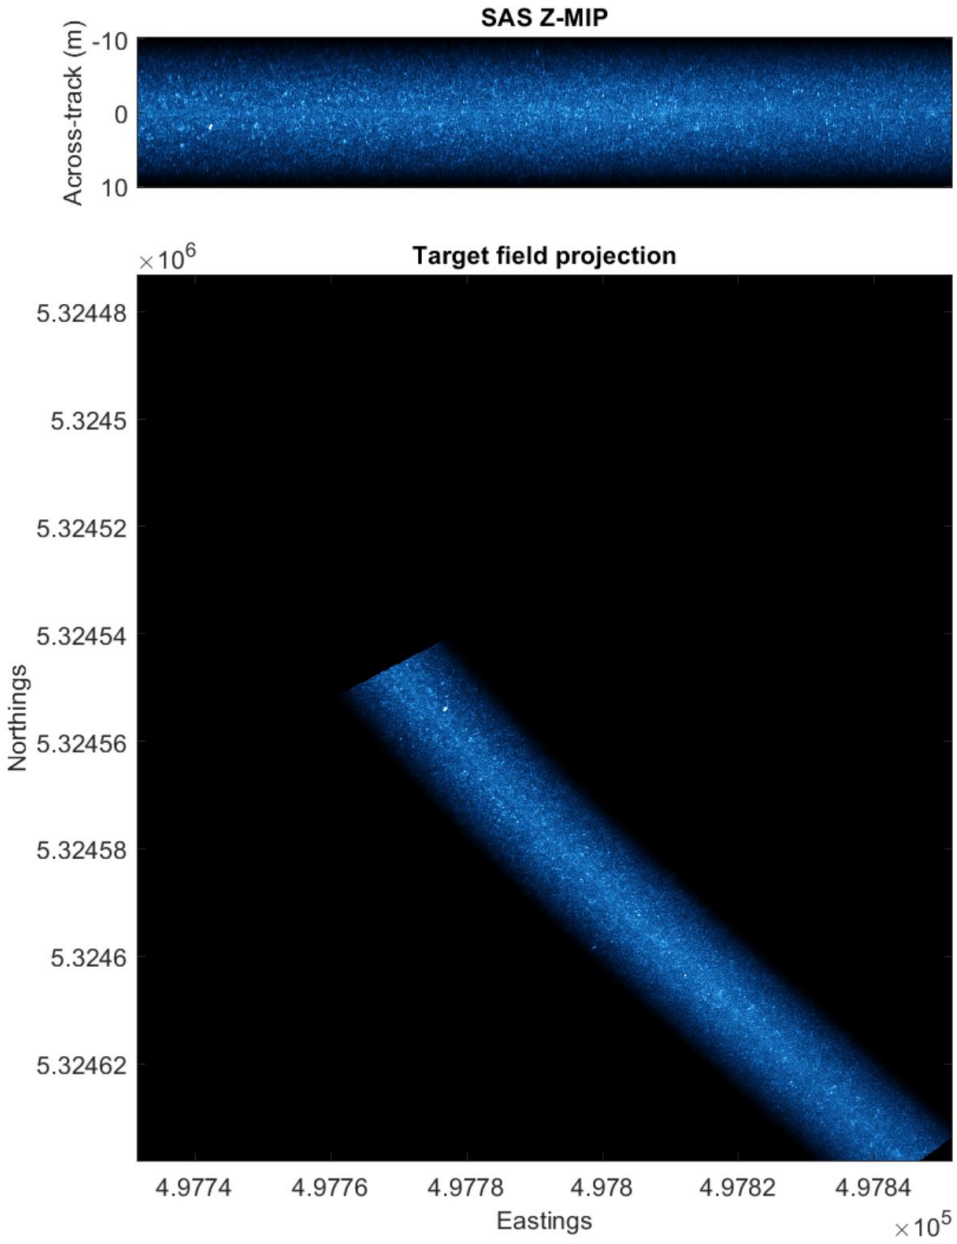

Figure 100: Layer data stored for a JSF file (top) and following projection onto target field grid (bottom).

Correlation pre-processing

Prior to performing cross-correlations for alignment estimates between scan projections the data must be pre-processed to prevent correlation peaks from being biased by unwanted features. Potentially biasing features include 1) low-frequency trends associated with beam patterns rather than sediment features, 2) very bright non-stationary scatterers, and 3) artificial features with high spatial frequency content like scan cutoff points, such as the end of the sonar path in the dataset projected in Fig. 100. Each of these problems are addressed by the pre-whitening (low-frequency trend compensation), local normalization (bright non-stationary features), and tapering

(artificial data cutoff) steps in shown in Fig. 99. A before and after comparison for gridded data after it has undergone this pre-processing step is shown in Fig. 101.

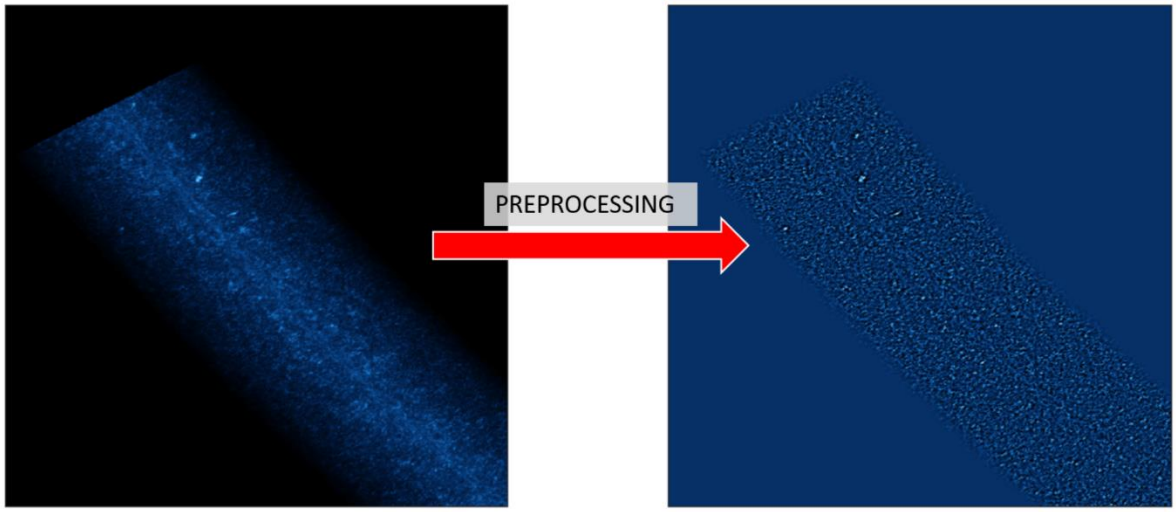

Figure 101: the effects of pre-processing on gridded MIP layer data.

The pre-processing has the effect of turning small sediment features into a high spatial-frequency barcode that gives proper alignment between alternate scan patches a high score vs. correlation background noise.

Masking and scan association

The naïve approach toward alignment would be to correlate each gridded scan within the target region with all other scans. If several hundred scans have been acquired then this could result in tens of thousands of correlation operations, many of which are pointless because of a lack of overlap between scans.

To trim un-necessary pairwise correlations from being performed, binary masks for each scan are generated that have pixel values of 1 for all regions of valid sonar data, and zero elsewhere. Overlap between scans can be determined by multiplying masks for different scans and seeing if the result has any non-zero pixels. The drawback to this approach is that overlap may be detected even if the overlapping region is too small to generate a significant correlation peak. To improve the likelihood that overlap detections will result in enough common sediment area to generate a decent correlation value, the morphological erosion operator is applied to the binary masks prior to multiplication. Erosion essentially performs a boundary shrink around the valid sonar region. Examples of a mask and an eroded mask for the gridded scan in Fig. 100 are shown in Fig. 102.

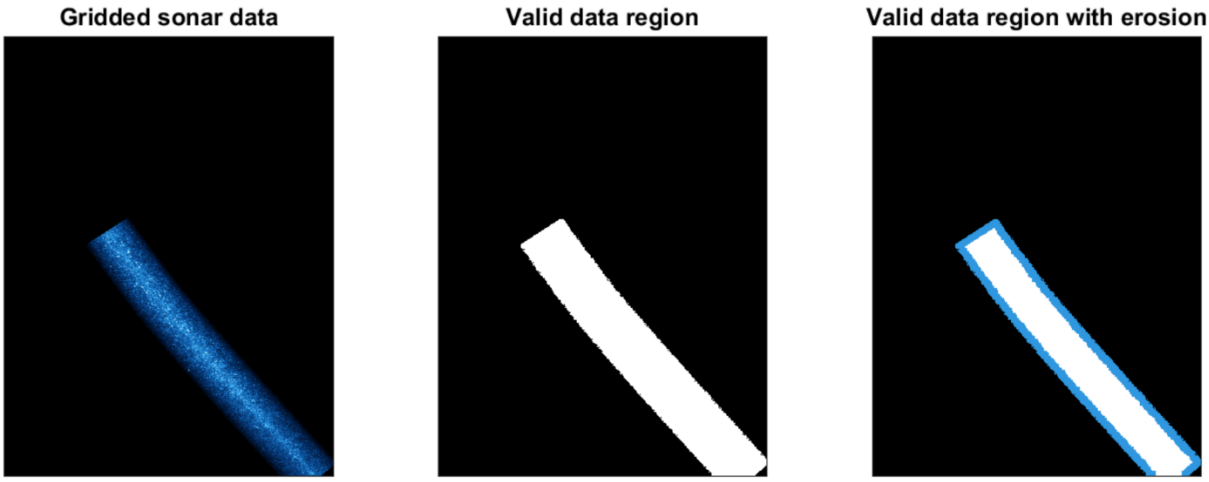

Figure 102: Gridded sonar data (left) and the binary mask representing the valid data region (middle). The white region in the figure on the left represents the valid region used for overlap detection after the blue bounding area has been eroded.

In the present case the erosion radius has been set to 2 meters. This value was chosen based on an estimate of a 2 RMS positioning error for the system so that if overlap is detected there should be a high likelihood of having at least a 2-meter wide common region between scans with detected overlap.

Correlation and translation estimates

Correlations between images *i* and *j* are computed using two-dimensional Fourier transformations, i.e.:

$$\boldsymbol{CC_{i,j}(x,y) = \mathcal{F}^{(-2)}\{\mathcal{F}^{(2)}\{S_i(x,y)\} * conj(\mathcal{F}^{(2)}\{S_j(x,y)\})\}} \quad \textbf{Eq. 57}$$

In (57), $\mathcal{F}^{(2)}$ and $\mathcal{F}^{(-2)}$ represent the forward and backward 2D Fast Fourier Transformation operations. The standard normalization component, $\sqrt{\sum|S_i|^2 + \sum|S_j|^2}$, which is typically placed in the denominator of (57) to arrive at a correlation coefficient value ranging from -1 to 1, is omitted in the present case because it is almost guaranteed that only a small fraction of the scan regions are overlapping, and this normalization factor will normalize using the spatial information from the whole scene rather than just the unknown region of overlap. A quality estimate is critical for subsequent navigation refinement, however, so to get a quality estimate from the correlation peaks the ratio of the peak versus the local absolute mean is calculated by dividing the correlation data by the absolute value of the correlation data averaged over an 8 x 8 meter uniform 2D rectangular window. Lastly, a search-radius is applied that zeros all correlation data outside of a specified radius (6 meters in the present case), limiting the region over which correlation peak maxima can be reported. This normalization and search limiting process can be described via (58):

$$\boldsymbol{CCnorm_{i,j}(x,y) = M(x,y) * \frac{CCnorm_{i,j}(x,y)}{\mathrm{F}_{8x8m}\{|CCnorm_{i,j}(x,y)|\}}} \quad \textbf{Eq. 58}$$

where $M(x,y)$ is a binary mask zeroing all cross-correlation data corresponding to x- or y-shifts greater than 6 meters in the x or y directions. The quality of the shift estimate is the peak value of *CCnorm* and the shift values are calculated by mapping the x- and y-indices of the peak value of *CCnorm* to displacements via the pixel resolution value d_{xy}. Figure 103 plots the Euclidean distance of the translation estimate vs. quality for all scans in one example field operation:

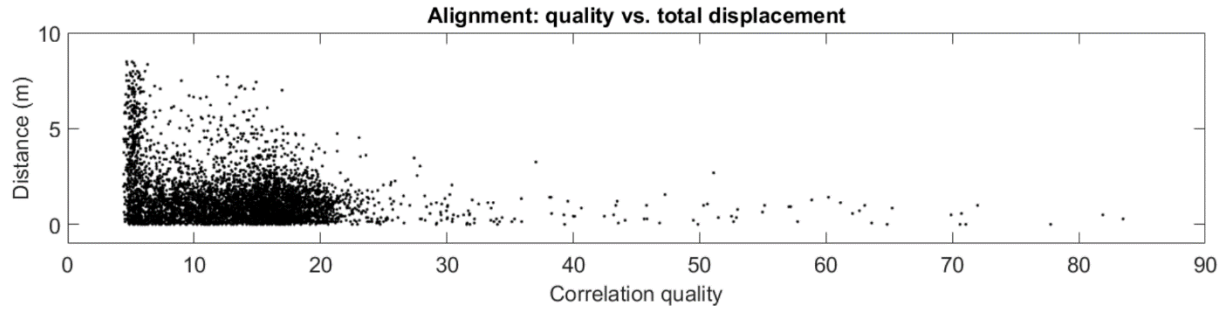

Figure 103: Total displacement vs correlation quality determined via equation (5). Quality values below 6 are randomly distributed over the search area, indicating the noise floor of the quality measurement.

From Fig. 103 it can be clearly seen that quality estimates never pass below a lower threshold, approximately 5 in this case, and are uniformly distributed in the region just above this threshold, a clear indication of the noise floor. Based on these observations the error function is then used to map the distribution of quality values to a range from zero to 1 using:

$$\boldsymbol{Qual_{i,j} = erf\left(\frac{CCnorm_{i,j}-5}{10}\right)}$$

Eq. 59

These are the weighting values fed into an iterative reweighted least-squares solver that attempts to determine the best translations for each sub-image that make sense of the shift observations. The assumption that each scan is best fixed by a translation is not an accurate assumption because the position error grows vs. distance travelled underwater. Nevertheless, the goal of applying this IRLS algorithm is to generate an improved set of weights that more accurately describe the quality of the displacement measurements: least-squares regression can give very poor results in the presence of outliers, and the IRLS process drastically reduces the effect of spurious outliers by appropriately penalizing their associated quality values. A plot of the correlation quality vs. IRLS quality is illustrative because it provides a much stronger separation between “poor” and “good” quality shift measurements (Fig. 104).

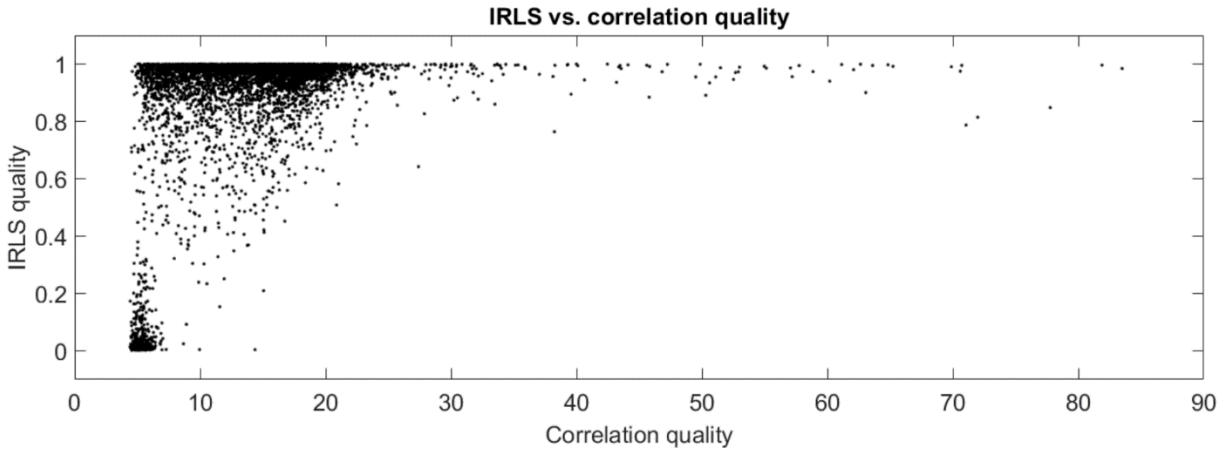

Figure 104: A comparison of correlation quality measurements and the distance measurement weightings arrived at via an IRLS solver, for which weights were initialized using the correlation qualities mapped using Eq. 59.

Solving the data driven navigation problem

To get an actual navigation solution from the translation estimates the displacements and quality estimates need to be associated with specific, time-stamped towfish locations from the previous navigation solution. This enables the displacements to be formulated as a series of relative translations between portions of the navigation solution, which can be easily incorporated into a solver. The first step of mapping the translations to an actual renavigation solution is therefore associating successful translation measurements to specific time stamps and vehicle path locations in the scan segments contributing to the measurement.

In the current approach this is done by computing the centroid of the mask overlap. The UTM coordinates of the centroid are matched to the points of closest approach in the contributing scans. The time-stamps of these points are then recorded and subsequently used during the navigation correction process. Figure 105 is an illustration of this time-stamp association process being performed.

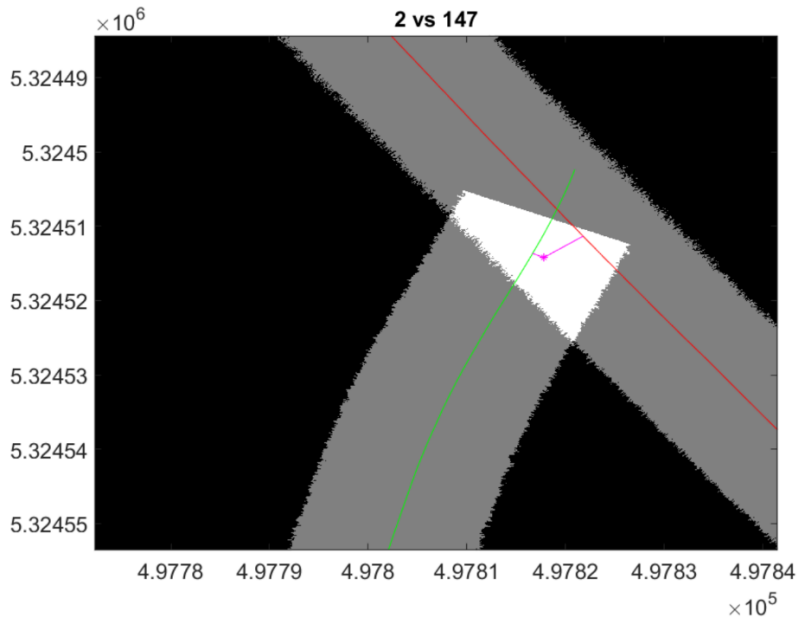

Figure 105: Displacement-to-timestamp mapping procedure: the overlap centroid (magenta star) is associated with the timestamp for the path locations of closest approach (magenta lines) for the contributing scan paths (green and red lines).

Following the association of all displacement measurements to pairs of time-stamps a large system of equations is formed in which x- and y- displacements are solved that match the measured displacements between time-stamps while strongly penalizing solution curvature independently in the x- and y- dimensions (for temporally continuous segments of the navigation solution). The curvature penalization allows the displacement measurements to provide gentle navigation deviations from the previous solution vs. being limited either to perfectly rigid translations (over regularization) or having to wander excessively to match every displacement estimate. Following the solving of these equations the smoothed displacement corrections in *x* and *y* are added to the original navigation solution to form the feature-driven navigation solution.

Feature-aligned mosaic generation

Upon generation of a feature-based renavigation solution the data can now be combined at "feature-scale" precision in a mosaic that will be exploited by the post-mission analysis tools discussed in the next section. Mosaicking is the process of combining these scans into a single image that:

- Enables simple localization in geo-rectified coordinates of any of the targets in the survey region.
- Reduces the number of images the user must sort through to find targets.
- Enables an intuitive understanding of the relative location of features both inside and outside of the target field.
- Potentially decreases speckle.
- Potentially improves the SNR of targets and the ability to detect them.

The "Set_MOSAIC" tool in the MuST software library provides the user with the ability to generate mosaics by specifying: 1) the JSF_table associated with the batch-processed files, 2) the renavigation corrections augmenting the navigation solution, 3) which layers to mosaic (a separate mosaic is created for each layer), 4) the grid resolution of the mosaic, and 5) the lower-left and upper-right corners of the grid region in latitude and longitude (decimal degrees). The tool converts the corners of the specified region into UTM and creates an evenly sampled grid. Files correspond to scans having swaths that overlap with the user-specified region are geo-rectified to the grid using the augmented navigation solution and are individually stored. Note that the SAS images have not been re-processed using the new navigation solution: thus far it has not been evident that any improvement in image focus is achieved by replacing the original renavigation solution with the feature-based solution.

Previous work has demonstrated that the manner in which multi-aspect data is combined can greatly affect the kinds of features accentuated in images [36]. Storing each scan individually facilitates combination of the rasterized scans using methods that accentuate different features. Presently, Set_MOSAIC generates three separate mosaics for each layer created from the set maximum, mean, and the *contraharmonic* mean:

$$\boldsymbol{S_{MAX} = MAX\{S_n(x,y)\}_{n=1}^{N}} \quad \textbf{Eq. 60}$$

$$\boldsymbol{S_{Mean} = \frac{1}{N}\sum_{n=1}^{N} S_n(x,y)} \quad \textbf{Eq. 61}$$

$$\boldsymbol{S_{CMN} = \frac{\sum_{n=1}^{N} S_n^2(x,y)}{\sum_{n=1}^{N} S_n(x,y)}} \quad \textbf{Eq. 62}$$

In (60 – 62) $S_n(x,y)$ represents the n^{th} rasterized scan of the set corresponding to a particular layer. The *maximum* operation accentuates the brightest features of targets; typically the broadside glints. This method also frequently provides the most consistent and seamless view of the sediment because the lower-level, outer swath regions of images are hidden by datasets captured in the same region that have higher signal levels.

The *mean* operation often gives the best despeckling results and also strongly emphasizes stable, isotropic features of targets. Often the mean will provide the best SNR for a given feature. The downside of the mean is that trends and data-gaps very strongly affect the mosaic, creating many false edges and artifacts.

The *contraharmonic* mean, a sub-case of the *Lehmer* mean, is a very effective tool for balancing some of the good qualities of the mean (e.g. despeckling) while countering the artifacts caused by over-weighting regions of scans with low signal level. The contraharmonic mean is essentially a weighted sum where the value of the weights are equivalent to the values being summed. Conveniently for the present situation, regions of low (or zero) signal are completely removed in this operation, however a large amount of de-speckling is still achieved. Figures 106 and 107 show sonar mosaics generated using the three different combination methods.

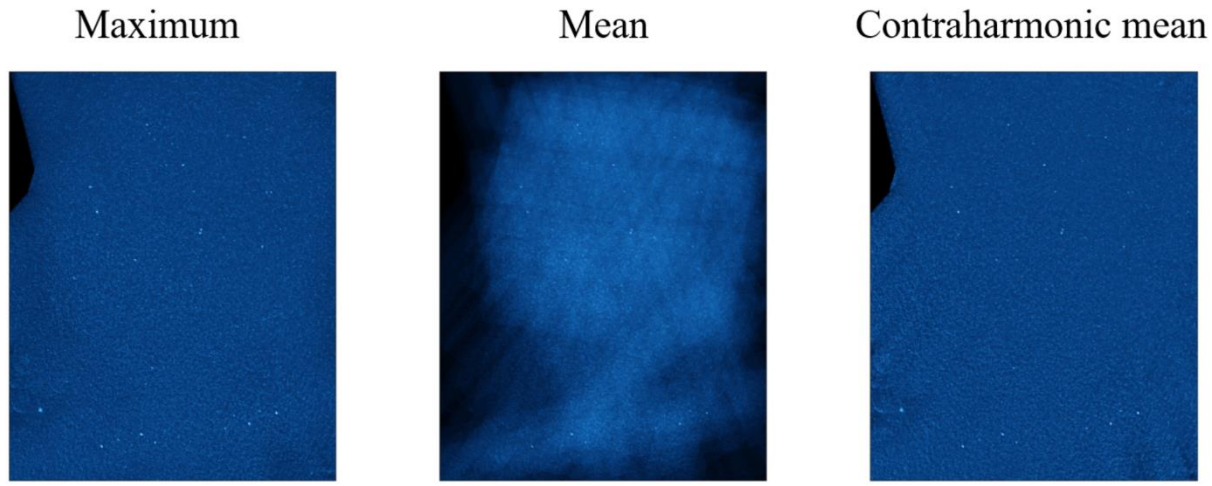

Figure 106: Mosaics of a survey region created using the *Max* operation (left), *mean* (middle), and *contraharmonic mean* (right). The maximum has the fewest image artifacts created by scan density and beam pattern trends, followed very closely by the contraharmonic mean. The arithmetic mean (center) exhibits many artifacts based on density of coverage and beam-pattern effects.

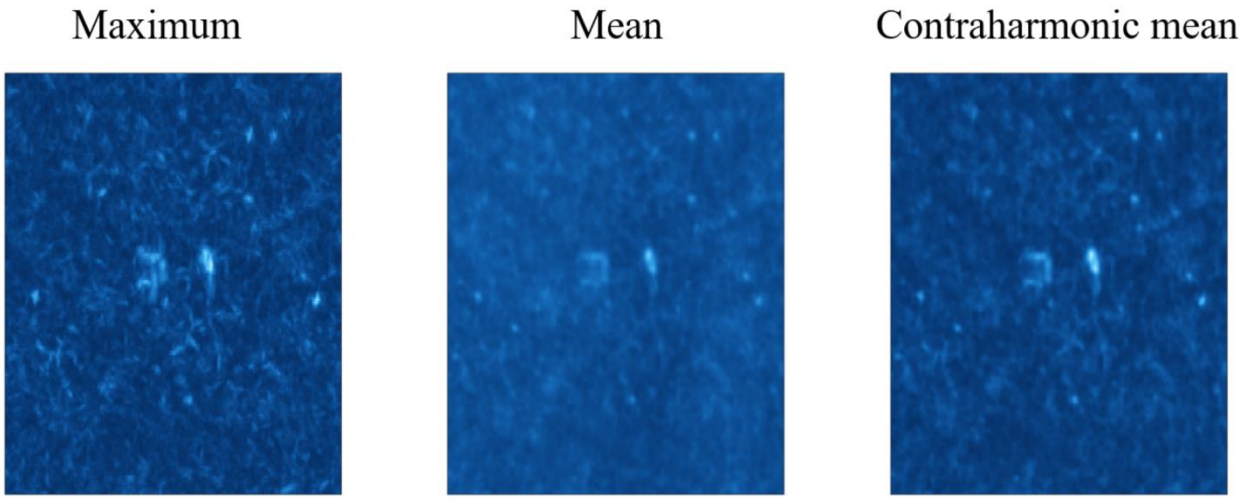

Figure 107: A zoomed view of some targets in the region of interest showing the effects of the different combination methods. The *max* operator produces images with highly speckled, almost feathery appearance. The mean and contraharmonic means create nicely de-speckled images.

Each mosaic type is preferential for different purposes and are accessible in the post-mission analysis tool. The set of mosaics corresponding to each sediment projection layer and each combination method are stored in a file that is accessed and used as the primary means of region-of-interest detection in the post-mission analysis tool which is discussed in the subsequent section.

Post-Mission Analysis

PMA tool main screen and ROI selection

In the present context, post-mission analysis refers to the observation and analysis of beamformed data to ascertain meaningful information furthering the objective of the acquisition mission: UXO detection and classification. To simplify this process software has been written that gives the end user extensive options for how data is visualized, targets are selected and labelled, and associated data is recalled and/or transformed to accentuate features of interest. The software makes use of the file-association tables described earlier along with mosaics described in the previous section. Figure 108 shows the user interface of the BOSS UXO post-mission analysis (PMA) tool, loaded with a mosaic created from the Sequim Bay 2021 UXO testbed survey.

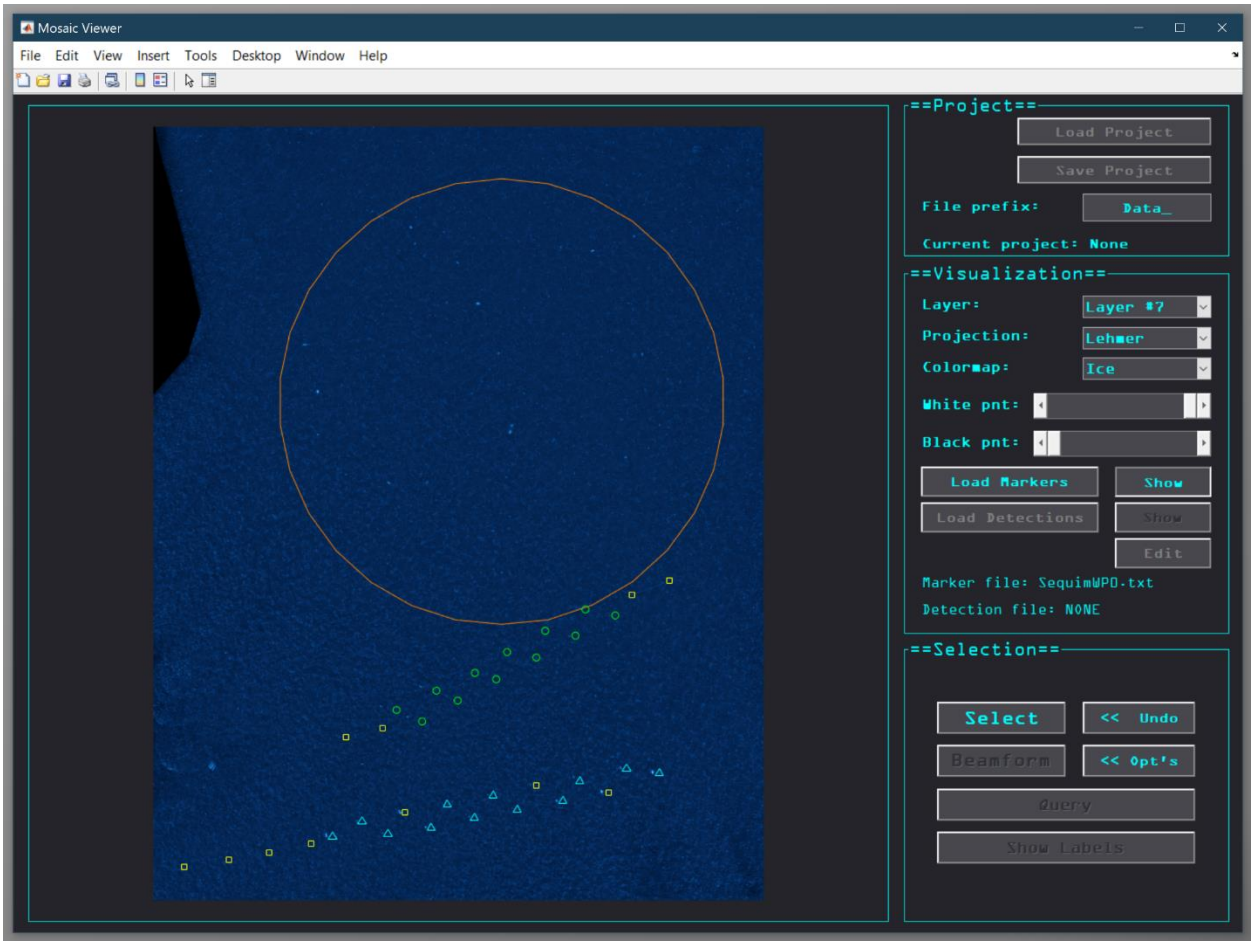

Figure 108: The BOSS post-mission analysis tool user interface.

The user interface is comprised of a large data visualization screen on the left and three sub-panels on the right pertaining to the current survey project (top), visualization of the mosaic (middle) and selection of regions of interest (bottom). Several features are still in development (multi-project flexibility and importing of detection data such as those created via the previously described Williams detector). The sediment layer can be selected via the “Layer” drop-down menu in the visualization panel. It defaults to the maximum layer number, which is always

equivalent to the projection through the entire sediment column. Layers 1 through *N*-1 (where *N* = 6 in the present case) correspond to the layers created during batch processing. The projection type (*Max*, *mean*, *contraharmonic*) can be specified under the "Projection" drop-down menu and defaults to "Lehmer" (the contraharmonic mean is a special case of the Lehmer mean). The user also is presented with colormap options and the ability to adjust the intensity thresholds via the black- and white-point sliders in the panel. The user can also load marker files of the same format used by MuSTLive, which can be superimposed on the mosaic. The default marker map for the 2021 Sequim Bay testbed survey is shown in Fig. 108: the boundaries of the testbed are shown in orange, and markers corresponding to calibration targets are visible beneath the testbed region. The markers can be toggled on or off using the "show" button to the right of "Load markers." The detections compiled after 4 days of surveying operations are shown in Fig. 109, along with several other modifications made to the mosaic visualization.

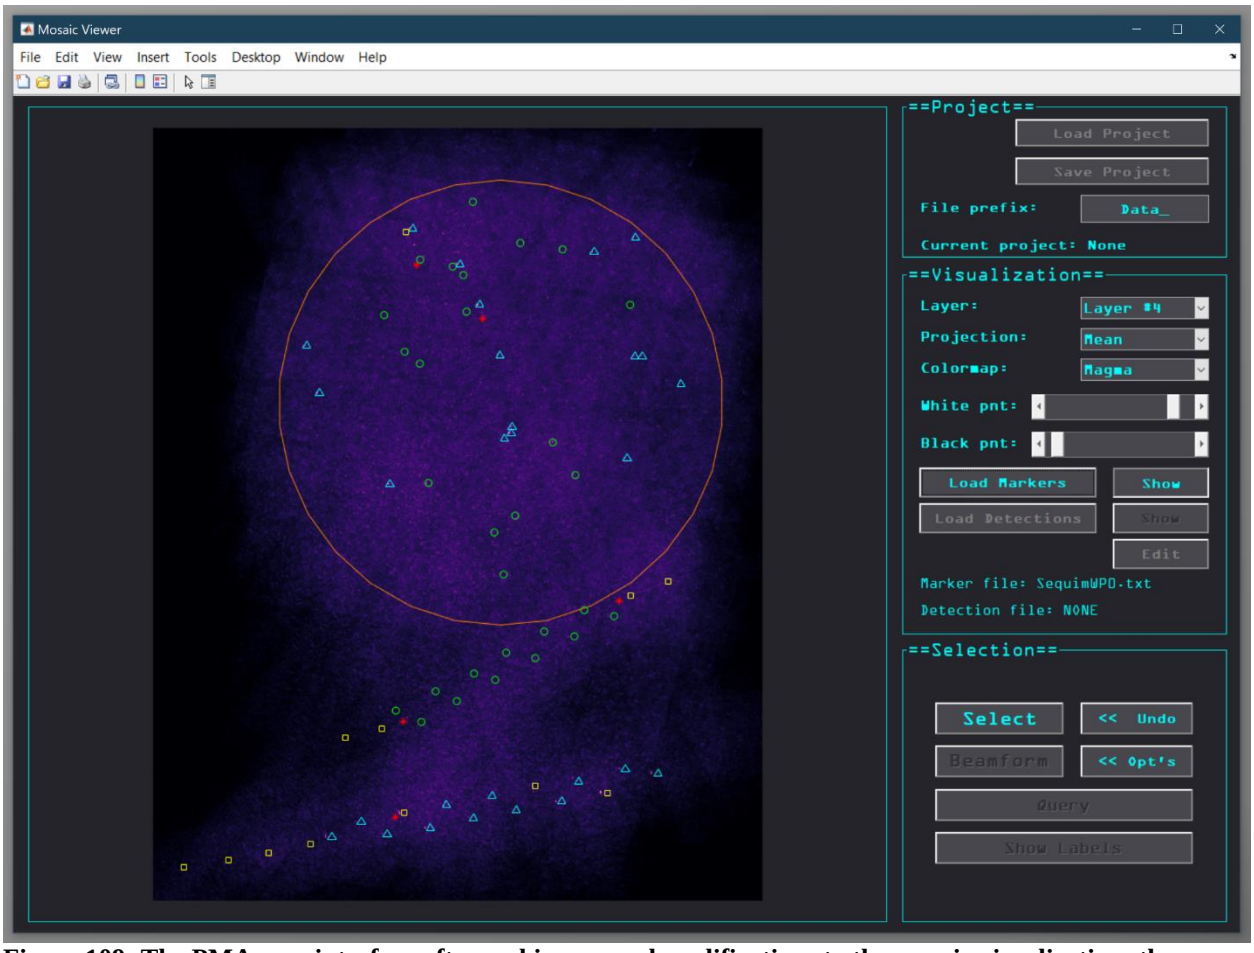

Figure 109: The PMA user interface after making several modifications to the mosaic visualization: the combination method has been set to "mean," the black and white points have been altered to improve contrast, the layer straddling the interface has been selected, the colormap has been altered and the markers from the last day of survey operations have been included.

The selection panel at the bottom allows the user to select and patch-process regions of interest. The user can zoom and pan the mosaic to find targets of interest and make selections by clicking on the "select" button. The cursor will become a cross-hair when hovering over the mosaic and

the user can click on a point of interest. As soon as the user makes a selection the “Beamform” option becomes available. The user can make as many selections as desired and selecting “beamform” will batch-process the set of selections. In the following example the mosaic has been zoomed to a region near the center of the testbed and four selections have been made. (The colormap has been returned to “ice”).

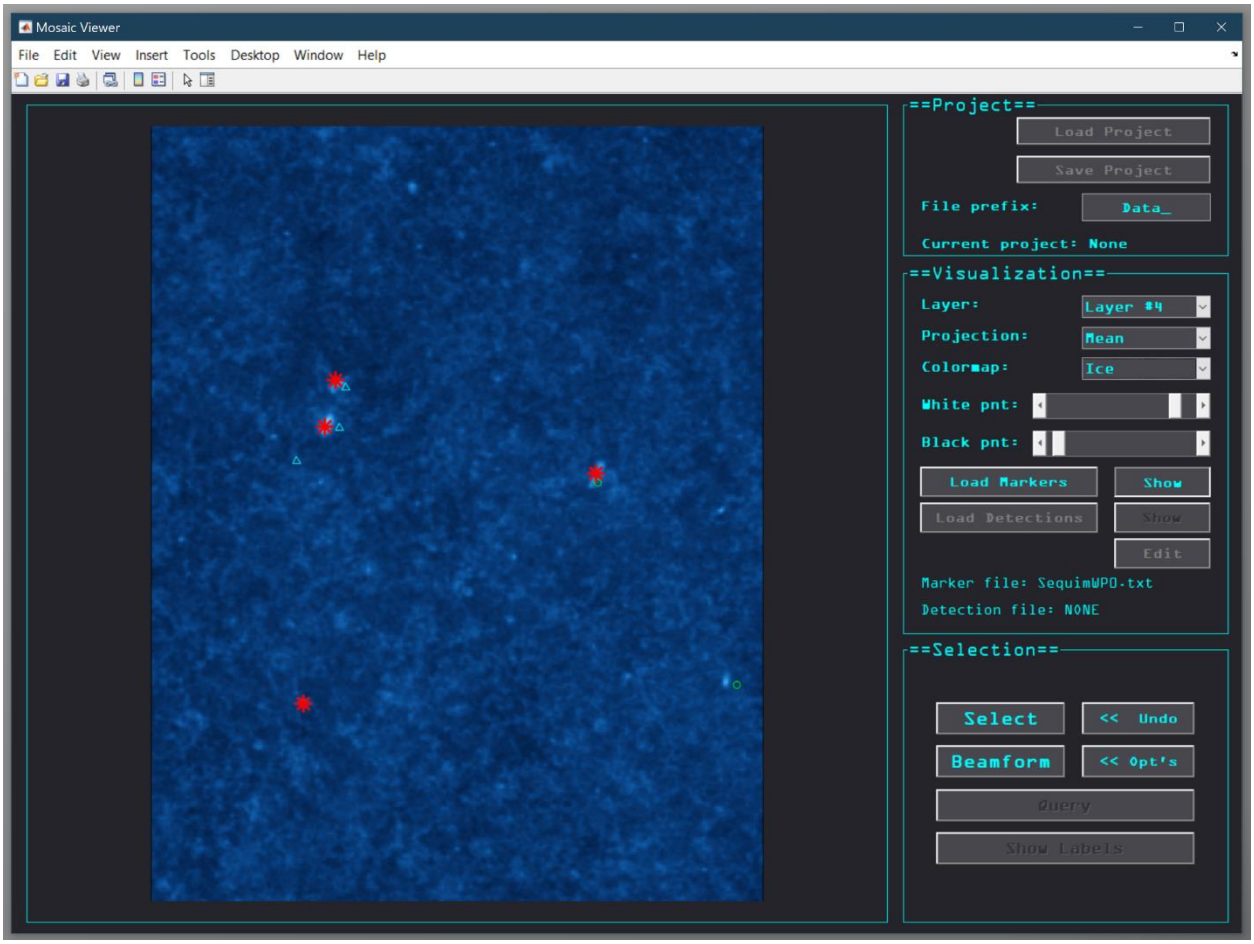

Figure 110: An example of target selection in the PMA tool. Selections are denoted by red stars. Four locations have been selected.

The user can modify beamforming parameters by clicking the “Opt’s” [options] button to the right of “Beamform.” The settings window opened by this action is shown in Fig. 111.

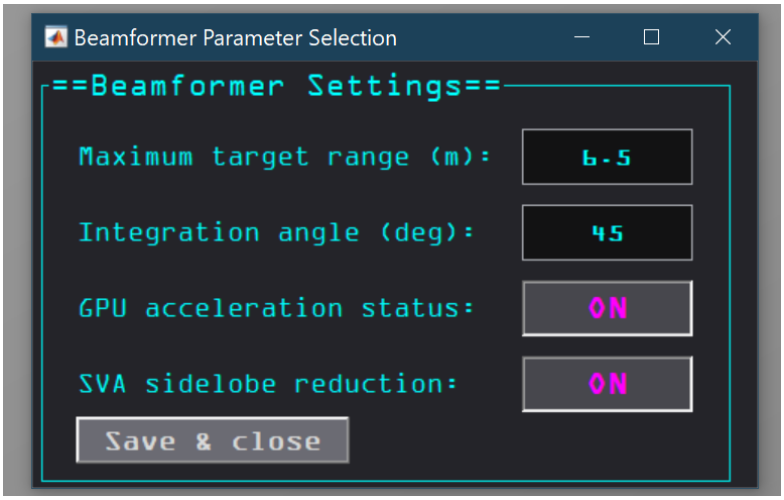

Figure 111: The beamforming options selection window.

Figure 111 shows the beamforming options available to the user. The "Maximum target range" field specifies the maximum track-line distance from the target that will enable a scan to be considered a "contributor" to the target of interest. Increasing the value will result in more observations, but the observations will be at longer range and lower grazing angle. The second option, "Integration angle," sets the beamformer "half-angle" in a manner identical to the same settings discussed previously for MuSTLive and the SAS post-processing beamformer. The user can also determine if the GPU should be used for beamforming to accelerate the process (it can be switched off if the user does not have a GPU accelerated computer), and spatially variant apodization (SVA) can be turned on or off by toggling the adjacent button as well.

ROI review tool

When the "Beamform" button is pressed in the main screen, the PMA program searches through the JSF table linked to the active mosaic and finds the contributing scans based on the search-radius specified in the options (default radius = 6.5 meters). The files are accessed and portions contributing to selected regions of interest are beamformed and assigned to detection class-objects that contain as properties beamformed imagery and metadata providing information about the regions and the scans contributing to the regions, labelling and class score information.

Currently all regions of interest are defined as 2 x 2 x 2 meter cubes, and are beamformed at 2.5cm voxel resolution (~500,000 complex voxel values per region-of-interest). The volumes are all beamformed directly into geo-rectified Eastings-Northings-Up (ENU) UTM coordinates. Due to having certain advantages for acoustic color processing and spatial masking, each scan is also beamformed into a path-centered *cylindrical* coordinate frame having coordinates {*range, along track distance, azimuthal angle*}.

Following completion of the beamforming operation the ROI review tool opens. Figure 112 shows the ROI review screen that opens after the four selections shown in Fig. 110 have been beamformed.

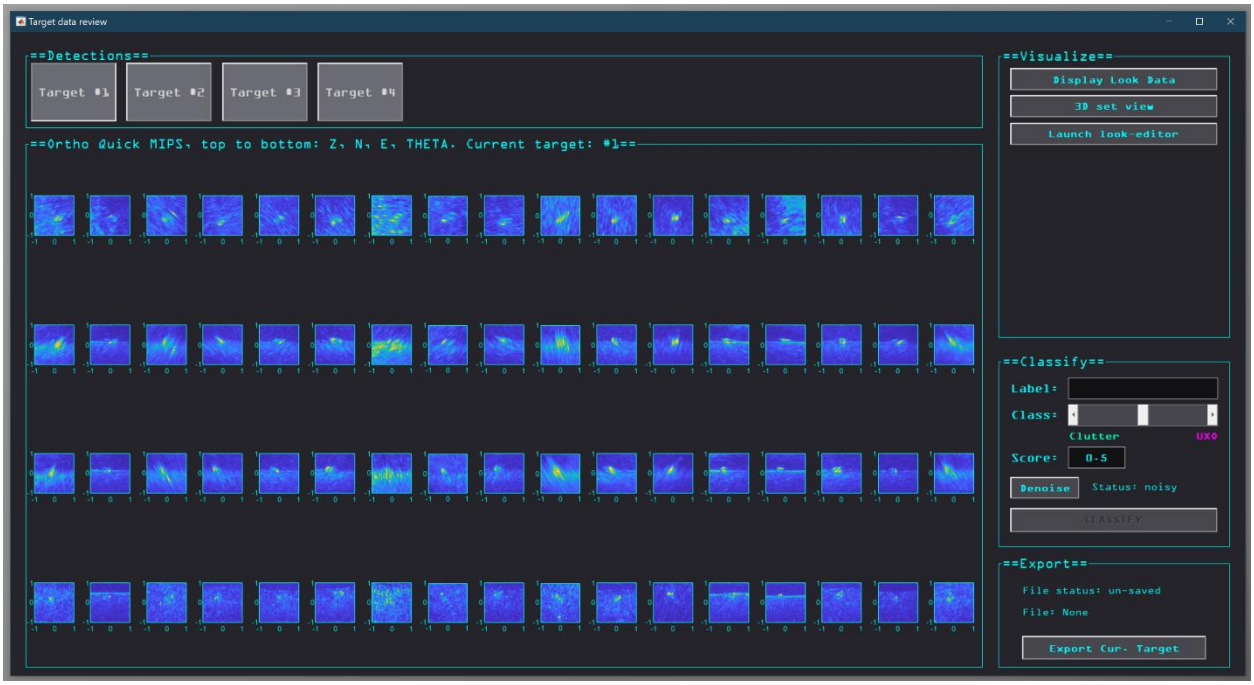

Figure 112: The ROI review screen created upon beamforming the selections shown in Fig. 110. This screen will be the launching pad for most of the user's ROI observation and manipulation and classification activities.

The ROI review tool shown in Fig. 112 gives the user access to most of the data visualization and feature enhancement tools developed as part of this project, and a number of subsequent sections will be devoted to explaining these features. The top row of the figure contains a row of selection buttons corresponding to each target that was batch processed based on the user's selections in the mosaic screen. The selection defaults to Target #1, the first selection made by the user.

The panel directly beneath shows maximum intensity projections along the 3 axes of the Cartesian coordinate volume image for each look acquired of the selected target. Therefore in Fig. 112 it can be seen that 17 individual observations were made of target #1. Each row corresponds to a projection through a different axis: Z (height, top row), Northings and Eastings. The bottom row is a projection through the azimuthal angle axis of the cylindrical coordinate image and its image axes are *range* (m) and *along-track displacement* (m).

The panels on the right of the screen allow the user to perform many functions related to visualization and classification of the data corresponding to the selected target. The middle panel, labeled "Classify," allows the user to determine a class score for the target (adjustable via a slider, 0 = Clutter with 100% certainty, 1 = UXO with 100% certainty). The user can add a label to the object (e.g. "UXO #1," "Cinder block," "Unknown clutter item" etc.) in this panel, and a place-holder is left for insertion of automatic target recognition algorithm (or algorithms) via the disabled "Classify" button. The user can also apply a spectral-denoising function that removes specular glints and filters out interface and volume scattering to improve the signal-to-noise ratio of structural acoustic features by pressing the "denoise" button. This function will

apply the denoising algorithm to *all* of the looks acquired for the current target, and the algorithm behind the filtering operation will be described in a subsequent section.

In the “Visualization” panel on the top right the user is presented with a number of different methods for viewing and editing look data. If the user selects “Display Look Data” a more complete version of the MIP projection set is opened in a new window, shown in Fig. 113. This figure contains the labelled projections shown in the bottom left panel of the ROI review screen, but additionally shows acoustic color plots (collapsed to two dimension by projecting through the wavenumber axis corresponding to azimuth) and, if the user has denoised the dataset, the denoised acoustic color spectra as well.

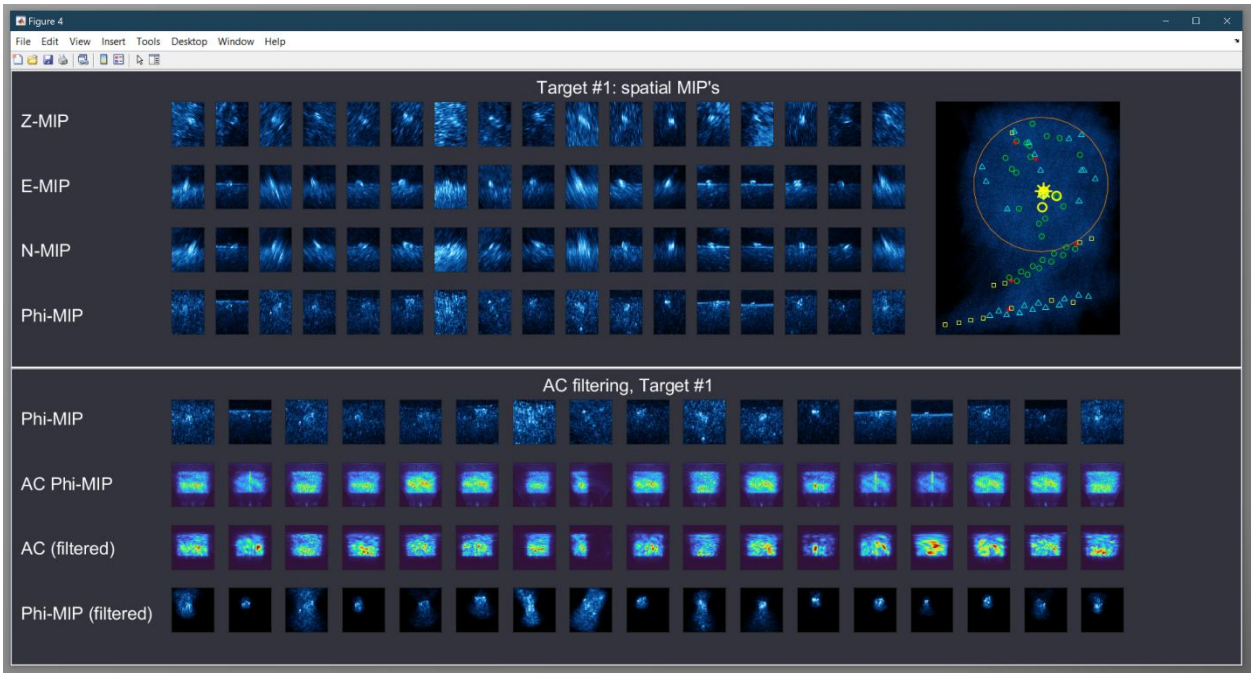

Figure 113: The look-data display window showing projections through the spatial domain (top) and filtered data (bottom). An inset view of the mosaic with the location of the current target denoted with a yellow star is shown in the upper right corner.

This window is useful for quickly reviewing the spatial and acoustic color data and reviewing the location of the target in the mosaic.

Selecting “3D set view” from the visualization panel opens up a multi-panel 3D maximum intensity projection display that allows the user to uniformly change the viewpoint of all of the different target looks in the set using the provided viewpoint manipulation box. The figure on the bottom right of the set is the magnitude-sum of all of the target looks and often produces a greatly improved view of the target, having lower speckle and a more complete reconstruction of target features. The viewpoint of any individual look rendering can be modified independently by clicking and dragging the cursor over the rendering.

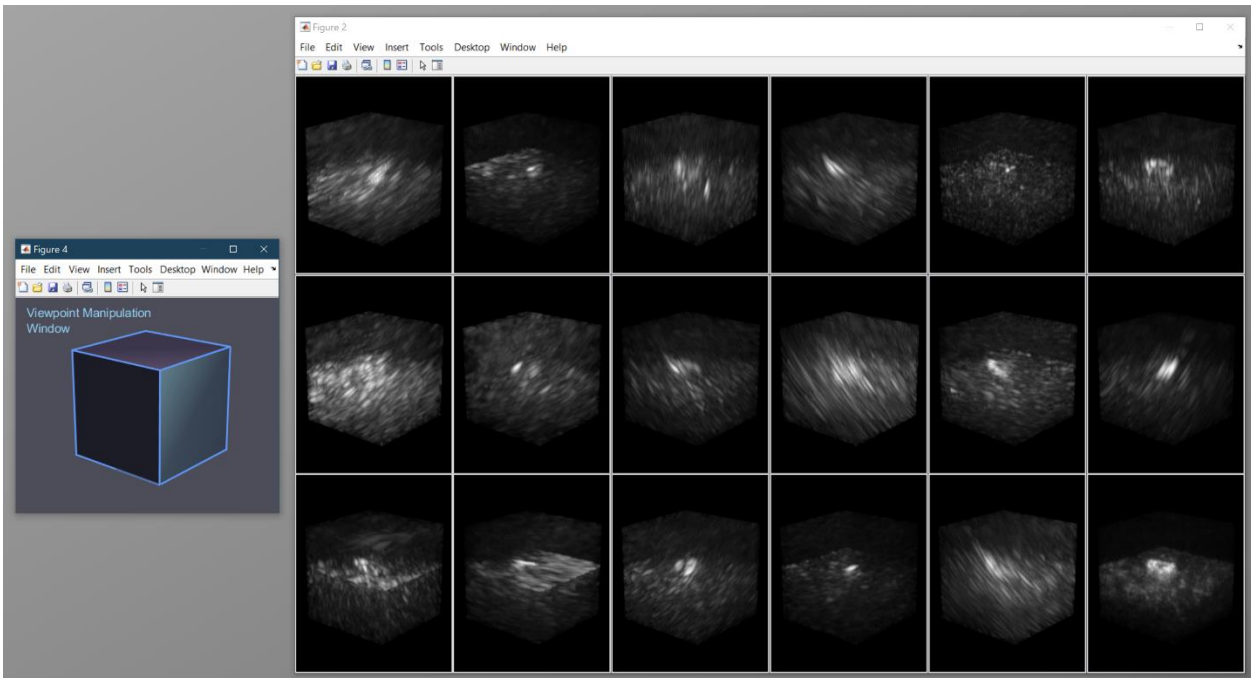

Figure 114: The 3D set-view visualization tool. Clicking and dragging the viewpoint box on the left simultaneously changes the viewpoint of the rendered looks on the right. The image on the bottom right is a magnitude-sum of all of the contributing target looks.

Finally, selecting "launch look editor" from the visualization panel opens up a new user interface that gives the user detailed control over the display and denoising operations of individual target looks. This will be described in the next section.

The Look Editor

Figure 115 shows the user interface for the look editor. The top panel allows the user to select the currently active look from a drop-down menu and displays the *.jsf* file from which the data was beamformed. Beneath this selection panel the user has a number of options related to data display and, significantly, length measurement tools and custom data denoising.

Figure 115: The Look-Editor user interface, containing options for look selection, visualization, denoising and length-scale measuring.

If the user selects “Launch MIP viewer” from the Tools panel, the following display links itself to the editor panel. This is like a zoomed version of the rendering corresponding to the current look that was part of the set shown in Fig. 114. The MIP viewer provides more visualization features, such as the ability to apply dynamic range compression and limiting and colormap selection.

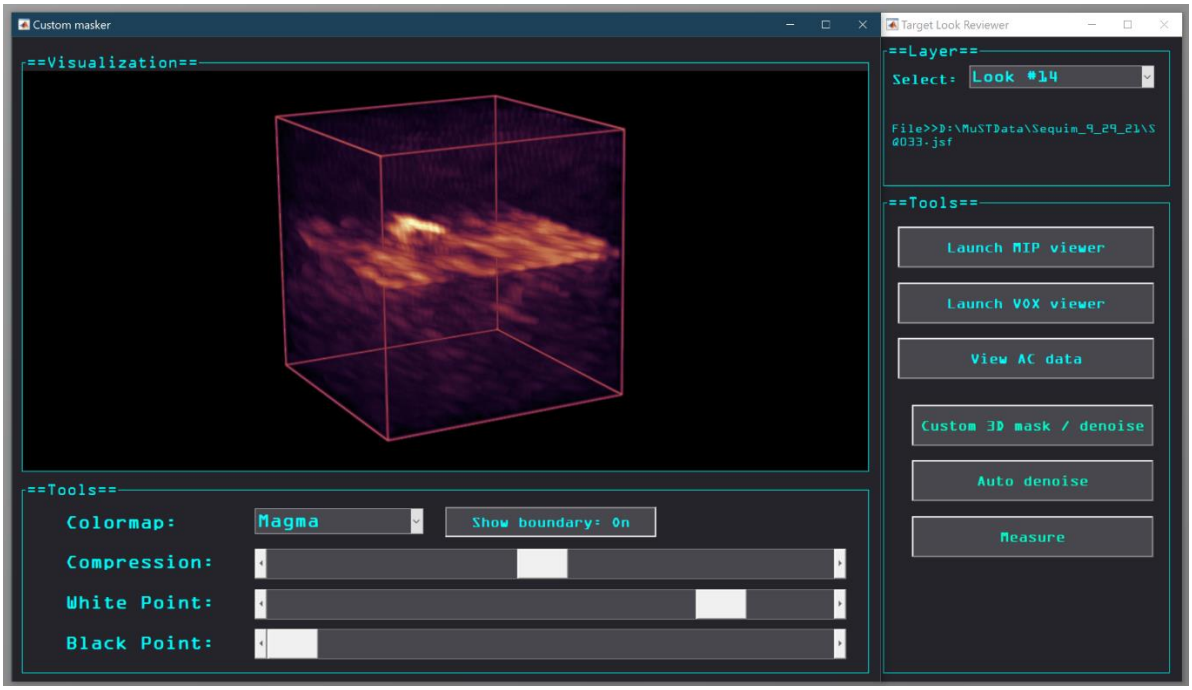

Figure 116: The MIP viewer opened upon selecting "Launch MIP viewer."

Selecting "Launch VOX viewer" creates a rendering using alpha-blending via VOXview rather than a maximum-intensity projection.

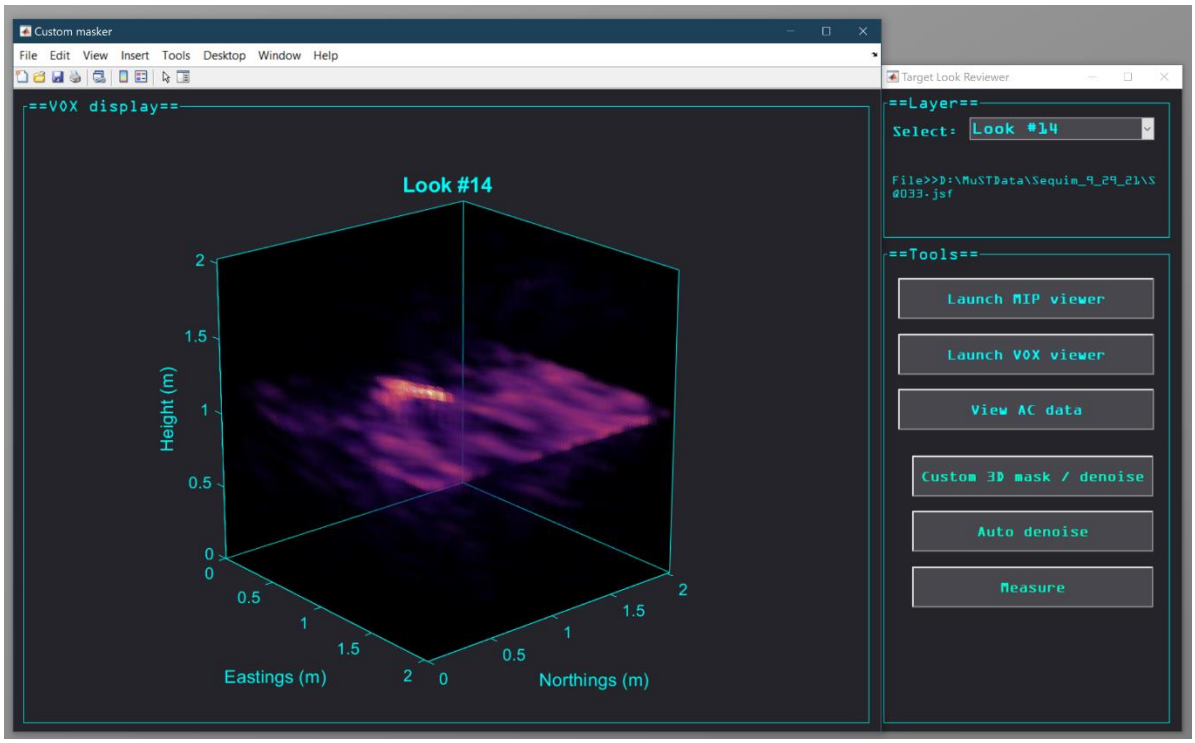

Figure 117: The alpha-blending rendering opened by selecting "Launch VOX viewer."

Acoustic color data corresponding to the selected look can also be viewed by selecting “View AC data.” Selecting this option opens up the linked pair of windows shown in Fig. 118.

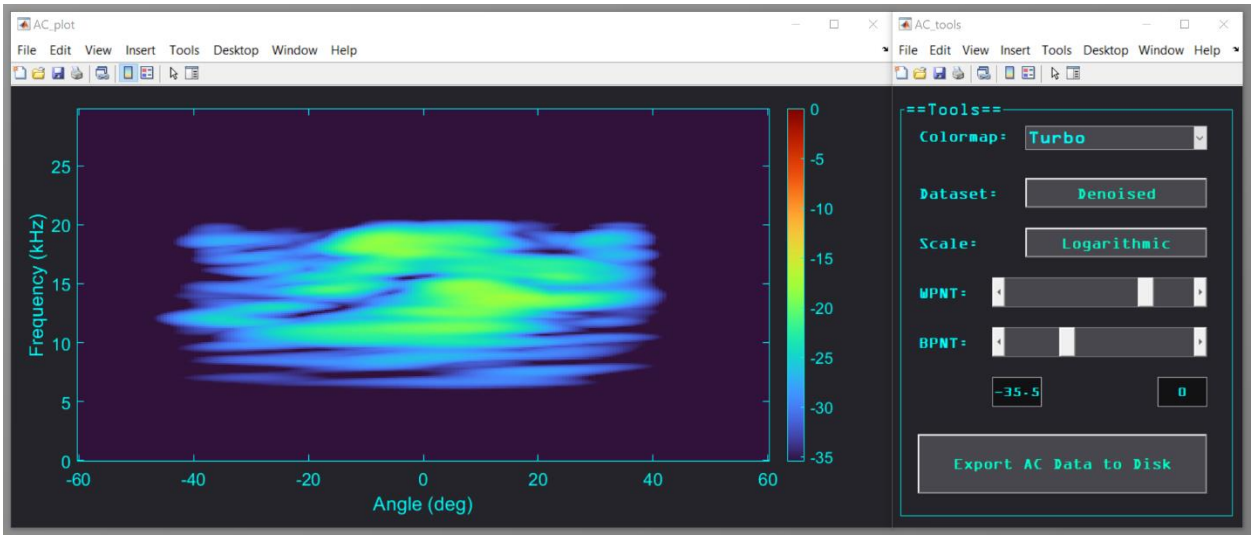

Figure 118: The acoustic-color visualization tools opened by selecting “View AC data.”

A variety of options are available to the user for visualizing acoustic color data: the user can chose whether to display data on a logarithmic or linear scale, and the thresholds of the visualization can be selected with the white-point and black-point sliders (set as a default to 35.5 dB dynamic range). The user can directly enter the desired dynamic range thresholds by editing the numerical fields below the sliders. The user can also swap back and forth between the “denoised” data (default, if it exists / i.e. if the user has already performed either a custom or batch denoising operation), or the unfiltered data. Figure 119 shows the original un-filtered version for comparison, selected by toggling the “Dataset” button to “original.”

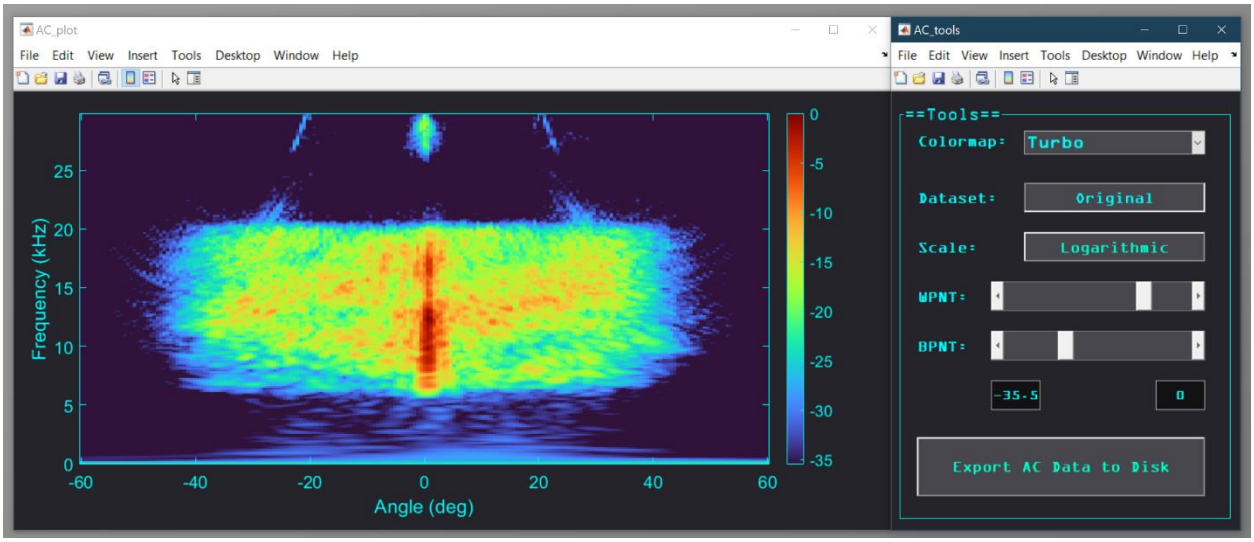

Figure 119: The original acoustic color plot prior to filtering, selected by toggling the “Dataset” switch to “Original.”

Note that in Fig. 119 there appears to be a bright, narrow specular glint near 0 degrees that is completely missing from the "denoised" version. This would intuitively imply that the filter is removing critical target features, however this is *not* the case: the narrow spike in the original, unfiltered signal is actually the contribution from the coherent component of the specular glint off of the sediment: this look at the ROI is nearly at nadir. The specular glint is visible in the ROI visualizations in Fig.'s 117 and 118 and manifests as a rippled flat surface right at the sediment / water interface. This can be verified by creating a custom mask, which will be demonstrated shortly. Lastly, the user can select the "Export AC Data to Disk" button to save the acoustic color data as a two-dimensional matrix with supporting metadata that can be read by other programs for classification, plotting purposes, or adding to a database. Upon clicking the button the user is prompted to select a directory and assign the file a name.

Selecting the "Custom 3D mask / denoise" option opens the following screen which allows the user to create a custom three-dimensional mask around the target to filter out volume scattering noise and nearby clutter objects. This option is really only necessary in cases where the automatic filter fails to produce good results, however it is a useful tool for performance verification as well. Figure 120 shows the window open by selecting this option.

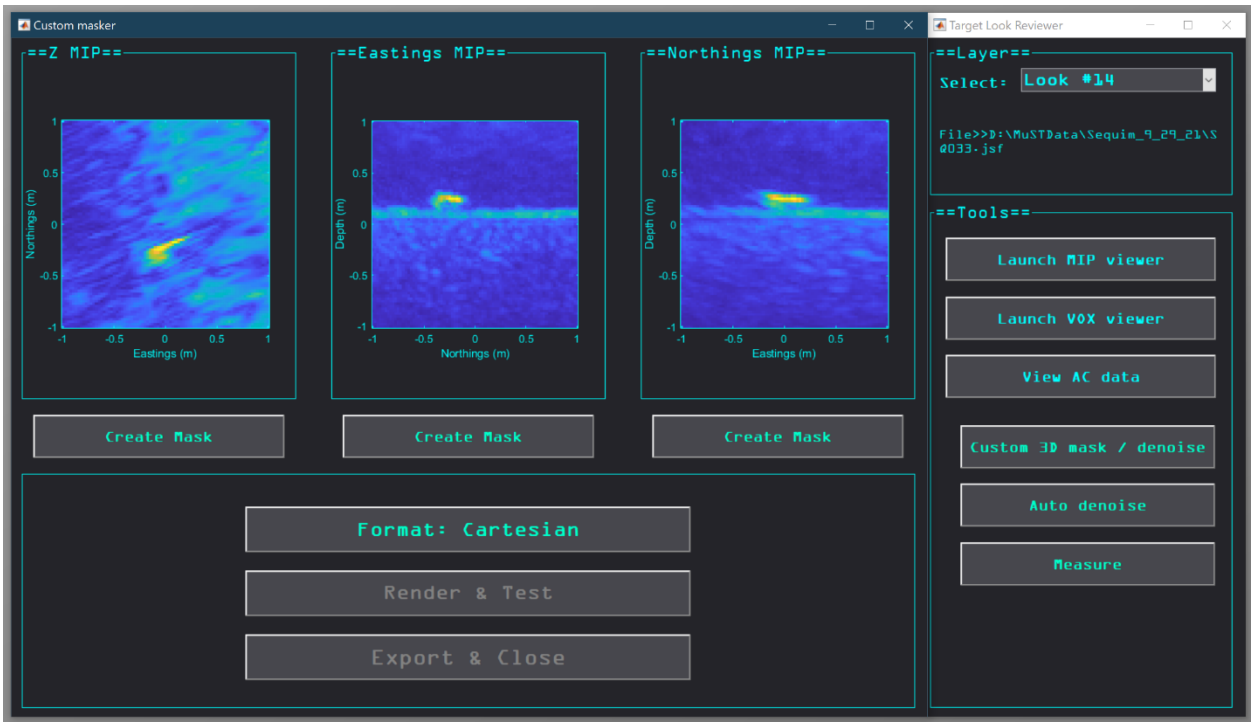

Figure 120: The custom denoising window pane. The user defines a boundary for each window which is combined to create a three-dimensional spatial mask around the target.

When the user selects "Create Mask" under an image the image will highlight and the user will be able to define a polygon by selecting corners. The top toggle switch in the bottom panel swaps between a Cartesian volume and a cylindrical-coordinate volume. In many (but not all) cases the cylindrical coordinate representation is more useful for masking, and Fig. 121 shows masks created in cylindrical coordinates.

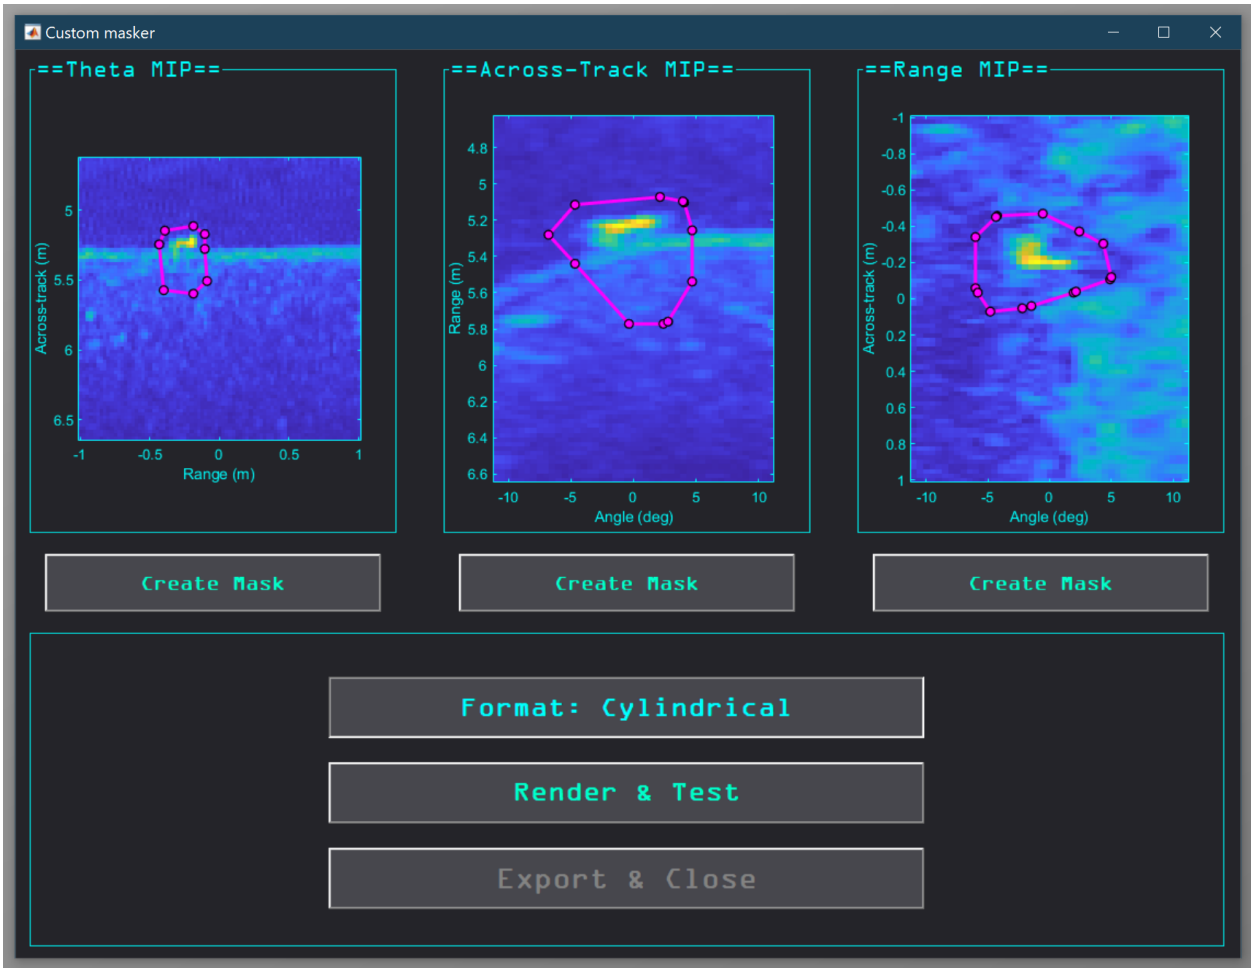

Figure 121: Mask polynomials created in cylindrical coordinates. Note the labels above the images show the dimensions through which the maximum is being found. The "Render & Test" button has also been activated.

When masks have been defined for each dimension then the user can select "Render & Test" to check the effects of the mask on the spectrum, as shown in Fig.'s 122 and 123. In Fig. 122 the filtered spectrum is shown in the upper left plot, and the comparison (unfiltered) spectrum is sown in the bottom left plot. The 3D data is shown with the user-defined three-dimensional mask shown as a translucent 3D polygon in the figure on the right. Note that the spectrum is normalized by the brightest value in the filtered spectrum, so most of the noisy spectrum clips, however the bright surface glint feature is visible down the center.

Toggling the comparison set button on the lower right allows the user to compare the spectrum with the current denoising solution – in this case the auto-denoising solution. This comparison is shown in Fig. 123.

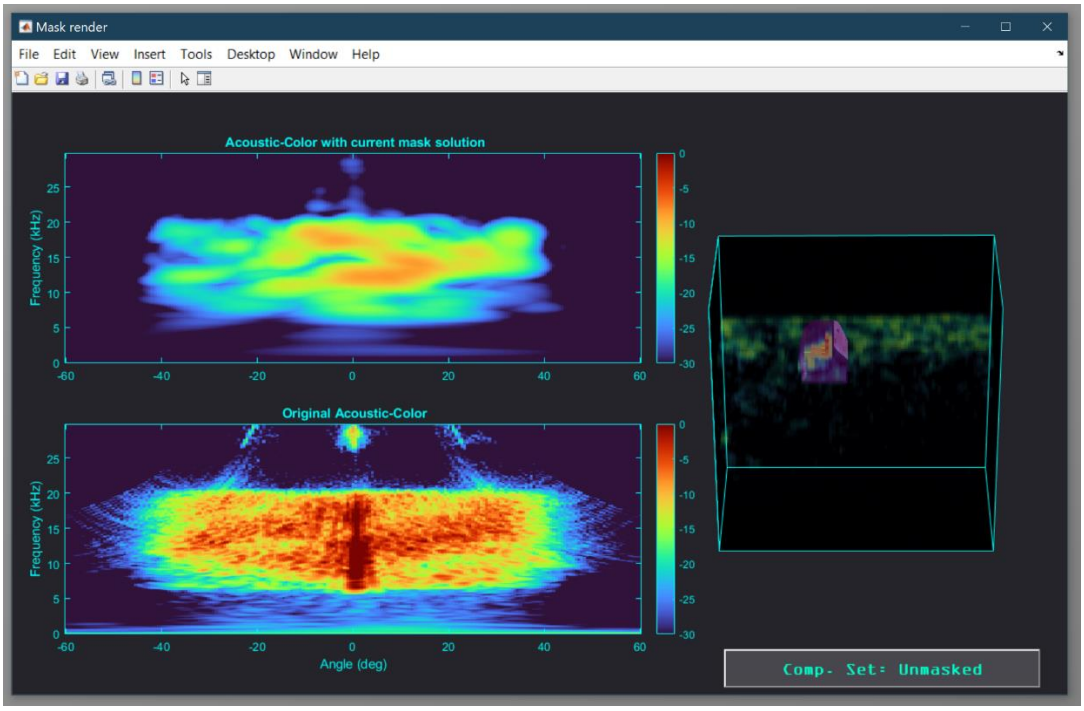

Figure 122: The acoustic-color spectra for the masked object (top) vs. the original, un-masked object (bottom). The spatial mask is shown as a transparent 3D polygon on the right.

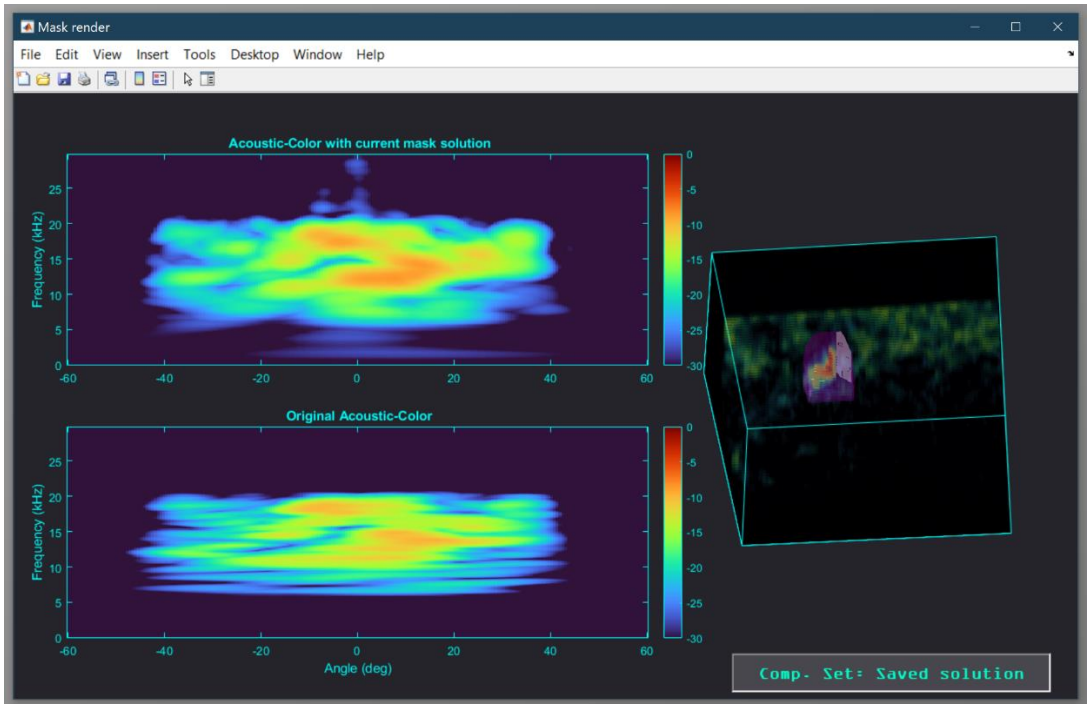

Figure 123: The custom solution (top) vs. the previously saved solution (bottom). The figure perspective has also been altered on the right.

If the user is dissatisfied with the current filter performance this window can be closed and the user can create a new mask. When the user is satisfied with the performance of the mask the

denoised data can be saved and associated with the current look by pressing “Export & Close” in the custom masker main screen (Fig. 121: the “Export & Close” button becomes activated when a masking solution has been rendered and tested.)

The two remaining buttons in the Look-Editor are “Auto denoise” and “Measure.” Auto-denoise is a single instance of the batch auto-denoising tool available in the ROI review tool. If the user creates a custom mask but desires to revert to the automatic version this can be accomplished by pressing “Auto denoise.” Alternatively, the Look-Editor can be opened prior to batch denoising the dataset, and pressing the “Auto denoise” button will apply the standard automatic acoustic-color denoising solution.

The last option available to the user is “Measure.” Clicking this button opens the window shown in Fig. 124.

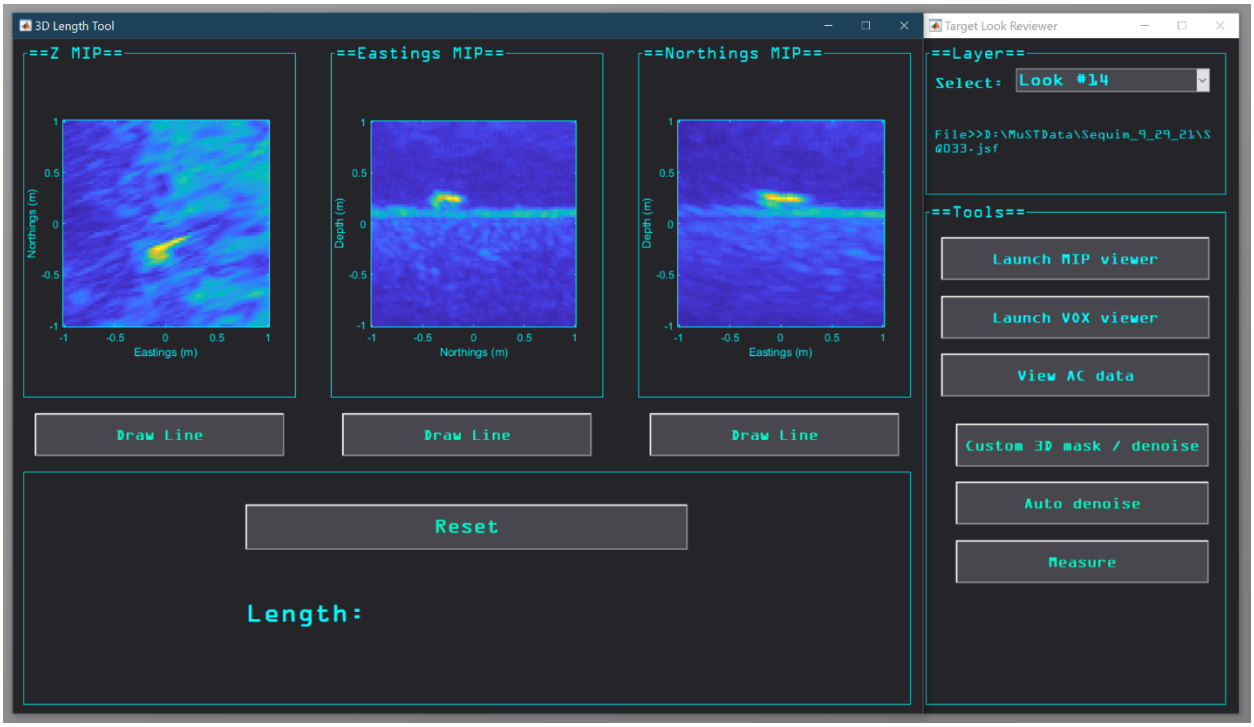

Figure 124: The length measurement tool opened by clicking “Measure” (left panel). The interface is very similar to the 3D masking interface.

Specifying the end-points for a line in a 3D volume is not a task with an immediately obvious solution so the approach taken in this tool is to give the user the three orthogonal views through the Cartesian volume. By selecting “Draw Line” under any of the views, the user’s curser will turn into a cross-hair over the selected view and the user can create a line by clicking and dragging the cursor. The created line form the constraints for the location of the endpoints in the dimensions of the projection. The user can determine the coordinates in the missing dimension by clicking and dragging in either of the two remaining windows. The values of the endpoints in the dimension not defined in the first selection are selected from this second user defined line, which creates a 3D line, and the “Length” field becomes populated as shown in Fig. 125.

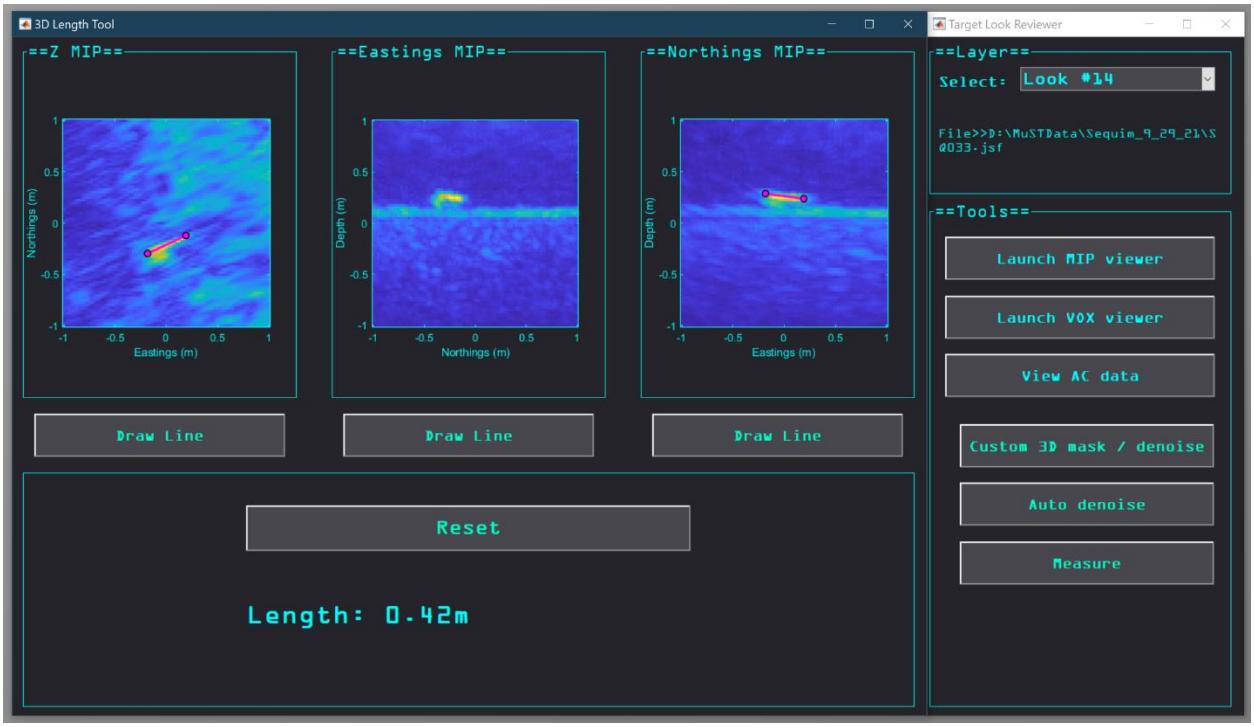

Figure 125: 3D length scale measurement example: A line is defined via the Z-MIP (left) and Northings-MIP (right). The Z-MIP line was first selected by the user and defines the line endpoints in Eastings and Northings. The Northings-MIP provided the Z-dimension values for the endpoints. The length of the target is approximately 42 centimeters.

Selecting "Reset" at any time clears the current endpoint values and the length measurement.

Exporting Data

In the lower right corner of the ROI review tool (Fig. 112) the user is given the option to export an organized structure containing the data and metadata associated with the currently selected target.

Figure 126: The data export panel in the ROI review tool.

When the user clicks "Export Cur. Target" a prompt will appear requesting the user to select a directory and provide a file name. Following exportation the "File status" and "File" fields are

updated as shown in Fig. 126. Note that if data has been denoised, either individually using custom masks or automatically using the auto-denoiser algorithm, these denoised acoustic-color spectral plus all label and classification information are also saved in the exported datasets.

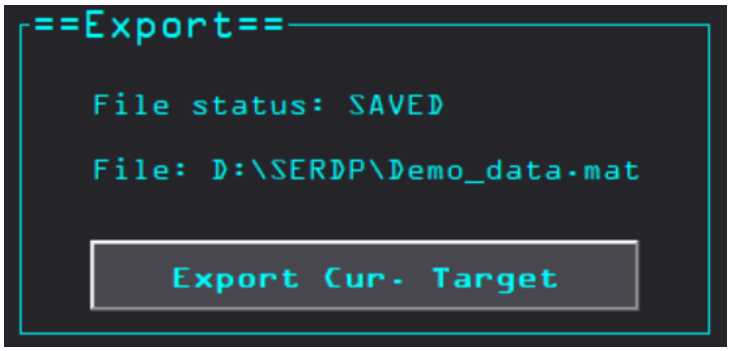

Figure 127: An example of the field update caused by exporting data corresponding to a particular target.

The best format for storing data associated with a detection is currently an open question. The present project opted for biasing on the side of over-completeness vs. sparseness, for the amount of information packed into each exported detection file: detection files contains information about the target and information about how the target was beamformed so that the data could be re-created from raw data given the information in the metadata associated with the detection. Opening a detection file in Matlab reveals a data structure with the following fields:

- *QP:* "Query point:" the location of the target
- *Files:* A cell array of the *.jsf* file names used to beamform the target looks
- *parms:* The beamforming parameters (e.g. integration angle, max search radius etc.) used to beamform each look
- *images:* A cell array of complex data and metadata associated with each look
- *label:* The user-created label for the target (e.g.: "155mm Howitzer shell")
- *class_score:* A value between 0 and 1 defined by user or ATR algorithm containing the classification score
- *UID*: A 32 digit unique ID number. This number will prevent ambiguity about which target is being specified when a large database of targets is assembled.
- *denoised*: A Boolean value that is set to "true" if data has been denoised. This value indicates the presence of a new field of data within the *images* cell array of look data structures.

The *images* cell array is an array of target structures containing information about each target look. This has a lot of information that is exploited by data visualization tools and spectral feature enhancement algorithms:

- *bsX*: ROI box size (meters) in the X (eastings) dimension
- *bsY*: ROI box size (meters) in the Y (northings) dimension
- *bsZ*: ROI box size (meters) in the Z (up) dimension
- *dR*: The voxel edge size (meters)
- *bXYZ*: The coordinate locations of each voxel of the ROI in UTM coordinates

- *img3*: The complex (i.e real + imaginary component) 3D look data produced by the patch beamformer
- *cyl3*: The complex 3D look data produced by the patch beamformer, beamformed in local *cylindrical* coordinates
- *cyl_axes*: A structure containing the axis and sample spacing definitions for the cylindrical coordinates plus projection information required to interpolate the cylindrical axis data onto Cartesian coordinates and vice versa (used for acoustic-color feature enhancement and masking operations).
- *w0*: The center frequency of the acoustic band of the data processed to form the image
- *c*: The local sound speed measured when the look was acquired
- *amax*: Integration angle maximum used for beamforming
- *contributors*: A list of the pings from the associated JSF file that were beamformed to generate the current data volume
- *xyzO*: The UTM coordinates of the centroid of the ROI box
- *ArrayXYZ*: The UTM coordinates of each array element during target acquisition (necessary for predicting the spectral coverage)
- *nAlt*: Nominal vehicle altitude above the sediment as estimated by the MuST altimeter at the time of acquisition.
- *nMind*: Ping index associated with the minimum distance to the ROI centroid
- *denoised*: A data structure, present if the image data has been denoised, containing filtered and unfiltered 3D acoustic color data, axis information, and the complex filtered data in the cylindrical coordinate spatial domain.

A detection exploration library has been developed as part of the MuST software toolbox that allows the user to explore detection data saved in the previously described format. This data exploration toolset provides a variation of the "Look-Editor" (e.g. Fig.'s 115 – 125], along with additional visualization tools that will be further discussed in the present Methods section of this document under "Data Visualization."

PMA: querying processed ROI's

The user can re-access all of the look data created in a PMA session, even if the ROI tool is closed by querying the processed location. After the data clicked "Beamform" and the selections shown in Fig. 110 were beamformed, the red stars denoting selections were changed to circles, the color of which are dependent on the classification score of the target (Yellow ~ = class score of 0.5, green = class score of 0, red = class score of 1). This is shown in Fig. 128. Note that the "Query" button in the PMA panel on the bottom right is now active.

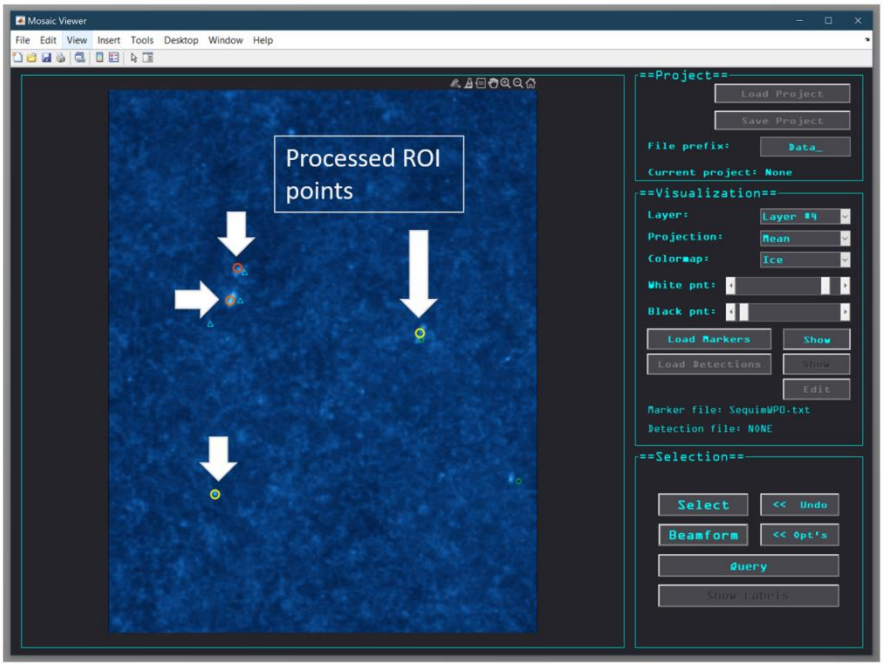

Figure 128: The processed ROI points are now displayed as colored circles. Note that the "Query" button is now active.

When the user selects "Query" in the PMA main screen the cursor will again become a cross-hair in the display screen. The user can select one of the query points by clicking on it, and this will bring open the ROI screen again (e.g. Fig. 112) and the user can perform all of the previously discussed operations provided by the ROI tool. Filtering and export information (i.e. whether the target has been exported) is all retained and reviewable in the ROI tool.

Summary

This section has explored the many capabilities provided to the end user by the PMA tool at the risk of creating the impression of an overly-complex tool. In reality the PMA tool is quite simple, having only three layers of interaction:

- PMA main screen: view mosaic data, select and query points of interest
- ROI review tool: multi-look data visualization, classification and data exportation
- Look-Editor: examine acoustic color data, make measurements, define filters

The apparent complexity only comes from the wide array of tools at the user's disposal within these layers. Many aspects of the PMA tool are still under development, for example, incorporation of an automatic target classifier, definition of project files for reloading previous project session data, and incorporation of modifications from the experience of users.

V. Structural Acoustic Feature Extraction

Overview

Certain robust, identifiable features for targets are not necessarily most obvious or most stable in the spatial domain. One alternative domain used for feature detection is the "acoustic color"

representation, in which the magnitude of the frequency response is measured as a function of target ensonification angle. This representational space, which was discussed previously in the background section of this document, has been used effectively for a number of classifiers developed specifically to tackle the UXO classification problem. As is generally the case, performance of classifiers using acoustic features – in the acoustic color *or* image domains – benefits from improvement in the signal-to-noise ratio (SNR) of features. Methods for improving the SNR of acoustic color features have been developed elsewhere (see, e.g. [6] and the associated interim report), however the present case, which involves three dimensional data products with non-uniform image statistics in the three spatial dimensions and highly non-uniform coverage in the spatial wavenumber domain, requires careful consideration.

A number of alternative approaches were explored in this project but the final algorithm settled on for automatic data denoising is a two-stage algorithm that combines a variant of the wavenumber-moment filter (the "*kM*" filter) developed in [6] with a spatial domain specular glint filter. A flowchart of the algorithm is shown in Fig. 129.

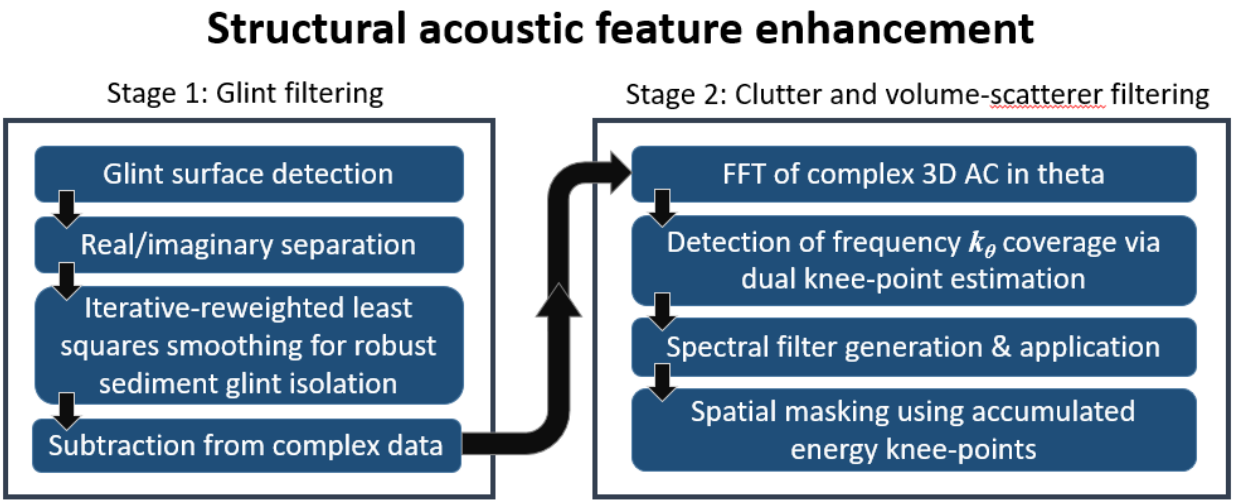

Figure 129: A flowchart of the structural-acoustic feature enhancement algorithm developed in this project. Stage 1 (left) removes the coherent portion of specular glints in the spatial domain, and stage 2 transforms the 3D acoustic color data created from the de-glinted image into the frequency - k_θ domain, where filter boundaries are auto-detected and applied.

The stages shown in Fig. 129 will be explained in more detail in subsequent sections.

Glint filtering

The de-glinting routine developed for target acoustic color enhancement is significantly different than the filtering method developed for large-scale batch processing. It is very important to avoid any sort of local intensity normalization in acoustic color processing because this can distort the spectral features associated with the structural acoustic response of targets. De-glinting therefore needs to be subtractive rather than multiplicative, i.e. the algorithm should isolate the spectral component and subtract the glint contributions directly from the complex backscattering data rather than null the regions having specular scattering via multiplication with a normalization function which will also scale any features of the target that lie in the same scattering plane.

The first step of the de-glinting process is simply to determine if a glint exists. This is done by checking where the target is relative to the sonar (i.e. if it is within some narrow angle relative to nadir), and search for a delta-function near the estimated sediment boundary. If the sediment is within the specular glint region and a glint is detected, the height plane of the sediment detection and a buffer of pixels surrounding it are extracted. A low-order polynomial model for the phase variation across the interface is estimated and a unity-magnitude complex phasor with a phase argument equivalent to the conjugate of the estimated phase is multiplied with the data to create slow spatial variation in the values of the real and imaginary components of the specular glint. The imaginary and real components are computed using a non-linear, iterative reweighted least-squares smoother that iteratively de-weights regions of rapid spatial signal variance (e.g. targets embedded in the interface). The resulting residual is multiplied with the phase screen and subtracted from the original data to remove the specular glint. This process is graphically illustrated in Fig. 130.

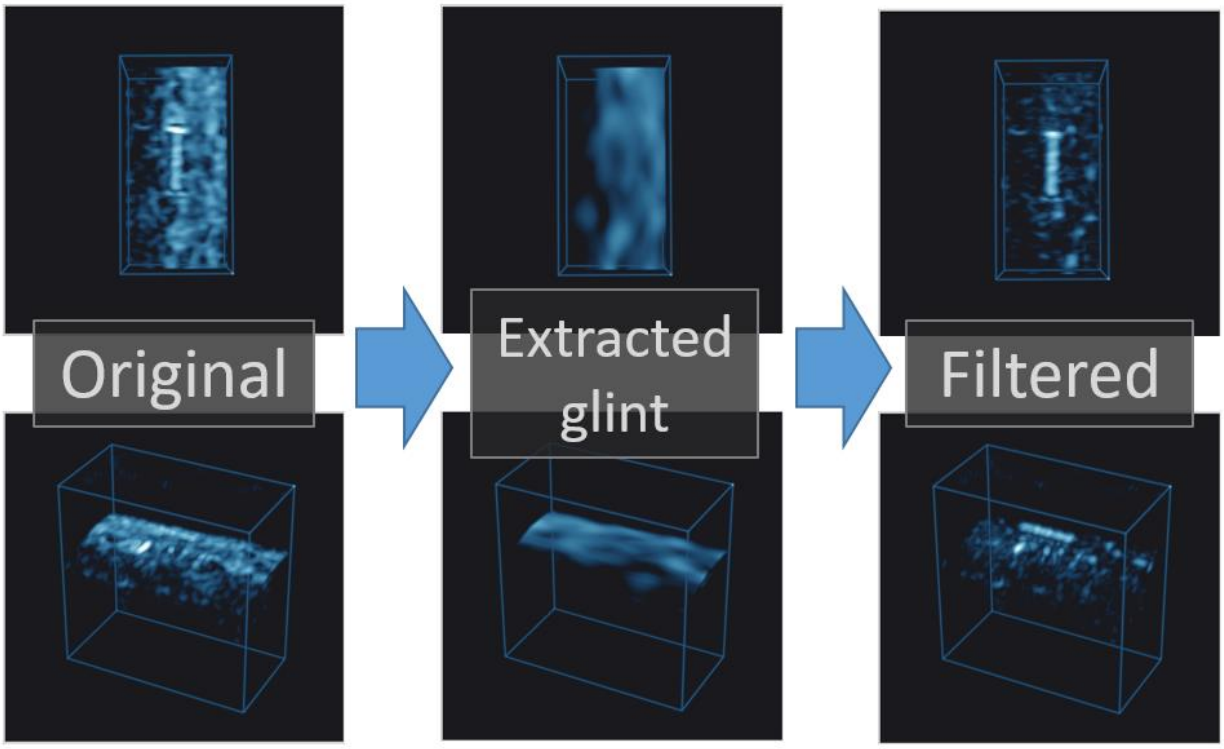

Figure 130: Visual depiction of the de-glinting process used for structural acoustic feature enhancement. The sediment glint feature is detected and isolated via iterative reweighted least-squares smoothing and subtracted from the original dataset. Two views of the original dataset are shown on the left, the extracted glint feature is shown in the middle, and the filtered dataset is shown on the right.

In the current script associated with the filter a flag exists that will prompt the filtering routine to 1) always apply the filter, or 2) only apply if the target is in a region predicted to contain sediment glints. The glint filter hasn't been found to have negative effects on the acoustic color data products generated for a region of interest, however it can slow down the filtering operation by several seconds as the present implementation hasn't been optimized for execution speed.

Volume noise reduction

The incoherent volume scattering component forms the dominant portion of the signal noise in most regions of interest. The next few sections describe how the *kM* filter of [6] was adapted to the present imaging scenario.

Origin shifting and transformation to the ω-k$_\theta$ domain

It was observed in [6] that transforming raw circular synthetic aperture target data into the ω-k_θ domain via a 2D Fast Fourier Transform conveniently condenses signal power while distributing noise power if, as a pre-processing step, the range-migration curves corresponding to the target location were flattened via application of an angularly varying delay function. In the ω-k_θ representation a wavenumber filter can be generated that very effectively improves the signal to noise ratio of structural acoustic features that become visible upon a subsequent application of an inverse Fast Fourier Transform in the k_θ dimension.

In the present case, the data is captured by linear, rather than circular scans and raw data is *not* immediately amenable to this type of filtering; however in a post-beamforming context it is possible to apply this same type of filter – provided beamforming has been performed with adequate spatial sampling and the correct transformations are applied as a pre-processing step.

The first step in transforming target data captured by the BOSS to the now three-dimensional {ω, k_θ, φ} domain (i.e frequency, moment, azimuth domain) where wavenumber filtering can be applied is to shift the centroid of the object of interest in the 3D spatial domain to the origin of the beamformed region-of-interest (ROI). This shift accomplishes two necessary things: it minimizes the phase gradients of the features associated with the target in the spatial-spectral representation, reducing interpolation artifacts, and secondly it condensing acoustic target information in the ω-k_θ representation in subsequent transformations, improving the filtering capability of the wavenumber filter.

This origin shift is performed by zero-padding the 3D data array by a factor equivalent to 75% of the size of the ROI in each dimension, then convolving the magnitude of the ROI with a 3D Hanning function having a size equivalent to a quarter of the size of the ROI in each dimension. Following filtering, the peak is detected and shifted to the first index of the array: the typical convention for the origin of FFT operations in Matlab (i.e. a unit impulse at this index has zero phase slope in all dimensions following an FFT. Generally, most FFT routines in most languages consider the first index of an array the “origin” in this sense). The large degree of zero-padding is to allow for the unlikely case that the target of interest may abut the edge of the ROI. It is virtually impossible to perform a spectral interpolation operation on data with spatial phase gradients that are critically sampled, and the excessive zero-padding prevents any of the original signal from having near critically sampled phase after origin shifting. Note that this processing is being performed on the *cylindrical coordinate* representation of the target data, not the Cartesian coordinate representation, because it aligns the synthetic and real aperture integration angles with the orthogonal spherical coordinate dimension θ and φ.

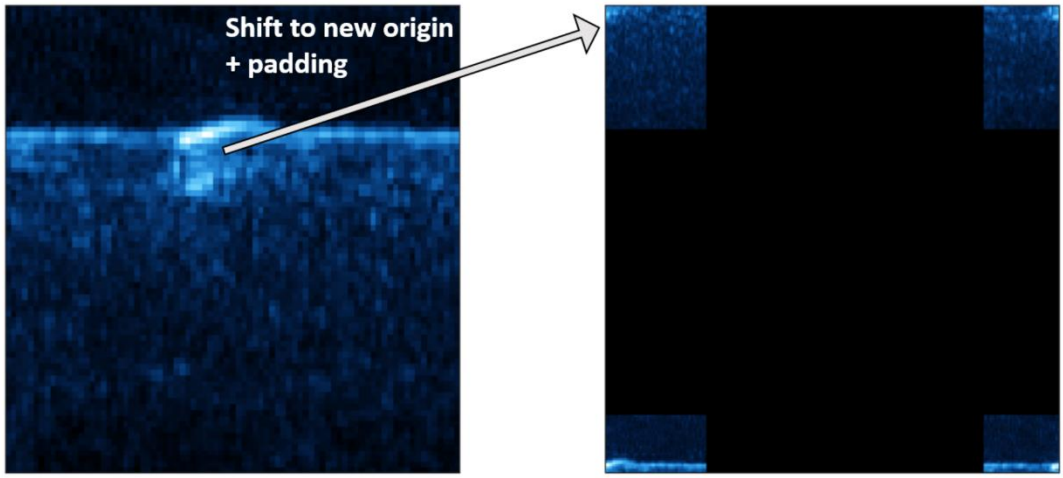

Figure 131: MIP through cylindrical φ (phi) axis showing a UXO target before (left) and after the origin shift + padding operation (right).

Subsequently, the 3D FFT of the padded and shifted ROI is computed. The result closely resembles a sampled region of the spectral hemisphere corresponding to the 3D target frequency response (e.g. Fig. 11 of the final report of [6]). In the present case, however, the orientation of the spectrum is relative to the synthetic aperture, and is not completely analogous because of the φ-dimension of the cylindrical coordinate ROI. As was described in the first interim report of [6], the data captured from a linear scan can be transformed into "CSAS" format by mapping from k_r and k_{AT} (the wavenumber axes corresponding to the range and along-track dimensions respectively) to ω-θ directly via a Cartesian-to-polar transformation. (This interpolation forms the zero-padding requirement in the spatial domain: rapid phase gradients in the spectrum must be surpressed for this to be successful). Following this interpolation a Fourier transform in the θ dimension is performed to convert the transformed data into the ω-k_θ domain – where wavenumber filtering can be applied. This CSAS format transformation process is graphically described in Fig. 132.

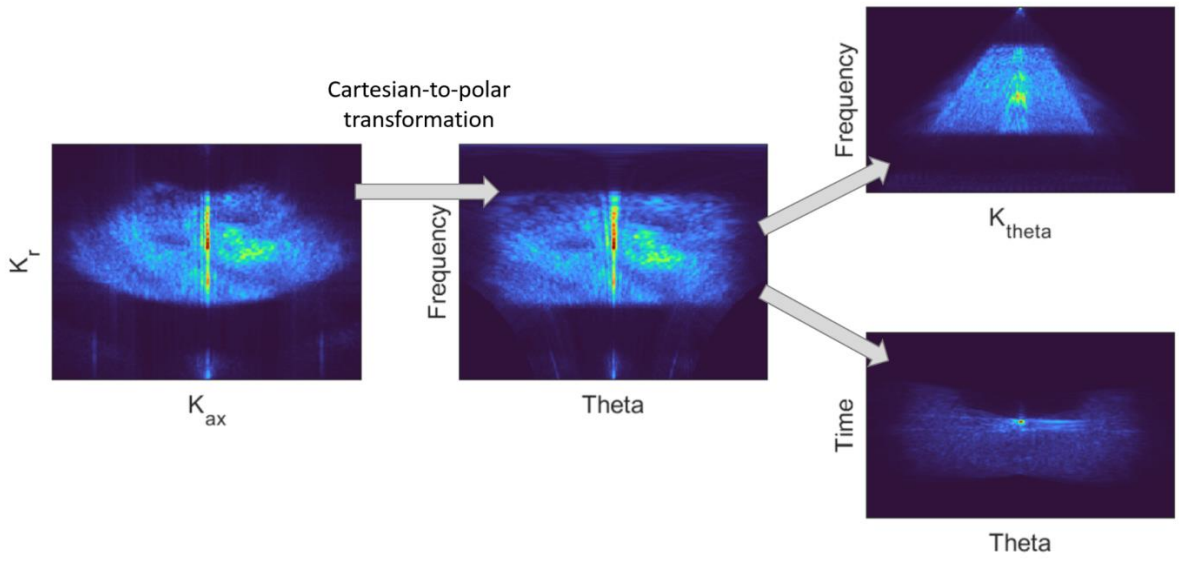

Figure 132: Mapping from a 3D image spatial spectrum to CSAS-format data. The left-most image shows a MIP of the 3D FFT of the cylindrical-format image cube (projected through real-aperture angle φ) showing the sampled segment of the wavenumber spectrum resembling an annulus. This is mapped to the frequency vs. angle domain via a Cartesian-to-polar transformation (center). The FFT can be computed in the look angle (theta) dimension to arrive at the ω-k_θ representation (where a wavenumber filter is generated and applied) on the top right, or in the CSAS-like time-angle representation (bottom right), which shows the impulse response vs. look angle for the target.

Wavenumber filtering

Following the transformations described in the previous section, the axes of the 3D data are (k_φ,ω,k_θ). Noise is not strongly separate from the signal in k_φ because it is the array dimension and therefore it is has less spatial bandwidth relative to the other bands. Therefore a MIP is performed through k_φ to show the predominant signal distribution in ω-k_θ for the construction of a wavenumber filter.

As described in the first interim report of [6], a ω-k_θ filter $M_{\Theta_1\Theta_2}(\omega, k_\theta)$ is a spectral mask created to preserve the regions between angles Θ_1 and Θ_2, defined as:

$$\boldsymbol{M_{\Theta_1\Theta_2}(\omega, k_\theta) = \begin{cases} 0 \; for \; \Theta < \Theta_1 \\ 1 \; for \; \Theta_1 \leq \Theta \leq \Theta_2 \\ 0 \; for \; \Theta > \Theta_2 \end{cases}} \qquad \textbf{Eq. 63}$$

where:

$$\boldsymbol{\Theta = \tan^{-1}\left(\frac{k_\theta}{\omega}\right).} \qquad \textbf{Eq. 64}$$

The degree of filtering is determined by the spacing of angles Θ_1 and Θ_2, and it is worth noting that one of the reasons for applying the origin shift is to cluster the distribution of target energy in the ω-k_θ representation to minimize this spacing. The optimal values for angles Θ_1 and Θ_2 are automatically selected via a knee-point detection process. For the MIP $\boldsymbol{S_{MIP_k\varphi}}(\omega,k_\theta)$ of the transformed signal $S(k_\varphi,\omega,k_\theta)$, the keystone transformation:

$$S_{MIP_k_\varphi}(\omega, k_\theta) \rightarrow S_{MIP_k_\varphi}(\omega, \tilde{k}) \,, \qquad \textbf{Eq. 65}$$

$$\tilde{k} = \left(\frac{\omega}{\omega_0}\right) k_\theta \,, \qquad \textbf{Eq. 66}$$

is performed, which aligns acoustic MIP energy over Θ. The root-mean-square is then computed to determine the distribution of target energy versus $\tilde{k}$:

$$\overline{S}(\tilde{k}) = \sqrt{\frac{\sum_{i=1}^{M}\left(S_{MIP_k_\varphi}(\omega(i), \tilde{k})\right)^2}{M}} \qquad \textbf{Eq. 67}$$

The distribution has two levels; a background level from random scattering and a region of increased signal due to scattering from the target. The region over which target energy is distributed is determined by finding the index corresponding to the "center" of the target distribution by aggressively smoothing $\overline{S}(\tilde{k})$ and finding the peak and accumulating outward in both directions. This forms a set of curves with well-defined knee-points.

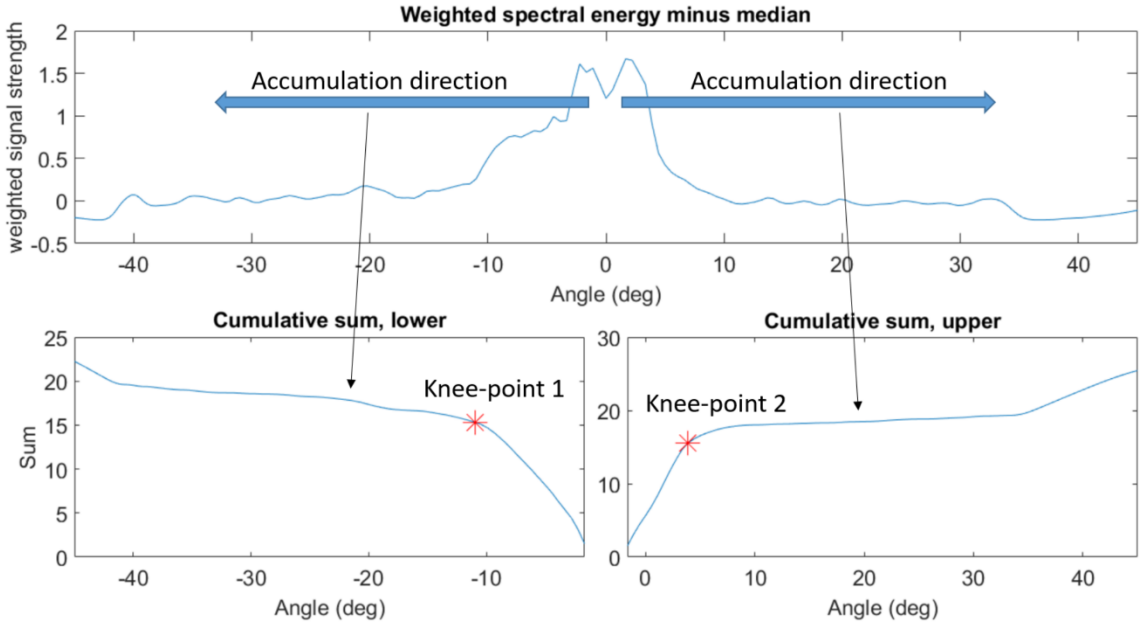

Figure 133: Knee-point detection. The weighted spectral energy is integrated outwards from both side of the estimated center-point and the knee-points of the integration are used as the filter boundaries.

A standard knee-point estimation technique that finds the intercept of two linear functions minimizing the sum of squared residuals is used to determine the knee-point locations. The indices associated with the knee-points are used to calculate Θ_1 and Θ_2 via:

$$\Theta_{1,2} = \tan^{-1}\left(\frac{\tilde{k}_{1,2}}{\omega}\right). \qquad \textbf{Eq. 68}$$

Using these two values the filter $M_{\Theta_1\Theta_2}(\omega, k_\theta)$ is constructed and multiplied with $S(k_\varphi,\omega,k_\theta)$ to denoise the signal. In practice a set of error-functions is used to smoothly taper the preserved spectral region, as is done for spectral masking in [6] and [36], e.g.:

$$\bar{\boldsymbol{S}}\left(\boldsymbol{k_\varphi}, \boldsymbol{\omega}, \boldsymbol{k_\theta}\right) = \boldsymbol{M_{\Theta_1\Theta_2}}(\boldsymbol{\omega}, \boldsymbol{k_\theta}) * \boldsymbol{S}\left(\boldsymbol{k_\varphi}, \boldsymbol{\omega}, \boldsymbol{k_\theta}\right) \quad \textbf{Eq. 69}$$

where:

$$\boldsymbol{M_{\Theta_1\Theta_2}}(\boldsymbol{\omega}, \boldsymbol{k_\theta}) = \frac{\left(erf\left(\alpha(\Theta-\Theta_1)\right)+1\right)*\left(erf\left(-\alpha(\Theta-\Theta_2)\right)+1\right)}{4} \quad \textbf{Eq. 70}$$

In Eq. 69, $\bar{S}$ with the overscore represents the filtered version of the original data S, with the coordinates of the current transformation being defined by the variables in parenthesis. In Eq. 70 α is a tunable mask sharpness parameter, Θ is substituted for compactness but can be expressed in terms of ω and k_θ directly via Eq.'s 68 and 66. Graphically the masking is shown in Fig. 134.

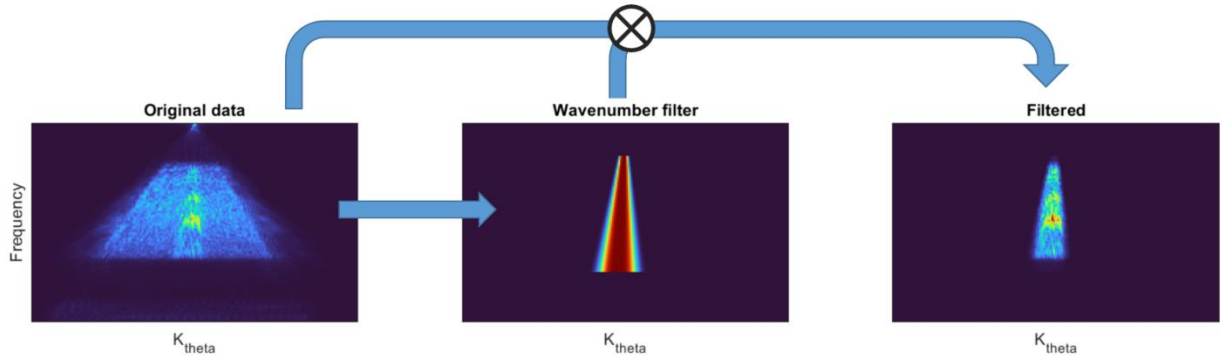

Figure 134: Application of the wavenumber filter. The left figure shows a MIP through the raw data, $S(k_\varphi,\omega,k_\theta)$, which is used to define the coverage of the wavenumber filter via the knee-points shown in Fig. 32. The filter (center) is then directly multiplied with the data.

A de-noised acoustic color data projection $\bar{\boldsymbol{S}}_{\boldsymbol{MIP_k_\varphi}}(\omega, k_\theta)$ can be recovered by applying an inverse Fourier transform to $\bar{S}\left(k_\varphi, \omega, k_\theta\right)$ in the k_θ dimension and projecting through k_φ. This was shown in the second interim report of this project. Note however, that if the filtering process stops at this point then no time-gating has been applied. This is significant, because the structural acoustic features associated with a target have a practical time duration associated with then (e.g. for elastic resonances, related to the quality factor), and beyond this time duration including additional range data only increases the noise of the acoustic-color representation. To deal with this problem a time-gating algorithm with temporal thresholds determined by energy integration and knee-point selection, similar to the wavenumber filtering approach, is again employed. To determine the temporal thresholds an inverse 3D FFT is applied to the masked data to put it into "CSAS" format to produce angle and time axes (see Fig. 132, bottom right):

$$\bar{\boldsymbol{S}}(\boldsymbol{\varphi}, \boldsymbol{t}, \boldsymbol{\theta}) = \boldsymbol{\mathcal{F}^{-3}_{k_\varphi,\omega,k_\theta}}\left\{\bar{\boldsymbol{S}}\left(\boldsymbol{k_\varphi}, \boldsymbol{\omega}, \boldsymbol{k_\theta}\right)\right\} \quad \textbf{Eq. 71}$$

In this format maxima of the envelope are computed across the φ and θ dimensions to create a time signal showing the maximum values of the signal level:

$$\overline{\boldsymbol{S}}_{\boldsymbol{MAX}(\boldsymbol{\varphi},\boldsymbol{\theta})}(\boldsymbol{t}) = \boldsymbol{MAX}_{\boldsymbol{\varphi},\boldsymbol{\theta}}\{|\overline{\boldsymbol{S}}(\boldsymbol{\varphi},\boldsymbol{t},\boldsymbol{\theta})|\} \quad \textbf{Eq. 72}$$

An example is shown in Fig. 135:

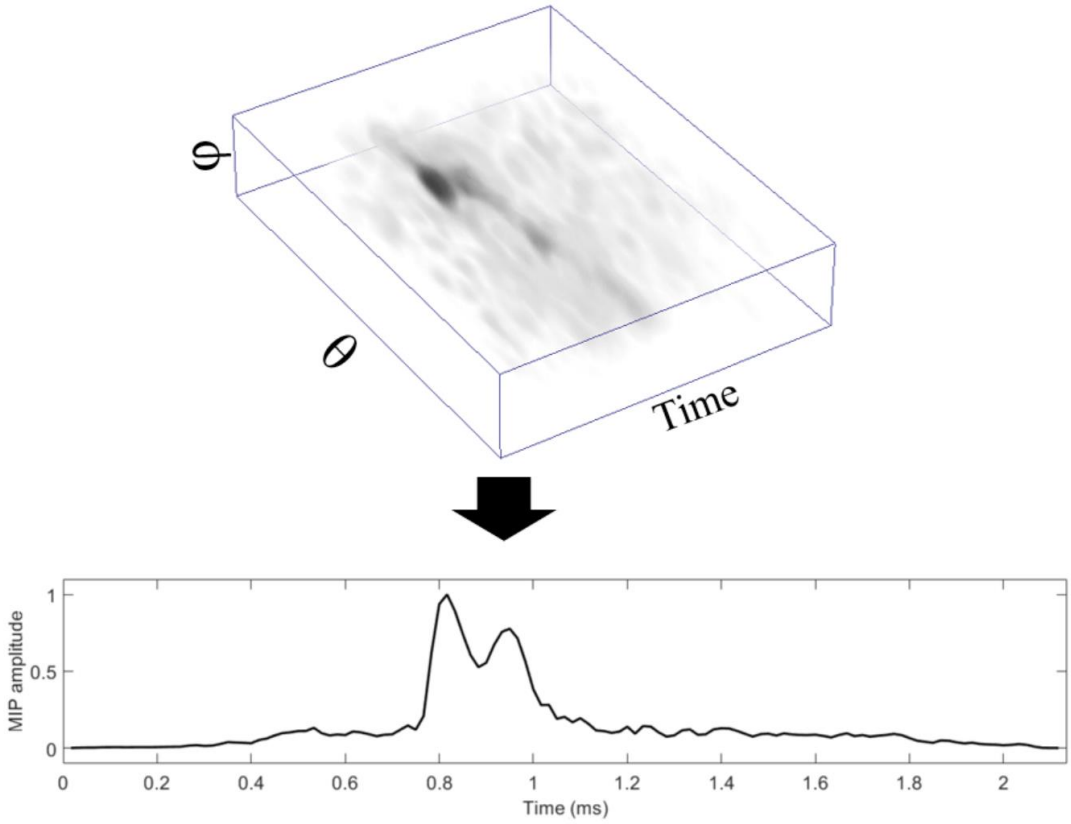

Figure 135: A maximum intensity projection of the filtered "CSAS format" data $|\overline{\boldsymbol{S}}(\boldsymbol{\varphi},\boldsymbol{t},\boldsymbol{\theta})|$ (top) is performed along θ and φ to create the time series on the bottom, used to determine the temporal gating thresholds for the data.

The centroid of the data could be estimated and the outward-accumulation again used to estimate the knee-points, however often when dealing with low SNR datasets the signal value doesn't have enough prominence relative to the background to make the knee-points unambiguous. (This is not true in the present example but can occur when small targets are embedded in the sediment volume). To increase the prominence of the valid signal region, increasing the sharpness of the knee-point, the background is estimated by performing a robust fit of a quadratic polynomial to the signal via iterative-reweighted least squares. The quadratic curve is then subtracted from the signal as shown in Fig. 136 and the signal is smoothed to determine the approximate centroid of the high-signal region.

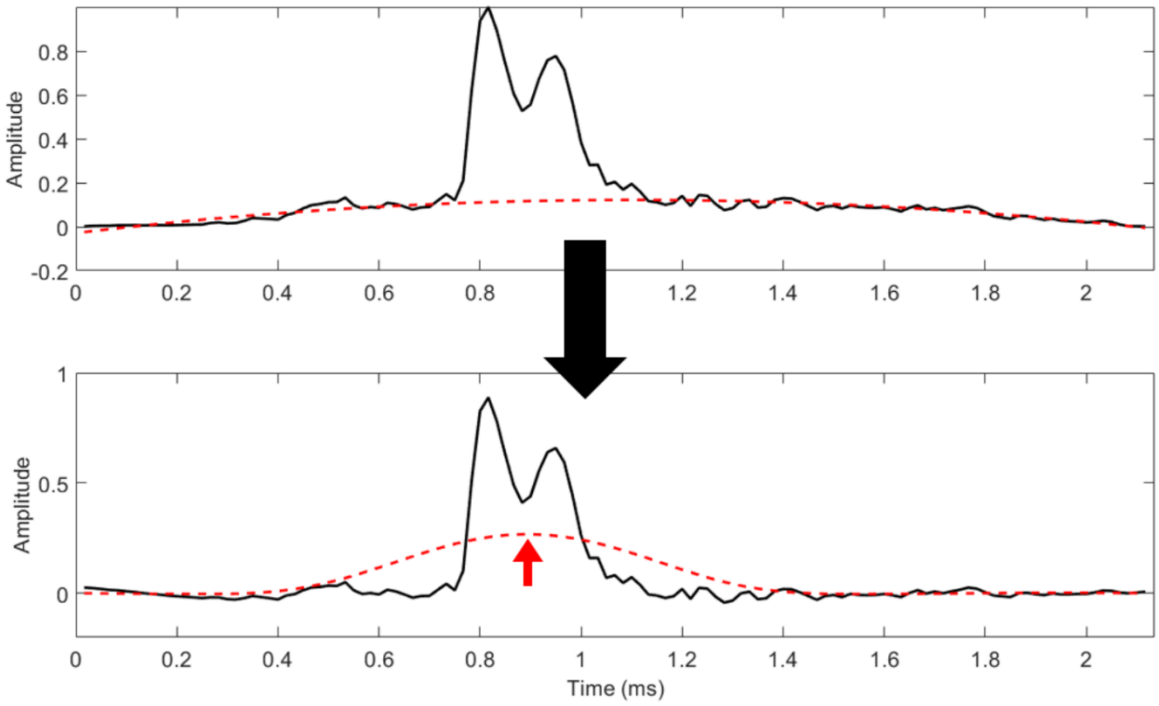

Figure 136: The projected data (solid black line, top) and robust quadratic fit (red dashed line top), and the compensated projection (solid black line, bottom) and the smoothed compensated line (red dashed line, bottom). The red arrow points to the peak of the smoothed curve. This point determines the center point for outward accumulation.

Identically to the process previously shown in Fig. 133, curves are generated by accumulating outward and finding the knee-points of the curves. The range-gates are determined by the location of the upper and lower knee-points with an adjustable pad-value that is currently set to ~1/3 the length of the distance between the knee-points and a new mask is created via the dual-error-function method again, e.g.:

$$\bar{\bar{S}}(\boldsymbol{\varphi}, t, \boldsymbol{\theta}) = \boldsymbol{M}_{\tilde{\tau}_{1,2}}(t) * \bar{\boldsymbol{S}}(\boldsymbol{\varphi}, t, \boldsymbol{\theta}) \qquad \textbf{Eq. 73}$$

where:

$$\boldsymbol{M}_{\tilde{\tau}_{1,2}}(t) = \frac{(erf(\alpha(t-\tilde{\tau}_1))+1)*(erf(-\alpha(t-\tilde{\tau}_2))+1)}{4} \qquad \textbf{Eq. 74}$$

$$\tilde{\tau}_{1,2} = \tau_{1,2} \pm 0.3(\tau_2 - \tau_1) \qquad \textbf{Eq. 75}$$

Note that in Eq. 73 the filtered and range-gated signal is denoted with a double overscore, $\bar{\bar{S}}$. The final implementation of this filter included a third filtering operation along the φ dimension which operates in the exact same manner as the range-gating operation, i.e. the time axis can be swapped with the φ axis and Eq.'s 72-75 recomputed for application of a third mask in the φ dimension. At this final point the filtered signal can be transformed into any other domain via the appropriate transformations and the relevant projections made for visualization or classification (e.g. through φ). Generally this filtering process worked extremely well when a target was

compact. Failure cases tended to be when targets had low signal-to-noise ratio distributed features, which cause breakdown in the selection of knee-points due to an inability to establish which features constitute clutter and which features constitute signal. Examples will be shown in the results section.

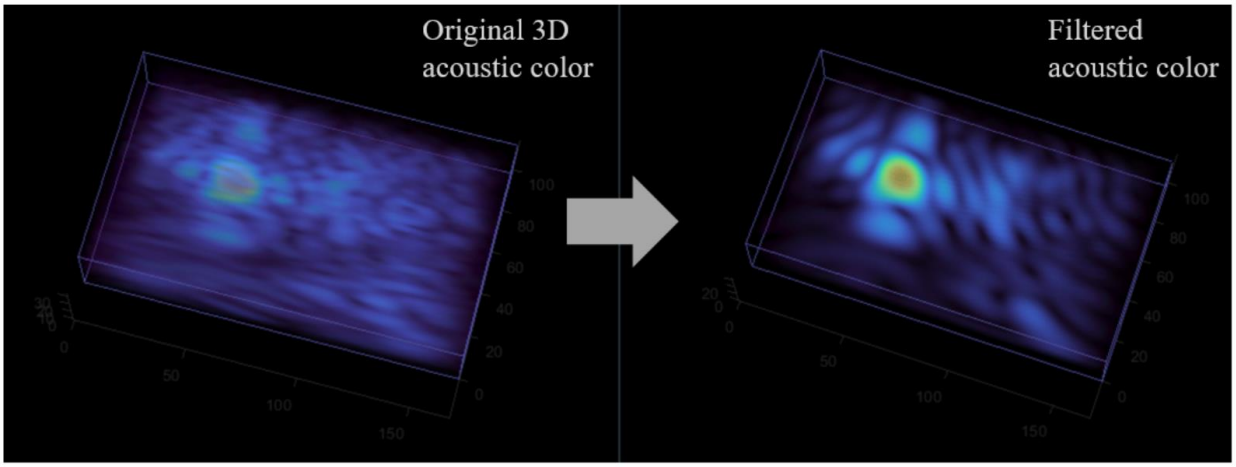

Figure 137: The 3D acoustic color dataset used to create the projection shown in Fig. 135 before (left) and after (right) the application of the structural acoustic enhancement filter described in this section.

VI. Additional visualization considerations and tools

MIP vs Alpha renderings

As discussed in the background section of this document, one of the technical objectives of this project was to improve the standard methods used to visualize BOSS data which, via a simple literature review (e.g. [1-5]) is almost entirely limited to orthogonal MIP projections of single look data. Before venturing further it ought to be pointed out that orthogonal MIP projections are the visualization method of choice for a good reason: general experience with noisy single-look acoustic volume data shows that MIP's are 1) the technique that most reliably shows the significant features of the volume with the highest signal-to-noise ratio and 2) are extremely fast to compute when performed along the orthogonal views of the volume. The universally acknowledged downside of MIP's are the loss of 3D spatial relationships [37]. This is certainly true in the single-look case, less true when the combination of projections through different volume axes are viewed as a set, which is typically why MIPs are shown as an image pair (e.g. in MuSTLive) or a triplicate. Even more optimally, MIP images are viewed with animated rotation, or augmented with distance scaling and range-attenuation, which provides additional clues regarding the spatial relationships between features. Fig. 138 is an example of a set of rotated MIPs augmented in this manner, created using a custom augmented MIP rendering tool created as part of this project.

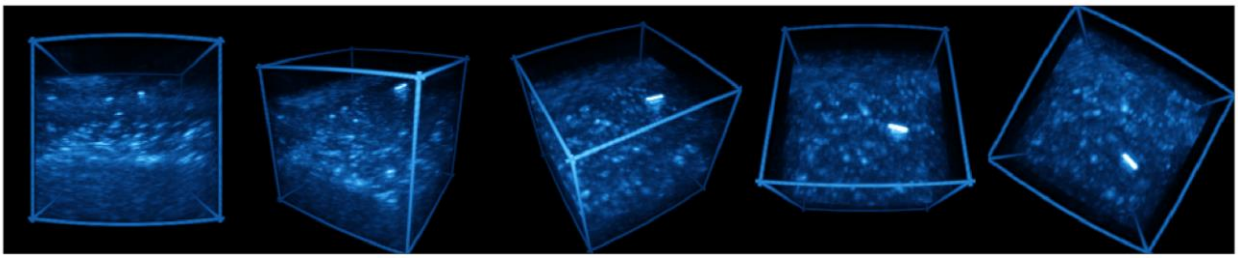

Figure 138: Augmented MIP views through a volume containing a UXO facsimile target. The MIP volume is rendered from the perspective of a user-defined camera location. Attenuation as a function of range combined with perspective scaling and multiple looks creates a greatly improve sense of the 3D spatial relationships between scatterers.

Volume rendering via alpha-blending, which calculates the attenuation of rays passing through voxels with predefined transparency values, is a popular alternative to MIP rendering. The occlusions caused by transparency tend to be very helpful for visually preserving spatial relationships – at the cost of potentially obscuring the most important features of the volume. Particularly with noisy volumes, alpha-blending can be problematic for this reason: transparency thresholds have to be set to allow the most significant features to be visualized, which often involves placing the transparency threshold above the level of the noise. There is often a tradeoff therefore, in volume rendering noisy data: spatial relationships can be preserved at the cost of losing data due to thresholding (alpha blending), or spatial scales can be lost with the benefit of preserving a higher dynamic range (MIP rendering). This tradeoff is demonstrated in Fig. 139, which compares MIP and alpha-blending techniques used on spheres embedded in noise.

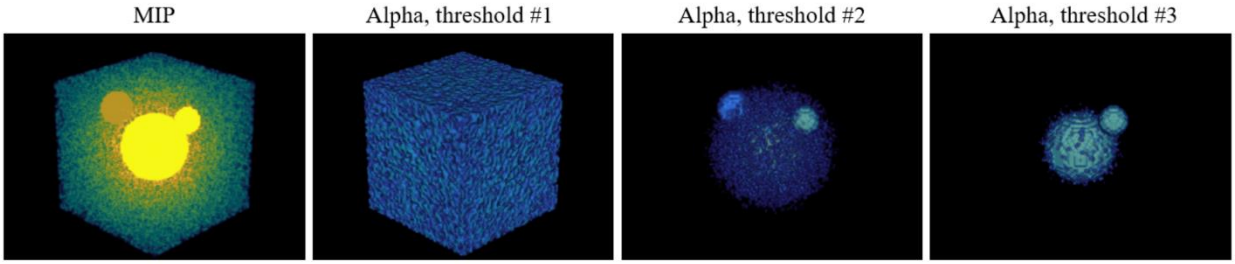

Figure 139: An illustration of the tradeoff between feature preservation and spatial relationship preservation for spheres embedded in non-uniform noise. The MIP rendering on the left clearly shows the key features of the volume, including the three spheres and spatial variance of the noise, but all spatial scale is lost. In contrast, alpha-blending (right 3 figures) preserves the spatial relationships of the spheres, but no single threshold captures the major features of the volume.

The practical implications for BOSS data visualization is that alpha-blending and MIP renderings are often best used in pairs; the alpha-blending can be used to indicate to users the relative spatial scale placements of features and visualize lower intensity scatterers, and MIP renderings allow the most significant scatterers to be easily spotted even in high levels of noise. Examples of MIP and alpha-blending pairs are shown for a ship and for the target depicted in Fig. 138 in the following figures.

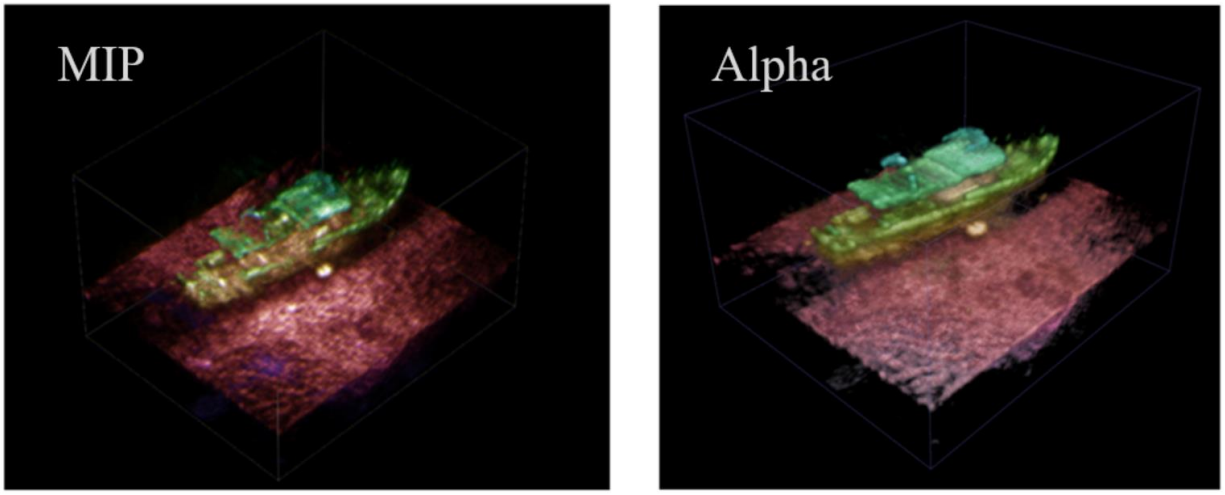

Figure 140: Side-by-side renderings of a multi-scan volume mosaic containing a fishing vessel. The MIP rendering on the left shows structural details inside of the hull, whereas the alpha-blended rendering on the right is more useful for visualizing the external features of the vessel.

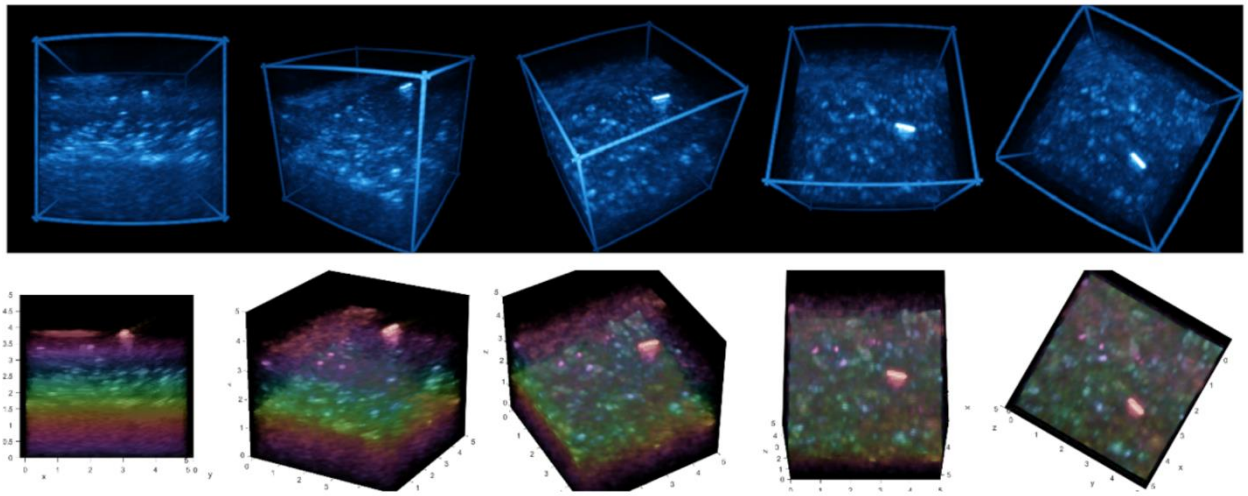

Figure 141: Augmented MIP's (top) of a multi-layer sediment volume containing a target, juxtaposed with color-by-depth alpha renderings generated in VOXview.

Note that both Fig. 140 and 141 include renderings that employ a color-by-depth strategy. This strategy has proven invaluable for identifying the relative displacements of various features in the height dimension and it is frequently used in volume renderings where vertical discrimination is important. This type of visualization encoding has been provided to users working with the SAS_view data display tool as demonstrated in Fig.'s 92 and 93. In general, both the PMA tool and SAS_view allow the user to choose either alpha or MIP renderings when selecting a volume visualization method, and the rendering is often accompanied with dynamic range compression and thresholding options; the reader is referred to the previous sections on the PMA ROI review and Look-Editors, and the SAS_view tool for further detail.

Unincorporated data visualization tools

A number of additional region-of-interest (ROI) multi-look tools have been developed that have not yet been incorporated into the PMA tool but are part of the MuST BOSS Toolbox and are accessible as command-line functions. These include tools for wavenumber coverage visualization, color-by-aspect rendering, 3D spatial-spectral data visualization, and linked multi-mode volume rendering. These tools are briefly described in subsequent sections.

k_coverage*(<<IM_dats>>)*: Spectral coverage rendering

In some applications it is beneficial to know how much wavenumber coverage has been created in a multi-aspect scan of a target; for example, it can be used to determine how much trust can be placed in a given classification score by depicting how much wavenumber diversity the classifier has to work with. Wavenumber coverage can be calculated by looking at the angles of the array phase centers to the target as a function of position along the synthetic aperture array. By examining all of the angles captured for all of the different looks acquired at a target the total three-dimensional wavenumber coverage can be visualized. (This also is an example of a reason for storing the array locations as metadata along with the complex backscattering values).

This analysis is performed by the function *k_coverage(*<<input>>*)*, where the input to the function is any ROI data structure created using the "Export Cur. Target" option in the ROI review tool (e.g. Fig. 126 & 127). The output of this function is shown in Fig. 142.

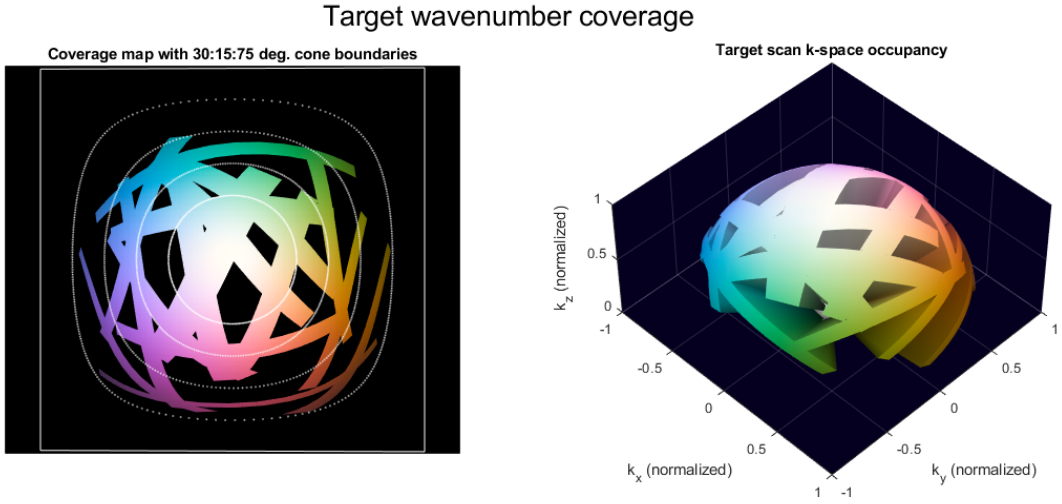

Figure 142: The wavenumber coverage displays created using the *k_coverage*() tool for an example ROI.

In Fig. 142 the wavenumber coverage is shown in spherical coordinates on the left, with conical boundaries corresponding to {30, 45, 60, 75} degree angles from nadir overlaid on the coverage map. The figure on the right shows the wavenumber coverage directly in wavenumber coordinates $k_{x,y,z.}$ A color-by-aspect scheme is superimposed that is both adopted in color-by-aspect renderings of multi-look fusions, and useful for matching coverage features between the spherical coordinate and Cartesian coverage maps.

Note that in practice the wavenumber coverage is not strictly limited to the regions defined in Fig. 142 because the angle of incidence changes depending on where the scatterer lies in the

volume. Refraction through a sediment interface for objects suspended in the sediment volume could be an additional confounding factor.

CBA_ROI(): Color-by-aspect rendering

Color-by-aspect (CBA) rendering [27] is a useful tool that maps spectral location to color, simplifying differentiation between backscattering mechanisms and features captured by different acquisitions. Color-by-aspect rendering is normally performed by operating on a coherently combined dataset, applying spectral windows, associating rotating colormaps to the spectral windows based on the backscattering, and fusing the results in a single RGB image. In the present case the data is not coherently combined and an approximation to color-by-aspect rendering is generated by determining color for individual looks based on the heading angle of the point-of-closest approach, and color-map saturation via the angular separation between nadir and the ray to the target at the point of closest approach (scans over the target don't have a well defined azimuthal backscattering angle and are rendered with reduced saturation).

The *CBA_ROI*(<<IM_dats>>, <<Power_mean_value>>) tool generates color-by-aspect volumes that can be rendered via a color MIP or VOXview using a color-by-aspect scheme. The first input to the function is the standard data structure exported from the PMA ROI review tool, and the second input sets the exponent value of the power-mean used to combine contributing looks. Lower values emphasize anisotropic scatterers, higher values emphasize brighter scatterers. Examples of CBA images computed with different power means for two different targets are shown in Fig. 143.

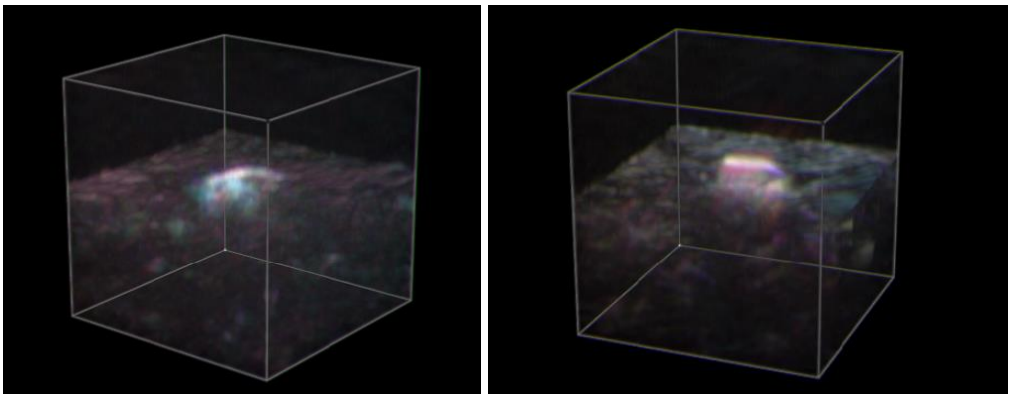

Figure 143: Two ROI rendered using CBA coloration. The left target is using a power mean exponent of 1, the volume on the right is rendered using a power mean exponent of 2.

CBD_ROI(): Color-by-depth rendering

Color-by-depth (CBD) rendering is an alternative to color-by-aspect rendering in which hue is determined by voxel depth rather than ensonification angle. This type of visualization is helpful to separate scatterers based on height-plane. Note that though this capability is included in the PMA tool the CBD_ROI(<<IM_dats>>, <<Lehmer_mean_power>>, <<Rotation rate>>) function gives the user control over the sub-look combination method (Lehmer-mean exponent value), and color rotation rate. Examples of several ROI showing various proud and buried objects detected in the Sequim Bay 2021 UXO testbed surveys are shown in Fig. 144. As in all previous ROI snapshots the edges of the imaged volume are 2 meters in length.

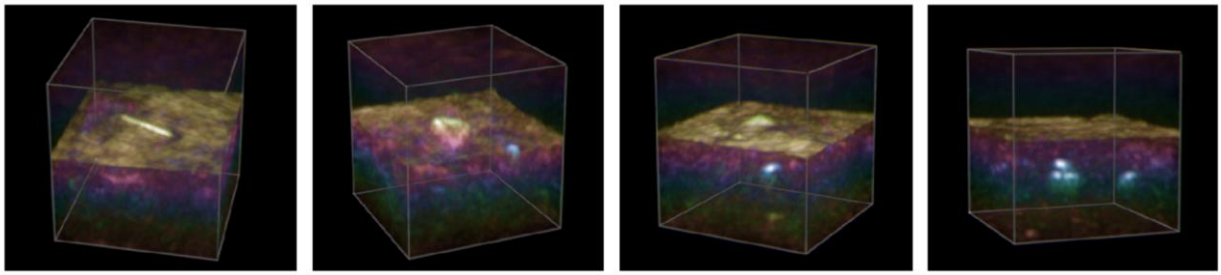

Figure 144: Color-by-depth ROI volume visualizations rendered using color-MIP's. Objects are detections found during the 2021 Sequim Bay UXO testbed survey.

RenderSpect(): Spatial-spectral data rendering

The *RenderSpect*(<<IM_dats>>, <<Looks>>, <<visType>>, <<denoised>>) allows the user to view the spatial spectra of the different looks contributing to an ROI detection and gives the user considerable flexibility in how the data is viewed. The user can select the specific looks to include (as an array of look indices, or the string 'all' if all are to be included in the visualization). The user can determine if MIP or alpha-blending is used for volume rendering, and select whether or not the data visualized is the original data or the denoised data with enhanced structural acoustic features. This is useful for identifying recognizable structural acoustic features that are not well represented by any individual acquisition. An example of this is shown in the figure 145, in which the spatial spectrum from a few independent looks are compared to the aggregate.

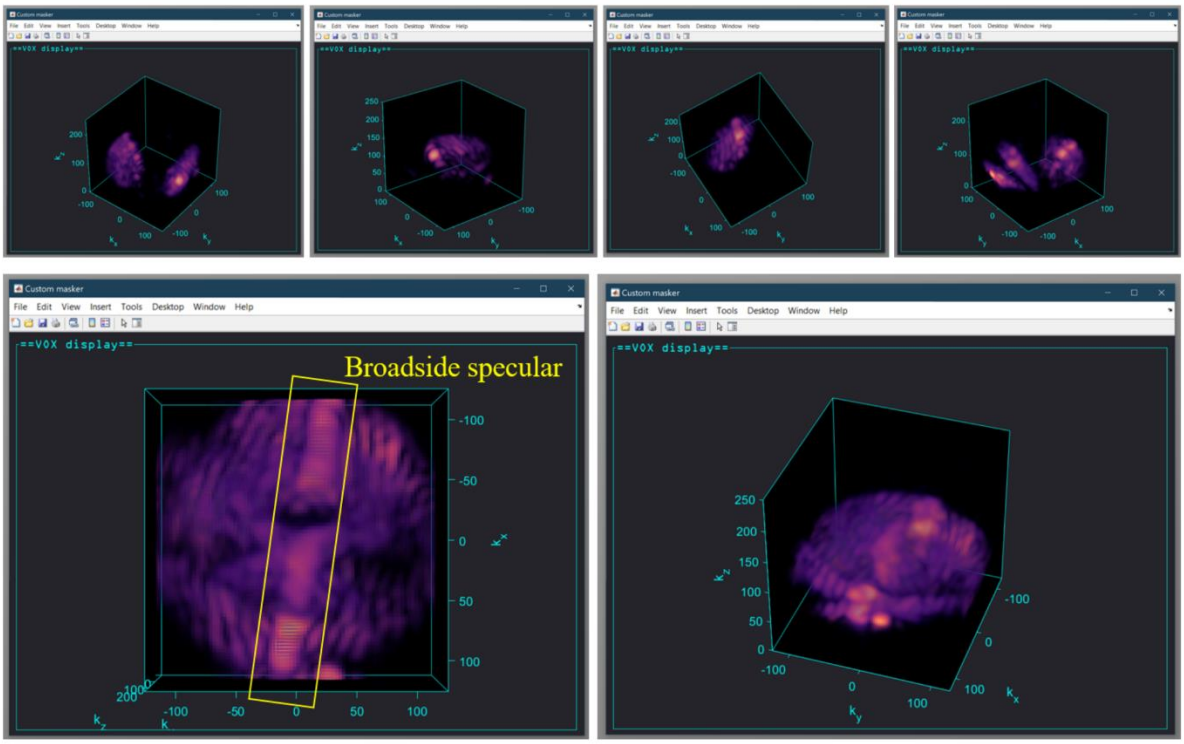

Figure 145: RenderSpect showing spectral data for 2, 1, 1, and 3 separate scans (top row) and the ROI aggregate (bottom). Note that the aggregate creates a piece-wise reconstruction of features such as the broad-

side specular and interference fringes from corner scattering that are not well captured by any single scan within the set.

A variety of visualization options are presented if MIP rendering is selected. Figure 146 shows the aggregate MIP with and without denoising. The central spike created by the nadir glint is readily visible in the noisy version.

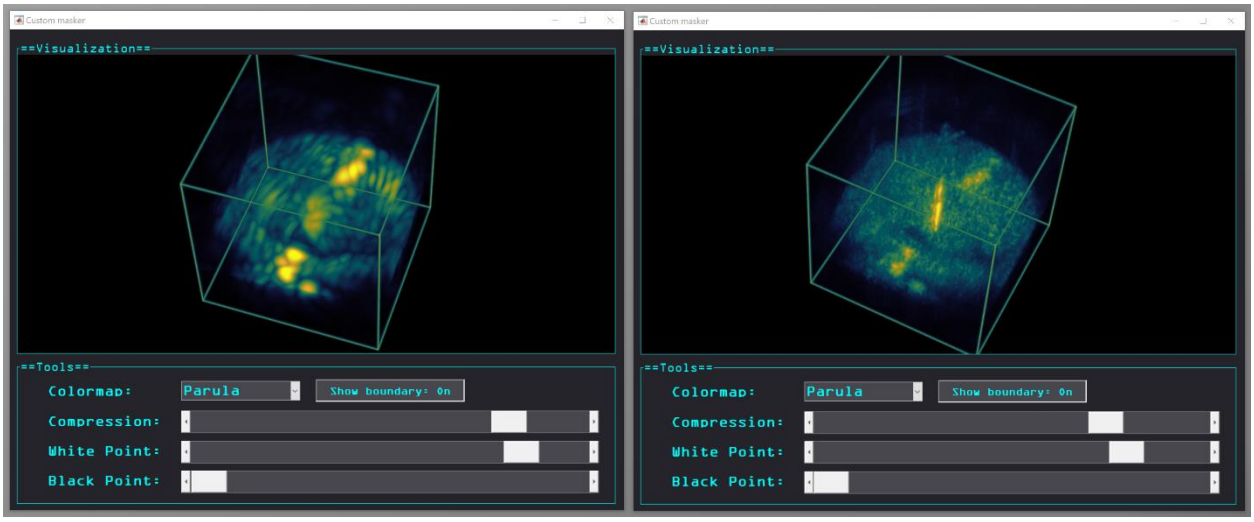

Figure 146: MIP views of the denoised aggregate spectrum (left) and the original noisy aggregate (right). Note the bright vertical spectral line extending upwards from the center of the spatial spectrum: this spike corresponds to the sediment specular glint. This feature is completely absent in the spectrum on the left.

ROI_TriView(): Linked multi-modal volume rendering

The different advantages and disadvantages of MIP and alpha-blended renderings were previously discussed. The ROI_TriView(<<IM_dats>>, <<Looks>>, <<IsoValue>>) tool allows the user to display both types of visualizations plus an iso-surface rendering (i.e. the surfaces carved out by a pre-defined threshold value) from a linked perspective. As with RenderSpect, the user can define which looks to combine, and can also define the surface value for the iso-surface rendering. An example of this type of visualization is shown in Fig. 147.

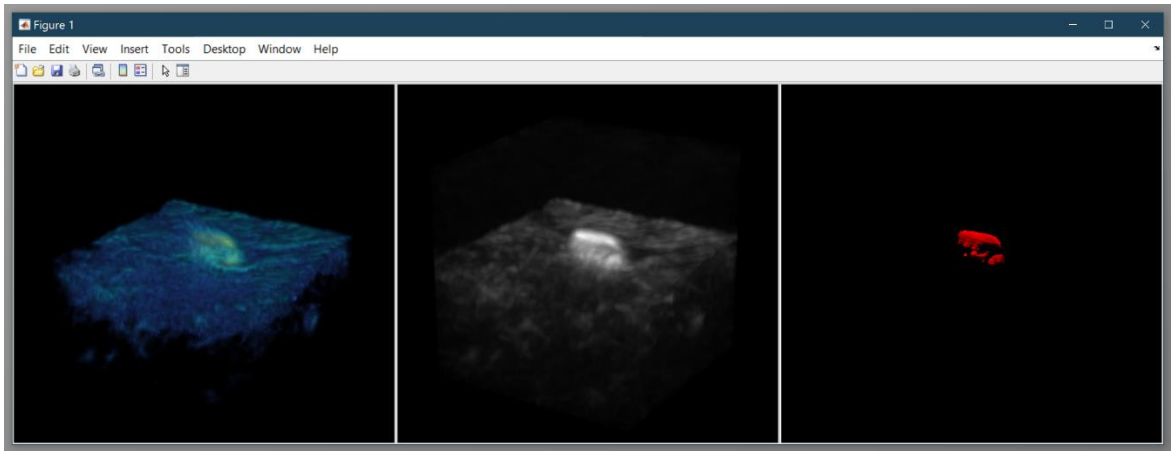

Figure 147: The ROI_TriView visualization showing an alpha-blended rendering (left), MIP rendering (middle) and iso-surface rendering (right) viewed from the same perspective.

Results

I. Overview

In the Materials and Methods section of this document the simulation tool and various BOSS sonars used for data collection were described alongside the tools that were developed for analyzing the data acquired by them. The *Results* section of this document will first cover the BOSS data simulator and 3D FFBP beamforming performance tests. This will be followed by an overview of the data acquired by the ONR AUV system. Subsequent datasets acquired from BOSS systems were collected using the Multi-sensor Towbody (MuST). MuST BOSS acquisitions that will be described in this document include:

- Initial MuST data acquisition test (8/2019)
- First MuST Sequim Bay survey (9/2019)
- MuST GS4 INS integration test (6/2020)
- Second MuST Sequim Bay survey (9/2020)
- MuST IXBlue INS integration test (4/2021)
- Third MuST Sequim Bay survey (9/2021)

Many of these surveys were system tests and did not include scans of UXO or science targets (e.g. INS integration tests). Nonetheless data acquired from these surveys were useful for evaluating BOSS capabilities and the developed post-processing tools, so results will be included along with the UXO-specific data acquisitions.

II. BOSSim and 3D FFBP beamformer results

Overview

The BOSSim simulator was used to create data for testing the 3D FFBP algorithm and compare its performance with generalized back-projection (BP, also referred to as "time-delay-and-sum"). To test the *generality* of the beamformer, the simulated data was subject to perturbations in all 3 translation dimensions (surge, heave, sway), and all three orientations (roll, pitch, yaw). For relevance, design parameters for the MuST system were used: hydrophone size, spacing, and transmitter size, spacing, and firing order. The full simulation parameters can be found in Table 2 of the *Materials and methods: BOSS data simulation* section. Quantification of performance was done by comparing the SNR, spectral coverage, and resolution of points generated by both algorithms, where the generalized backprojection results are considered the standard for comparison because of its general usage and robustness. Resolution comparisons between FFBP and BP are performed in the *x*, *y*, and *z* dimensions by taking slices through the image volumes and looking at -6dB down points. Since targets were distributed uniformly in *x*, SNR comparisons were computed by comparing the median vs. peak values for fixed *y* and *z* values containing distributions of point scatterers. Beamforming completion times are also compared, since an improvement in computational efficiency is a key factor in the development of the FFBP algorithm.

Since both methods are time-domain algorithms that rely heavily on multiple interpolation operations, quality comparisons will be made between *cubic spline* and *linear* interpolation methods: the latter method estimates the value at a sampling location by performing a linear fit using only its immediate bounding points, whereas the former performs a piecewise 3rd order

polynomial fit using multiple surrounding data points to try to more accurately estimate the value at the sampling location at the cost of computation speed. Comparisons between spline and linear interpolation will be performed for both beamformers to show how interpolation parameters can affect the point-spread function. Lastly, the spatial spectrum of the beamformed signals will also be compared.

Simulation results: data slices

Figure 148 shows the array and distribution of scatterer points, along with the planes for which volume slices are extracted from the 3D beamforming results. The planes are at x = 5 meters, y = 0 meters, and z = 5.5 meters, intentionally placed to bisect points from the three-dimensional distribution that lie either on a plane or along a line.

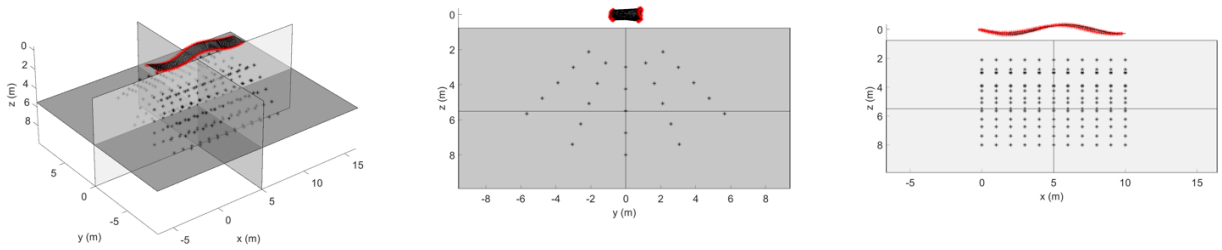

Figure 148: Visualization of the simulated system track and point scatterers, along with the slicing planes sampled from the volume that will be used to compare the performance of standard back-projection with 3D FFBP.

Figure 149 shows a perspective-MIP rendered through the data-cube generated by the FFBP algorithm, indicating that the points beamformed in a well-behaved manner and according to the distribution that can be expected given the parameters in Table 2.

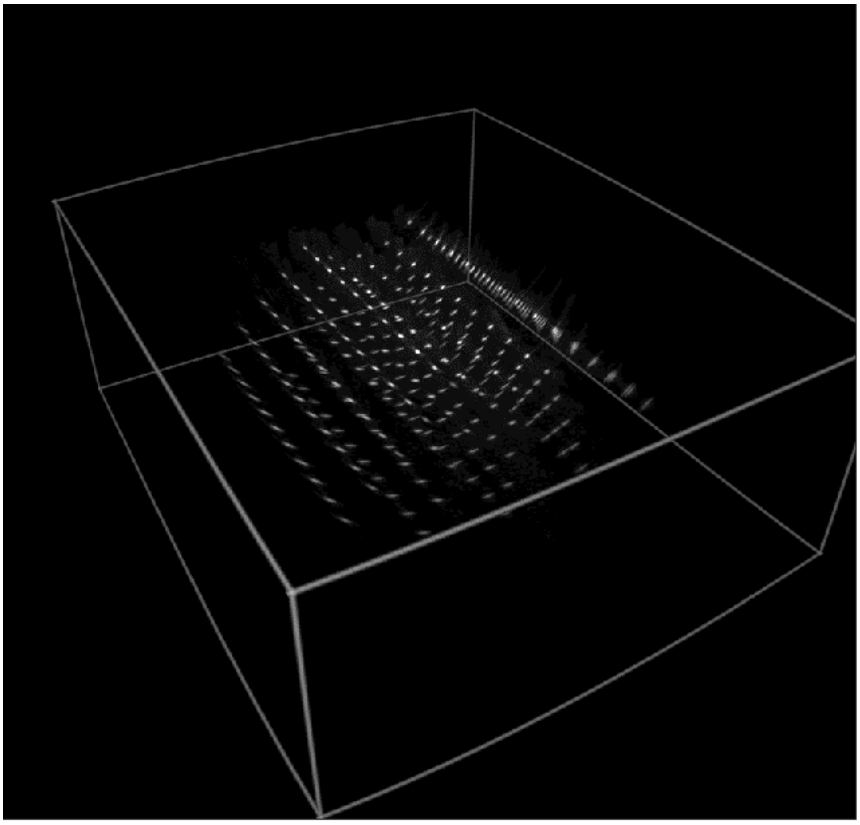

Figure 149: A perspective-MIP rendering of the 3D image generated by applying the 3D FFBP algorithm to the data generated using BOSSim and the simulation parameters listed in Table 1. The points are well behaved with low side-lobes and the distribution follow the expected radial distribution around the sensor trajectory. The amplitude is plotted in a linear scale normalized by the value of the brightest point.

Figure 150 shows the previously specified volume slices extracted from the volumes generated by the FFBP and BP algorithms. In both cases, the *cubic-spline* function was used as the interpolant, and the magnitude information was normalized by the brightest point. The images were displayed using a log-scale color axis with a 50dB total dynamic range – purposefully selected to make side-lobes and interpolation artifacts visible.

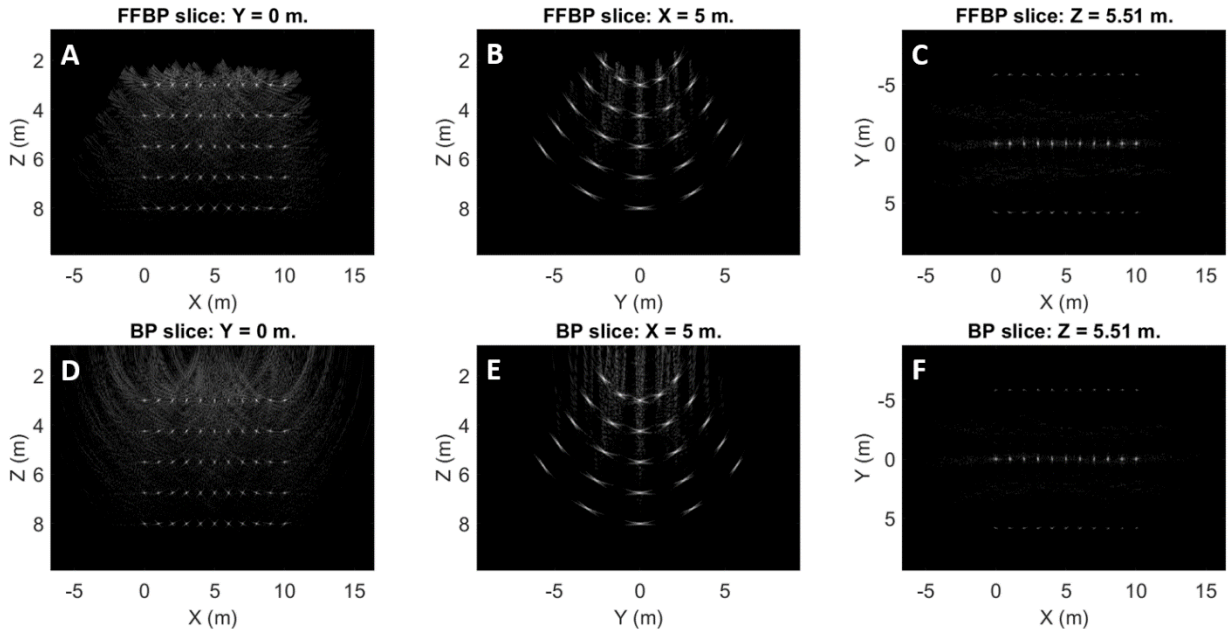

Figure 150: slices through the image cube using FFBP (top row) and back-projection (bottom row), imaged using a log-scale color axis and 50dB dynamic range to accentuate interpolation artifacts. Color-scale is 50dB relative to the brightest point in the image.

Qualitatively, both algorithms produce well behaved point spread functions, with interpolation artifacts at levels ~40 dB below the brightest scatterers. The major source of these artifacts are not the beamforming algorithms themselves but the slightly under-sampled nature of the array in the along-track dimension due to the alternating transmitter configuration used for the simulation. There are no significant deviations between beamforming algorithms other than that the interpolation artifacts have a more limited range in the FFBP algorithm due to the imposed integration angle limitations in the along-track dimension (-60 to 60 degrees). From slices (B) and (E) it can be seen that these interpolation artifacts extend in the along-track dimension only, rather than the across-track dimension.

Simulation results: SNR, BP vs FFBP

SNR comparisons will be made as follows: the median of the log-scale values will be calculated across the sample span to estimate the noise floor in dB. The median is selected because its value is robust to outliers and is only trivially effected by the presence of the points. The signal value will be found by averaging the value of all peaks within 10 dB of the brightest peak. This range was selected empirically to correctly capture the peaks in the data.

Figure 151 shows a line sampled at $y = 5$, $z = 5.51$, $x = -1$ to 10.5 meters through the FFBP data. This is equivalent to sampling the row of points at $z = 5.51$ in Fig. 150.A over the specified range. The noise level is plotted as a bold solid line and the signal level is plotted as a dashed line.

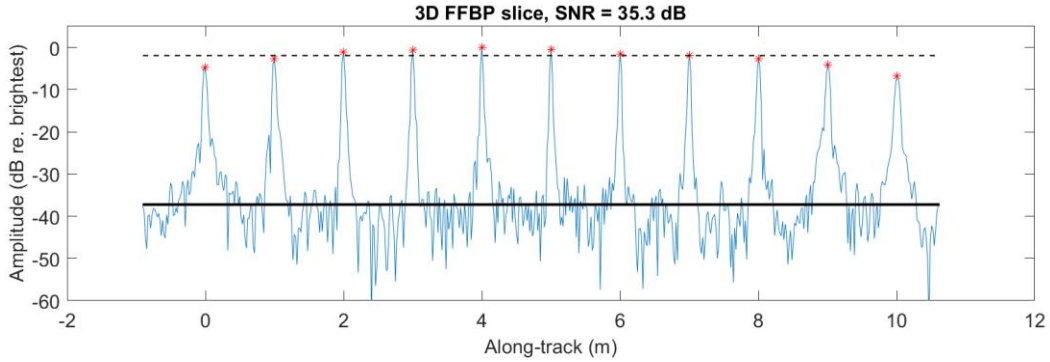

Figure 151: A data slice through the 3D FFBP image plotted in Fig. 150.A at the line *z* = 5.51 meters. The peaks are highlighted with a star, the median signal level is shown with a dashed line and the noise level is shown with a bold line. The signal-to-noise level is measured to be 35.3 dB.

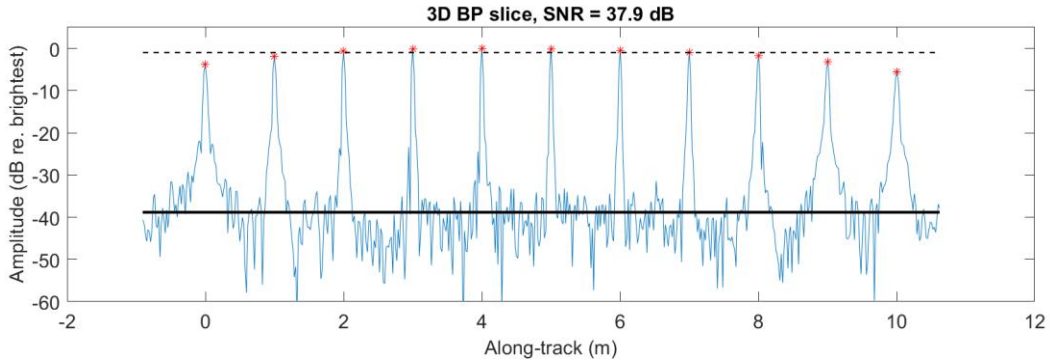

Figure 152: A data slice through the 3D back-projection image plotted in Fig. 150.D at the line *z* = 5.51 meters. The peaks are highlighted with a star, the median signal level is shown with a dashed line and the noise level is shown with a bold line. The signal-to-noise level is measured to be 37.9 dB, about 2.5 dB greater than for the 3D FFBP case.

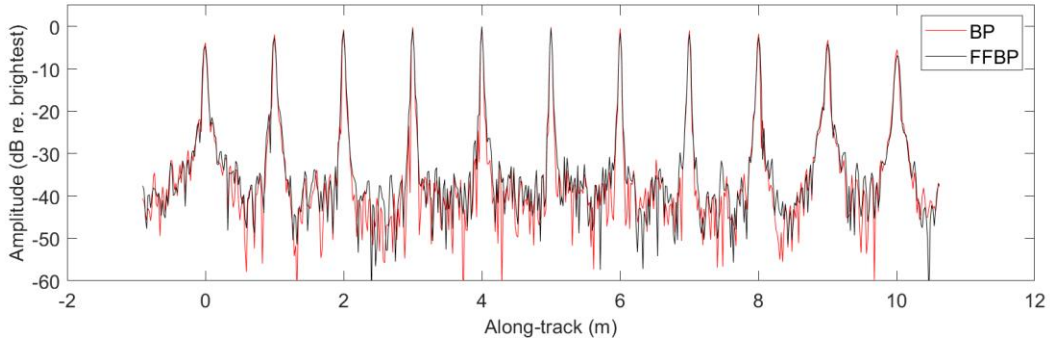

Figure 153: BP (red) and FFBP (black) profiles plotted on top of each other, showing how similar the general structure and levels of the noise and point spread functions are.

This process was repeated for all five rows of points visible in the volume slice at y = 5. The SNR for the BP and FFBP algorithms are plotted vs. each other in Fig. 154.

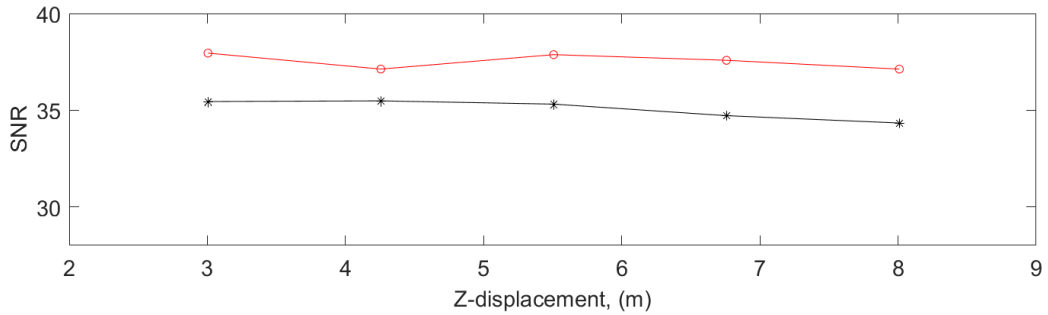

Figure 154: SNR estimates for back-projection (red, circles) vs. FFBP (black, stars). Vertical axis limited to exaggerate the differences. BP tends to have ~2 to 3dB higher SNR relative to interpolation-artifact noise floor.

Across all ranges the FFBP algorithm produces acceptably similar results to BP, varying only by a few dB at maximum. In both cases the SNR is high enough that beamforming artifacts will not be the limiting factor in target analysis for real data: the signal-to-clutter ratio is anticipated generally to be much lower than the ~35 dB that the FFBP algorithm maintains across ranges.

Simulation results: Resolution

The data are nearly critically sampled in the *x* and *z* dimensions, and to compare resolution values sinc-interpolation is used to up-sample the data and determine the -6dB width of the point-scatterers in SAS processing (*x*) dimension. As with SNR estimation, peak-detection with a thresholded margin relative to the maximum value is used to find the location of points. The widths of the peaks when -6dB down relative to the peaks are averaged for all points at the specified range and plotted for both BP and FFBP in Fig. 155.

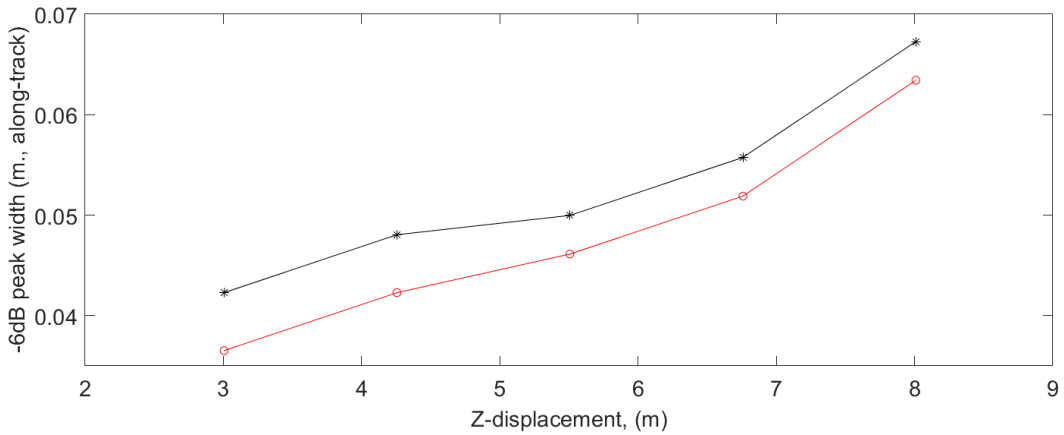

Figure 155: Peak widths vs. range, back-projection (red, circles) vs. FFBP (black, stars). The decreasing resolution is due to the fact that the resolution is aperture limited not beam-width limited in the simulation.

The resolutions of both back-projection and FFBP degrade as a function of range, which is not typically expected for SAS systems, however in this case the wide angles of integration mean that most of the targets in the scene do not have full apertures contributing to their point-spread functions. This problem is exacerbated with range, resulting in the resolution reduction. This

isn't an indication of a malfunctioning algorithm, rather it is expected given the length of the aperture in the simulation and the width of the integration angles.

The resolutions from the different methods are different, however by comparing the relative resolutions it can be seen that the FFBP algorithm has at worst only about 15% lower resolution. The relative resolutions vs. range are plotted in Fig. 156:

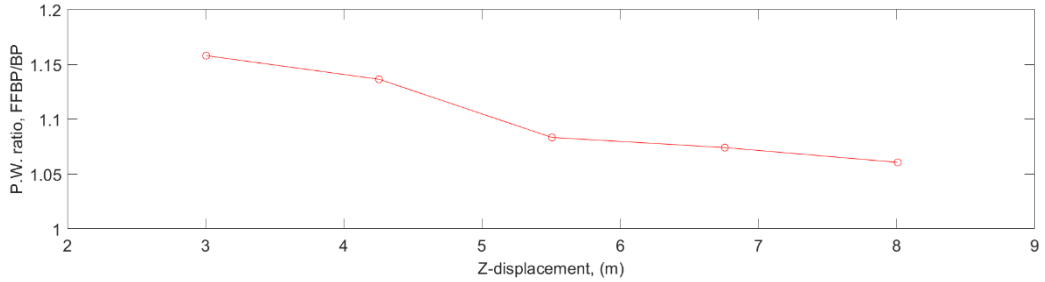

Figure 156: Peak width ratios, FFBP / BP, vs. range. FFBP resolution is, at worst, only about 15% lower than back-projection. Qualitatively, this difference is nearly impossible to detect in images.

From Fig. 156 it can be seen that the resolution of the FFBP algorithm appears to converge towards BP. There is a simple explanation for this: the back-projection algorithm was defined with no angular integration limits, and thus the wavenumber coverage (and resolution) was much higher at close range. At longer ranges, the point-spread-functions generated by both algorithms become more limited by the length of aperture, not a pre-defined integration angle, and thus the two algorithms produce more similar results. The integration angle difference can be seen more clearly in the spatial spectra shown in the following section.

Simulation results: spectral aliasing

Fig. 157. shows the spatial spectra for the FFBP volume slices shown in Fig. 150.A and 150.D. Overlaid on the spectrum are the boundaries imposed by the along-track integration-angle limitations specified in the FFBP beamforming algorithm (+/-45 degrees along-track) and the signal bandwidth (5 – 35 kHz, cosine-tapered).

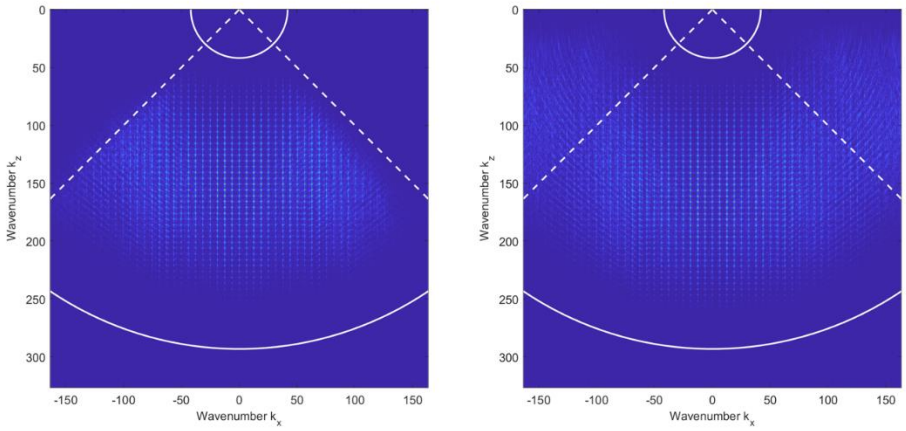

Figure 157: Spatial spectra of 150.A (FFBP, left) and 150.D (BP, right). The spatial spectrum of the image beamformed by FFBP stays within the bounds specified by the integration limits (denoted by the white

dashed lines), and the signal bandwidth (solid white curves). In contrast, the BP routine, which had no integration angle limits imposed beyond that of the aperture itself, shows energy distributed beyond the integration angle limits.

Visible in Fig. 157, the energy of the spatial spectrum stays within the boundaries imposed by the FFBP algorithm's integration limits (as it should), but the BP algorithm, which had no integration angle limitations, shows significant energy beyond +/- 45 degrees, some of which will inevitably be aliased given the along-track sampling rate of the system.

Also note that the signal power does not approach the boundaries defined by the signal bandwidth. This is due to the use of a Tukey parameter of 1 (see simulation parameters, Table 2), which indicates the use of a fully cosine tapered signal.

Simulation results: interpolation comparisons

Figures 158 and 159 show resolution and SNR comparisons between 3D FFBP algorithm implementations using spline vs. linear interpolation as the interpolant.

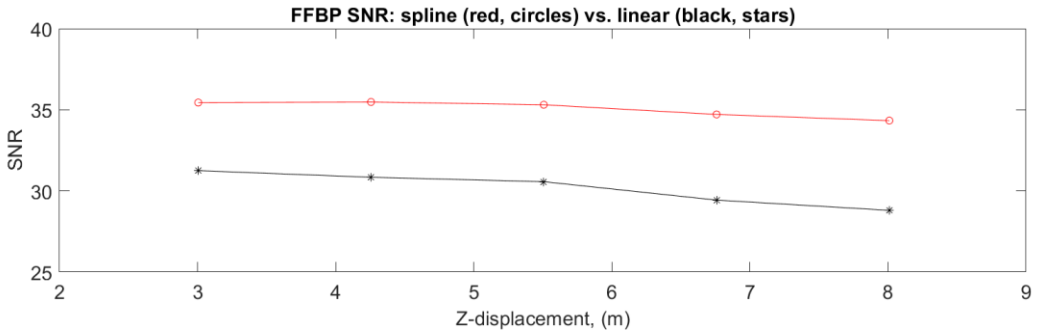

Figure 158: SNR comparisons between FFBP implementations using spline (red, circles) interpolation and linear (black, stars) interpolation. The interpolation can make a 4 to 5 dB difference in SNR.

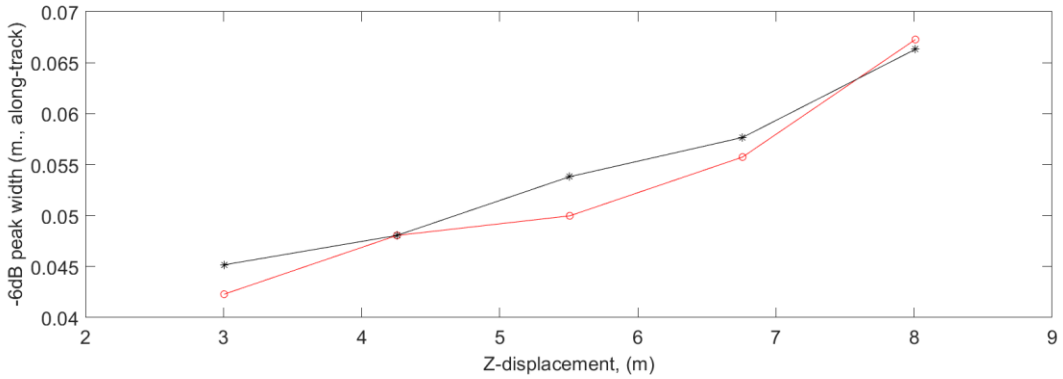

Figure 159: Peak width comparison in the along-track dimension for FFBP implementations using spline (red, circles) and linear (black, stars) interpolation. Unlike with SNR, there doesn't appear to be a huge difference in the resolution results for the different interpolation methods.

It can be seen that interpolation method has a stronger effect on the signal-to-noise ratio than the resolution. The fact that linear interpolation primarily effects SNR, which in practical applications is not limited by the beamformer but by other factors, implies that linear interpolation is an acceptable alternative to spline and it is much faster and easier to implement.

The GPU accelerated beamformers implemented in the MuST BOSS toolbox all use linear interpolation.

Simulation results: computational efficiency

Table 3 shows a comparison of computational times for FFBP and time-delay-and-sum. Due to the excessively long times predicted for BP (e.g. > 10 days), computation times were estimated by computing single slices through the volume and multiplying by the total number of slices required to populate the volume. The point-spread-function analyses performed in the previous sections also exploited single-slice beamforming to allow beamformer comparison. Each configuration was executed five times and the values of the mean and standard-deviation are shown in the table.

	Completion time (s)	Standard dev. (s)
Back projection (linear)	859,476	1856
Back projection (spline)	924,690	14362
FFBP (linear)	344	26
FFBP (spline)	892	10

Table 3: Back-projection and FFBP completion time tables for different interpolants. Each operation was performed six times and the means and standard deviations are shown in the table.

The back-projection tests were completed on a Windows 10 laptop with 64 Gigabytes of memory and 6 Intel i7 CPU cores running at 2.2 GHz. No additional hardware acceleration beyond Matlab's native multi-core acceleration for FFT and interpolation operations was used. All computations were performed in 64-bit precision on complex datasets (i.e. 16 bytes per point). The memory usage of the computer was monitored and it was determined that the memory ceiling was never approached during either of these beamforming operations, and the bottlenecks from a computational standpoint for the standard back-projection routine were the interpolation and phase multiplication stages. From Table 3 it can be seen that FFBP is multiple orders of magnitude faster than standard back projection. Though the spline version of FFBP is slower than the linear version it is still almost 1000 times faster than back-projection.

Note what is being claimed in Table 3 in regards to the performance relationship: it is *not* being claimed that a fast implementation of back-projection couldn't be developed on a computer if professional programming practice is applied (i.e. *c*- or *Fortran* programming applied to the direct data stream), or other acceleration hardware is utilized (e.g. GPU's). Rather it is being claimed that FFBP produces nearly the same results but with significantly fewer computations. This is particularly convenient for Matlab implementations because it makes beamforming and data exploration vastly more feasible using this widely used data analysis tool.

Simulation results: summary of FFBP performance

Application of the 3D FFBP algorithm described in the *Materials and Methods: Beamformer Development* section of this document to BOSSim data and comparison with standard back-projection indicate that the algorithms produce nearly identical results but the computational

efficiency gains are enormous. For this reason 3D FFBP (in various implementations) was used exclusively for the production of the 3D BOSS beamforming results in this final report.

III. Field test results: ONR AUV-mounted BOSS

Overview

In January 2019, field data collected from a prototype BOSS system mounted on an AUV was given to APL-UW for analysis. The data from this BOSS system provided a valuable tool to test the 3D FFBP algorithm on data that had the noise, calibration issues, and trajectory perturbations that accompany real-world systems. The sediment in this region was sand, and the object of interest was a 3 meter buried pipe. At either end of the pipe were placed, on the sediment surface, spherical markers indicating the location of the pipe.

3D FFBP beamforming results

Figure 160 shows the results of the real-aperture body coordinate frame beamforming operation. The block-size for the algorithm was set to 128 pings, thus there are 128 ping real-aperture frames:

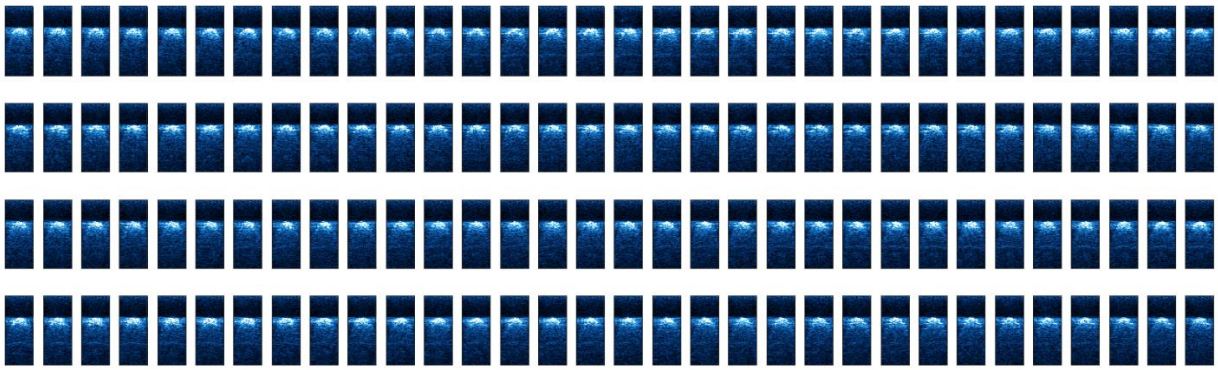

Figure 160: Frames generated from one block of pings. These frames will be used to make the first stage spherical coordinate sub-apertures in the 3D FFBP algorithm.

In this 3D FFBP implementation, $K = 16$ and thus the frames in Fig. 160 combined to generate eight low resolution three dimensional spherical coordinate sub-apertures. These sub-apertures are shown in Fig. 161:

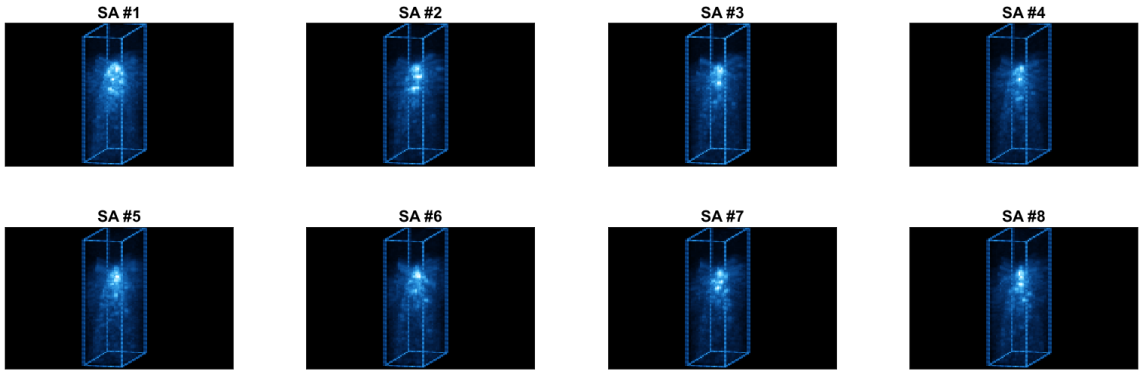

Figure 161: Stage-1 low resolution spherical coordinate sub-apertures generated from the frames visualized in Fig. 160.

Following the construction of the first stage of low-resolution spherical coordinate sub-apertures, aperture pairs are successively combined to generate half the number of 3D images, each having twice the resolution. The process of combination is shown in Fig.'s 162 – 164:

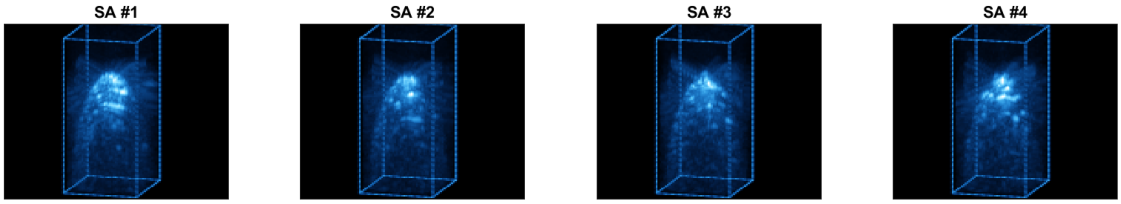

Figure 162: Stage-2 spherical coordinate sub-aperture images. The pair-wise coherent combination of the images in Fig. 32 has resulted in half the number of apertures having twice the resolution.

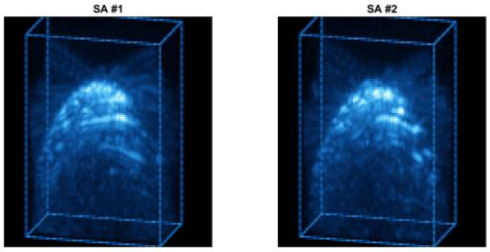

Figure 163: Stage-3 spherical coordinate sub-aperture images.

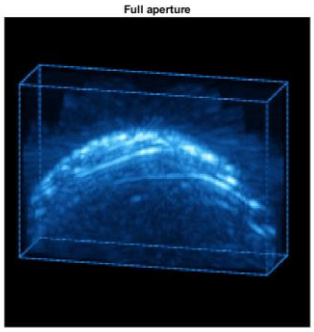

Figure 164: Final stage (stage-4) spherical coordinate sub-aperture. Depending on the block-size and value for *K* the 3D FFBP algorithm may require fewer, or more, stages to generate the final sub-aperture corresponding to a single ping block.

Note that the sediment interfaces visualized in all of the sub-apertures and sub-aperture stages appear to lie on a curved manifold. This is a natural result of imaging the data in spherical coordinates. When the data is finally projected to the Cartesian image domain the sediment no longer has this curved appearance. This is shown in Fig. 165, in which the final Spherical-to-Cartesian transformation has been applied:

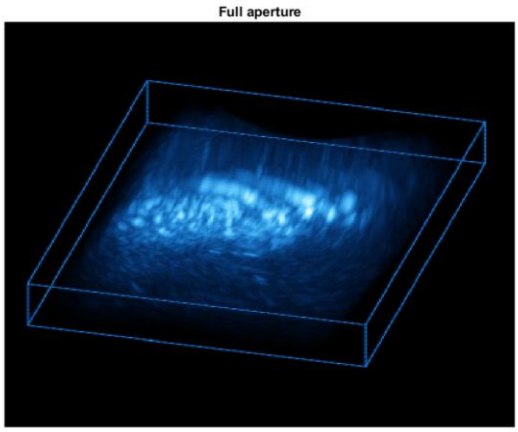

Figure 165: Cartesian data projection. Note that the sediment no longer appears to lie on a curved manifold.

The final data-product that can be used for target detection and mapping is a combination of these final stage Cartesian coordinate images.

Mosaic results

As discussed previously in the *Methods and Materials* section, three dimensional images can quickly drain memory resources as their size increases. A modestly sized, uncompressed 1000 x 1000 pixel gray scale image would require 1 Megabyte (MB) of memory if each pixel represented an 8-bit grayscale level. In contrast, a 1000 x 1000 x 1000 voxel image cube would require a full Gigabyte (GB) of memory at the same bit-depth. If each pixel represents a single-precision complex number (i.e. 8-bytes per pixel), then that 1000 voxel-on-edge cube would require 8 GB's of memory – a significant chunk of memory on many modern systems. If voxel sizes are small (e.g. 0.02 cm on edge) then the image, only 20 meters on edge, may represent just a small portion of the scan.

When this AUV data was being analyzed the MuST Toolbox software had not yet been developed, and methods for mosaicking volume data and creating multi-ping maps were still being explored. In these initial tests, mosaics were generated from Z-MIP's through Cartesian-coordinate sub-aperture image cubes. For each sub-aperture image cube generated in Cartesian coordinates, a geo-rectified Z-MIP was generated. The stack of images from all sub-aperture images were combined to form a Z-MIP mosaic representing the scanned region. In this manner, the salient targets in the sediment volume could be geolocated efficiently, and the approach formed a template for the layer combination method implemented later in the MuST Toolbox mosaicking tool. An example of a Z-MIP mosaic for two scans (one of which incorporates the data from the previous figures) is shown in Fig. 166. Note the highly non-linear trajectory which causes strong oscillations of the scan in the across-track dimension. Despite this trajectory nonlinearity, the images still show good focus.

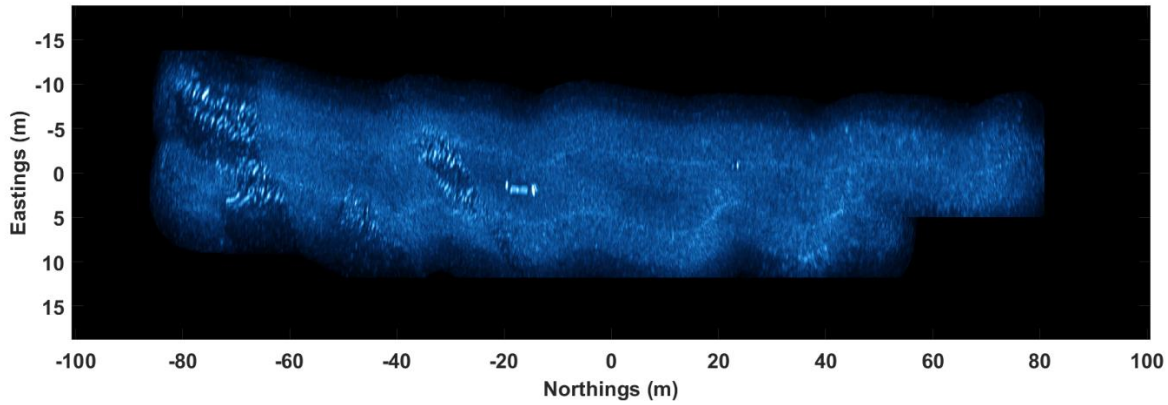

Figure 166: Z-MIP mosaics of two scans from the AUV-mounted BOSS system. The target, a buried pipe, is distinctly visible, along with the end-markers, at ~2 meters east, -18 meters north.

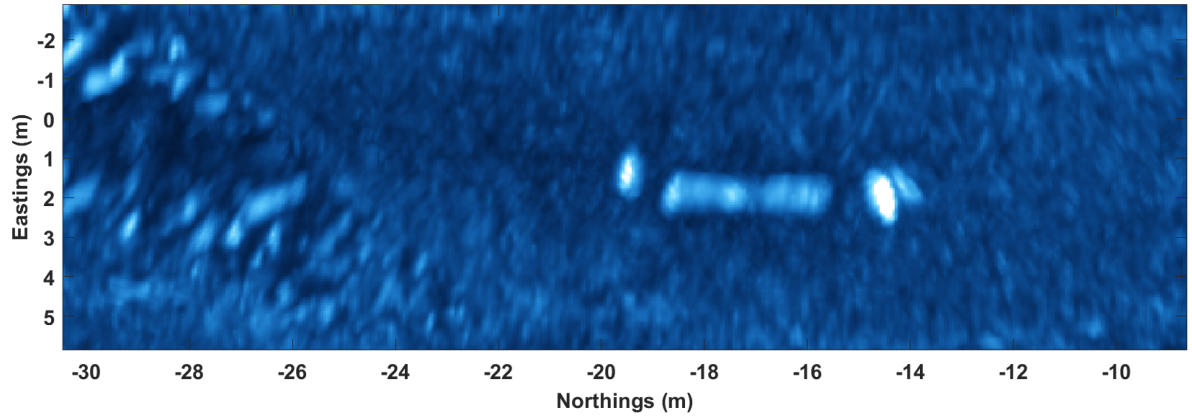

Figure 167: Zoom-in of the buried pipe target and the end-markers. A ripple field is visible on the left side of the image.

Acoustic camera imagery (e.g. Fig.'s 42 – 44) was also generated from the data. The following figures are examples from a movie showing the construction of a mosaic and the acoustic camera view corresponding to the spherical coordinate sub-apertures being used to generate the Z-MIP's contributing to the mosaic:

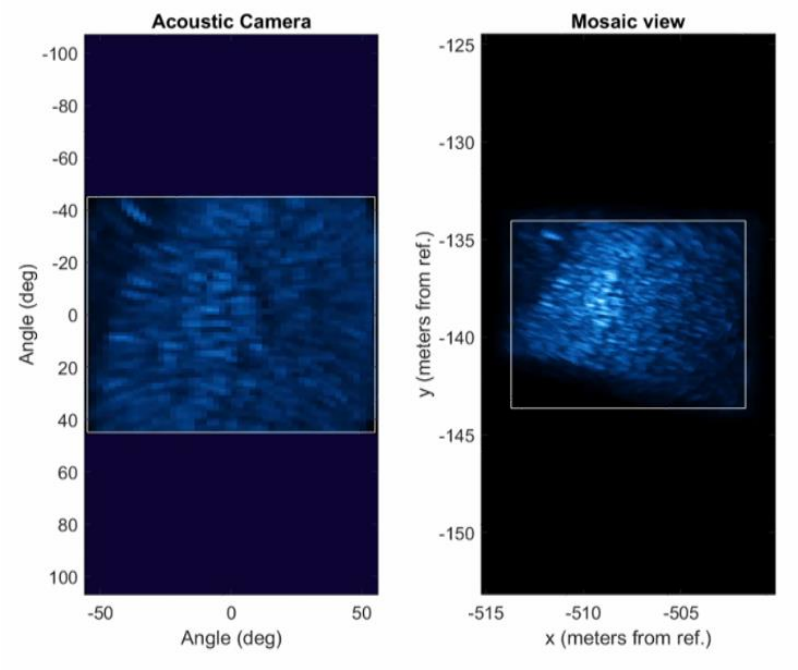

Figure 168: Frame 1 of an acoustic camera / Z-MIP generation demo. The Acoustic camera view is on the left, with angle in the synthetic aperture dimension being in the vertical axis, and angle in the real-aperture dimension being the horizontal axis. The Z-MIP mosaic imagery (as it is generated) is shown on the right with a white box surrounding the most recent addition to the mosaic.

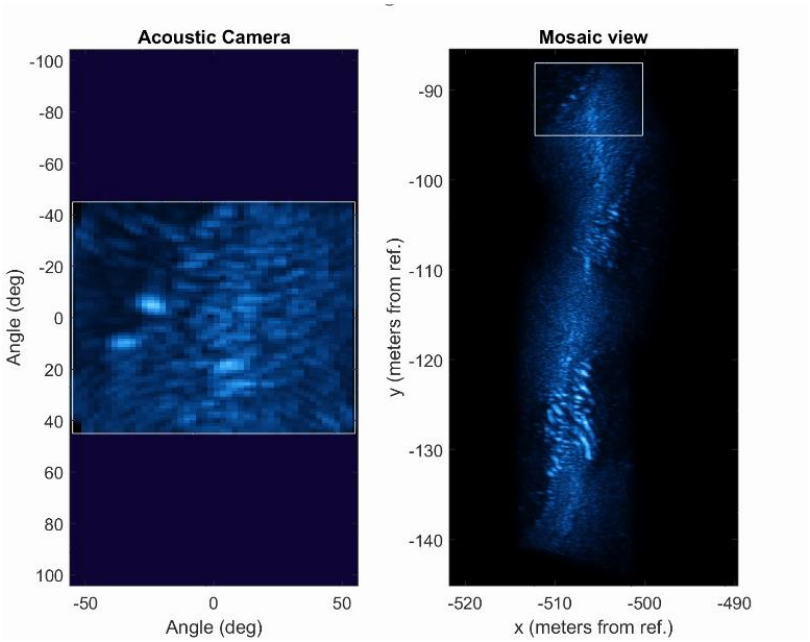

Figure 169: A subsequent frame of the mosaic, showing the progression of the mosaic.

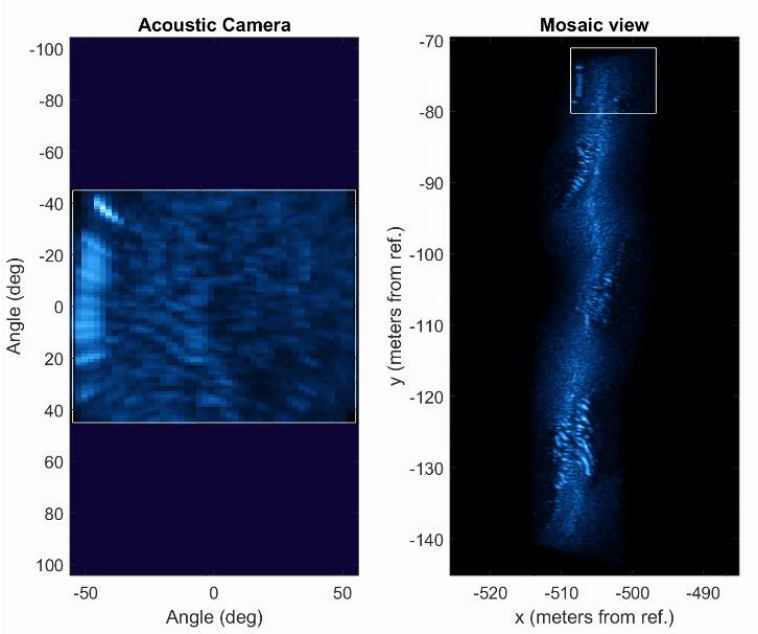

Figure 170: A frame showing the pipe target in the acoustic camera image (left) and its contribution to the mosaic, on the right.

Note that both modes of visualization were later abandoned in MuSTLive in favor of the scrolling *Z*- and across-track MIPs shown in Fig.'s 50 and 51. Reasons for this included the geometrical consistency of the projections (i.e. shapes didn't bulge or bend due to being viewed in spherical coordinates), and the resolution of map scaling issues (i.e. the continually increasing size of the map in the approach above caused many practical display problems).

Autofocusing

An extremely significant result that came from analysis of AUV BOSS data was the startling discovery that the inertial-navigation data packaged with the AUV SAS was entirely sufficient for achieving well-focused SAS imagery. This was found to be the case despite the fact that the AUV platform was relatively unsteady and experienced non-trivial levels of heave and sway due to wave action. The original technical approach included development efforts for autofocus algorithms that would use detected sediment features and backscattering statistics to estimate navigation corrections that would enhance image focus to make up for the perceived insufficiency of the navigation aids that normally accompany SAS systems. In fact, data-driven "micro-navigation" techniques for refining the navigation estimate produced by SAS navigation aids, (e.g. DPCA processing [38]) are the norm for virtually all SAS applications. In the present case these do not appear to be necessary. An explanation for this can be found in terms of the aperture length involved with processing BOSS data. Most SAS systems are meant to operate with maximum ranges of several hundred meters, and have wavelengths on the order of several centimeters or less, consider e.g. the HiSAS 1030, which has a nominal center frequency of 100 kHz and 30 degree total beamwidth. At 200 meters, fully coherent SAS processing with the HiSAS would require a 107 meter long synthetic aperture with navigation accuracy on the order of $\lambda/8$, or ~2mm accuracy. In contrast, for an AUV flying at 3 meters and imaging targets buried less than a meter, the synthetic aperture associated with an integration angle of +/- 35 degrees

would be only ~5.6 meters long, and the accuracy requirements reduced by nearly an order of magnitude due to the long wavelengths used in BOSS sonars. The end result is a *significantly* relaxed navigation constraint. For field tests employing the MuST towbody, the very high stability of the platform resulted in highly reliable SAS imaging even using the MuST INS system which contained only MEMs gyroscopes. (In contrast, most SAS systems exploit optical gyroscopes). The MEMs INS was eventually abandoned in favor of a higher precision INS but not due to poor imagery, rather, due to long term navigation accuracy degrading mosaic quality.

The lack of necessity for autofocusing meant that the research efforts initially planned for autofocusing could be redirected toward data-driven navigation alignment, and the developed algorithms for mosaicking are described in the *Feature-based renavigation* portion of the *Materials and Methods* section.

AUV BOSS summary

The AUV dataset provided by the Navy was an excellent resource for evaluating an implementation of the 3D FFBP algorithm with real data and trajectories subject to strong perturbations. As can be seen from the previously shown examples, the algorithm worked extremely well. Targets were well focused, with the limiting factor clearly being the resolution in the real-aperture (across-track) dimension. The synthetic aperture dimension, by virtue of its much longer aperture, exhibits much higher resolution.

IV. Field test results: MuST First Deployment in Lake Washington

Overview

The first official test of the "integrated" MuST system, combining the FOCUS 3 towbody with the EdgeTech 2205 and Buried Object Scanning Sonar systems, was performed from 2nd – 8th Aug. 2019 in the north end of Lake Washington. There were several goals of this experiment, the primary being to perform a basic systems test determining array and transmitter performance and optimal operational parameters such as sample rate and scan altitude. Additionally, the performance differences between transmitter configurations (VLA vs. alternating Neptunes) needed to be determined, and the sufficiency of a MEMS INS for navigation needed to be assessed. Another goal was to test some live-streaming software, (a very early pre-cursor to MuSTLive), that was developed to monitor the BOSS data stream and perform real-aperture beamforming.

Two targets were placed prior to the deployment of the MuST sensor: a 30 by 60 cm. solid aluminum cylinder (the "slug") and a hollow 30 by 60 cm. aluminum shell (the "cylinder"). These were positioned in a muddy region of the sediment in the northern part of Lake Washington. The slug sank into the mud however the shell stayed at or near the surface. Both targets had an attached cable allowing them to be retrieved following the conclusion of the sensor test. Determining whether or not these targets were visible to the eBOSS sensor was another important goal of this field test.

The ship from which the target was deployed was the APL-UW owned RV Robertson.

Data Collection, 8/2/2012

On the first day of data-collection the goal was to verify that the sonars were operational, and to perform a safety scan of the survey area to determine if there were any hazards in the watercolumn that needed to be avoided. The BOSS sensor was activated, however the sediment was never directly visible to the sensor because the pulse-repetition was maintained at 80 Hz (max range = (c/2)/80 ~= 9.5 meters) but the sensor altitude was never below 10 meters. It was found, however, that echoes from previous pings were strong enough to be visible in the data, and backscattering from the sediment was still visible.

Live data acquisition and hydrophone health diagnostics performed via Python scripts which connected to the BOSS computer via TCP/IP was tested, however it was found that BOSScontroller.exe was not transmitting a sensor data package containing the critical information to perform data analysis due to a programming bug in BOSScontroller. This was verified by EdgeTech, which made a patch and pushed the corrected BOSScontroller.exe program to APL-UW prior to the next day of experiments. This responsiveness was characteristic of APL-UW's interaction with EdgeTech throughout all of the field experiments.

Because live data streaming proved ineffective this first day, sensor health and hydrophone characterization was performed by making a remote connection to the BOSS computer's shared hard drive and pulling the BOSS datafiles ("*.jsf*" files). Loading several files into Matlab for analysis revealed an important error that would plague the system for the rest of the experiment, but would be resolved by EdgeTech prior to the Sequim bay experiments: all four BOSS hydrophone array panels were swamped by uniform noise across the array. The noise level was extremely high, making even the specular glint from the sediment barely visible. The figures below are data snips over approximately 2000 pings, captured from array channels 1 and 25, which represented the SNR span for the best and worst functional sensors on the array:

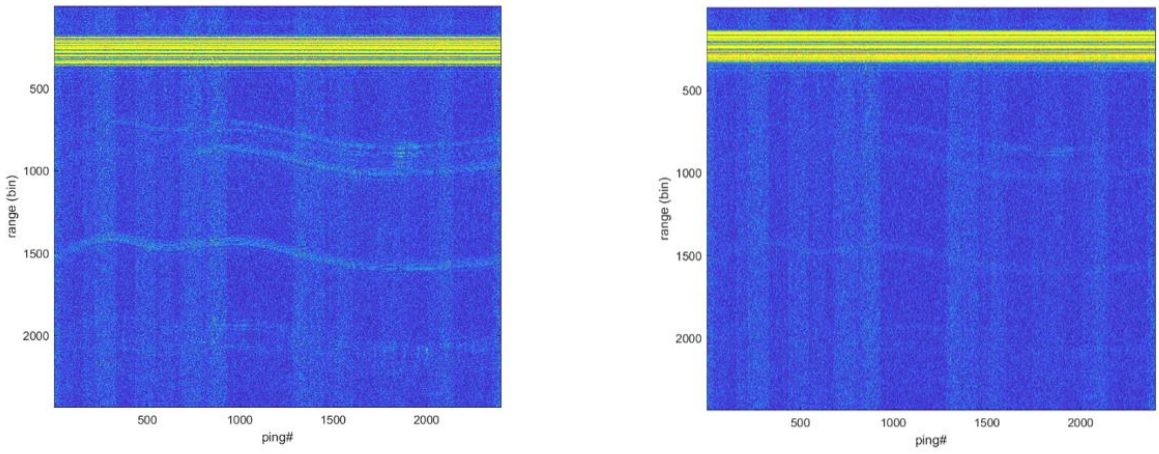

Figure 171: Data from channel 1 (left) and channel 25 (right) showing signal degradation due to panel noise.

On closer inspection it was found that the noise was nearly identical across all panels, with slight fluctuations in level. Figure 172 shows an example of the envelope of the raw data of a ping vs. time and channel number.

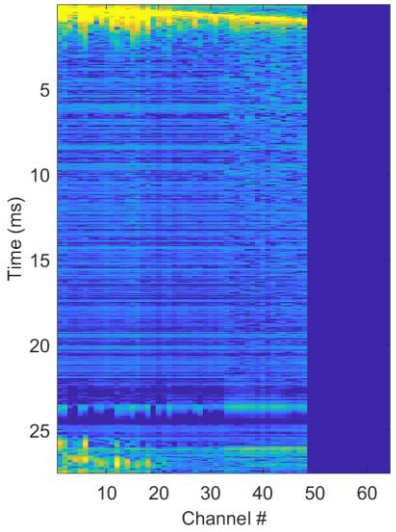

Figure 172: An example of the raw data captured from the BOSS system showing sensor noise (horizontal striations). The starboard-side array was malfunctioning and thus shows no signal (channels 49-64). The bright signal dominating the first couple milliseconds is the direct path signal from the source to the array. Pulse compression artifacts have wrapped around to the time series beyond 25 milliseconds. The specular glint signal from the sediment arrives at about 24 milliseconds.

In Fig. 172 the array noise is visible as the series of horizontal striations that dominate the signal throughout most of the time-series. In Fig. 172 the data has been matched-filtered, improving the SNR vs. the level visible in Fig. 172, but the level is still quite high, especially in comparison to the specular glint from the sediment.

To de-noise the data (at the cost of some signal) the noise level was estimated by averaging the signal from the watercolumn (where panel noise dominated) over several thousand pings to get relative channel noise sensitivity estimates. For all subsequent pings, the panels were first corrected for noise sensitivity, and the average across channels was subtracted from the array. Channel data before and after this de-noising process is shown in Fig. 173, followed by a profile estimated with and without de-noising, in Fig. 174.

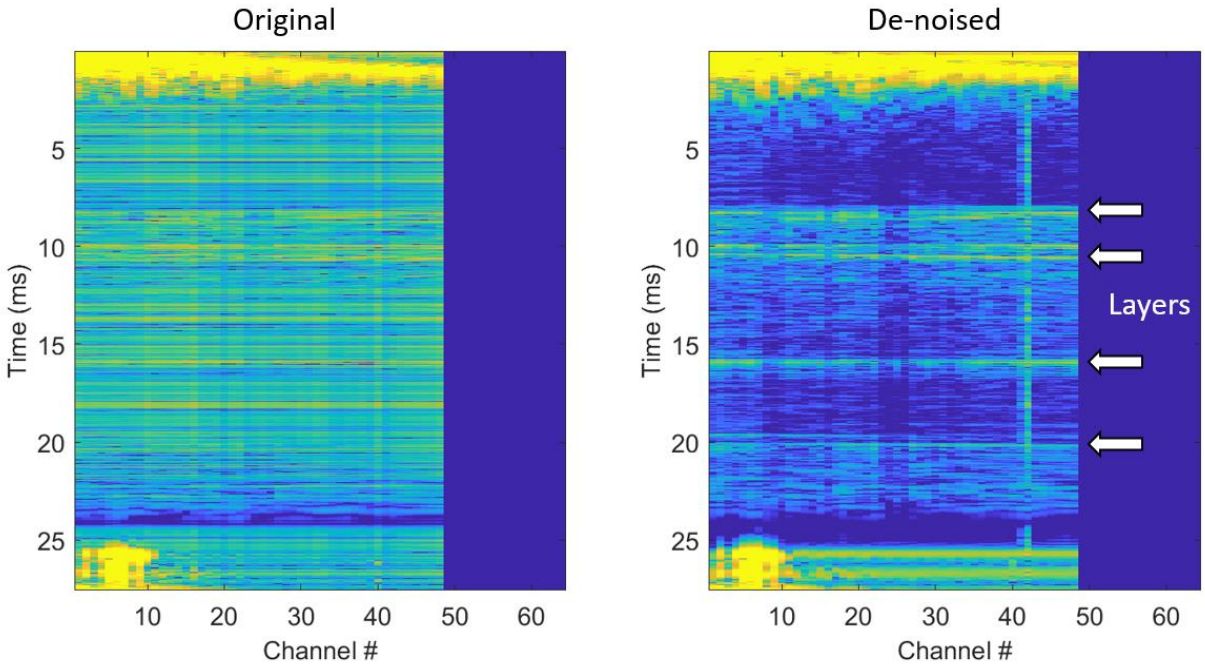

Figure 173: Noisy sonar data before (left) and after (right) de-noising. This ping was acquired on 8/8/19 (day 4) but is used here because the sensor is at low elevation and the masking of the sediment layers is easy to see.

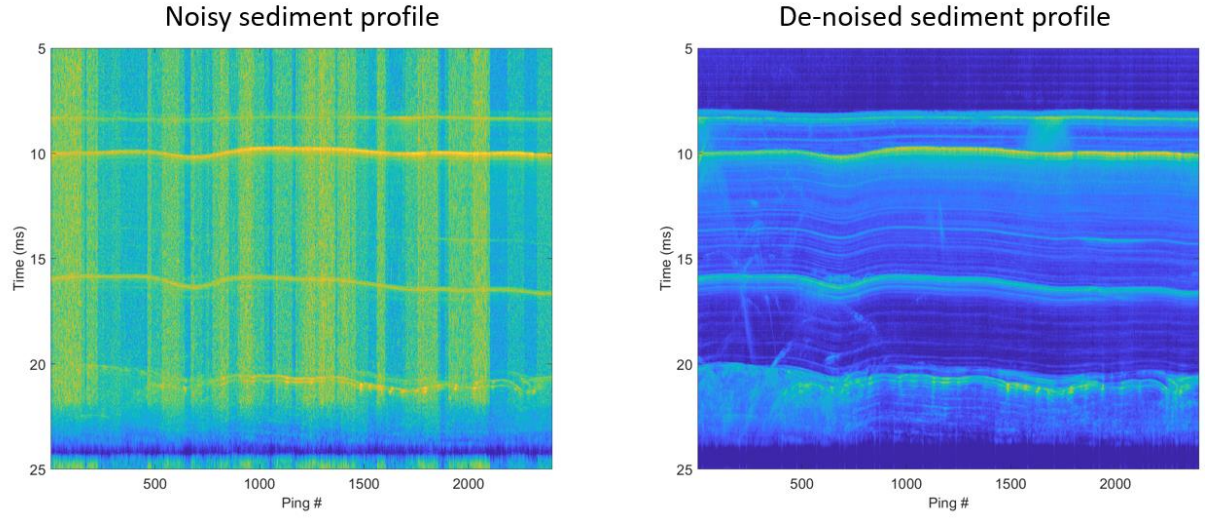

Figure 174: Noisy sediment profile (left) and de-noised profile (right) found by taking an average of the mean across the channels for the pulse-compressed data. Images are plotted using the same 40dB log scale color axis.

As seen in Fig.'s 173 and 174 the number of features visible in the sediment layers is greatly increased.

In addition to having noise issues in the array, the INS was found to give problematic position and orientation estimates. Figure 175 shows an example of a set of position estimates output from the INS for a particular time period during which the MuST system was deployed. The

black line shows the estimate of vehicle position output by the INS and the red dots are locations at which the MuST is surfaced and the GPS is receiving a signal.

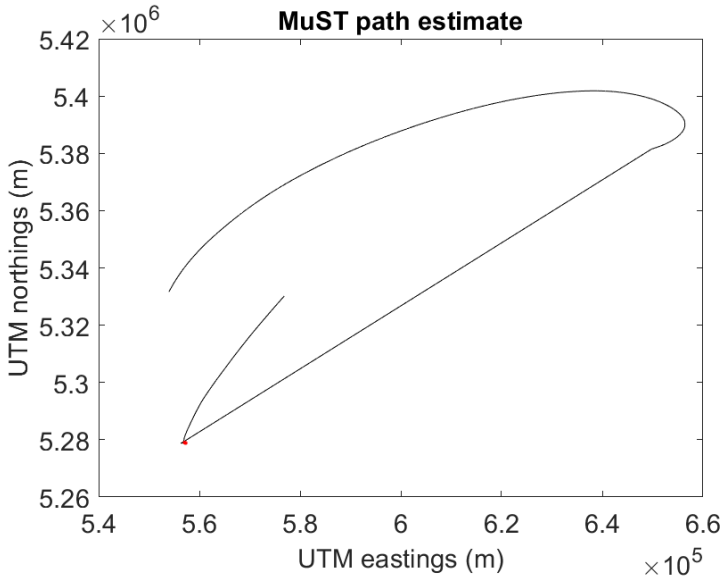

Figure 175: The MuST towfish position in UTM coordinates for a scan as estimated by the INS' internal Kalman filter. The bounding box required to enclose the path estimate is greater than 100 *kilometers* on edge, despite the track in reality only being on the order of a kilometer. The red points near the bottom left show the locations at which a GPS signal is received. The INS estimate is good when the GPS signal is received but rapidly diverges from having physical meaning when the GPS signal is lost.

Additional results seem to indicate instability in the tilt estimates. The following plot compares the INS' estimate for the roll vs. an estimate made by simply estimating the orientation of the gravity vector using the raw linear accelerometer data.

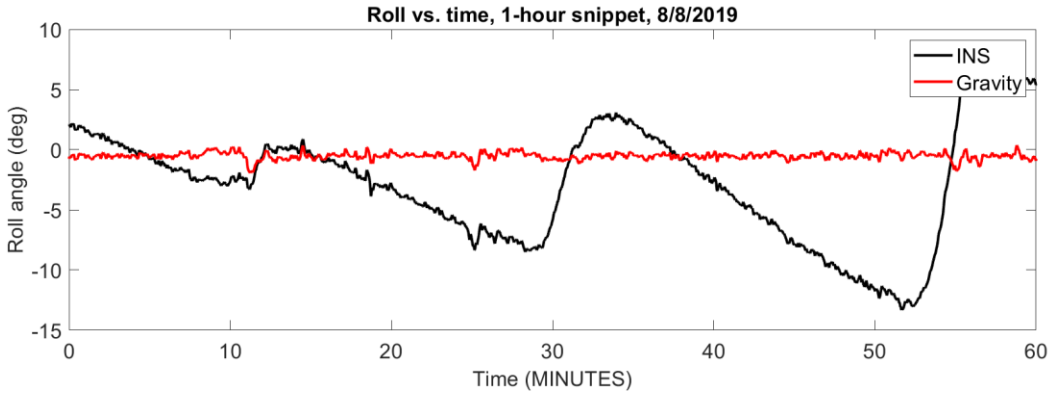

Figure 176: An estimate of roll angle using the effects of gravity on the raw linear-accelerometers in the INS, and the roll estimate actually output by the INS as estimated by its internal Kalman filter. Over the timespan depicted the towfish was not noticeably rolling, and the INS filter is giving clearly erroneous results.

Figure 177 shows the difference of the estimate from gravity and the INS' roll estimate using its internal Kalman filter. The difference is quite large, diverging by over 10 degrees at some points.

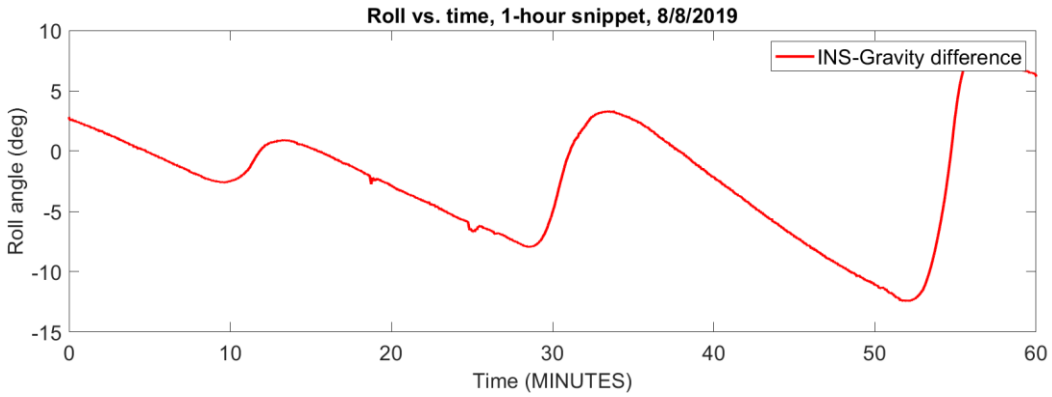

Figure 177: The difference between the INS' estimate for roll and the simple post-processing estimate found using the orientation of the gravity vector measured via the linear accelerometers.

Tilt estimates based on gravity can be erroneous if a vehicle is moving rapidly enough that there are acceleration sources that compete in level with gravity; in the case of the MuST, however, the platform is very slow moving and this special case never occurs. Discussion with the pilot of the towfish indicated that the vehicle was not experiencing significant roll, which further indicated that there was a problem with the internal Kalman filter used by the INS to estimate the vehicle's orientation.

The combination of orientation and translation errors completely prevented the MEMS INS from being useful from a beamforming standpoint for the entire Lake Washington field experiment and that portion of the project was put on hold to focus on testing raw data acquisition capability. Due in part to the issues with the INS but even more so the special beamforming modifications required to deal with the panel failure problems, the data from this set has never been SAS-beamformed.

At the end of the first day of experiments it was determined that there were no significant water column hazards either in the vicinity of the target or in the run-out region where the towfish would be ascending or descending prior to scanning the target.

Data collection notes, Aug. 5^{th} – 8^{th}

The first change made on the next test day was to update the BOSScontroller.exe program using the patch that EdgeTech pushed to APL-UW. Once the towfish was deployed and the BOSS system turned on it was found that the correct sonar data packages were being sent over TCP/IP, and that data could now be live-streamed from the system. Several live-streaming data diagnostics and beamforming tools were tested, however these remained in a fairly un-developed state. The main tools consisted of:

- A Python-based *.jsf* grabber that would assemble a BOSS data file of a predefined size from the live-streamed data, to be analyzed in Matlab.
- A Matlab script that performed real-time channel data visualization for diagnostic purposes (i.e. to see if channels / panels were active, dead, or in the process of dying). This Matlab code communicated with a Python script that was connecting to the BOSS via TCP/IP and acting as a data hub.

Concurrently, Matlab code that used its own internal TCP/IP connection routines was being developed for processing live-streamed BOSS data, however this software was not mature enough to be used until the Sequim Bay tests.

During the second day of tests, the alternating Neptune configuration was used and the sample-rate was reduced to 40 Hz to double the range of the system, reducing reverberation noise and increasing the maximum range of the system to ~18 meters. In later surveys, 40 Hz became the standard "close-range" sample rate used for imaging low-profile wrecks or pipes, though 60 Hz was used for UXO surveys.

Additionally, the towfish made several low passes over the target region. During one scan, however, the live-streaming data diagnostics tool showed the signals from the BOSS blinking and then disappearing, being replaced completely with coherent panel noise. The MuST was brought back aboard the R/V Robertson and it was found that the panels were non-functional. The experiment was aborted early and the vehicle returned to the dock for further electrical diagnostic tests by the engineers. Through the efforts of the engineers, two panels were restored to functionality on the following day (6th, August 2019) and a large amount of data was captured over the target field using the alternating Neptunes configuration. The engineers took the 7th of August to restore another panel, and during this time the live data viewing software was improved. On the fourth and final day of the Lake Washington field tests (8th Aug. 2019), three-panels were operational, and scans were conducted over the target field using both the VLA and the alternating Neptune transmitter configurations.

Data sets

Due to the unreliable output of the INS system the data was not beamformed, however a series of profiles were generated in which scattering from targets is visible. This section plots figures of the profiles that have currently been processed from the Lake Washington data sets. The best quality data was captured on the fourth and final day of the experiments, and provides the bulk of the examples shown below.

Figure 178 shows the signal level vs. channel and frequency, incoherently averaged over several thousand pings. The direct-path signal to the array from the source was gated out, and the averaged signals represent the combined response of the sensor and random scattering from sediment and clutter.

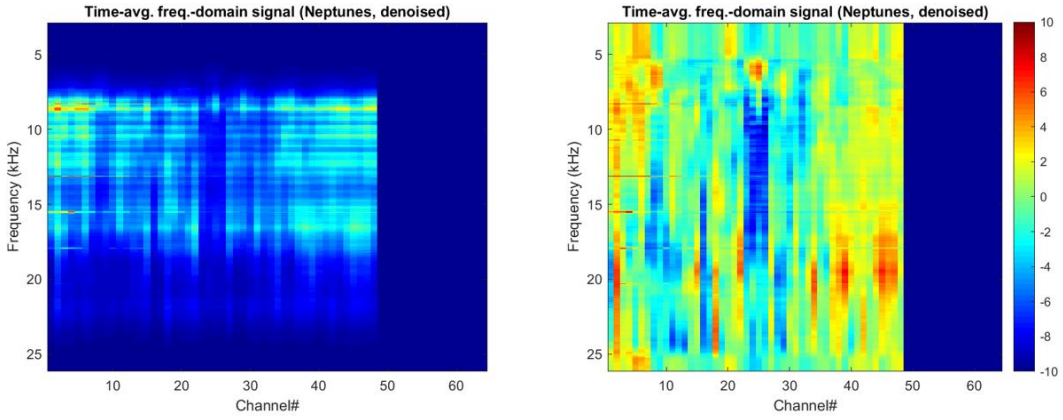

Figure 178: Sensor response averaged over several thousand pings (left) and a relative channel sensitivity calibration in dB relative to the median response across the array.

Note that though Fig. 178 doesn't show a true sensor calibration, which would later be determined via the APL barge calibration facility using upgraded BOSS array panels, (*Materials and Methods: MuST eBOSS Calibration*), the information contained was still useful for generating a filter that regularizes the sensitivity of the array and reduces artifacts when performing real-aperture beamforming. A sensitivity correction filter was generated in this manner and applied routinely to all data captured from the MuST system prior to processing the Lake Washington dataset. Note that this kind of sensitivity correction was estimated and applied *after* removing the coherent panel noise (e.g. Fig. 173, 174). Otherwise the estimates were biased by electronics noise, which didn't have the same channel sensitivity characteristics.

Figure 179 shows a side-by-size comparison of a real-aperture, single-ping image visualized in polar coordinates (left) and Cartesian coordinates (right). Note that flat surfaces such as sediment layers appear as curved surfaces in polar coordinate images, due to the variation in range of points on the surface relative to the origin of the coordinate system.

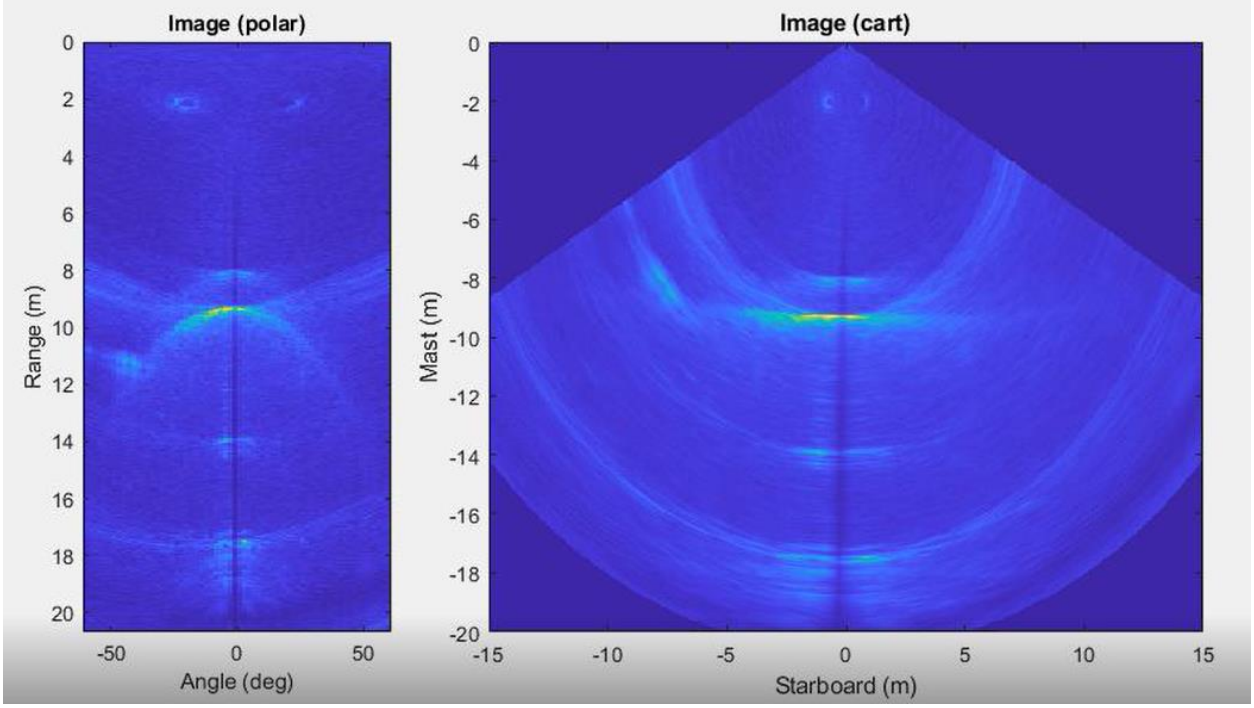

Figure 179: A single ping real aperture image visualized in polar (left) and Cartesian (right) coordinates.

In Fig. 179 a null near 0 degrees / 0 m. starboard is visible in both representations. This null is the result of the de-noising process. If the de-noising process had not been applied then the image quality was almost too poor to be useful. Figure 180 shows an example of a real-aperture image before and after de-noising, showing the central spoke that results if the de-noising process is not applied.

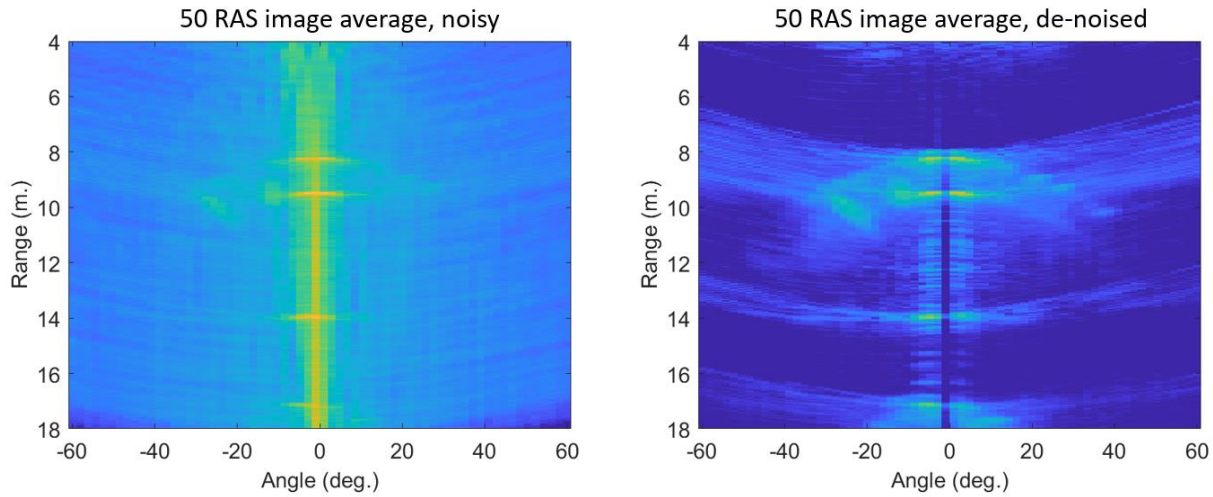

Figure 180: RAS images without de-noising, showing the spoke (left), and with de-noising, showing the null beam (right). Images are envelope averages of 50 RAS images, and plotted on the same 40dB dynamic range log-scale axis. Note that the null-beam isn't a perfect null because the broadside has phase curvature that is not modelled by a coherent sum across panels. Hence, a small amount of residual signal remains in the broadside direction.

By comparing averages of real-aperture images generated using the VLA transmitter with averages of real-aperture images generated using the alternating Neptune configuration, it was found that the omnidirectional Neptune sources enabled a wider swath to be collected. The following figure shows the averaged real aperture images used for comparison of transmitter configurations:

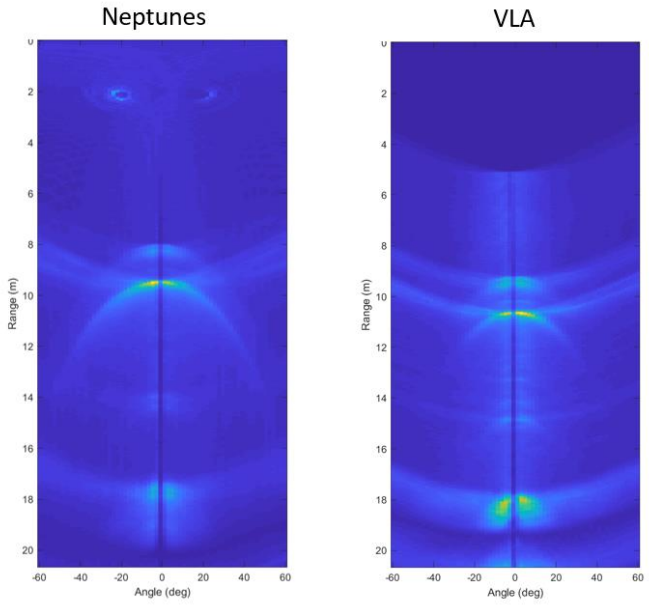

Figure 181: Averaged real-aperture polar images for the alternating Neptune configuration (left) and the VLA configuration (right). Linear scale.

In the Neptune configuration the sediment interfaces are visible over almost the entire 120 degree field of view of the polar image, vs. only about 50 degrees with the VLA. The difference in swath width is greater than a factor of two, however broad angles can still be imaged with the VLA, albeit with a lower SNR.

Sediment profiles beneath the sensor were found by applying pulse compression and de-noising to a set (in this context a set corresponds to the pings contained in a single 1 GB *.jsf* file: approximately 2400 pings), and calculating a MIP across channels. To enhance target signals, in some cases the pings were beamformed in the real aperture dimension and a MIP was computed across a defined set of angles. The following figures are profiles captured during day 4. Note that the overlaid red lines are the FOCUS 3's estimate of towfish altitude. The altitude sensor from the towfish was unreliable during these tests, hence the frequent drop outs visible in the plots. When the altitude estimate is valid, however, it corresponds accurately with the location of the specular glint from the top layer of the sediment.

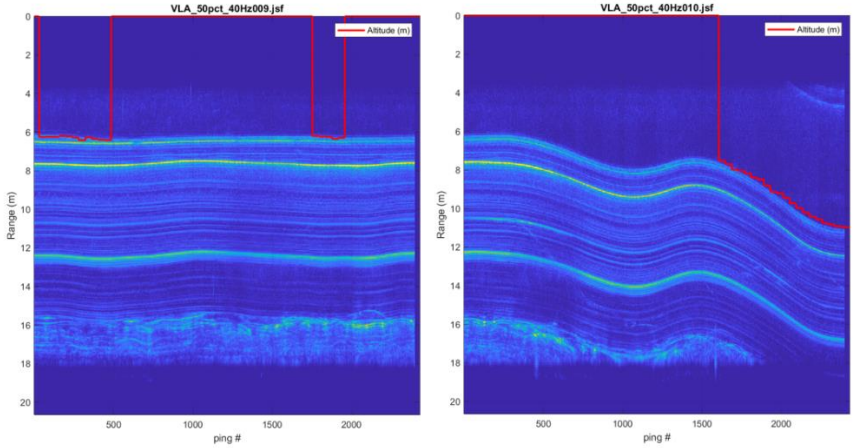

Figure 182: VLA profile, 8/8/2019, 8:14 AM (PST). Target range = 20m, not visible in scan.

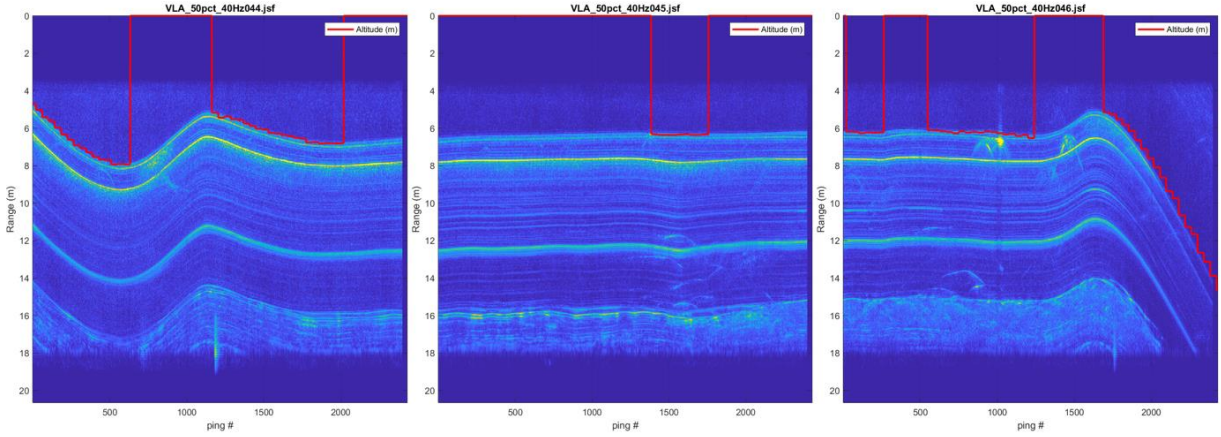

Figure 183: VLA profile, 8/8/2019, 8:50 AM (PST). Directly over pipe (target arcs visible in far right panel).

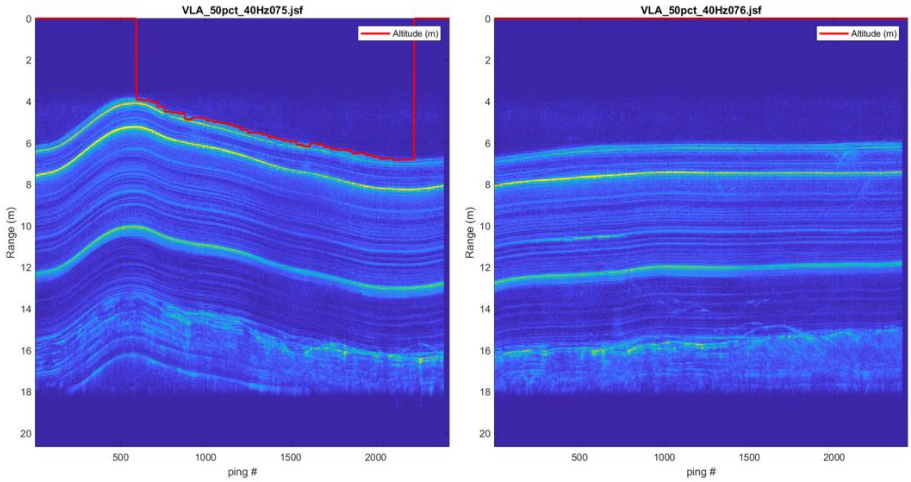

Figure 184: VLA profile, 8/8/2019, 9:20 AM (PST). The survey was nominally positioned to place the slug target +5 m. to the west of the north / south sonar survey line. A trace at ~7 meters near ping #2000 in the figure on the right may correspond to the target.

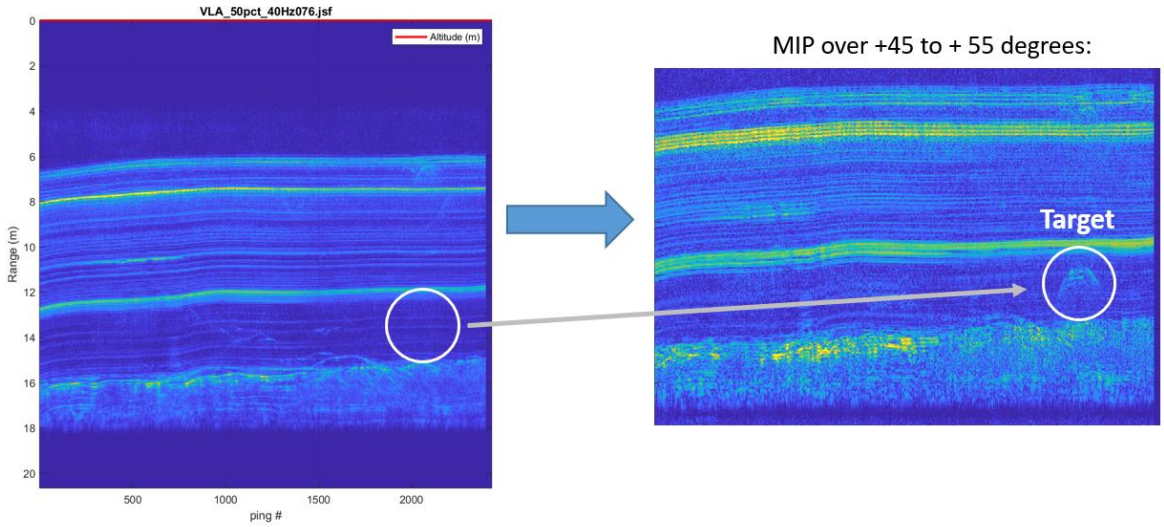

Figure 185: VLA profile, 8/8/2019, 9:20 AM (PST). Figure on the right shows a MIP through the angle-span containing the target (slug), making the target arcs visible. Note that the "layers" over this angle range are actually side-lobe artifacts.

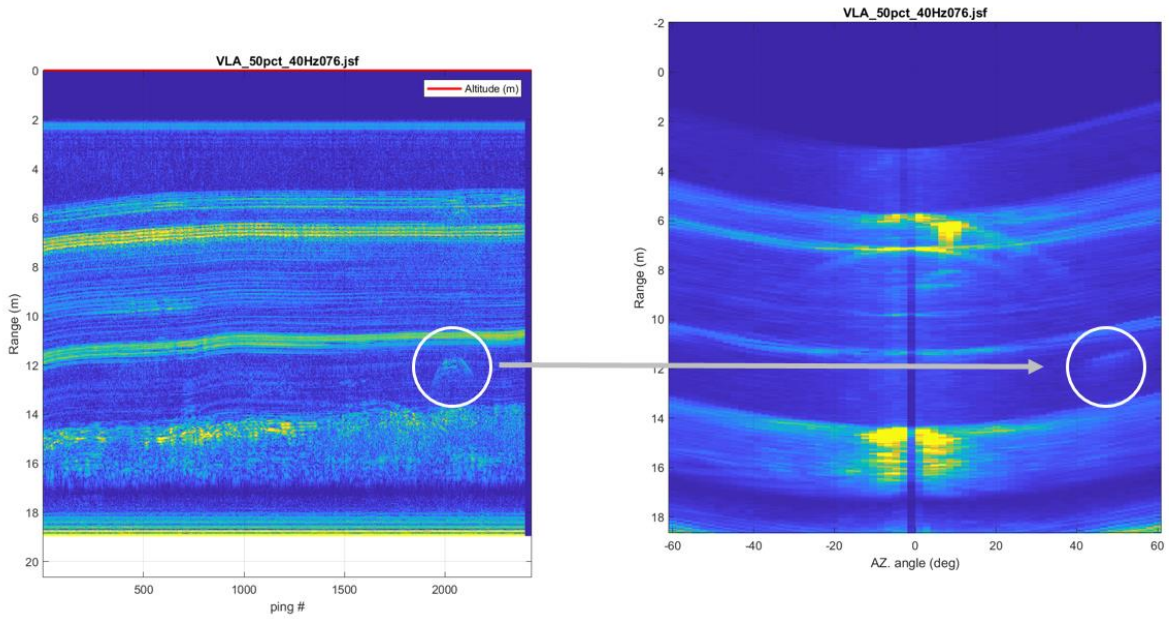

Figure 186: VLA profile, 8/8/2019, 9:20 AM (PST). A comparison of the profile scan showing the target with an average of the real-aperture images containing the target. The target is visible around +45 degrees (starboard). The high sidelobe level is an artifact of the window function used to weight the array during beamforming, poor uniformity of sensor response across the array, and de-noising process. Note that, due to VLA directivity the signal level from the target is low enough that poor placement could have put the target in a sidelobe that would completely drowned the target signal in noise. The wider directivity of the Neptune sources would prevent this type of problem from occurring as frequently.

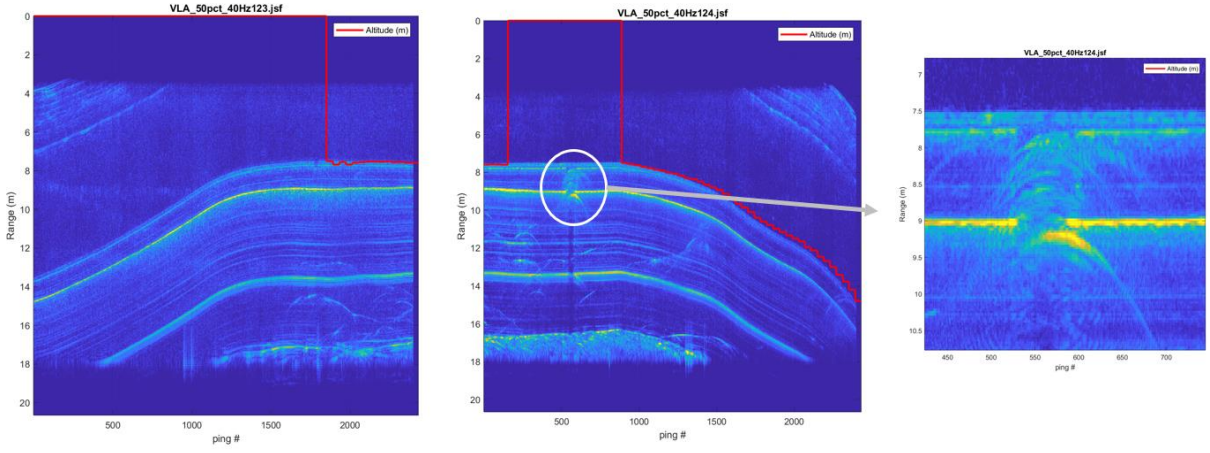

Figure 187: VLA profile, 8/8/2019, 10:10 AM (PST). Directly over target (slug). The target arcs corresponding to the slug are enlarged in the panel at right. Note that the target appears to have sunk down to the second sediment layer, and is also casting a shadow. Range aliasing (reception of backscattering data from previous pings) is visible in the top left and top right of the left and middle panels as well. This type of problem exists due to the rapid ping rate. The ping rate could be decreased, however this would cause major undersampling in the along-track direction, hampering SAS processing that will be applied to the data in the future.

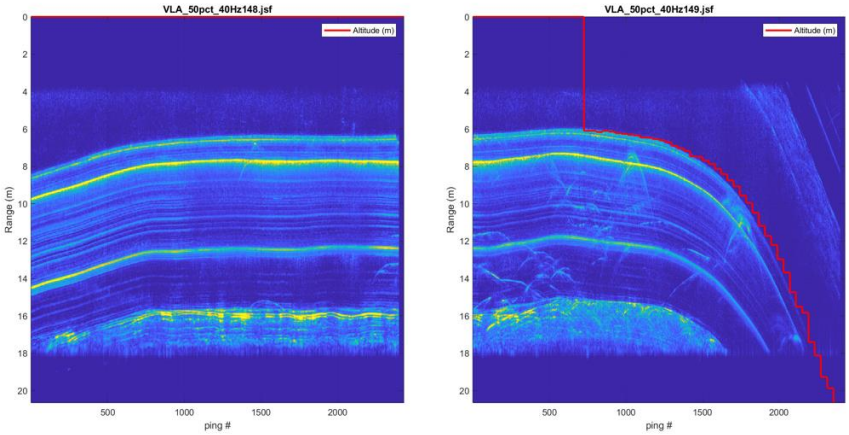

Figure 188: VLA profile, 8/8/2019, 10:56 AM (PST). Nearly directly over target (pipe).

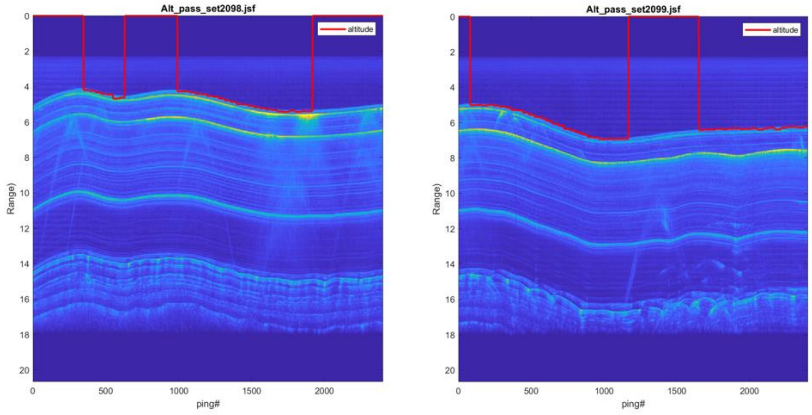

Figure 189: Alternating Neptunes, 8/8/2019, 1:15 PM (PST). This scan performed at different angle through target field. The scan is directly over the pipe and close to the slug.

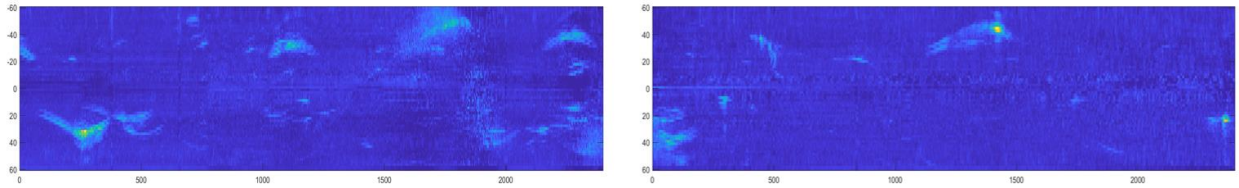

Figure 190: Alternative view of the data in Fig. 189. An adaptive filter was used to remove specular glints from sediment layers, leaving the rapidly evolving target arcs as a residual. A range-MIP through the real-aperture beamformed data was performed, showing the relative location of target and clutter target arcs relative to the track of the MuST. SAS processing, when applied in the future, should focus these arcs into targets and clutter. Vertical axes are angle, horizontal axes are ping number.

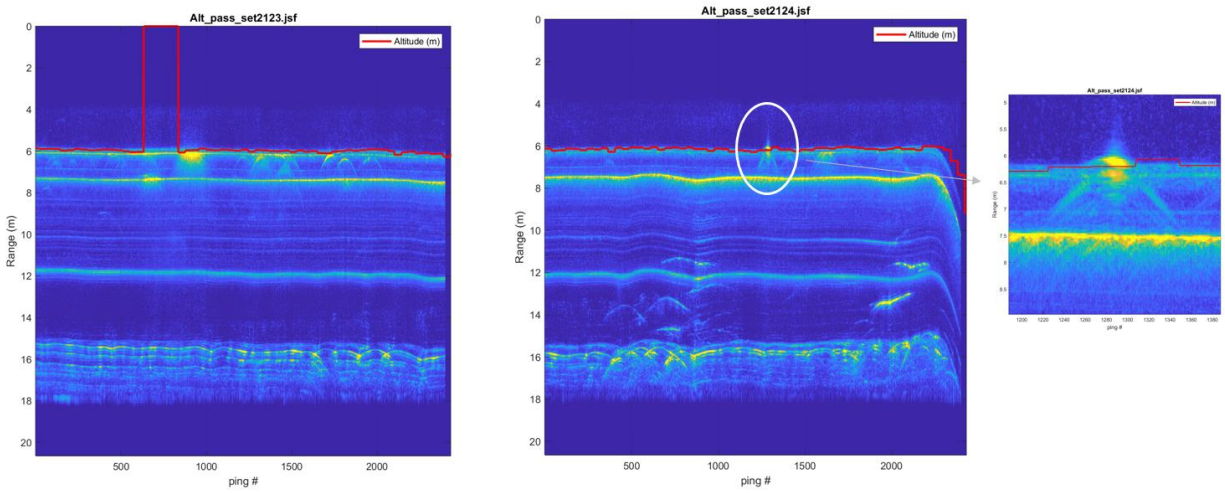

Figure 191: Alternating Neptunes, 8/8/2019, 1:40 PM (PST). Directly over pipe. Target is enlarged in the far right panel.

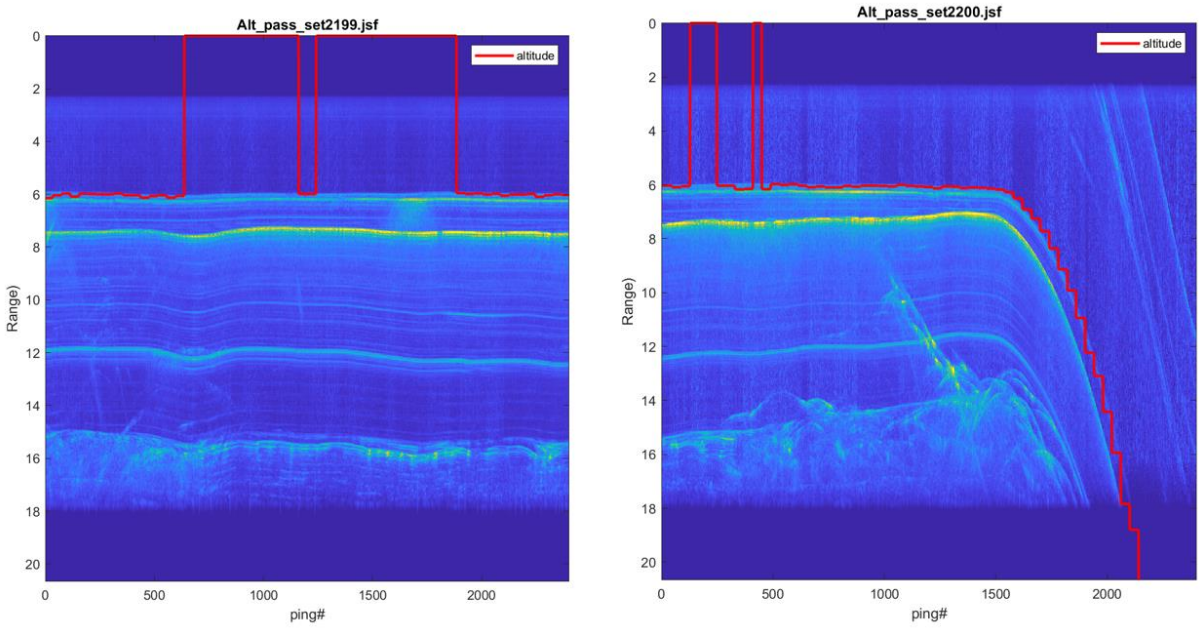

Figure 192: Alternating Neptunes, 8/8/2019, 3:07 PM (PST). Slug off to starboard.

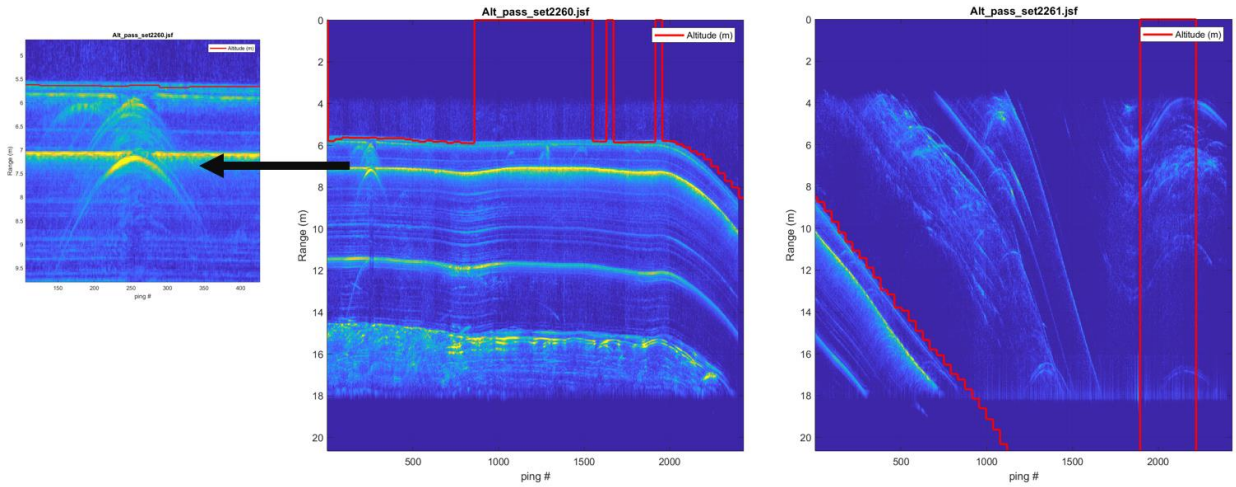

Figure 193: Alternating Neptunes, 8/8/2019, 4:09 PM (PST). Directly over slug. Channels 39-42 failed, in panel #3. The plot on the far right is an excellent example of range-aliasing from previous pings: as the vehicle climbs in altitude the backscatter data from the previous pings begin to fill in the region of the watercolumn.

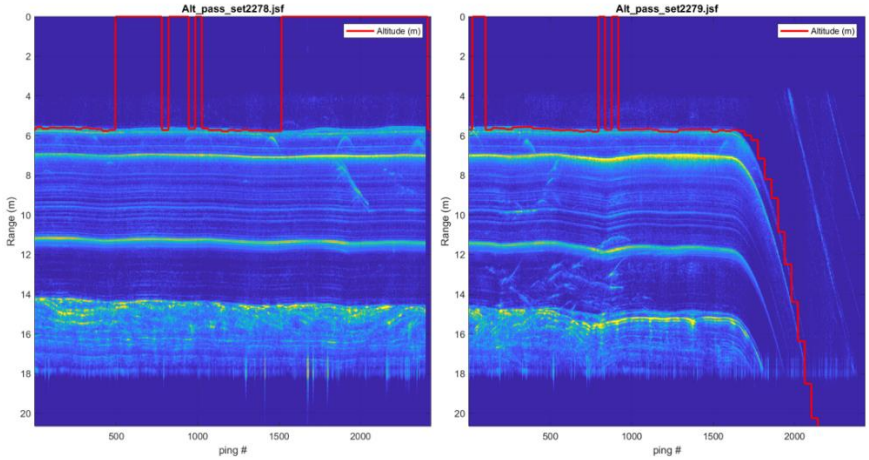

Figure 194: Alternating Neptunes, 8/8/2019, 4:36 PM (PST). Over slug. Final scan of the day.

In the previous figures it can be seen that, following denoising, quality appears good for both the VLA and Neptune transmitter configurations. In most of these scans it is possible to see sediment layers extending more than 10 meters into the sediment.

Despite a nearly complete failure on the part of the INS navigation filter, the manufacturer of the original INS allowed the system to be configured to output the raw signals measured by the linear accelerometers and rotation rate sensors embedded in the INS hardware, enabling a better solution to be estimated for the vehicle track in post processing. Results of the re-processing are shown in the next section and form the original basis of the renavigation algorithm incorporated later in the MuST BOSS toolbox. Note that for this original INS, angular orientation could *not* be assumed accurate (see, e.g. the roll state error shown in Fig. 177) and the angular values had to be estimated from the raw sensor data as well, creating a more complex non-linear optimization problem.

Inertial Labs INS and DVL navigation processing

During the interim period between the Lake Washington and Sequim Bay tests a renavigation algorithm was written to take all of the data from the various navigation and navigation augmentation sources on board the MuST and generate a more optimal navigation solution in post processing. As previously mentioned, the necessity to estimate orientations creates a non-linear problem which was solved using Newton's method. An example of re-navigation applied to the INS data is shown in Fig. 195, in which the re-navigation solution is compared to the original estimate from the INS:

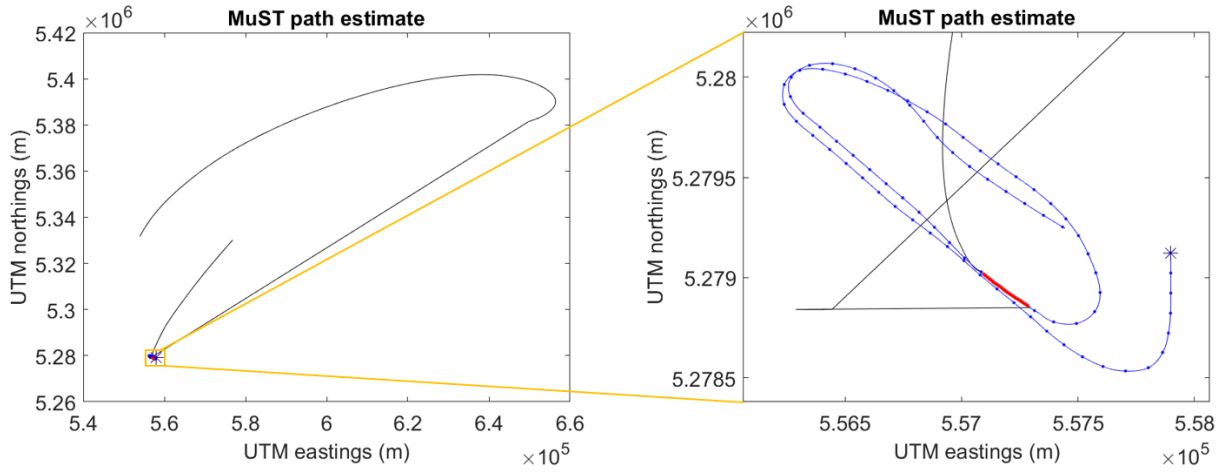

Figure 195: This figure shows the re-navigation solution as an inset to the INS estimate shown on left. The re-navigation position estimate is in blue, with blue dots spaced every 100 meters of track. The star marks the initial position of the towfish in the plot.

Note that the previous figure has been zoomed to allow the track to be visualized, and this ends up excluding many of the position estimates made by the INS systems internal software.

Due to ignorance regarding the problematic orientation states produced by the INS, it was not configured to output raw magnetometer data during the Lake Washington tests. When it determined that this would be useful for a complete navigation solution re-calculation, this option was selected for the Sequim Bay tests scheduled for the following month.

V. Field test results: Sequim Bay 2019 (MuST)

Overview

From Sept. 9 – 12th, 2019 the MuST system was deployed in Sequim Bay for its second official field test. During the interim between experiments the faulty panels had been sent to EdgeTech, where they were fixed and some adjustments were made. The new panels functioned much better than the previous set and, with the exception of a couple channels, operated without failure throughout the duration of the experiment. A large number of scans were acquired over target fields arrayed in both muddy and sandy sediment regions. The following figures show GPS ship positions for the separate days of the experiment, overlaid on a map of Sequim bay, and some pre-defined track lines for scanning the target field:

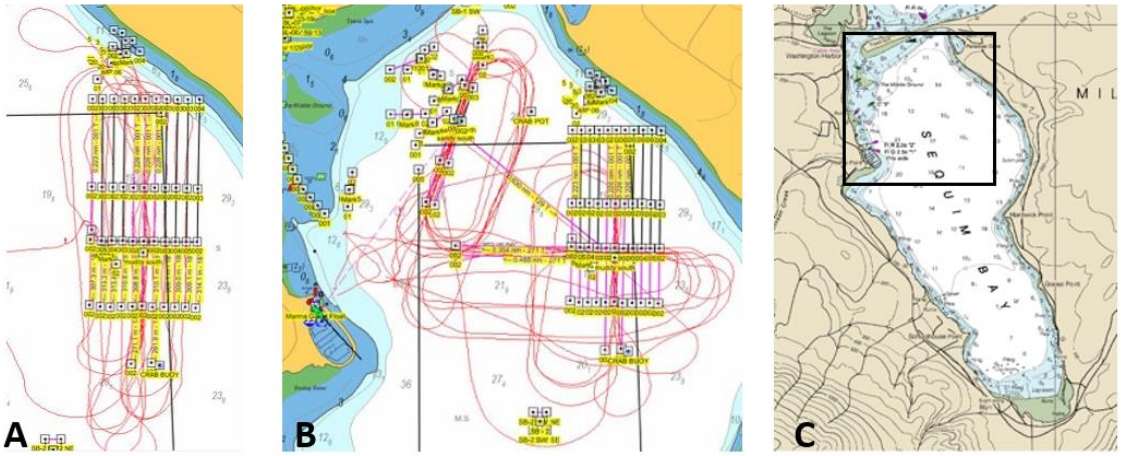

Figure 196: Track lines over the muddy target field and ship coordinates from day 1, 9/9/2019 (A) and day 2, 10/9/2019 (B). Scans were also conducted over the sandy region on day 2, visible as the set of race-track loops in the upper left of the figure. The region of the map shown in (B) is shown in its greater context in (C).

Figure 197: Ship coordinates for day 3, 11/9/2019 (left) and day 4, 12/9/2019 (right). The bold rectangle in the figure at right is approximately 500 by 1700 meters.

Target fields were positioned in two different environments within the bay: a muddy site and a sandy site. The target fields contained various UXO, science targets and clutter and pipes, which were arranged according to the diagrams in Fig. 198:

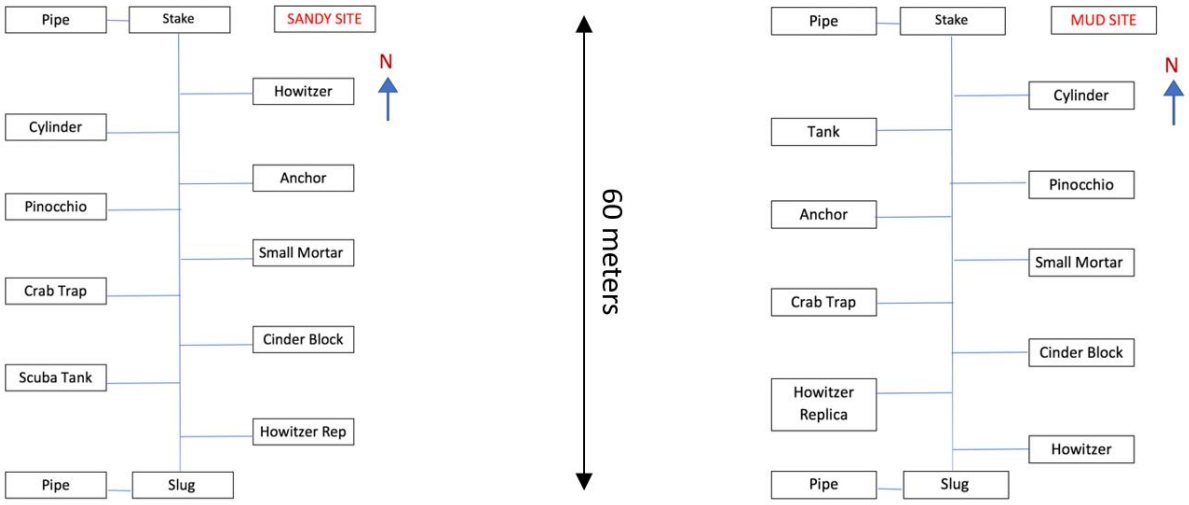

Figure 198: Target field item position maps for the sandy site (left) and muddy site (right). Targets were positioned 5 meters east / west of the center line for each field, and had 10 meters spacing in the north / south direction between other objects on the same side of the line. The total length of the fields were approximately 60 meters.

The APL-UW ocean engineering team compiled software to assemble all the pertinent navigation information into a single data stream that was broadcast via TCP/IP. Unlike during the Lake Washington experiments, in which navigation data was received through the BOSS computer and thus only when *.jsf* files were being output by the system, the navigations stream was now continuous, providing much better options for re-navigation. Additionally, the INS software was configured to export all raw sensor measurements, including each signal from the 3-axis magnetometer.

Live-streaming Data Visualization

The next few figures show the live data display panels that were used to monitor the output of the BOSS for the Sequim Bay experiment. Note the significant advancements to live data-processing that were accomplished between 2019 and 2021.

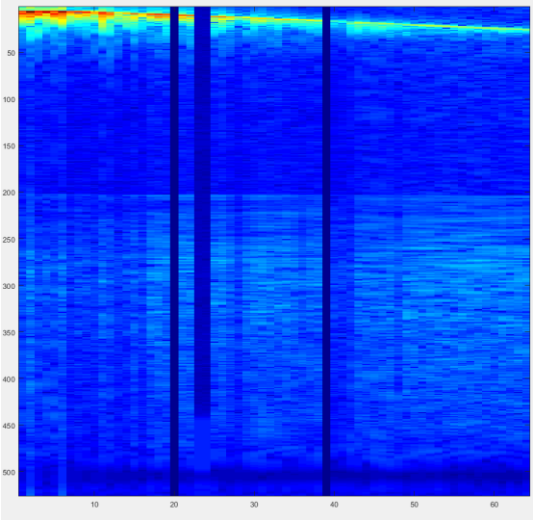

Figure 199: Live pulse-compressed channel data visualization panel. This window would refresh at approximately the ping repetition rate. Channel number is the horizontal axis. The port-side element is channel #1 and the starboard-side element is channel #64. Temporal sample bin is in the vertical dimension. The bright signal near the top of the matrix is the direct path, which peaks in amplitude near the port side, indicating that the transmitter configuration is the alternating Neptune configuration, and that this measurement is from a port-side transmission. Two channels, #20 and #39, are inactive, and two more, 3 & 4, have reduced sensitivity. On subsequent days of the experiment (Sept. 10 – 12), these panels were replaced with panels from EdgeTech that had an improved design and performance was improved.

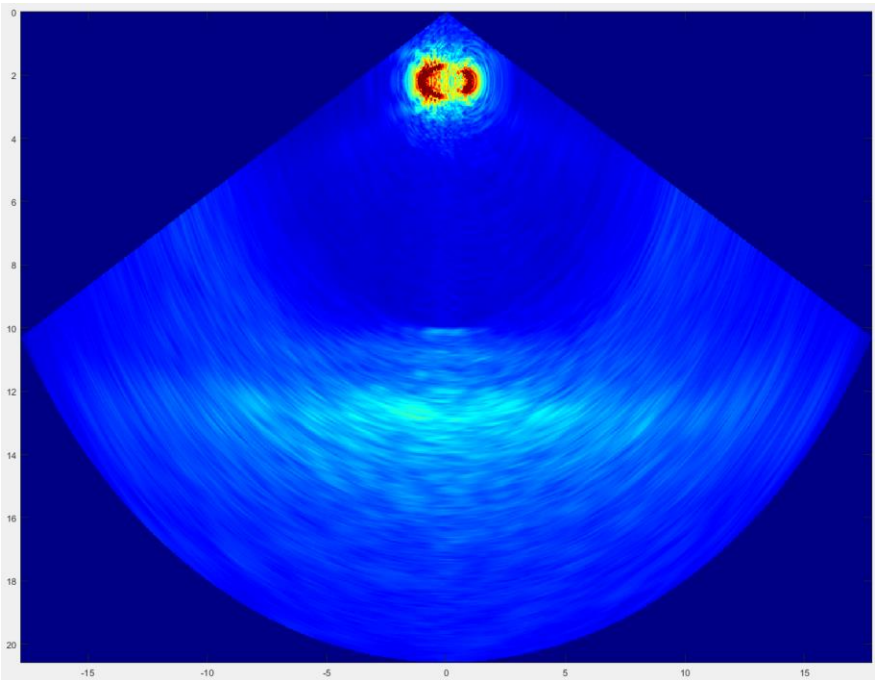

Figure 200: A screenshot of the real-aperture, single ping visualization panel. The bright artifacts near the image origin are beamforming artifacts from the direct-path signal to the array. In this image the sediment is approximately 10 meters below the sonar. The image is in Cartesian coordinates and approximately 20 x 35 meters in size. Only the highlighted region in the window contains beamformed acoustic data, however. This image is refreshed at approximately the pulse repetition rate.

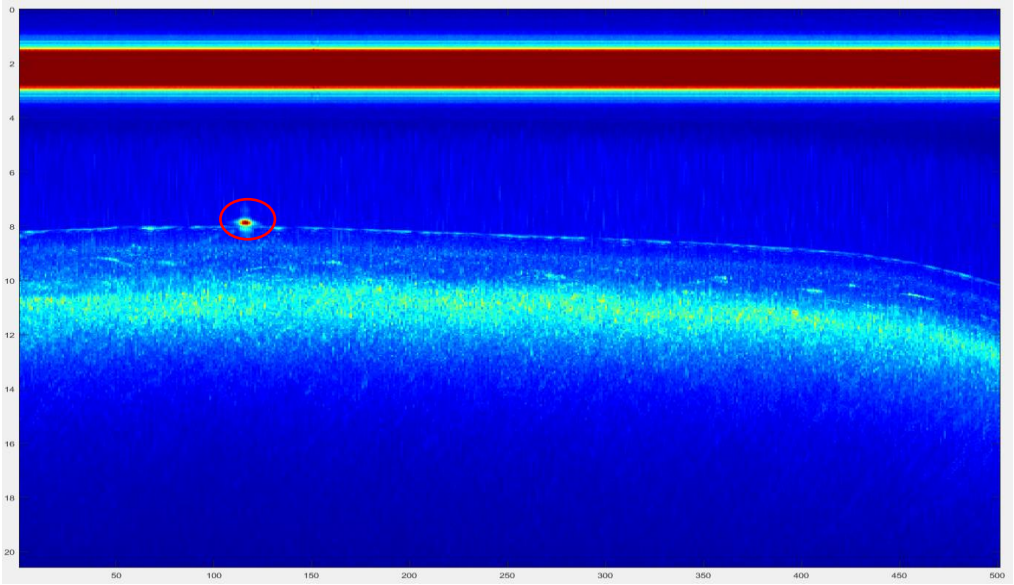

Figure 201: The scrolling sediment profile panel. This image is an updating queue of profiles captured using the central angles of the real-aperture frames. The vertical axis is range from the sensor, and the horizontal axis is ping (the numbers are reversed, new data appears at sample #500 and works towards sample #1). The profile scrolls from right to left as data is acquired A target near the surface is highlighted in the figure. The red strip near the top is the direct-path signal. The signal level is clipped to reduce the dynamic range of the image and make the lower level target signatures visible.

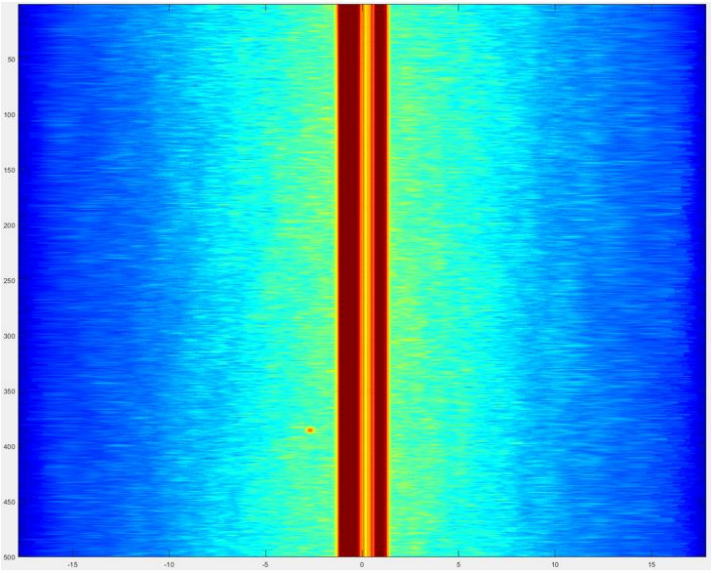

Figure 202: Vertical MIP view. This is a scrolling window that shows the azimuthal displacement in meters from the track of the vehicle, and scrolls from top to bottom. The horizontal axis is port/starboard displacement relative to the towfish, and the vertical axis is sample. The image is generated by computing a MIP in the vertical dimension through the real-aperture frames The target circled in Fig. 201 is visible in this figure at (-2, 380).

Using these data visualization panels it was possible to determine the health of the sensor and individual channels and, to some extent, determine when the target field was being crossed. The

lack of SAS processing limited the amount of information that could be assessed about individual targets, however.

Beamforming Results

Following collection of the data, SAS processing was applied in post processing. A number of changes had to be made to the 3D FFBP algorithm to make it compatible with the navigation data captured by the system. Figure 203 is a comparison of a range MIP through raw data with a Z-MIP through the data after it has been beamformed using 3D FFBP:

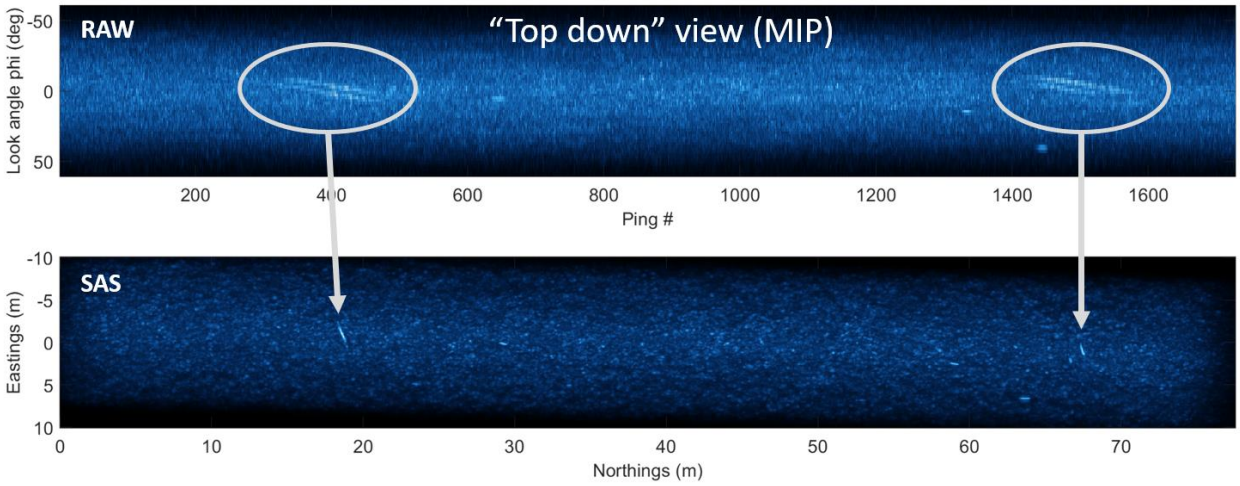

Figure 203: A MIP in range through the raw data (top) and through the data-cube generated by applying SAS beamforming. The ellipses in the top figure are the un-compressed target arcs corresponding to two targets in the bottom figure, denoted with the arrows. The target field is the field in the muddy sediment.

In Fig. 203, the enhanced resolution and SNR resulting from SAS beamforming is readily apparent, however the targets themselves are poorly visible, due to MIP processing displaying the contributions from the brightest targets across the vertical dimension. If this range is reduced to include just the region in the immediate vicinity of the sediment interface and region extending only half a meter below it, then the targets are immediately visible against the background:

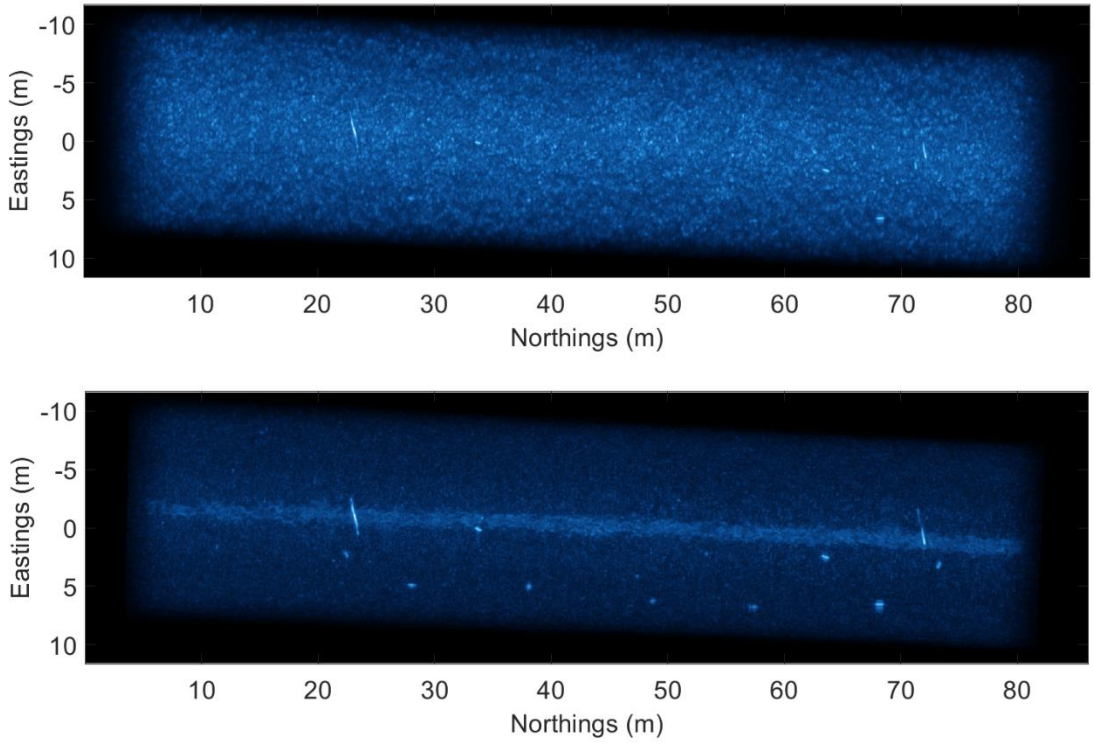

Figure 204: Z-MIPS through the entire beamformed volume (top) and just the top half-meter through the sediment, including the surface. When the range of the MIP is truncated, the target field becomes immediately obvious due to a reduction in clutter.

The reason for this becomes more apparent when a MIP through eastings (cross-track in this example, in which the towfish was moving north) is computed:

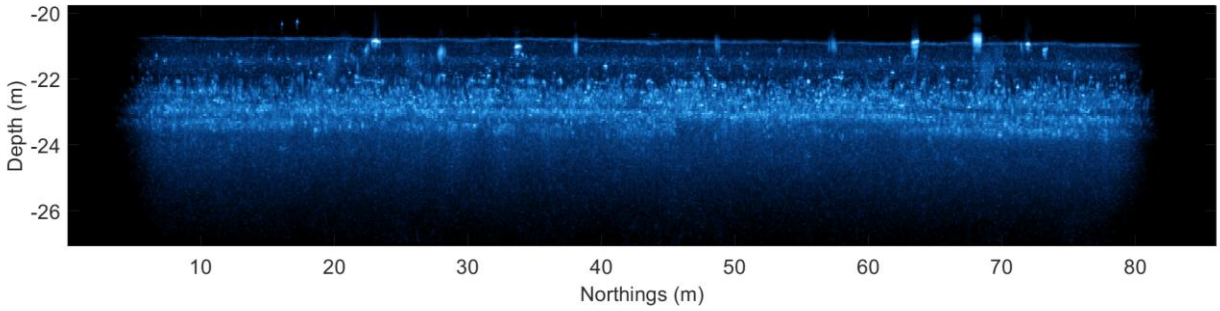

Figure 205: A MIP through the eastings dimension of mud field dataset shown in Fig. 203 and 204. A strongly backscattering sediment layer is visible ~3 meters below the water/mud interface that masks the target response when a MIP is calculated through the vertical dimension. Targets are the discrete bright scatterers clustered within a meter of the water/mud interface.

The only discrepancy between the target field observed in Fig.'s 204 and 205 and the target field described in the map (Fig. 198) is the lack of a crab-pot near the center of the target field. In fact, this target is visible, and represents one of the failings of MIP's and a difficulty of working with

volumetric data. Figure 206 shows the region in which the target should be present, a zoomed segment of the MIP image, and several slices of the volume in that region.

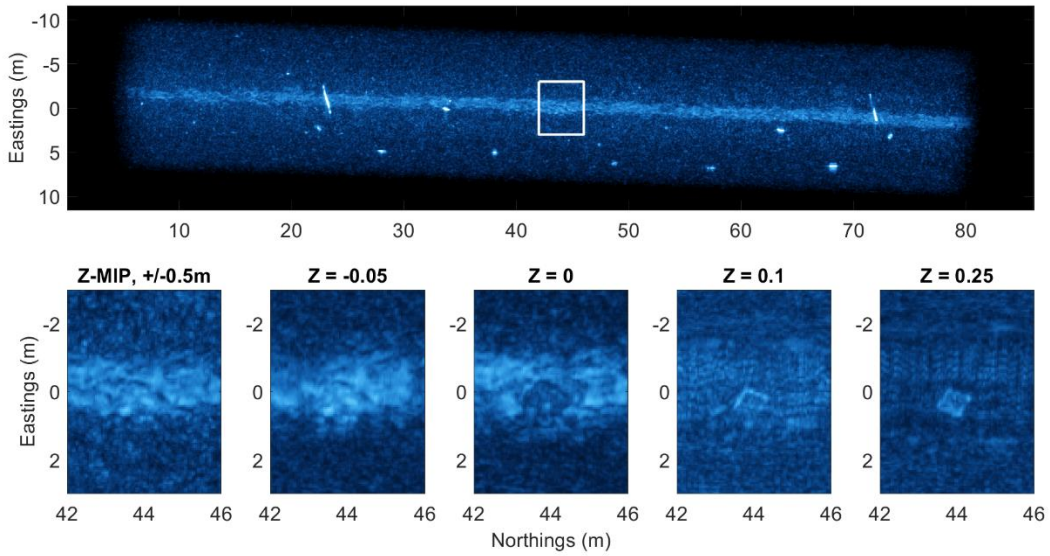

Figure 206: Crab-pot patch showing the Z-MIP and zoomed patch (top and bottom left), and iteratively changing the height plane of the slice through the volume. It can be seen that the crab-pot makes a disturbance in the sediment (middle plot), and is also visible in the watercolumn (bottom right two figures). Note that the dynamic range of the bottom right two figures has been altered to make the crab-pots visible. The low scattering strength of the crab-pot is why the object is invisible in the Z-MIP (bottom left).

It immediately becomes apparent that the target, a crab-pot, is present: it just happens to have a scattering strength so far below the scattering from the sediment that it is completely impossible to see in the MIP. The target does serve, however, as qualitative verification that the SAS processing is incorporating the navigation data into the beamformer correctly. As per the previous discussion on autofocus, if navigation data was inaccurate the crab-pot would experience the ordinary SAS focus defects, such as along-track blurring, that are common to cases when navigation error is corrupting image quality.

Two additional targets serve as excellent navigation and beamformer focus verification tools. At the southern end of the target field, two spherical floats were anchored near the pipe and suspended very near the sediment interface (i.e. several tens of centimeters). Often, floats defocus in SAS imagery due to drifting with the watercolumn. In the present case, however, ambient current oscillations appear to be stable enough that the targets focused to points. Not only are floats visible, but their anchors are visible, buried in the sediment. Figure 207 shows the location and zoomed MIP images of these targets:

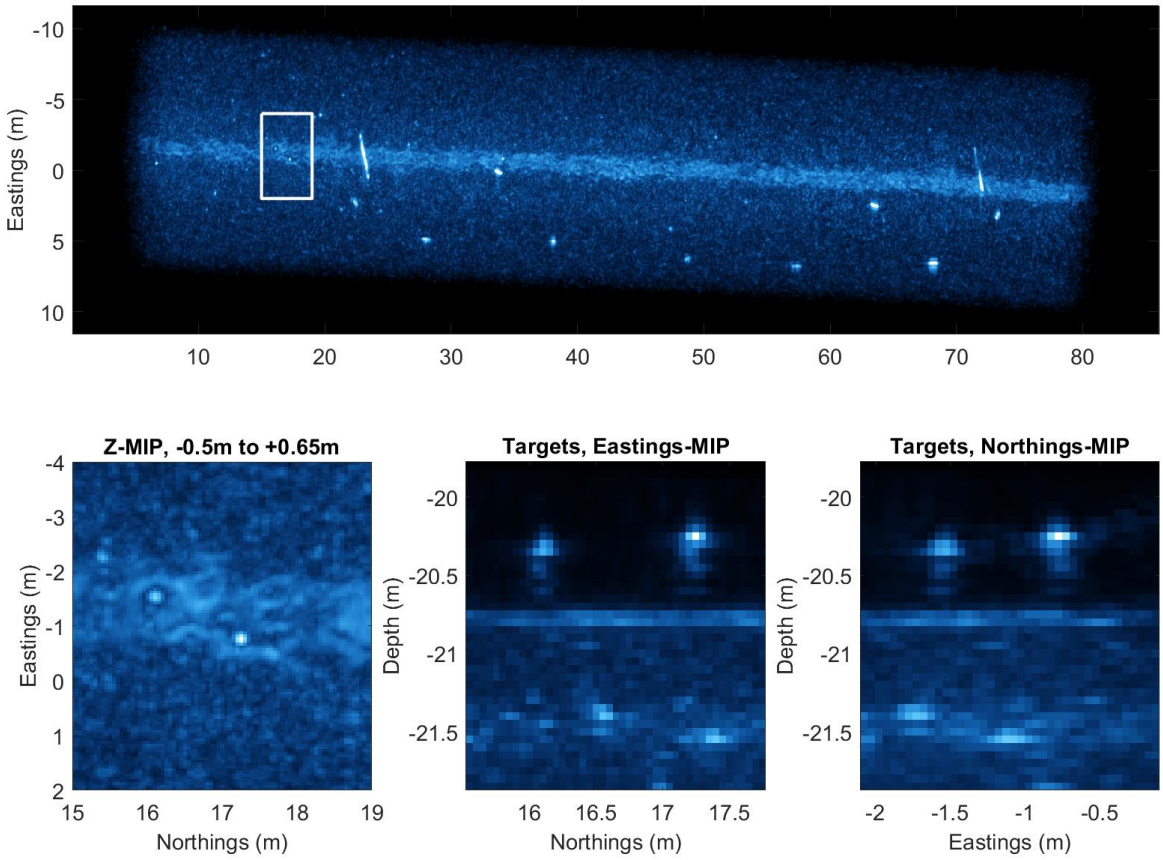

Figure 207: Target field end-marking floats. The bottom three images are MIP's through depth, eastings and northings, showing the targets and their anchors about half a meter beneath the sediment. The focus for these targets is very good despite being in the watercolumn.

In addition to targets, another scattering mechanism becomes apparent: the specular reflection from the mud layer. The specular reflection manifests as the bright strip that appears down the middle of scans and was the primary motivation behind the development of the de-glinting filter described in *Methods and Materials: De-glinting*. In soft mud sediment this strip is not extremely bright relative to the scattering strength of the targets. In sand, on the other hand, the specular reflection is very bright due to the higher impedance discontinuity between sand and water. This difference is visible in the example below which shows a Z-MIP through the sandy target field (compare with Fig. 204, top).

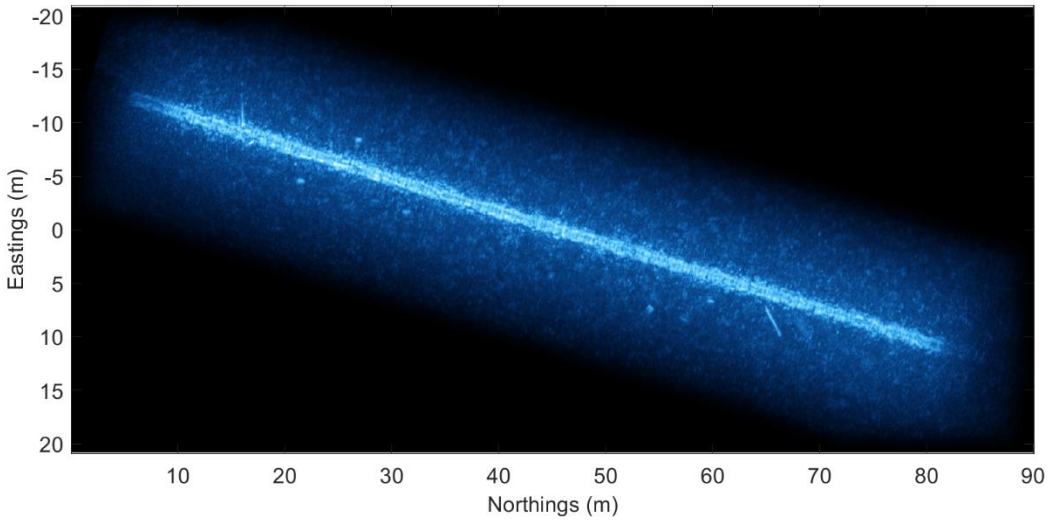

Figure 208: Target field in sand, showing a relatively strong specular glint compared to the mud patch shown in Fig. 204.

By comparing Fig. 207 (top) with Fig. 208 it can be observed that targets buried in mud at nadir can be easier to see in a Z-MIP than targets proud on the sand due to the brightness of this specular reflection.

Figure 209 is a side-by-side comparison of a dataset that has been beamformed using a single Neptune transmitter vs. the alternating configuration. The exact same data was used for the beamforming process, however data from one of the alternating transmitters was withheld. This method of resolution and performance comparison between single/double transmitter configuration was used because the towfish never traversed the exact same path using the VLA and alternating Neptunes. This method, in contrast, enables a direct comparison between configurations for many targets of opportunity.

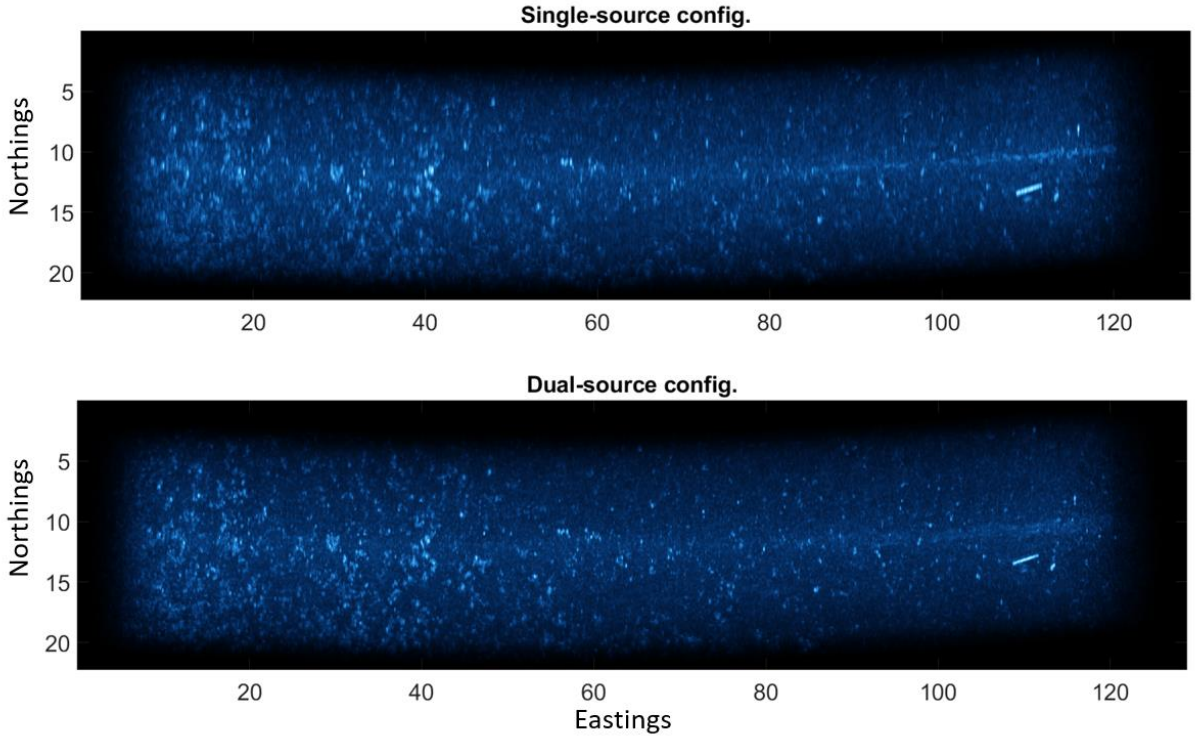

Figure 209: Single transmitter configuration (top) vs. dual transmitter configuration (bottom) using the Neptune transmitters. The figure on the top was generated from a dataset in which the data corresponding to port-side Neptune transmission was withheld, effectively making the data single-source.

The following plots are bisections of targets of opportunity found in the Z-MIP image.

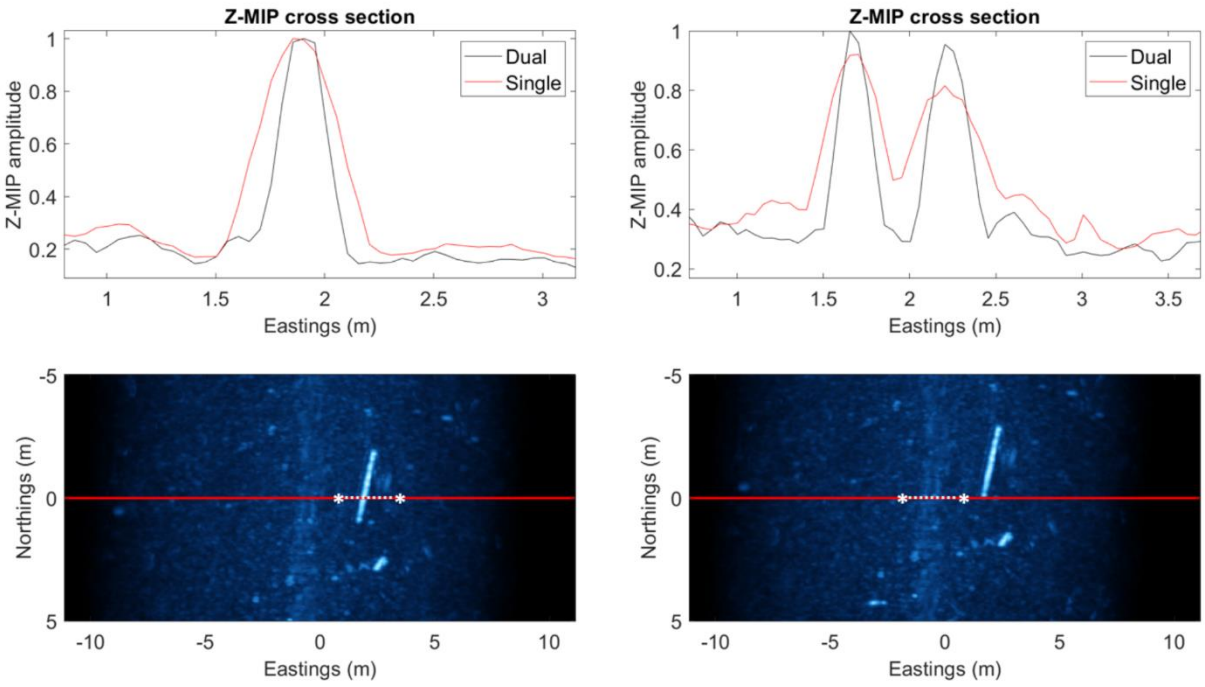

Figure 210: Amplitude values of two rows of the Z-MIP in Fig. 209 (rotated), showing the amplitude value of the row for the dual and single source configurations (top) and the location of the profile shown in the Z-MIP (bottom row plots).

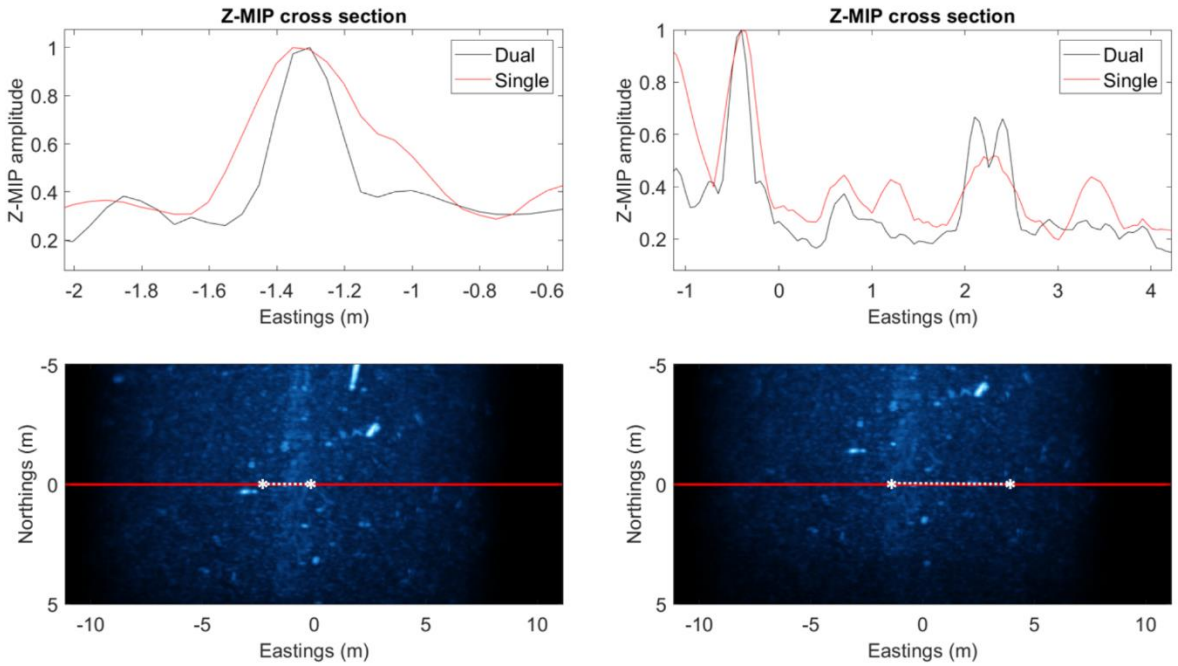

Figure 211: Profiles of two more targets of opportunity for the dual and single transmitter configurations (top), and the context of the slices (bottom). Note that in the profile on the top right there are two closely spaced scatterers that are resolved in the dual transmitter configuration but cannot be resolved in the single transmitter configuration.

On average, the resolution improvement due to the usage of the dual-source configuration was found to be a factor of approximately 1.7 times. This is quite significant, considering that the alternative for achieving similar resolution would be to make the array almost three meters long.

The following figure is a side-by-side comparison of the muddy target field captured using the VLA transmitter and the dual Neptune transmitter configuration.

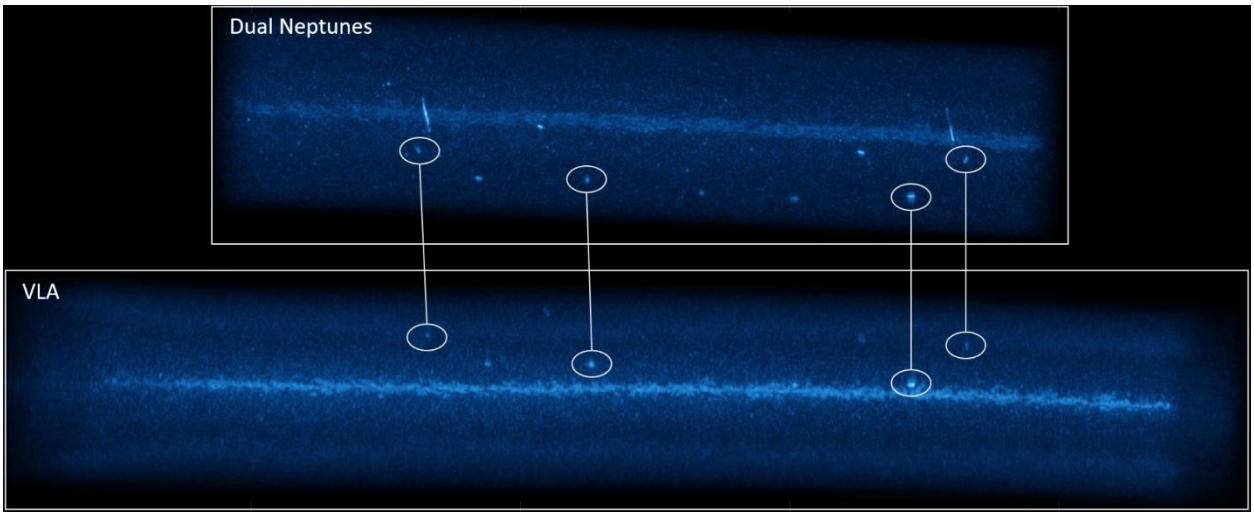

Figure 212: Transmitter configuration comparisons. The scan plotted on the top shows imaging results for the muddy target field using the dual Neptune configuration, and the bottom image shows the image generated using the VLA transmitter. Lines are drawn showing some of the common targets between the scans.

In Fig. 212 the towfish was moving north in both cases and the target field was off to starboard when dual Neptunes were being used, vs. off to port, when the VLA was used. It is interesting to note that the specular glint seems to be a more dominant contribution in the VLA image, and that, perhaps due to the tighter directivity of the VLA source, the regions near the extrema of the beam seem to be noisier in the VLA. Additionally, the pipes are not visible in the VLA scan, though this is most likely not a consequence of the transmitter being used but the fact that the broadside glint is lost due to the sensor not crossing the mainlobe of the backscattering response of that target feature.

BOSS Structural Acoustic Data: First Look

Structural acoustic feature enhancement had not been developed for BOSS processing when the 2019 Sequim Bay was collected and being analyzed, however it was found that, similar to lake Washington, off-nadir targets embedded in muddy environments generally have a much higher SNR than counterparts embedded in sandy sites, and unlike in sandy regions, decent acoustic color plots can frequently be generated without performing any extra filtering to reduce volume scattering, provided the region of interest doesn't contain clutter objects of competing scattering strength. An example of one scan over the muddy target field is shown in Fig. 213. The top figure shows a projection through a thin (i.e. ~50cm) projection through the sediment/water interface, and most of the targets in the target field are prominent in this projection, including the "Small mortar" (fig. 198). The bottom two figures show bounding boxes outlining regions-of-interest that are 1.5 x 1.5 x 2.25 m on edge. These bounding boxes have been drawn around several prominent objects in the target region (one low SNR buried object has been included as well). The regions of interest are labelled and have the following correspondence: 1 = Stake, 2 = Cylinder, 3 = Tank, 4 = 105mm HEAT, 5 = Cinder Block, 6 = buried unknown, 7 = Howitzer Shell Replica.

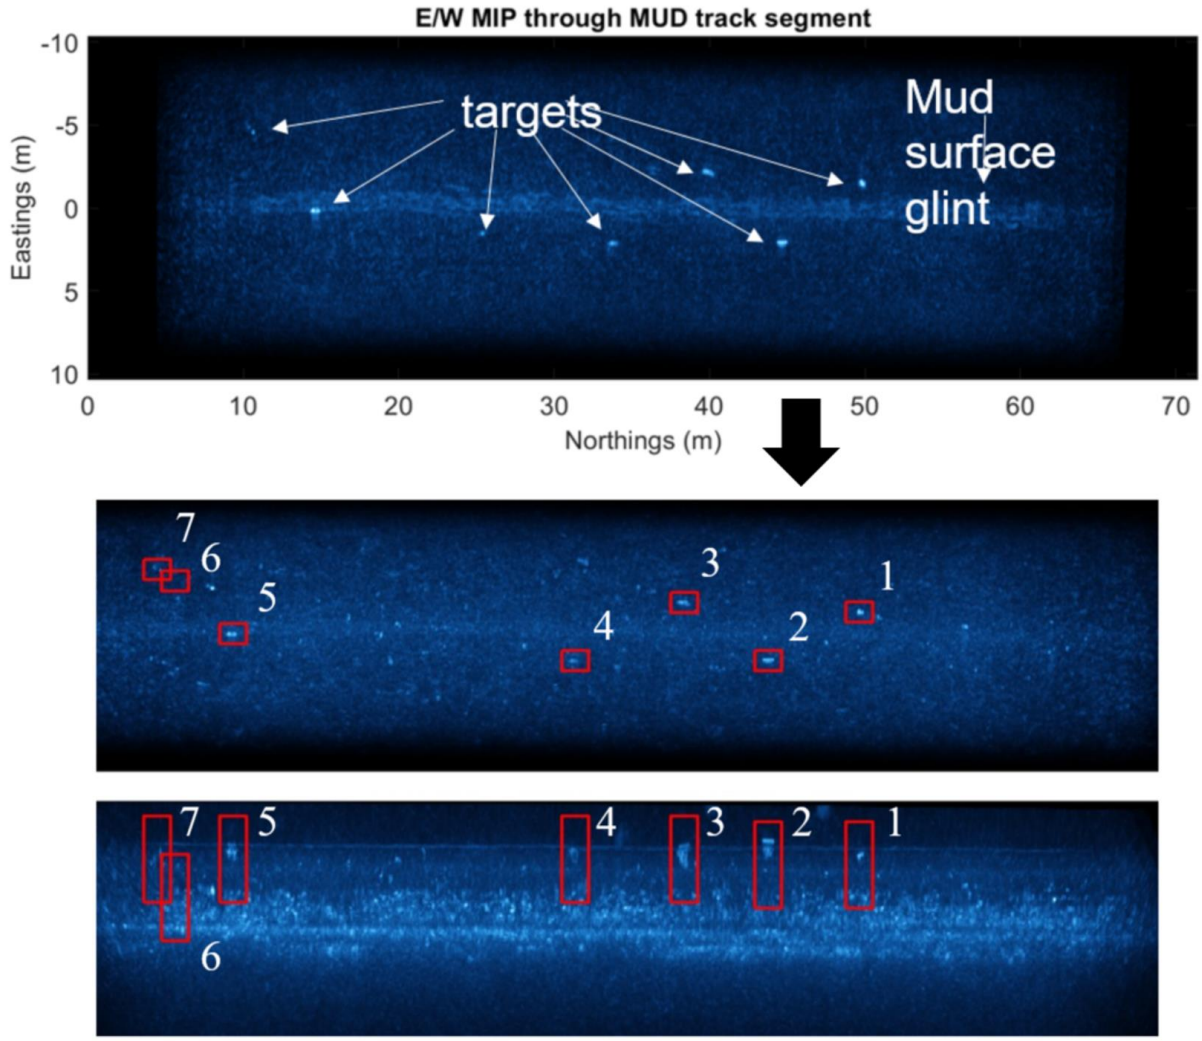

Figure 213: ROI's detected in the 2019 Sequim muddy target field. The top figure shows a MIP through a thin layer near the sediment / water interface and many targets are visible (compare with Fig. 207 where *all* of the targets are very prominent except the crab pot). The bottom two figures show a Z-MIP and Eastings-MIP with outlines around some regions of interest. Regions of interest are labelled as follows: 1 = Stake, 2 = Cylinder, 3 = Tank, 4 = 105mm HEAT, 5 = Cinder Block, 6 = buried unknown, 7 = Howitzer Shell Replica.

Figure 214 show the extracted regions of interest over several different rotation angles.

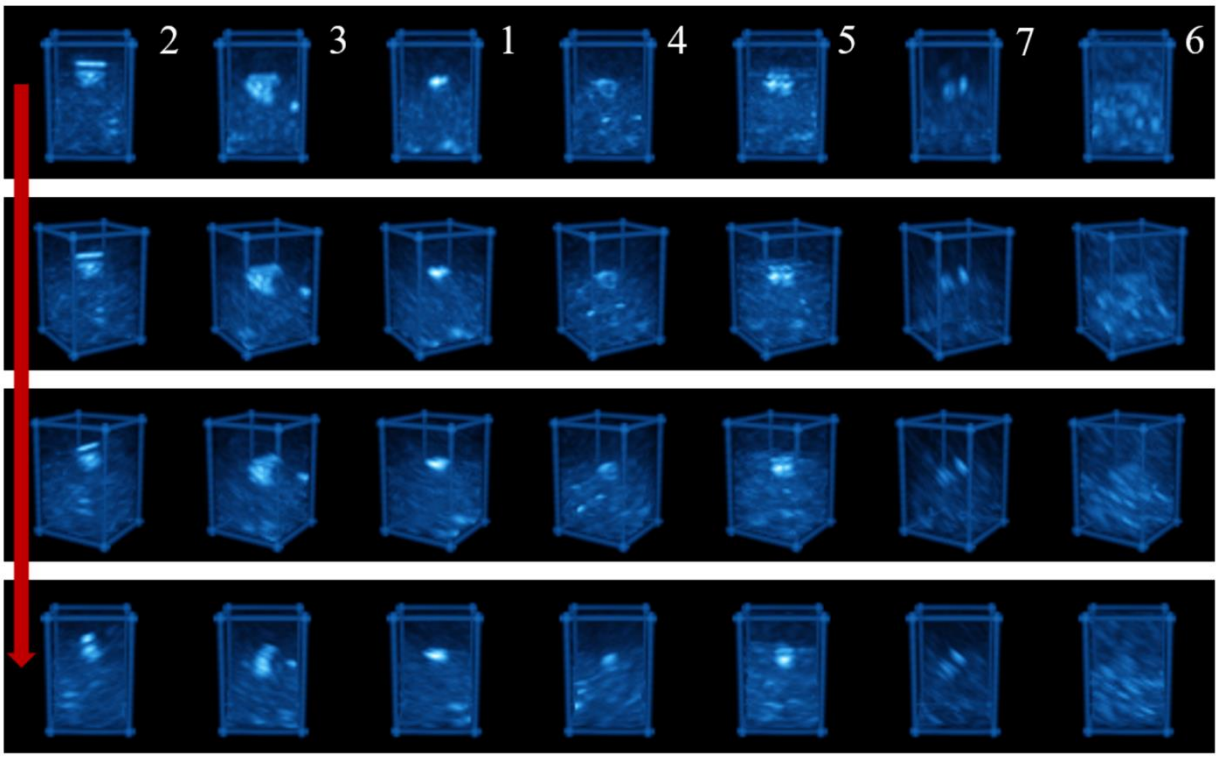

Figure 214: Rotations of MIP volume renderings corresponding to the ROI's outlined in Fig. 213.

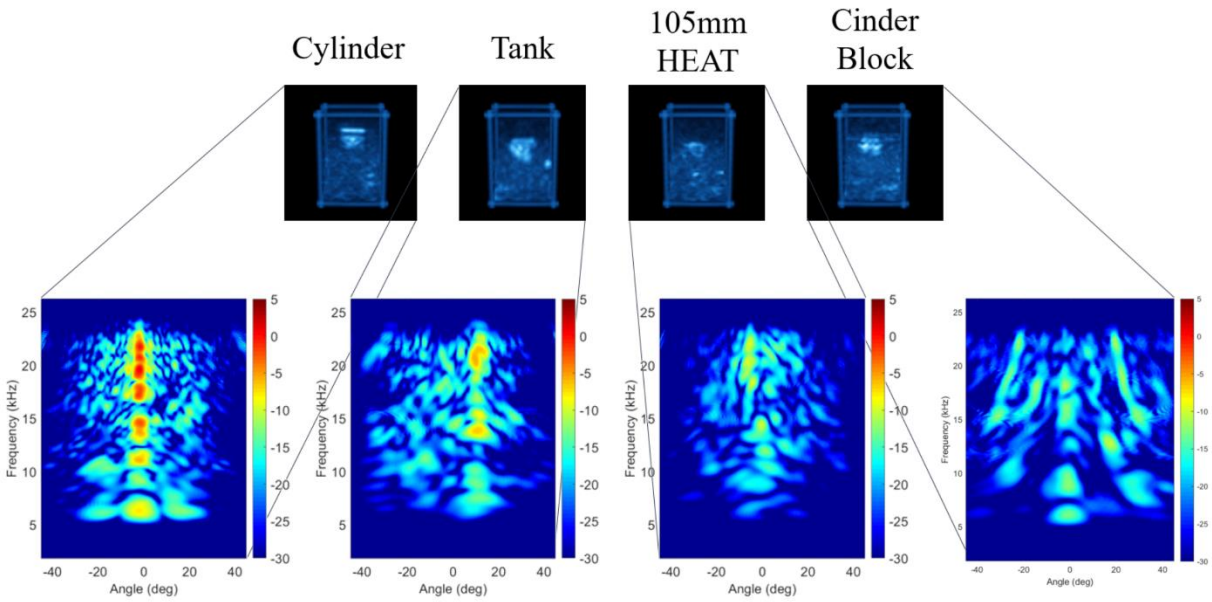

Figure 215: Acoustic color projections through the real-aperture dimension created from the hollow 30 x 60 cm aluminum cylinder (left), scuba tank (second from left), "Pinocchio" mortar (2nd from right) and cinder-block (right).

Figure 215 shows the acoustic color plots generated from several regions of interest containing no additional processing for signal-to-noise enhancement beyond some range-gating. These were

the first acoustic color plots generated from the BOSS mounted on the MuST system and gave very encouraging results due to the similarities (especially the pipe and cinderblock) with previous results generated in ONR / SERDP jointly funded BAYEX and CLUTTEREX experiments (see e.g. SERDP project MR-2505, "Acoustic Response of Underwater Munitions Near a Water-Sediment Boundary," PI: S. Kargl).

Summary

The 2019 Sequim Bay system test was successful both from a system integration and data acquisition standpoint and helped determine the direction future BOSS data analysis strategies would take. On the most basic level, the sensor functioned more reliably due to the improved receiver array panels sent by EdgeTech. Despite being simplistic compared to later developments, the crude live data acquisition and analysis tools proved useful during the entire experiment for applications ranging from determining when the target field had been crossed and highlighting key features of the sediment (e.g. layering and depth relative to the sonar), to performing array health diagnostics. The re-navigation algorithm worked well with the new navigation data format being assembled from the sensors and streamed to the BOSS data processing computers, and in post processing the beamforming worked extremely well, especially for the experimental dual-transmitter design. The successful integration of navigation data into the beamformer was verified by the compression of isotropic scatterers such as floats to points, and the good focus found on targets such as the crab-pot in Fig. 206.

For visualization it was shown that though MIP's are a useful tool for viewing 3D acoustic data they have drawbacks: when targets have extremely weak scattering response the signal-to-clutter ratio can be severely degraded if scattering from bright sediment layers is included along with scattering coming from the plane of the target field. This prompted the usage of projections through gated depth-ranges as a means of data visualization, and this method of visualization was incorporated into MuSTLive, SAS_view, and the MuST Toolbox SAS batch-processing software.

Lastly, the decision to upgrade the INS based on the poor underwater localization was significant as it led to INS model replacements, eventually resulting in the integration of a system which worked very reliably in the 2021 Sequim Bay experiments. Despite the upgrades, the D-GPS (Dual-GPS) unit that came with the original Inertial Labs INS was used as the absolute position reference sensor by all subsequent INS systems used throughout the rest of the project.

VI. Field test results: GS4 INS Integration

Overview

The purpose of this field experiment which spanned several days in June 2020, was to verify correct integration of the GS4 inertial navigation system and, while doing so, test the ability of the MuST BOSS to detect a 30 x 60 cm solid aluminum slug, a 30 x 60 cm hollow aluminum shell, UXO facsimile, and a straight length of nylon-coated ¼ inch steel cable. This was also the first field experiment in which a prototype for the realtime FFBP beamformer was tested, and results were extremely good. This was also the first dataset in which range-alias imaging

(*Materials and Methods: Range-alias imaging*) was tested. During the scan several targets of opportunity, including the Lake Washington PB4Y plane wreck, were scanned. This results section will show some of the results of these experiments.

Targets and cable

Several acquisitions from different angles were acquired over the target & cable deployment. Figure 216 shows an image of the target field captured by the MuST in which each target is visible *and* the nylon deployment and recovery rope connecting the targets is also visible. The visibility of the recovery rope is remarkable given that it has a very small scattering cross-section and low impedance-discontinuity relative to the water and sediment. A possible explanation for the visibility of the rope would be the presence of trapped air in the nylon braids, increasing the impedance discontinuity between the water and the rope.

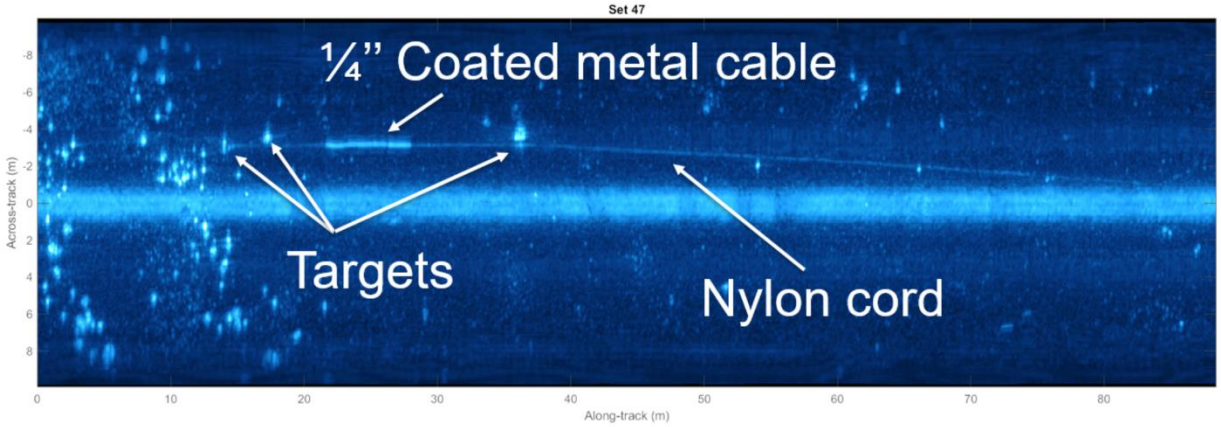

Figure 216: An image of the target and cable array deployed in Lake Washington for the 6/2020 INS integration test. Remarkably, in addition to the deployed targets, the nylon recovery rope is also visible.

Lake Washington PB4Y

In 1956 a PB4Y bomber crashed in Lake Washington. All crew were safely evacuated and a recovery attempt was attempted but resulted in failure [39]. The wreck currently sits at the bottom of Lake Washington and is a popular sonar image target, diving location, and was the subject of a large 3D photogrammetry imaging effort. The wreck was in the vicinity of the target array and a single pass over the plane captured at a pulse-repetition-rate of 40Hz was performed with the BOSS. Results were extremely good and the next several figures show some of the 3D images created from the wreck data. Due to the low frequencies and high resolution of the BOSS system many internal features of the fuselage are visible in addition to the gun nacelles and the actual gun barrels protruding from the rear turret.

A literature review (e.g. [1-5]) of BOSS-related research shows that the primary focus has been on classification and identification of UXO and other low-*ka* objects (e.g. mines). Most of the images created in this low-*ka* regime, unless enhanced via multi-aspect fusion, are blobs with occasionally recognizable features such as edges and specular glints. The scans of larger-scale objects like the PB4Y bomber are of interest because they illustrate that the BOSS system can

create very intuitive renderings of larger-scale objects that require little if any expertise or system familiarity to interpret.

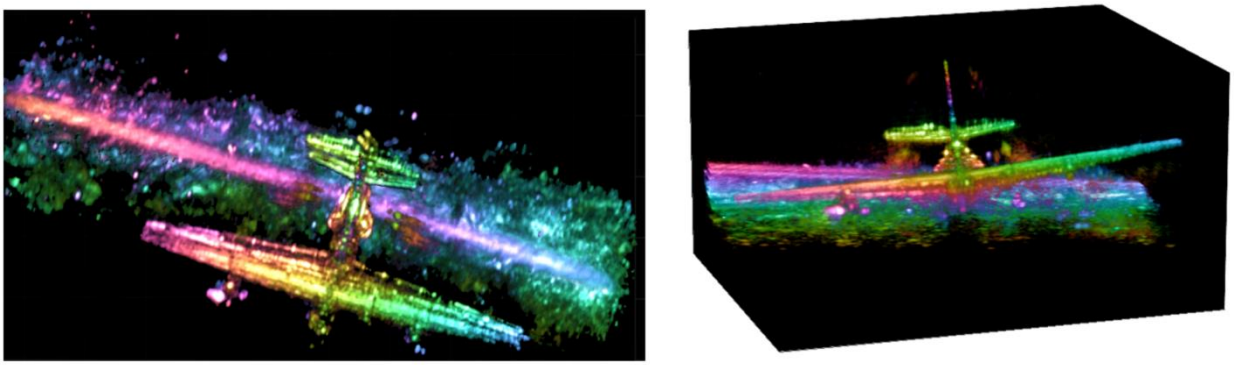

Figure 217: A top and 3D view of the Lake Washington PB4Y bomber rendered with alpha-blending and a color-by-depth colorization scheme. The path of the MuST went directly aft the port and starboard gun nacelles. The wings, tail, tail-fin, nacelles, and three of the engines (the fourth was lost during the failed salvage attempt) are readily visible.

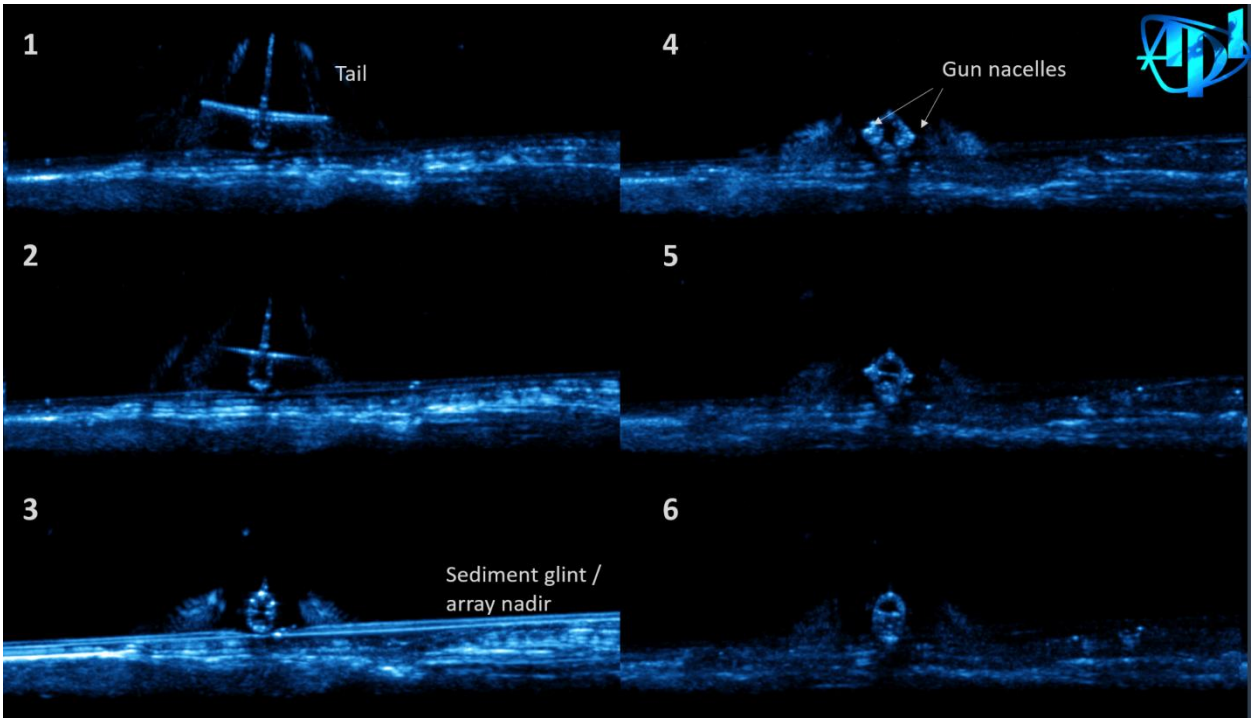

Figure 218: Cross sections of the plane in the along-track dimension moving from aft to fore, 1 – 6.

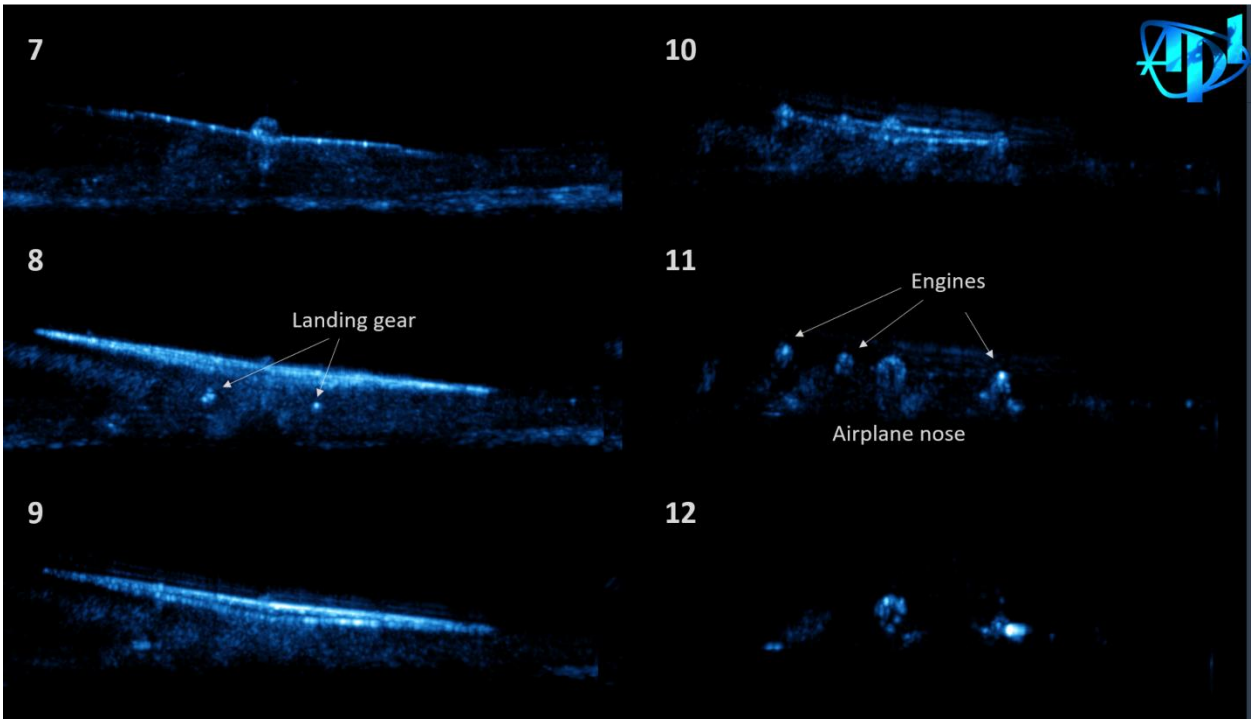

Figure 219: Cross-sections through the wing and forward. The cross-section of the wing and the landing gear beneath are visible. Most of the detail of the front of the plane is lost due to forward scattering.

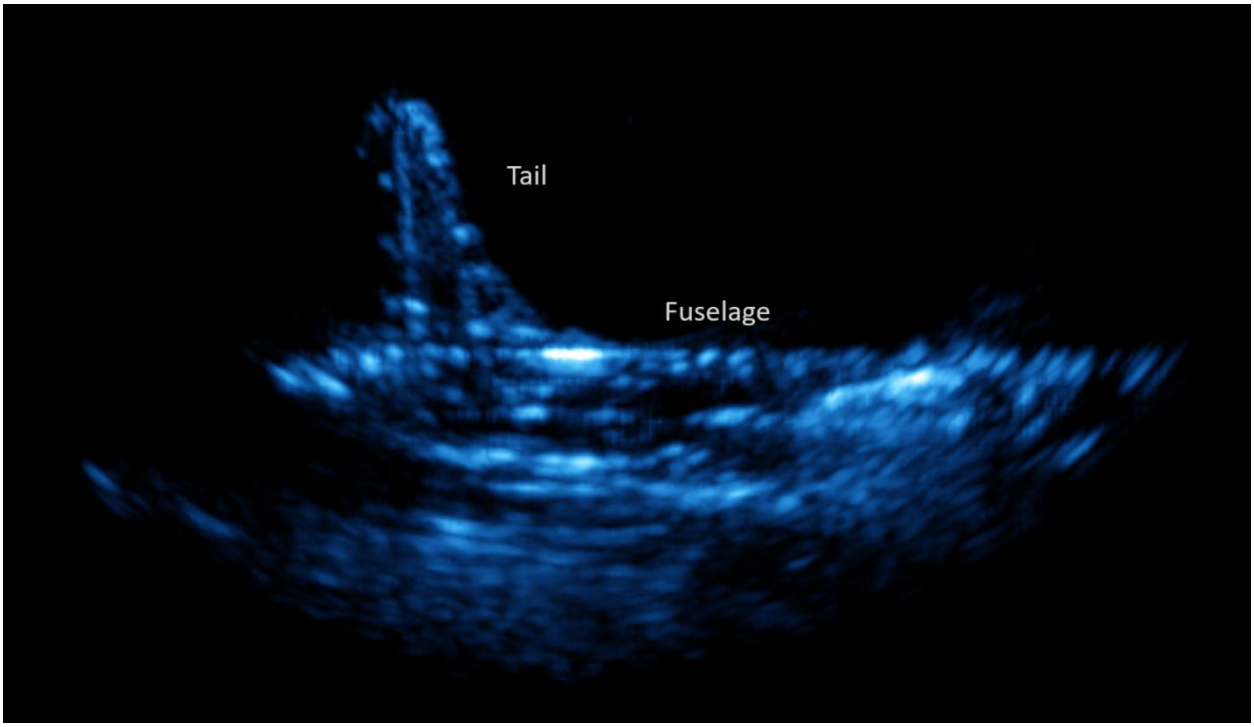

Figure 220: An across-track slice directly through the fuselage of the airplane. Resolution is considerably lower in this dimension yet many features of the plane are still very interpretable.

Range-alias imaging demonstration

The INS integration test was also the first test in which range-alias imaging was demonstrated, i.e. the ability to image data arriving from previous ping transmissions other than the current, extending the effective maximum range of the system. A high altitude scan of a shipwreck was used as beamforming verification because of the intuitive rendering produced by SAS processing. The system was pinging at 40 Hz and the target range was ~45 meters, placing it in "Zone 2" in the parlance of Fig. 86, which corresponds to beamforming echoes received from two pings previous to the current. Figure 221 shows an across-track and Z-MIP through the SAS reconstruction and several elevation slices showing deck features.

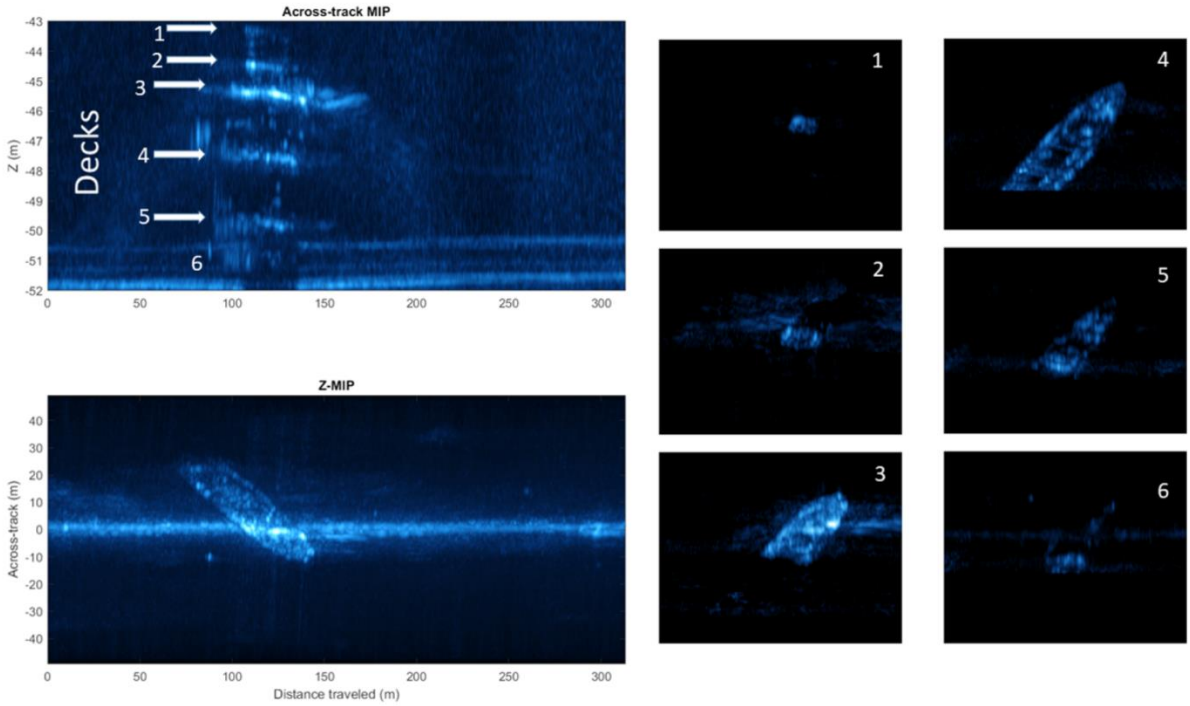

Figure 221: The first range-aliasing image created with the MuST BOSS sonar: echoes from the ship were in Zone 2 (see Fig. 86), i.e. 2 pings previous to the current. Several decks of the ship are reconstructed nicely: decks are annotated in the across-track MIP on the top and corresponding numbered horizontal slices of these levels are shown on the right.

VII. Field test results: 2020 Sequim Bay UXO trial testbed survey

Overview

In 2020, Pacific Northwest National Laboratories (PNNL) positioned 16 objects (both clutter and UXO) in unknown positions and in proud and buried states within a 100 meter by 100 meter testbed site. The testbed was located on a fine sand sediment in the northeast portion of Sequim Bay in Washington State.

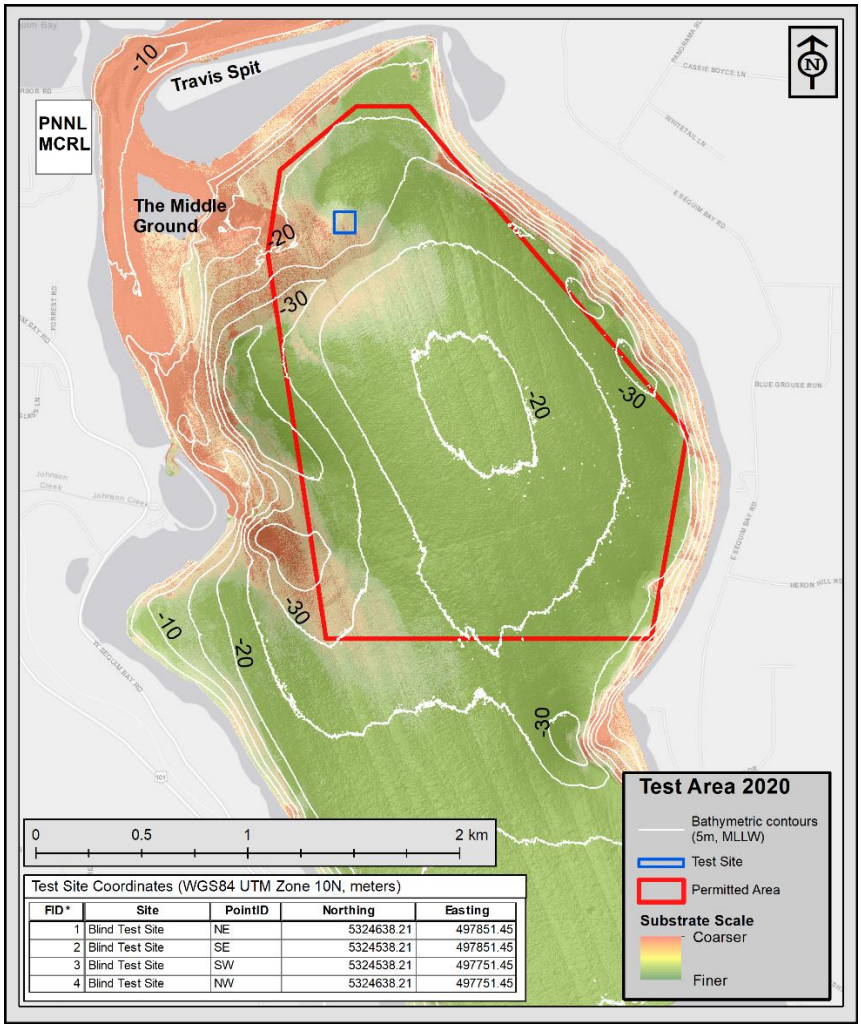

Figure 222: North Sequim Bay showing the sediment coarseness, bathymetry contours, and 100 x 100m box in which testbed targets have been arrayed.

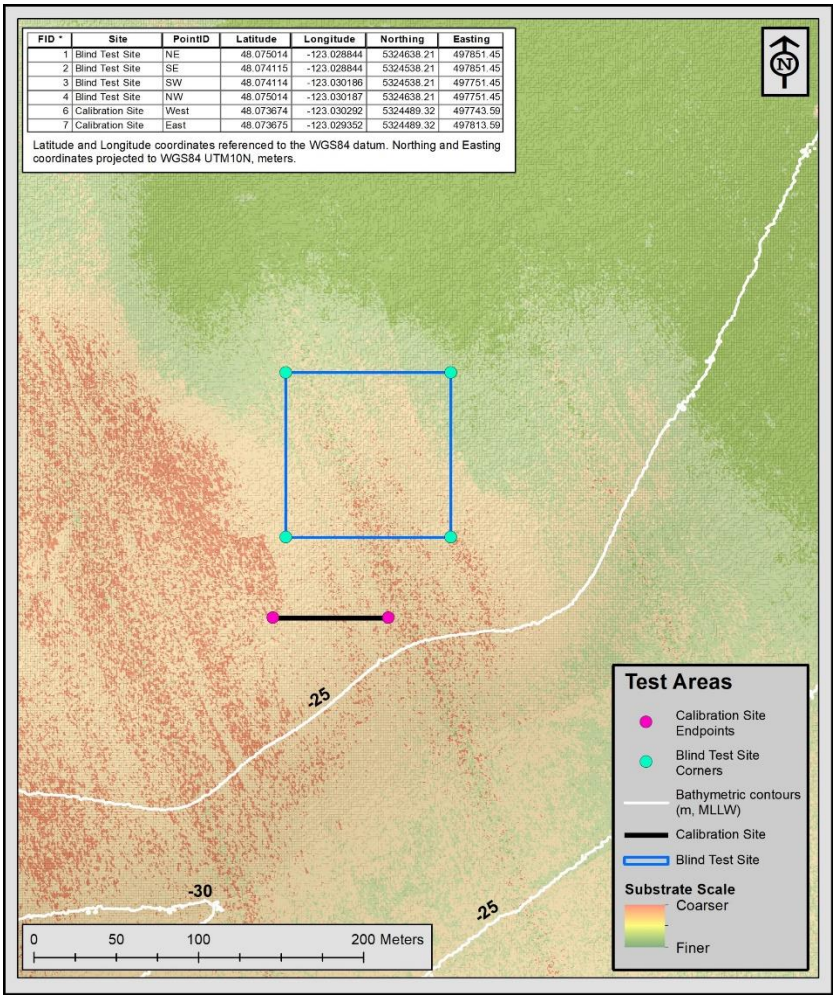

Figure 223: Zoomed in portion of Fig. 222 showing the testbed site and the location of the calibration line south of the target field which runs east-west. The sediment within the testbed region was dominantly fine sand, but it became muddier toward the north end of the box.

A list of the objects deployed in the testbed include science, UXO, and clutter objects:

Science targets

- 60 x 30 cm solid aluminum cylinder
- 60 x 30 cm hollow aluminum cylinder

Inert UXO and replicas

- 81 mm Mortar (M889A1) x 2
- 105 mm projectile (M80) x 2
- 105 mm projectile (HEAT) x 2
- 155 mm Howitzer projectile (M107) x 2
- 155 mm Howitzer projectile replica x 2

Clutter objects

- Anchor
- Cement block
- Scuba tank
- Crab trap

Photos of these targets are shown in Fig. 224.

Figure 224: All of the targets (excluding duplicates) deployed in the Sequim Bay testbed.

During the experiments, which lasted from 2020 Sept. 21 through Sept. 24, the Multi-Sensor Towbody surveyed the target field using the survey lines shown in Fig. 225. To increase the possibility of detecting broadside specular glints from targets of interest the towfish was scanned along three different heading angles oriented 60 degrees apart: 280, 340 and 40 degrees. Scan lines were spaced by 5 meters and arrayed to cover the extrema of the box and the calibration line.

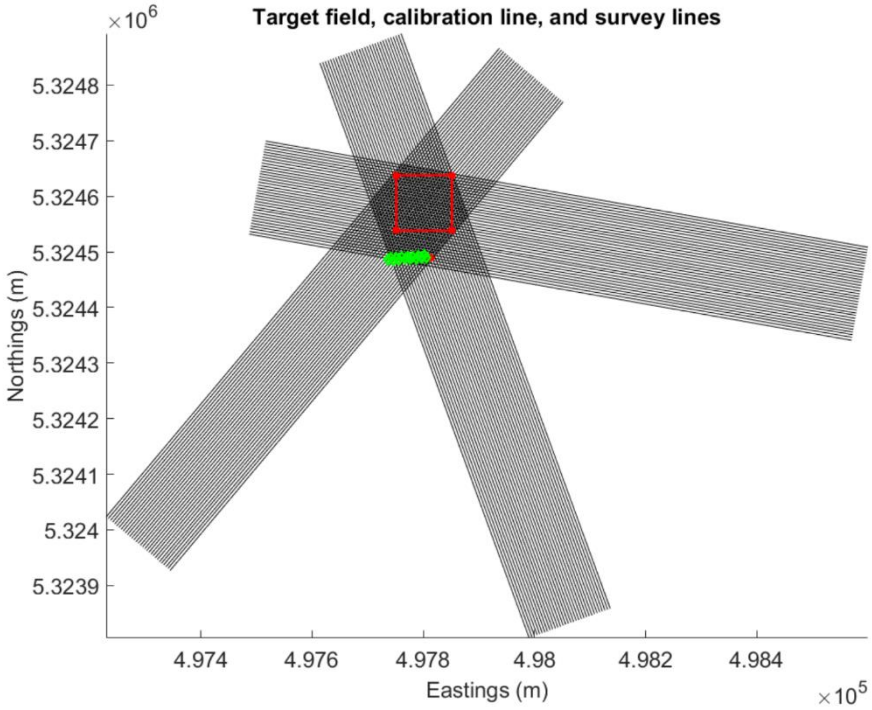

Figure 225: Target field (red square), calibration line (green points) and scan lines (black lines) used to survey the target field. The scan lines are along heading angles 280, 340, and 40 degrees.

The field test spanned four days, each of the first three days corresponding to one of the heading directions, and the final day was used to fill gaps from incompletions in the first three days. The INS integrated into the MuST in the previous Lake Washington trials was not yet a mature product and it suffered many problems during this experiment. For example, though it had a fiber-optic gyro (FOG) - a very sensitive rotation rate sensor - it experienced a constant ~20 degrees/hour heading drift that, even with a software patch, caused irregular and often poor position estimation performance vs. distance travelled. Furthermore, since the sensor was not north-seeking the heading estimate had to be calibrated each time the vehicle surfaced using the dual-GPS position sensor. This heading calibration seemed to have greater variance if the vehicle position was not stable due to wave action. Additionally, the drift vs. distance travelled seemed to increase rapidly if the vehicle moved in any manner other than a straight line.

To mitigate these problems the MuST operation team developed a strategy for minimizing position error in the navigation solution: the MuST would transit in a straight line on the surface to provide sufficient time to get a decent heading calibration, then dive at the last possible moment necessary to get a scan. Once the target area had been transited, the towfish would immediately resurface to get a GPS fix. This “minimal time down” approach decreased the error in the position estimate by reducing the amount of time the estimate could drift underwater. Following the conclusion of the field tests, the navigation data was re-processed using the renavigation algorithm described in *Materials and Methods: Renavigation*. Figure 226 shows the

total set of scans (with trajectories re-calculated via renavigation), conducted over the course of the experiment.

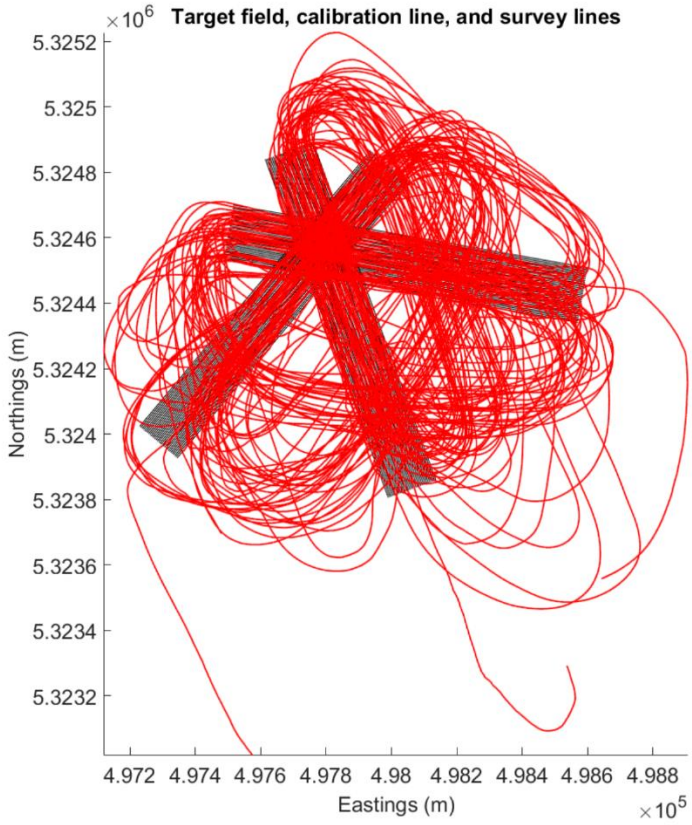

Figure 226: Towfish locations for all four days of field testing superimposed on the intended scanlines shown in Fig. 225.

The results shown in subsequent sections were gathered from the data captured during this experiment.

Renavigation results

Overview

This was accompanied by the first systematic development and testing of the feature-based renavigation algorithm described in *Materials and Methods: Renavigation* and this section will cover quantified results of the improvements resulting from the INS renavigation and feature-based renavigation stages. Two levels of renavigation have been developed (e.g. INS & feature-based) and therefore three navigation data-products are available for the given navigation data processing pipeline: raw (original), re-navigated, and feature-based re-navigated solutions. Ideally the level of accuracy should improve both from a precision standpoint (i.e. local variance) and from an absolute-accuracy standpoint with each stage of navigation refinement. The following sub-sections will first quantitatively compare the precision of the navigation methods then compare mosaics generated via each method for the Fall 2020 Sequim Bay field test with the ground-truth estimates of targets placed along the calibration line.

Quantitative precision comparison

The precision of the different levels of navigation refinement can be determined via examining the observed shifts between overlapping scan regions following image gridding. Figure 227 shows a scatter plot of the *x*, *y*, and absolute displacements vs. quality for the different navigation refinement levels.

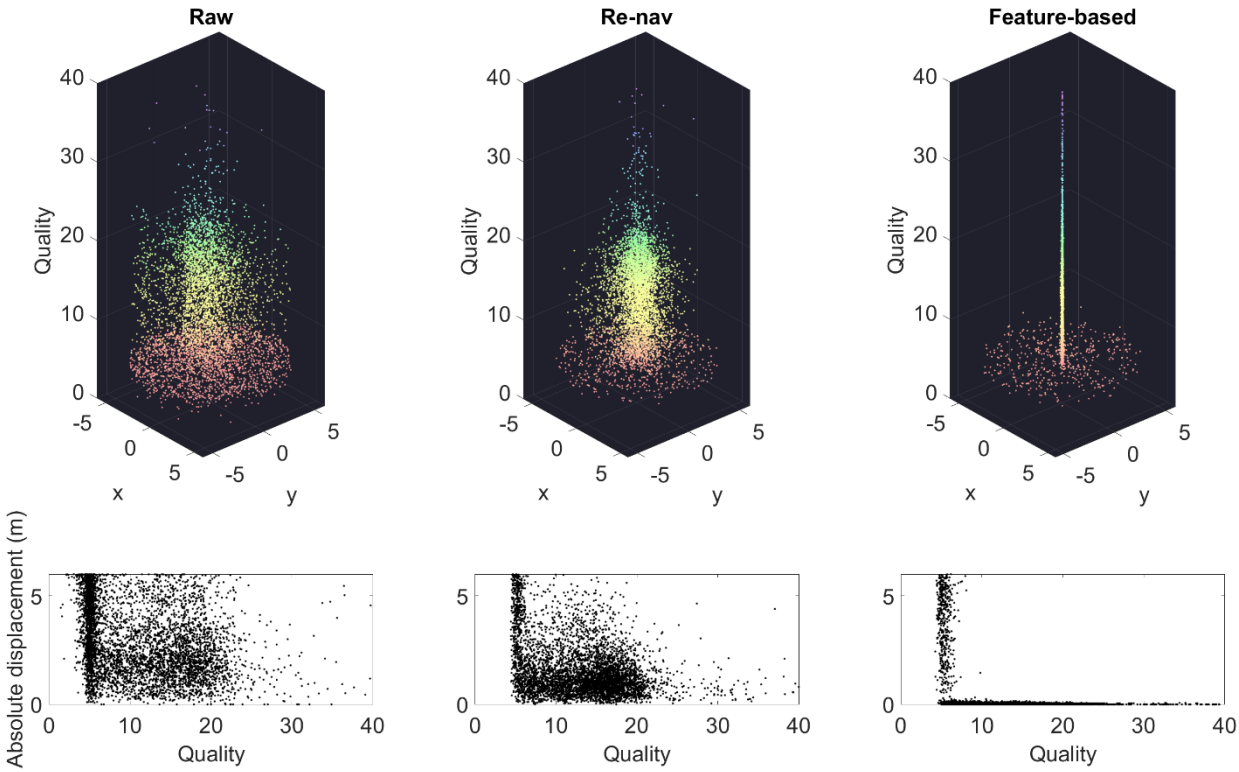

Figure 227: Alignment scatter plots showing measured *x* and *y* displacements vs. quality (top), and absolute displacements vs. quality (bottom). On left are the scatter plots for the raw navigation data, middle is after first-stage re-navigation and the figures on the right show the results after feature-based re-navigation.

From Fig. 227 it can be observed that the spread of displacements decreases considerably at each stage of re-navigation. It can also be seen that the number of high-quality displacement measurements is much lower for the raw navigation solution. This is possibly due to the fact that many scan pairs may have been mismatched due to inaccuracy in the estimate of overlapping regions. Furthermore, the correlation peak search radius was limited to 6 meters, potentially excluding many larger displacement values from the raw measurements.

Histograms of the absolute displacements are compared for the different solutions using displacements with quality values > 7 in Fig. 228.

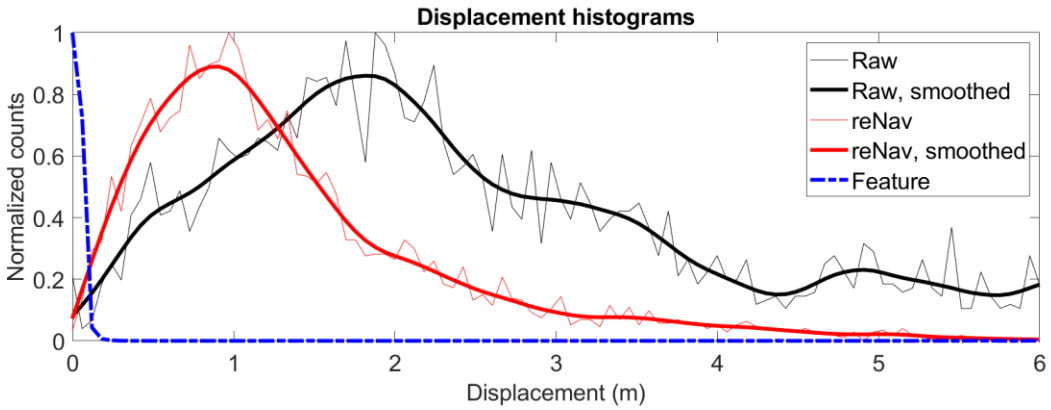

Figure 228: Alignment scatter plots showing measured *x* and *y* displacements vs. quality (top), and absolute displacements vs. quality (bottom). On left are the scatter plots for the raw navigation data, middle is after first-stage re-navigation and the figures on the right show the results after feature-based re-navigation.

The absolute displacement histograms illustrate some interesting properties of the different navigation solutions. For the raw navigation solution, the spread of the displacements is clearly not completely contained in the sampled region, and thus all subsequent efforts at determining displacement statistics will be skewed low. The peak of the RAW distribution is at approximately 2 meters, however the number of counts at the edge of the distribution is high enough to indicate a significant number of uncaptured displacements with very large values.

In contrast, the dominant portion of the navigation displacements measured for the re-navigated solution (red curve, Fig. 228) lie within the sampled region, and the measured statistics will be accurate. The peak of the histogram lies around 1 meter with instances decreasing rapidly above that level. Significantly, the number of very large outliers in this re-navigation solution is much lower than in the raw data.

Lastly, due its data-driven nature, the spread of the feature-based solution is extremely tight, with the peak displacement lying directly on zero due to the discretization of the displacement estimates. Table 4 is a table of the RMS, median-absolute-error (MAE), the radial threshold containing 95% of the samples, and the percentage of sample points with a quality value above the threshold using $Q_{threshold} = 7$.

	RMS (m)	MAE (m)	Radial 95% (m)	Quality %
Raw	2.81+	2.11+	5.3+	62
Re-navigation	1.70	1.10	3.45	83
Feature-based	0.22	0.00	0.06	85

Table 4: Distribution statistics for image alignment using the navigation solution with the three stages of refinement. For raw navigation data with no refinement, the statistics in the top row will all be biased low due to the lack of adequate sampling of the actual displacement distribution. The MAE value of zero for feature-based navigation is due to discretization of the displacement measurements (5cm resolution), resulting in a large number of 0m displacement measurements.

Though Table 4 indicates vastly improved performance for feature based re-navigation it is important to note that this type of navigation solution also has the most contingencies: it relies on accurate scan pair association and the statistics of the sediment and a distribution of enough significant scatterers on the surface to allow for good correlation. In contrast, the first stage of re-navigation, which merely reprocesses the raw data with different assumptions, also shows a major improvement vs. the raw navigation solution, but it is a simpler, faster, and potentially more reliable solution to achieve due to having fewer dependencies.

Mosaics and navigation accuracy

In contrast to alignment precision, absolute accuracy is more difficult to assess because it requires a set of reference points with known location in geo-rectified coordinates for comparison. In the fall 2020 Sequim Bay field experiments, a series of targets were placed along a calibration line by Pacific Northwest National Laboratories (PNNL) and their positions were estimated by determining the location of fixtures placed at either end of the line via a USBL transponder system. The USBL transceiver position was determined via GPS and lever-arms, and the end-points of the calibration line were determined via positioning relative to the USBL. A line was stretched between the end-points and divers positioned targets relative to the reference line between end-points. This method gives targets a position in absolute coordinates, however the associated position uncertainties were not supplied. The following mosaics have super-imposed the nominal positions of the calibration targets as provided by PNNL, however there is an unknown uncertainty in the actual location of the coordinates.

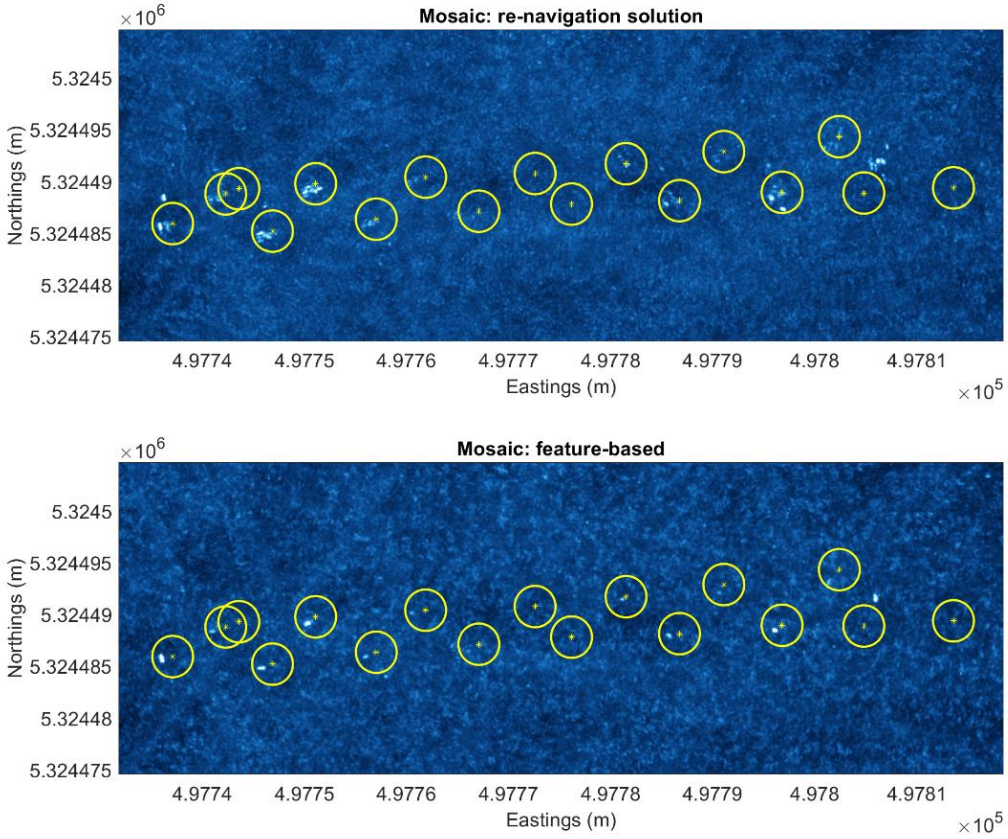

Figure 229: Gridded scan mosaics generated using the first re-navigation solution generated by re-processing raw navigation data (top) and generated using the feature-based navigation solution (bottom). The nominal locations of the targets in the calibration field are marked as yellow stars, and a 2 meter radius is drawn around their nominal locations.

Several important observations can be made from Fig. 229. The first renavigation stage (i.e. INS reprocessing) forms a fairly tight spread in terms of precision with the clusters typically being around one meter wide. Furthermore, most of the visible target instances lie within the 2 meter radial distance from the nominal locations of the calibration targets. The advantage of data-driven mosaicking is clearly visible in the bottom figure, however, in which the actual orientation of many targets can be determined due to the high level of focus. It can also be seen from Fig. 229 that there seems to be a systematic offset, or bias, between the PNNL locations and the mosaic solution: most of the targets in the mosaic appear to be about 1 meter to the south-west of their nominal locations. It also seems that at least one of the PNNL targets has shifted from their original location: the target at the extreme eastern end of the calibration line has moved about 8 meters west and 3 meters north from its original location. Lastly, the data-driven solution also focuses many features on the sediment which may have other implications for environment classification.

SAS post-processing results

Overview

In this section results of using the SAS post-processor engine described in *Materials and Methods: SAS post-processing* will be described. The results of modifying various settings will be reviewed along with computation times for different configurations, the effect of the de-glinter, and the dataproducts stored to the hard-drive.

Path-following SAS image volumes

As stated in *Materials and Methods: SAS post-processing*, the fast 3D SAS beamformer generates imagery in geo-rectified coordinates, meaning that all points in the beamformed volume were defined in UTM coordinates prior to backprojection. The image coordinates are not, however, rasterized (i.e. projected onto an equally spaced 3D gridded array in UTM coordinates) but are defined on a set of vertical planes that are rotated based on towfish heading. The altitude and central UTM northing and easting values for each plane are defined by the path of the towfish. The storage requirements for these path-following volumes tend to be much lower than for data beamformed onto voxels evenly gridded along Eastings, Northings and *Z*. In 2019 the SAS volumes were uniformly gridded in UTM, however in 2020 the path-following approach became standard for SAS post-processing with the MuST Toolbox. Figure 230 gives an example of the 3 orthogonal MIPs generated for one such path-following SAS volume as the sensor transited a buried sediment transition region in Sequim Bay. This volume rendering uses a ping queue size of 160 pings, an integration angle of +/-35 degrees, spatially-variant-apodization for sidelobe reduction, pre-whitening for spectral normalization and, significantly, de-glinting.

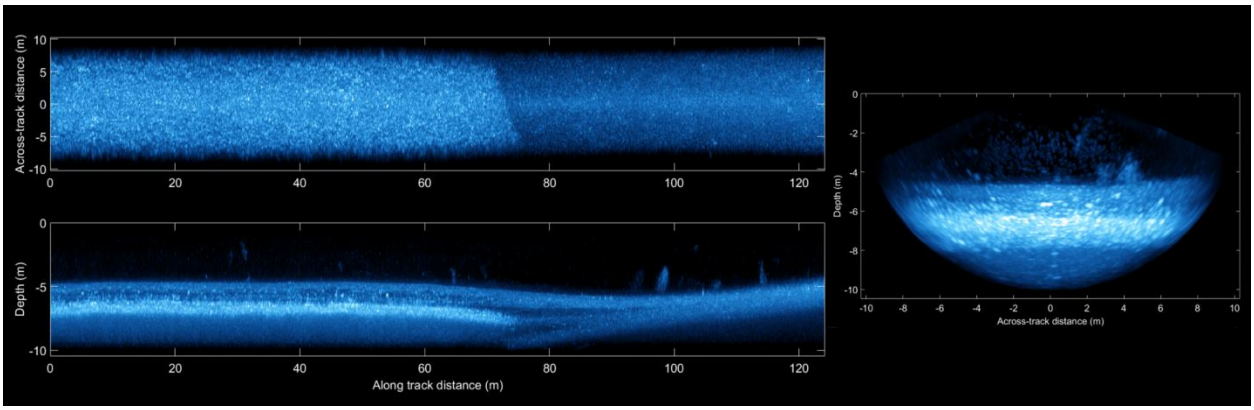

Figure 230: 3-direction orthogonal MIP views for a SAS volume generated using the path-following volume geometry. The region shown here is east of the target field and bridges a very distinct layer transition several meters below the surface layer. Several clouds of fish can also be seen floating above the sediment.

The next several figures will compare the results in Fig. 230 with the results created using different beamformer settings.

De-glinting performance and beamformer setting comparisons

Fig. 231 shows depth and across-track MIP's comparing the standard beamformer configuration of Table 1 but with and without de-glinting performed.

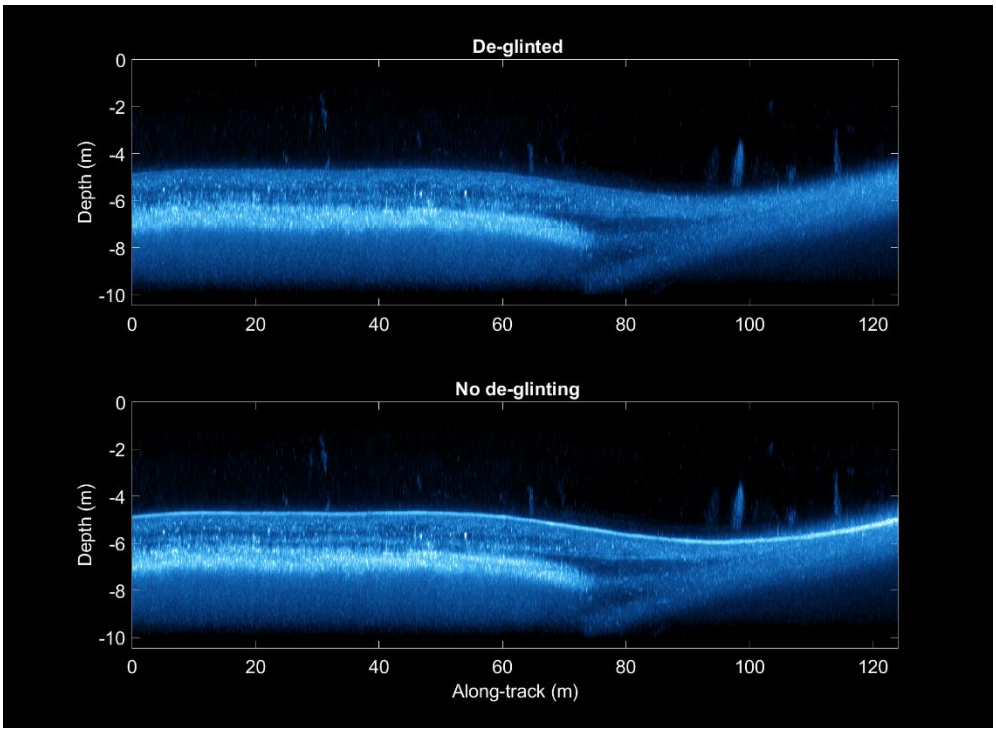

Figure 231: Profile-MIP for sediment transition data with (top) and without (bottom) application of the de-glinting algorithm.

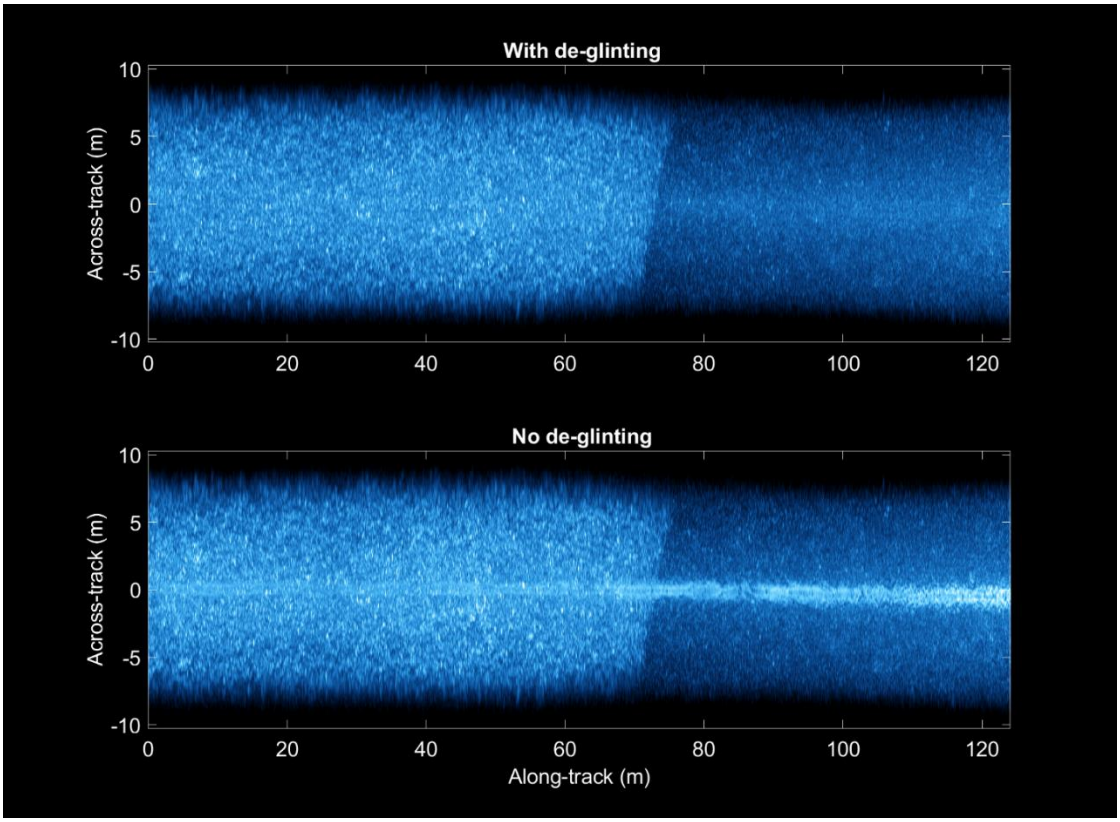

Figure 232: Z-MIP through the volume depicted in Fig.'s 231 and 232 with (top) and without (bottom) the de-glinting algorithm.

From Fig.'s 231 and 232 the influence of a specular glint on the imagery is clearly visible as a bright strip near nadir that occurs directly at the water-sediment interface. This is highly visible in the ordinary dataset with no de-glinting applied. It is important to determine whether or not the de-glinter removes the effect of targets in the near-nadir range. This would be a significant drawback to using the algorithm. Several near-nadir targets obscured by the specular glint were found in various datasets and the *before* and *after* de-glinting results are shown in Fig. 233.

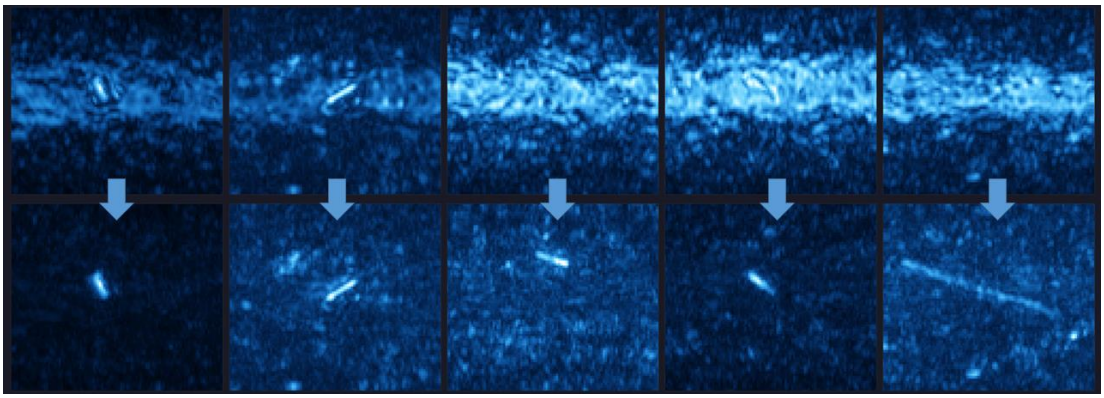

Figure 233: Various targets at nadir viewed with and without de-glinting. The target on the far right is the only buried target of the set, and is ~0.3m below the surface. It is completely obscured in the MIP image without the use of the de-glinting algorithm.

It can be seen from the examples in Figure 233 that the de-glinter greatly improves the detectability of targets that may be partially or completely obscured by the specular glint. Several of the targets (1, 3 & 4, numbered from left) are actually flush with the surface and though the de-glinting process is applied directly to the same voxels their prominence is increased rather than diminished.

Figures 234 and 235 show the effect of reducing the queue of pings used for synthetic aperture processing. In this case the ping queue size was reduced from 160 pings to 64 pings. The reduction to 64 pings decreases the computational burden but imposes a closer range at which the integration angle is queue-limited vs. ping-limited.

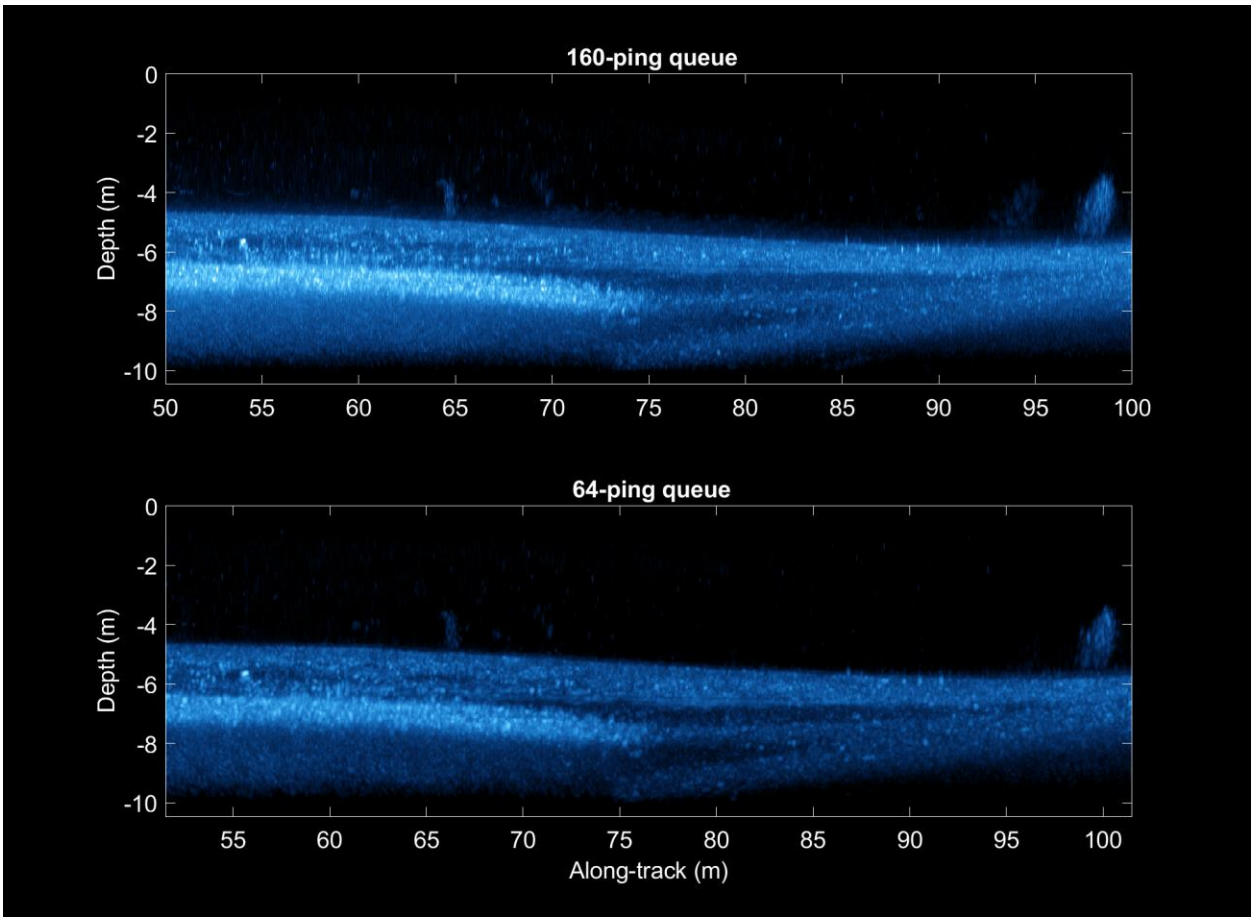

Figure 234: MIP profile comparing the default beamforming results with a 160 ping queue and a reduced 64-ping queue.

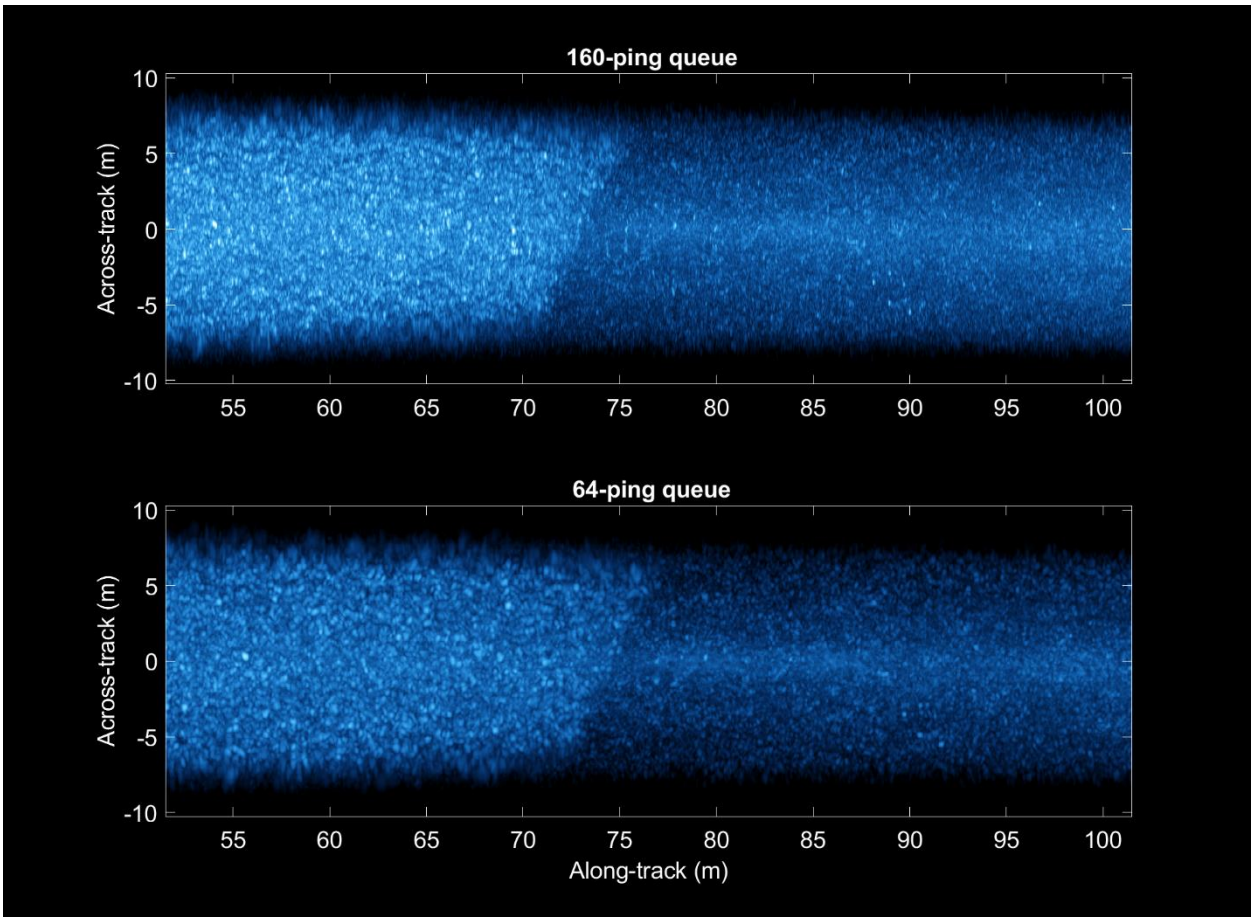

Figure 235: Depth MIP comparing the default beamforming results with a 160 ping queue and a reduced 64-ping queue.

The poorer resolution of the lower images in Fig.'s 234 and 235 is a result of the reduction in ping queue size from 160 pings to 64 pings. For a 35-degree half-angle for integration the beamforming becomes queue-limited at only 1.6m from the sonar.

Figures 236 and 237 show the result of a reduction of the total integration angle from 70 degrees to 10 degrees, a ping-queue size of 32 elements, and the de-glinting and spatially-variant apodization operations disabled. This "fast" beamforming configuration can be completed, from start to finish, in approximately the same amount of time as the data-collection is completed. The results are obviously inferior to the results with larger queue sizes, however it is still capable of performing object detection.

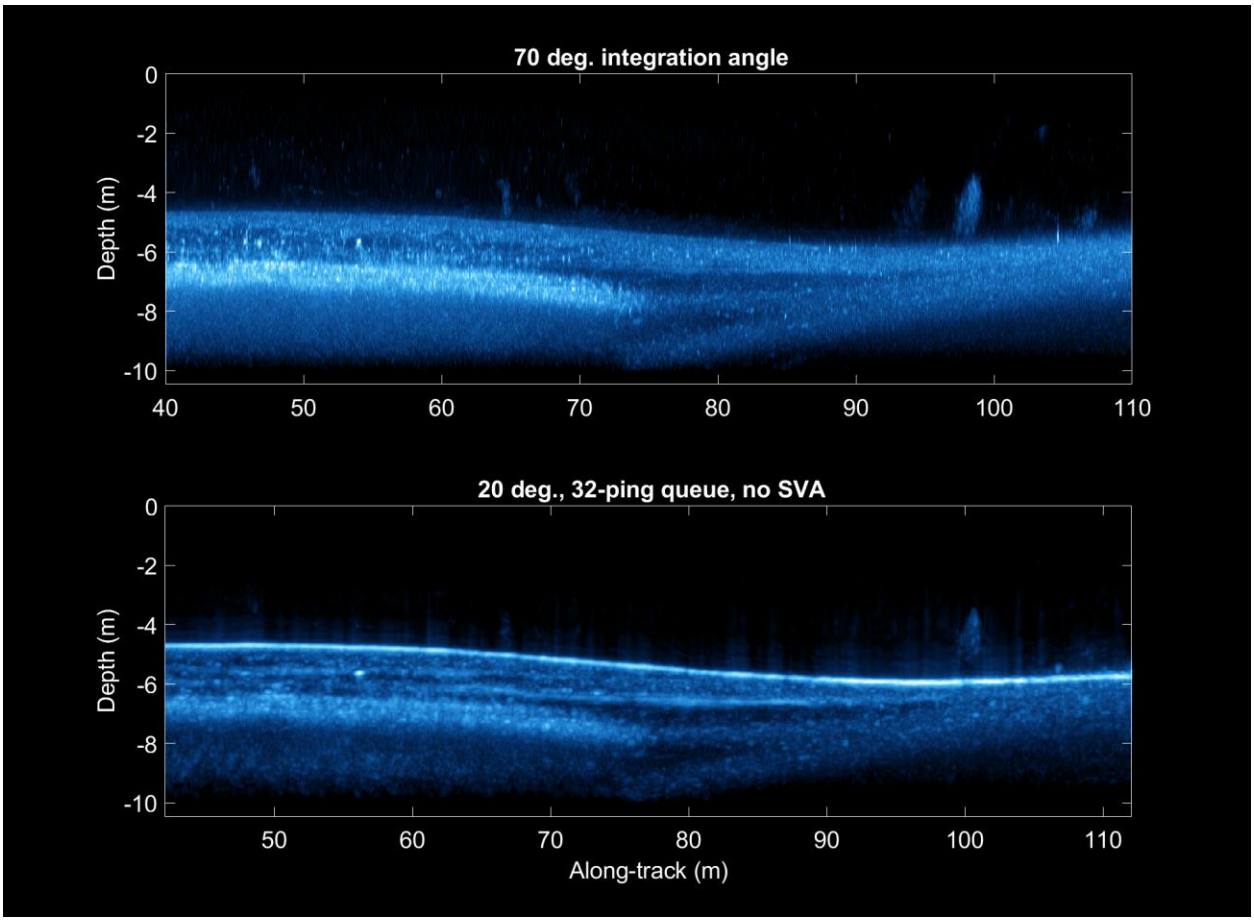

Figure 236: A comparison of the beamforming results with the standard configuration on the top and a fast configuration with a small queue size, integration angle, and no de-glinting or SVA sidelobe reduction is shown on bottom.

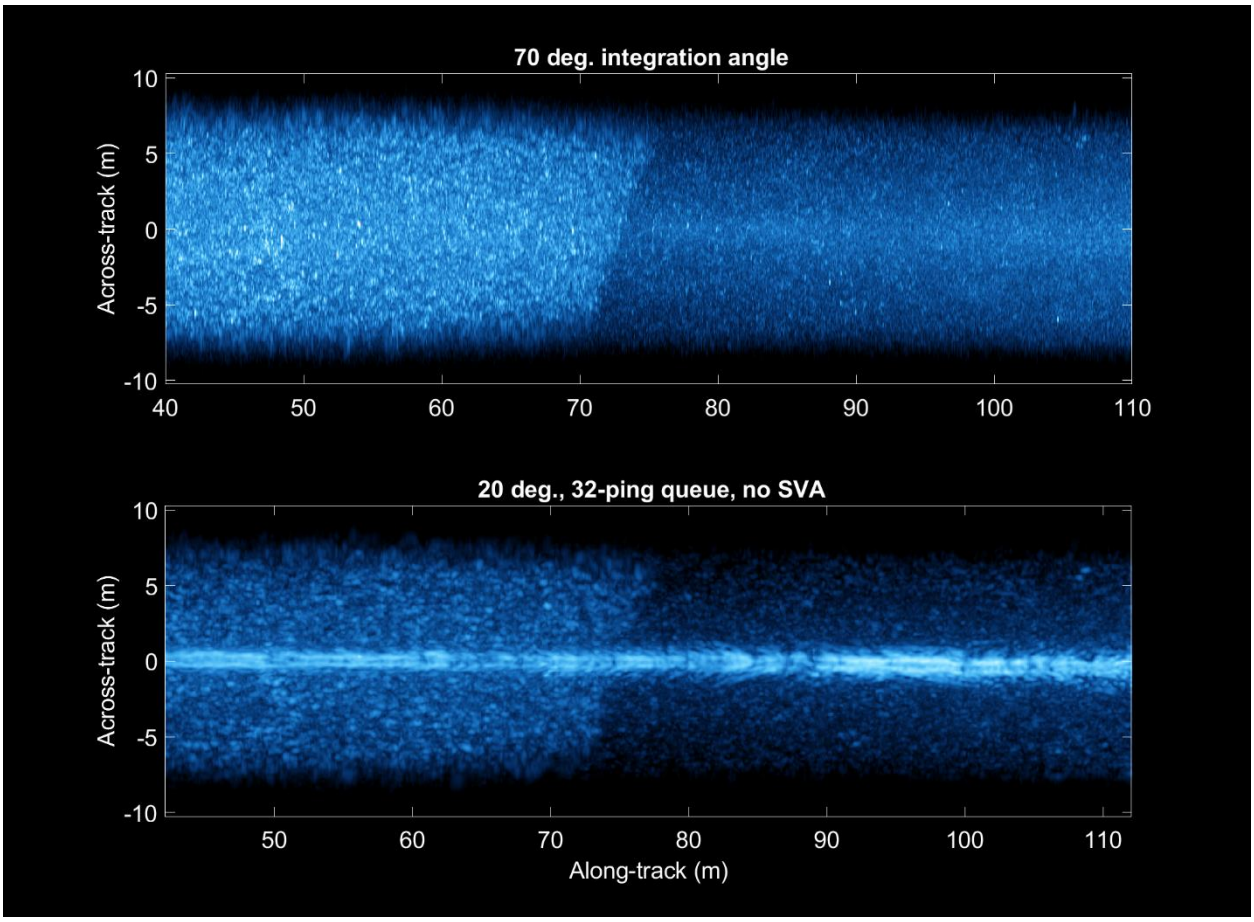

Figure 237: the Depth-MIP version of Fig. 236, showing comparing the standard configuration beamforming results (top) with the fast configuration (bottom).

An interesting feature of the fast-beamforming results appears to be that layers become more visible vs. the standard beamforming results. There are a number of possible reasons for this: by reducing the integration length the relative response of isotropic scatterers will greatly decrease relative to glint-like components resulting from acoustic impedance discontinuities. Alternatively, the SNR of some features will be enhanced due to the fact that aliasing is highly controlled by the tight integration angle.

Layer export

At the Sequim Bay 2020 experiment the batch-processing approach began to leverage layer-MIPs as a method for data storage. Figure 238 shows an example of a set of layers captured for a particular scan. A target is especially prominent in the layer centered ~40cm above the sediment, possibly the solid or hollow 30cm diameter 2-1 cylindrical aluminum science target.

Figure 238: Example MIP layer data stored for a given target with layer centers ranging from almost 90cm above the surface (top) to 90 cm. below the surface (bottom). This figure shows every other layer, 15 layers are stored in the current beamformer configuration.

Mosaics

The mosaicking tools described and leveraged in *Materials and Methods: Data gridding*, and *Materials and Methods: Feature-aligned mosaic generation* as well as the PMA mosaic analysis tool were only in the early stages of development during the post-processing phase of the Sequim Bay 2020 experiment, however many of the approaches taken were eventually adopted into the PMA tool. As described in *Materials and Methods: Data gridding*, layer data of the type depicted in Fig. 238 was loaded, gridded in UTM coordinates, and combined. The navigation

solution used for gridding was the feature-based re-navigation solution. The next several figures depict the results of that combination for the entire set of data spanning all four days of the Sequim Bay experiment, and corresponding to different MIP depths, using the *max* operation. The contra-harmonic mean (i.e. Lehmer mean with an exponent of 2) was not yet discovered as an improvement over the MAX operator for this mosaicking operation.

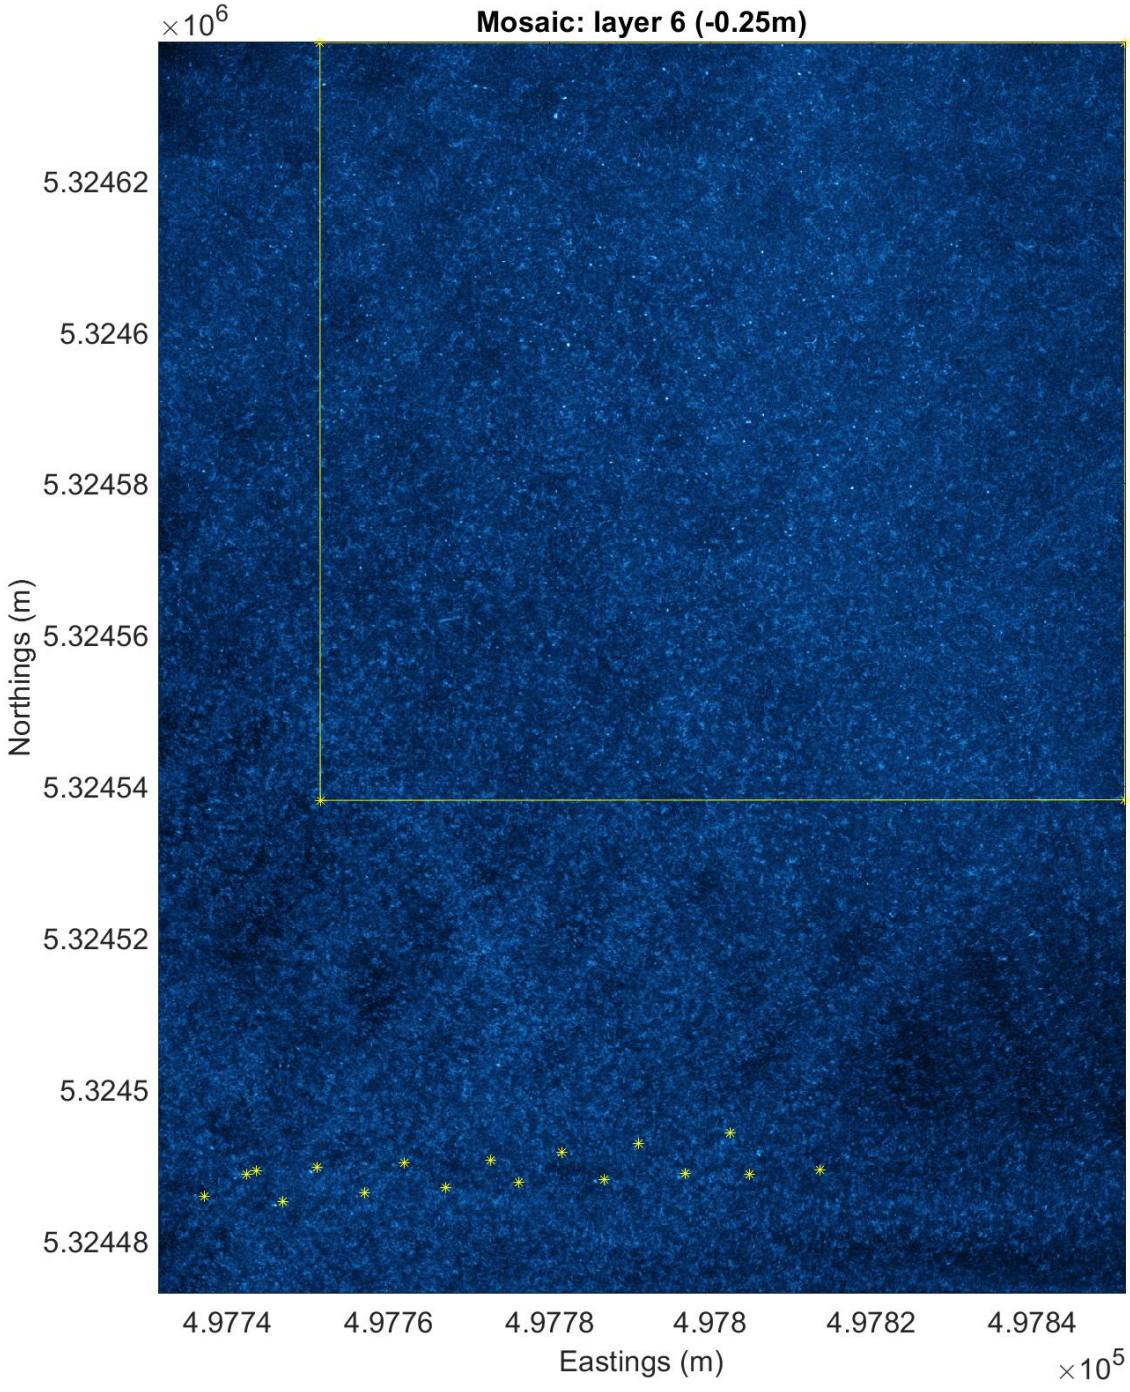

Figure 239: Sequim bay data mosaic, layer 6 (-0.25 m relative to sediment surface). The target area is outlined in the yellow box in the north-east corner of the image. The UTM coordinates of the targets in the calibration line are depicted with yellow stars and arrayed along the southern end of the mosaic.

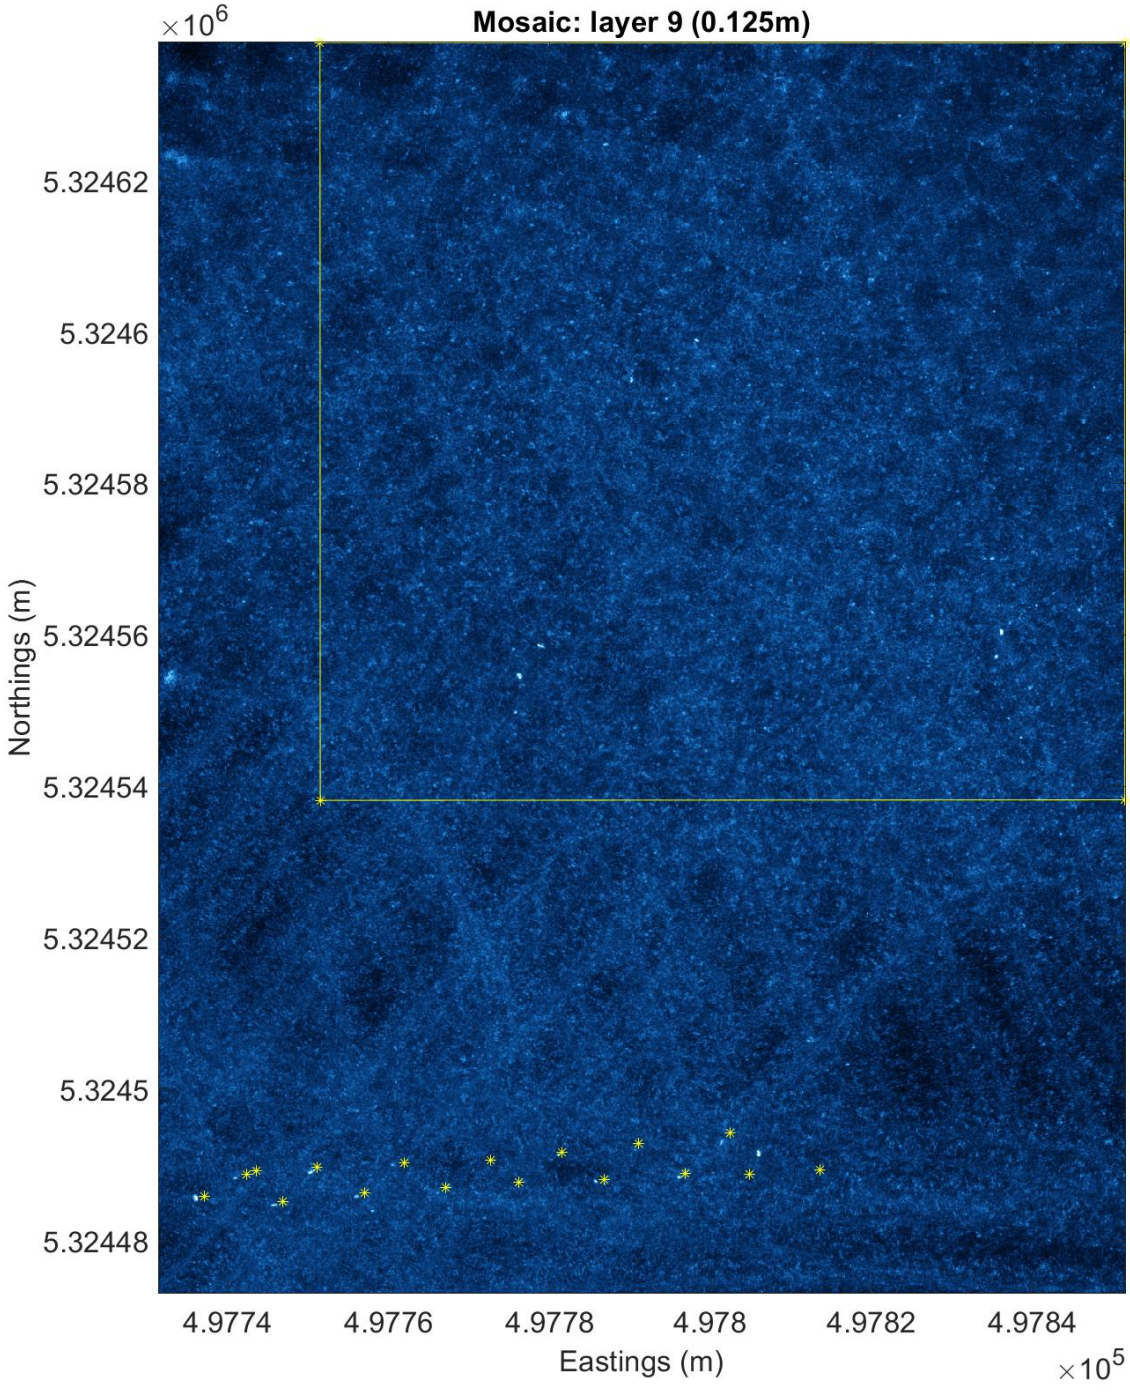

Figure 240: Sequim bay data mosaic, layer 9 (+0.125 m relative to sediment surface). The target area is outlined in the yellow box in the north-east corner of the image. The UTM coordinates of the targets in the calibration line are depicted with yellow stars and arrayed along the southern end of the mosaic.

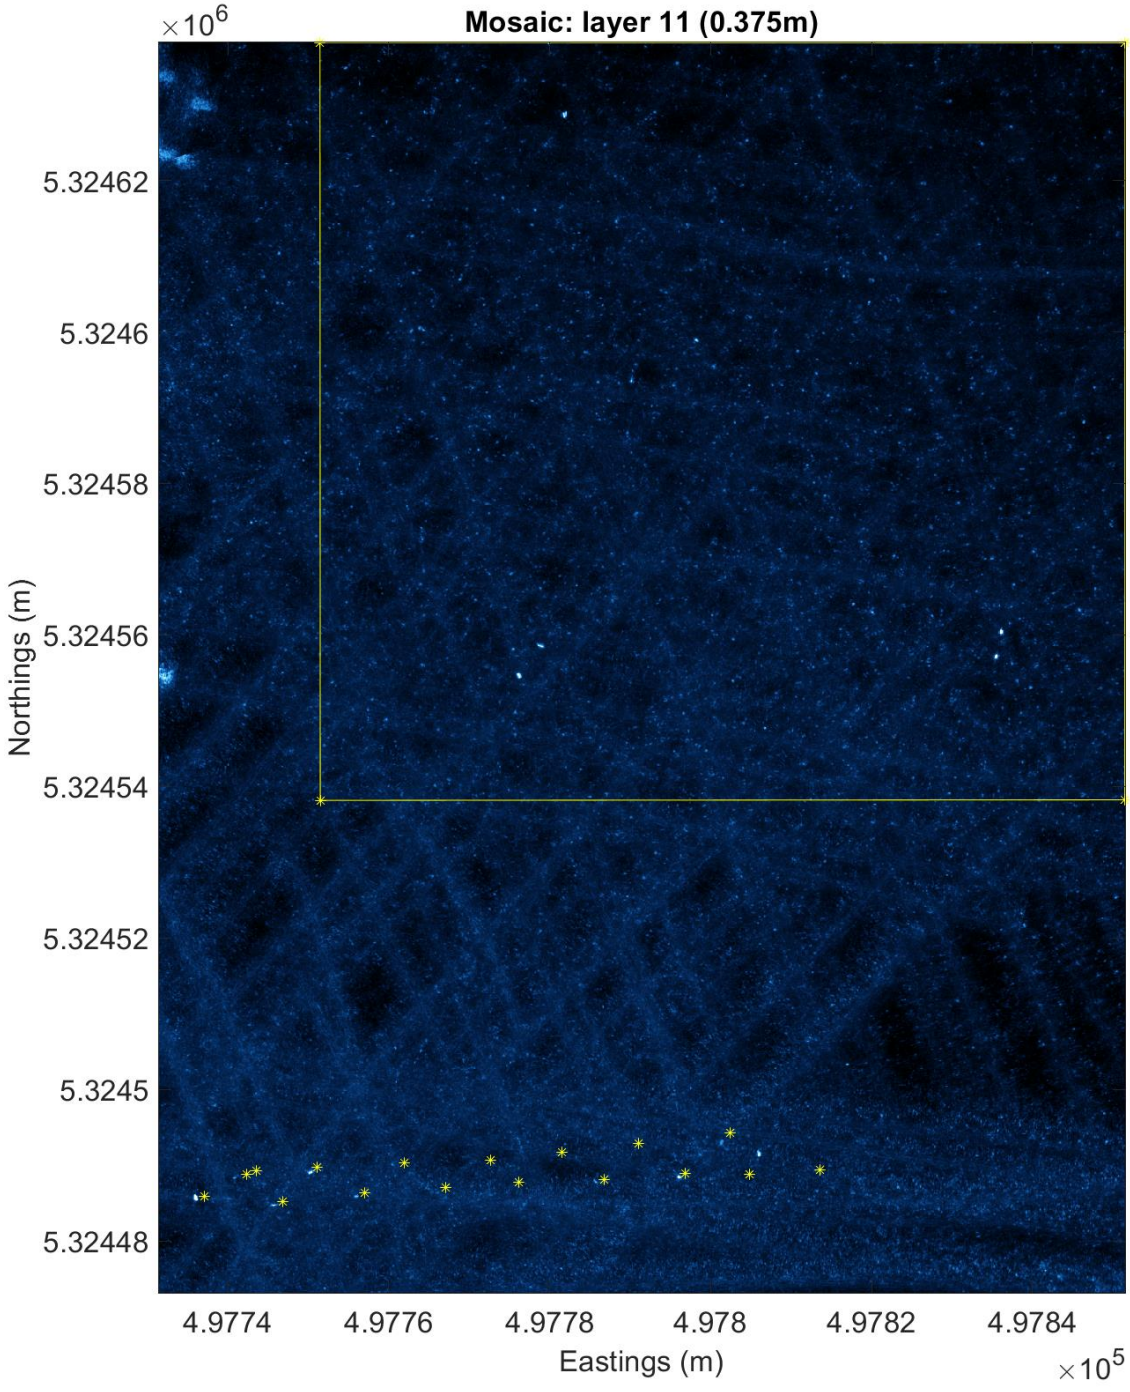

Figure 241: Sequim bay data mosaic, layer 11 (0.375 m). The target area is outlined in the yellow box in the north-east corner of the image. The UTM coordinates of the targets in the calibration line are depicted with yellow stars and arrayed along the southern end of the mosaic.

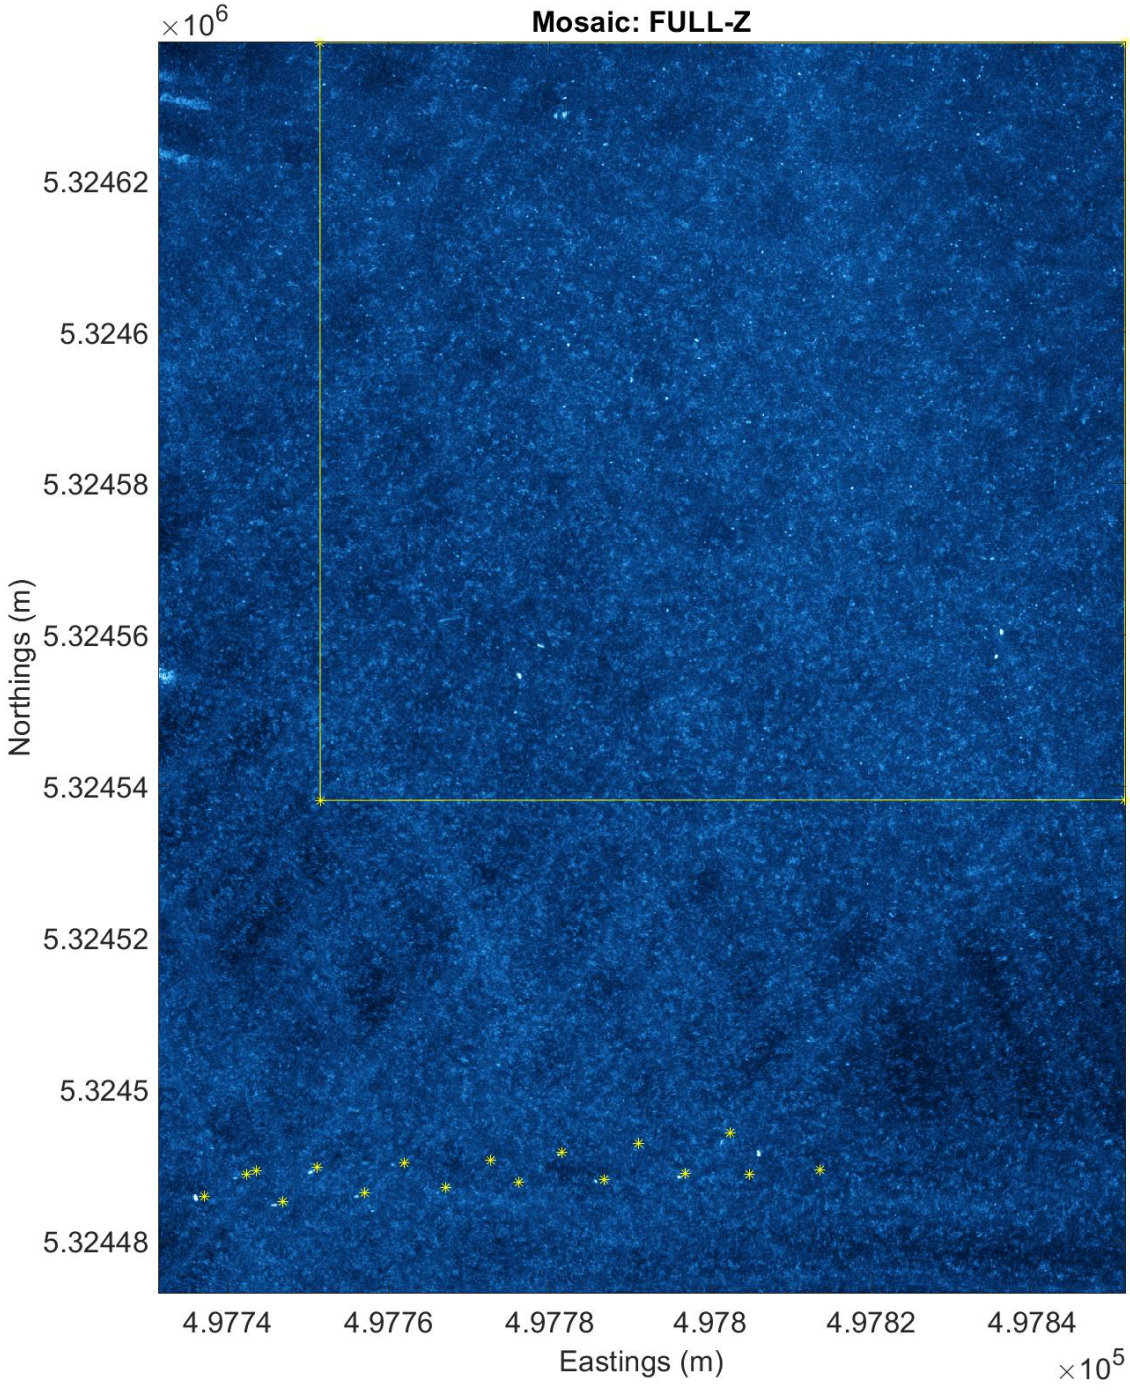

Figure 242: Sequim bay data mosaic generated using full Z-MIP images. The target area is outlined in the yellow box in the north-east corner of the image. The UTM coordinates of the targets in the calibration line are depicted with yellow stars and arrayed along the southern end of the mosaic.

A cursory review of Figures 239-242 shows that each layer has many different features and show different potential UXO detections. Most of the targets are not visible in layer 6, which is below the surface, however there is clearly a target that is visible in the north-west corner of the box that is invisible in the layer above. Additionally, near the south-west corner of the box a long slender buried object is visible. This is the same object that was shown in in the far-right plot of Fig. 233.

ROI patch processing

As in the current PMA framework, the post-processing tools used for Sequim Bay 2020 allowed regions-of-interest detected in the mosaics to be extracted and the data viewed using many of the ROI tools that were eventually incorporated to the current PMA tool. Examples include the look-data display window (*Materials and Methods: ROI review tool*, Fig. 113) and the 3D set-view tool (*Materials and Methods: ROI review tool*, Fig. 114). Fig.'s 243 – 250 show look-data displays for eight prominent targets that were scanned in the testbed.

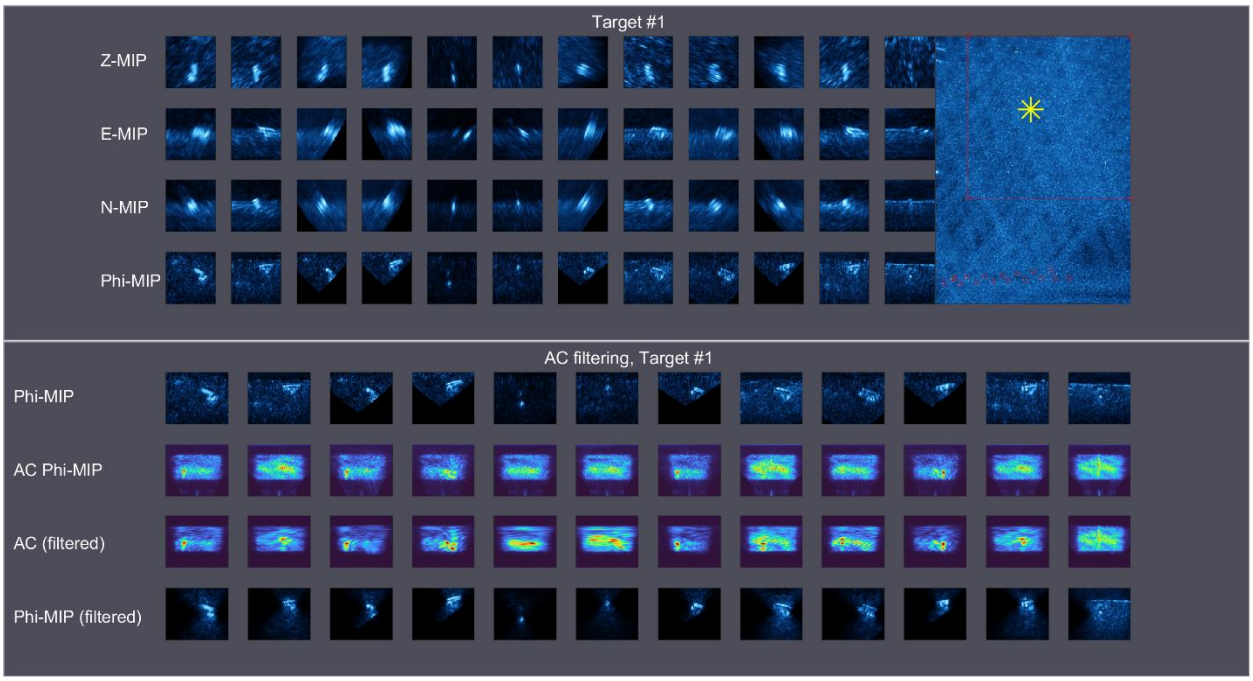

Figure 243: Detection set for a flush or mostly buried UXO in the target field showing different MIP's (Z, East & North) through the UTM gridded patch cube, a *phi*-MIP through the cylindrical patch, and the location of the detection in the target field (yellow star, upper right map). The set on bottom shows the 3D acoustic-color MIP's through *phi* and the accompanying spatial MIP's pre- and post-filtering.

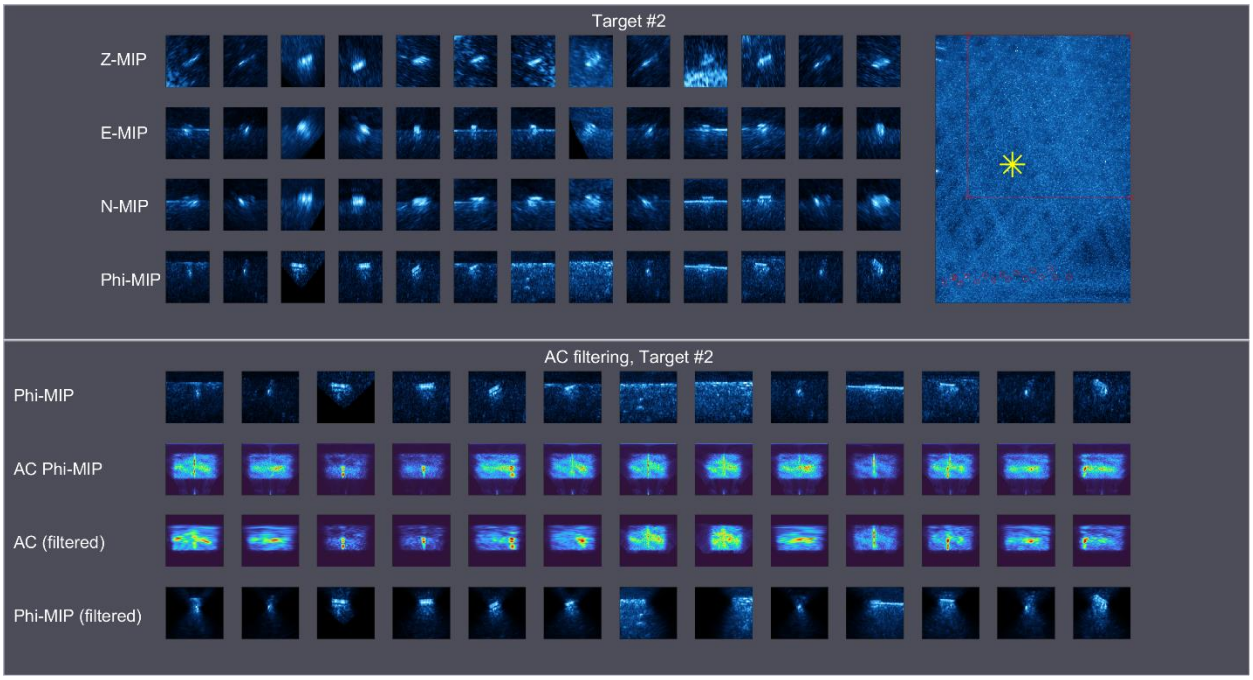

Figure 244: Detection set for the 30cm x 60cm solid aluminum slug.

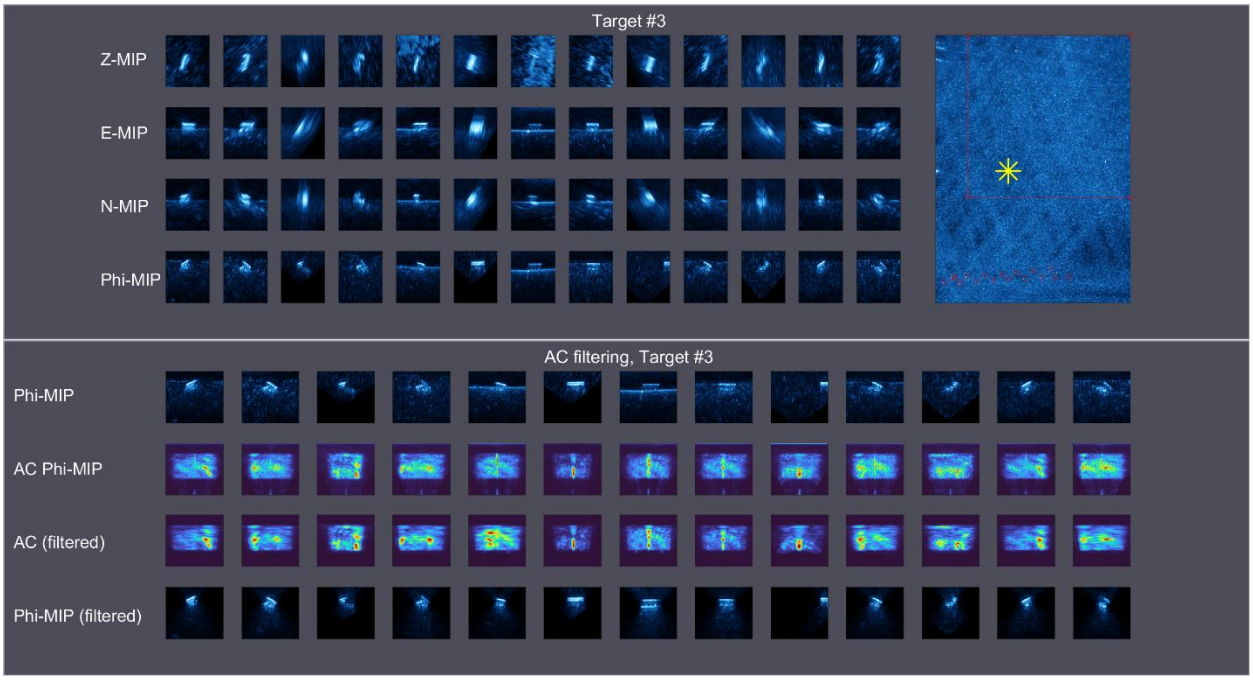

Figure 245: The detection set for a proud cylindrical object. The improvements from AC filtering are highly visible in this example.

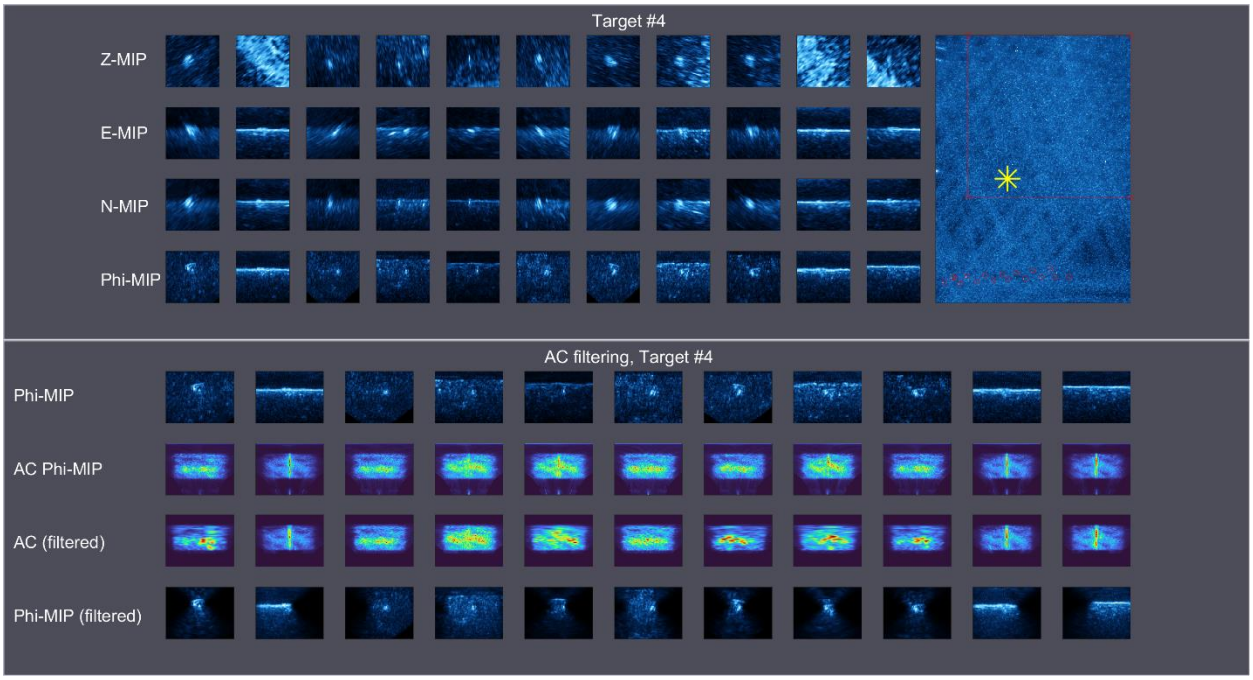

Figure 246: Detection hits for a smaller buried object that could possibly be a UXO. Several of the specular glints in this set were not filtered out by the AC filtering algorithm because the specular glint detection and removal filter described in *Materials and Methods: Structural Acoustic Feature Extraction* was not developed when these data visualizations were created.

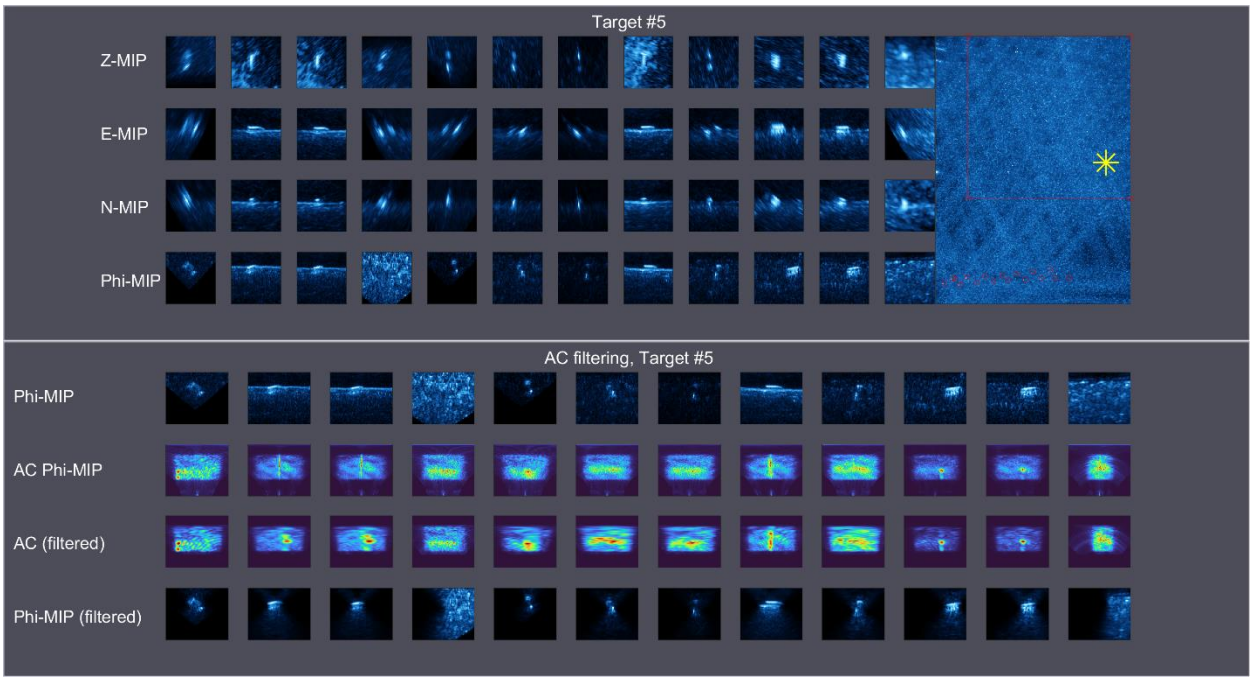

Figure 247: Detections for a UXO target.

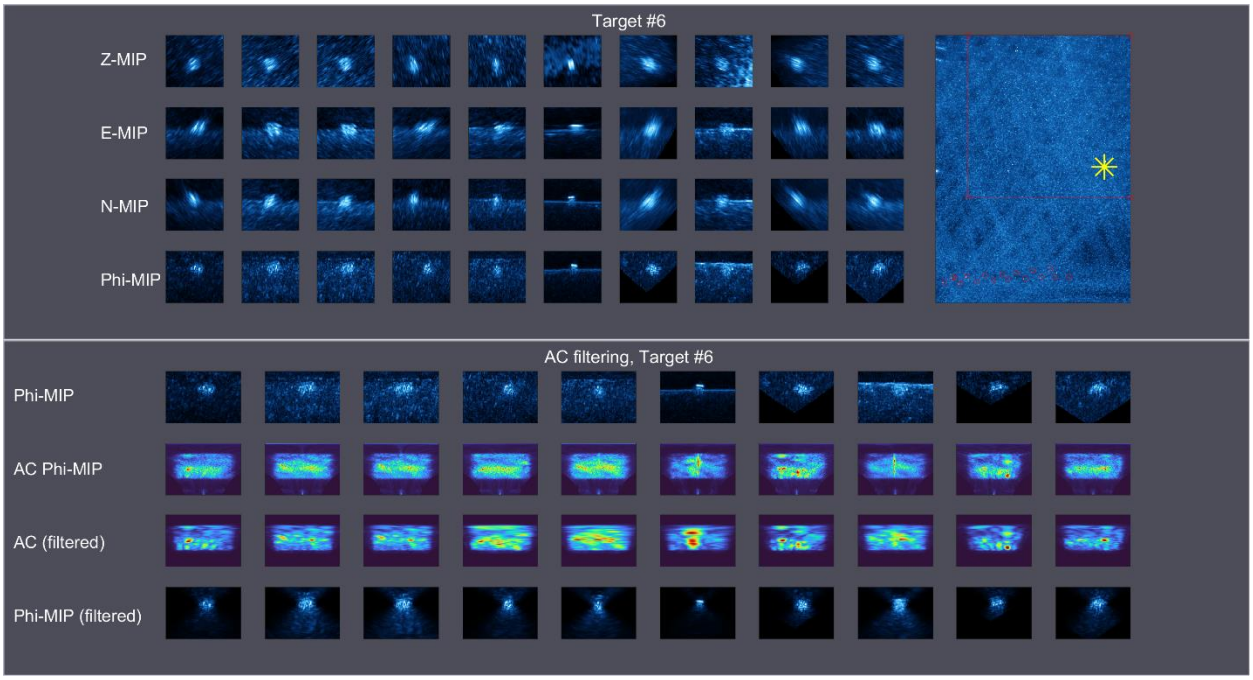

Figure 248: Detections and acoustic-color data for a cinder-block, flat-side up. The phi-MIP second from right shows what appear to be two distinct chambers in the object.

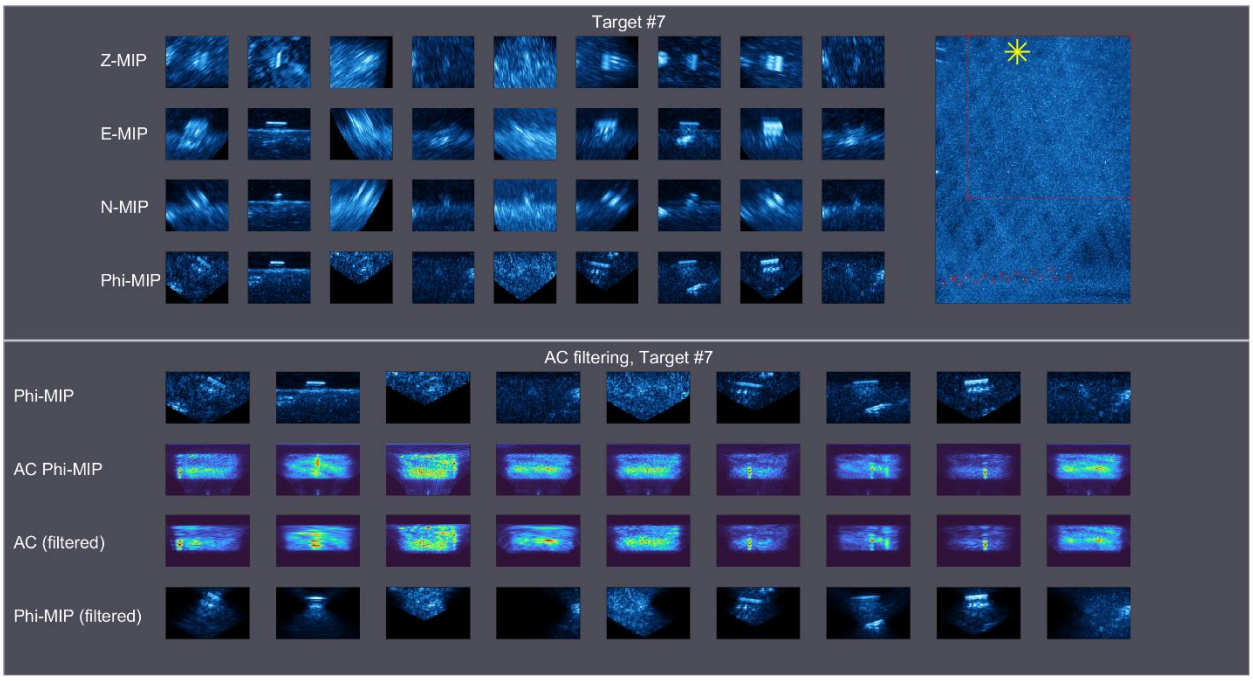

Figure 249: The hollow 2-1 aluminum cylinder. Another object was placed very near this object, interfering with the clutter rejection algorithm used at the time of analysis.

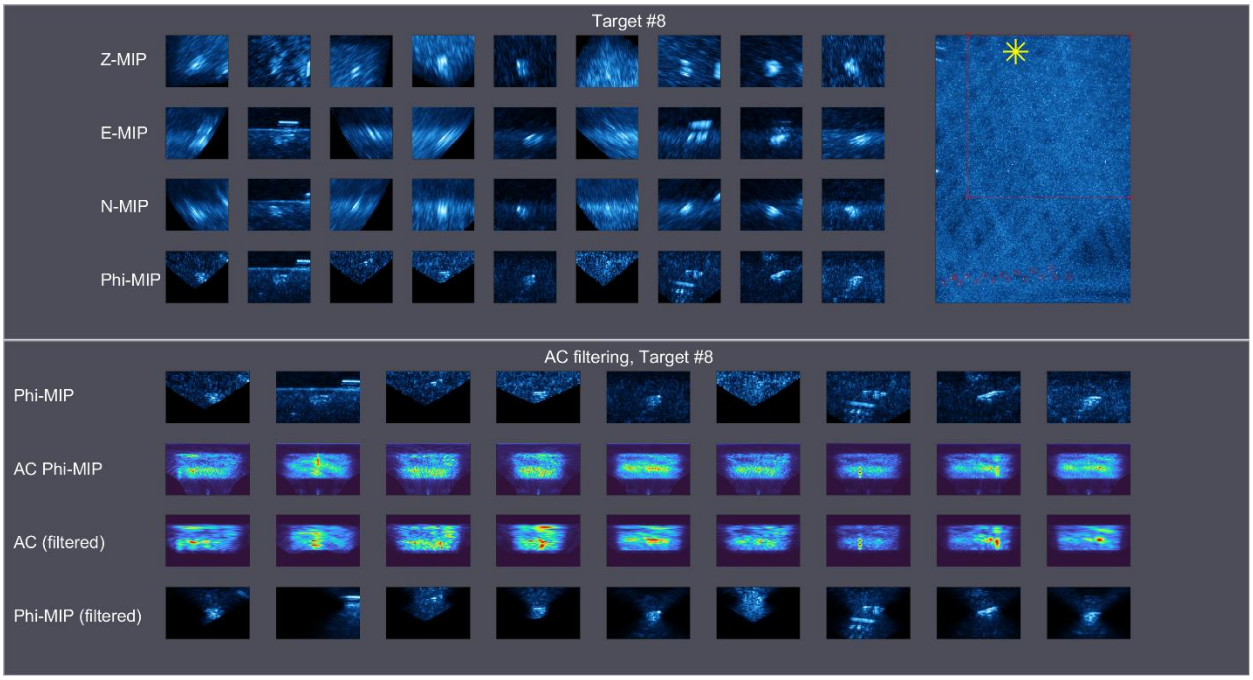

Figure 250: The detection set for the object adjacent the aluminum shell shown in the previous figure.

The previous eight examples of target multi-aspect scan sets were compilations automatically generated from a list of UTM coordinates corresponding to some of the prominent targets visible in the target field. Due to the tight alignment between scans and the organized post-processing framework that was to eventually evolve into the MuST toolbox, these datasets were extraordinarily simple to generate. The co-registration between looks was remarkably good, allowing for very simple multi-scan and multi-aspect feature comparison. There has been no curating of the results shown above and the "good" and "bad" results are included in the multi-look data visualizations. In some looks the targets are clearly prominent, and in others they can be almost invisible. The structural acoustic feature enhancement tool leveraged at the time of analysis clearly improved the acoustic color of the detected targets, however it was often hit-or-miss due, primarily due to the influence of specular glints; hence the development of the ROI glint subtraction process described in *Materials and Methods: Structural Acoustic Feature Extraction*. Additional SNR improvements were later realized by adding the auto-range gating feature and improving how the wavenumber filter boundaries were determined.

Technical Evaluation

The APL-UW created a classifier-aided detection and dig list for the blind survey indicating the likely position of surveyed targets and submitted this list for third-party review by the Institute for Defense Analyses (IDA). Figure 251 shows the ROC-curve created using a 2.5 meter detection halo around the locations supplied by APL-UW for estimated UXO locations. Of the nine targets placed in the testbed seven were correctly identified, two targets were missed and two false negatives were supplied. Both missed targets were 81mm UXO giving an indication of the UXO size scale at which detection performance begins to deteriorate. Despite completely

missing the 81mm UXO in the 2020 experiments, in the follow up 2021 Sequim Bay field experiment discussed next, many 81mm UXO were placed and over 90% of them were detected.

Note that though this testbed experiment were funded via MR18-5004 (“Multi-Sensor Towbody (MuST) for Detection, Classification, and Geolocation of Underwater Munitions”, PI: Kevin Williams [11]), the present project was integral to the post-processing and UXO detection list generation and therefore results are included here.

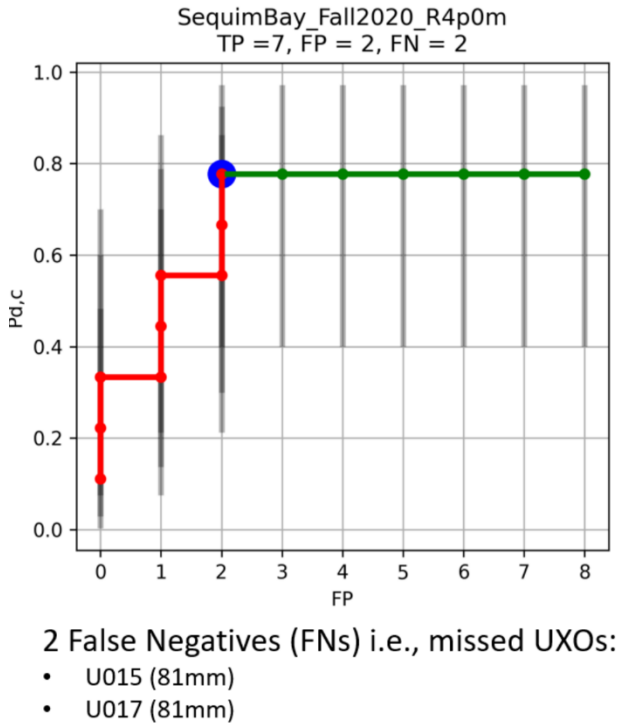

Figure 251: The UXO detection score for the 2020 Sequim Bay blind testbed target list created by APL-UW. The scoring of the detection list was performed by The Institute for Defense Analyses (IDA).

VIII. Field test results: 2021 INS integration tests

Due to heading drift issues in the INS model used for the 2020 Sequim Bay field trials a new INS was integrated into the MuST system in the spring of 2021. The primary purpose of this field test was INS performance verification, but it was also used as a springboard for testing various new features incorporated into MuSTLive. The new INS, which remains the model installed on the MuST at the closing of the current project, was evaluated to provide a consistent navigation error of half-percent distance travelled or better (i.e. 5 meters of error every kilometer travelled), a significant improvement over the previous model. Error statistics were determined by getting GPS fixes on the surface, towing the MuST system for different lengths underwater, out of range of a GPS, re-surfacing, then observing the difference from the GPS position. Fig. 252 shows the measured percentage distance-travelled error vs. distance travelled which, for lengths above 500 meters, is consistently better than half-percent distance travelled.

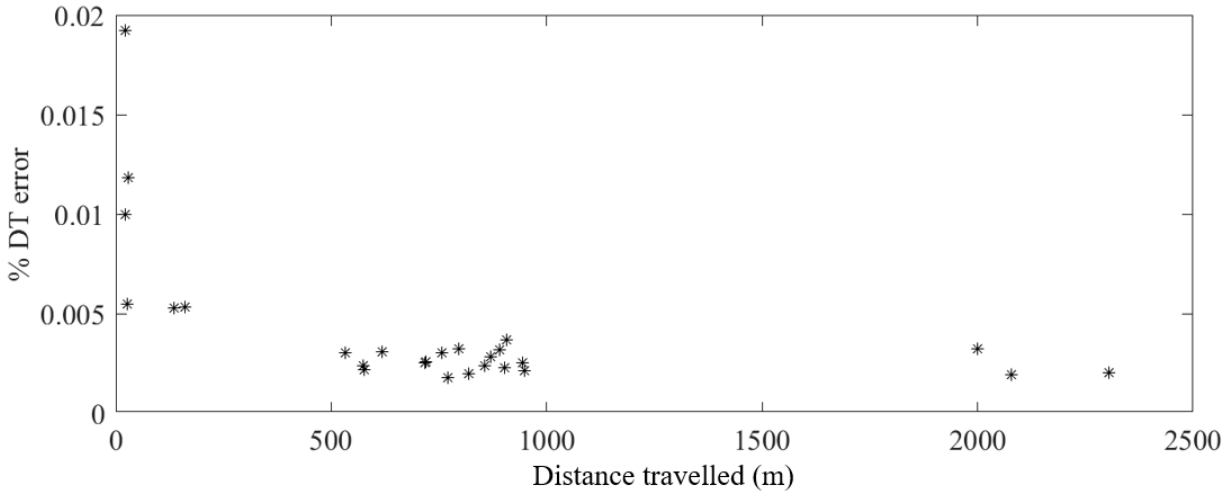

Figure 252: MuST Towbody INS error statistics (Percent distance-travelled error vs. distance travelled) compiled during the 2021 INS integration tests.

No targets were deployed prior to the integration test however a fishing boat and sunken barge were scanned from different angles and 3D mosaics were created by merging the different scans. Figure 253 shows a composite mosaic of a sunken barge and Figure 254 shows a composite 3D rendering of a fishing boat including a slice through the keel (single-scan image) with a 6-foot human polygon superimposed for a scale comparison.

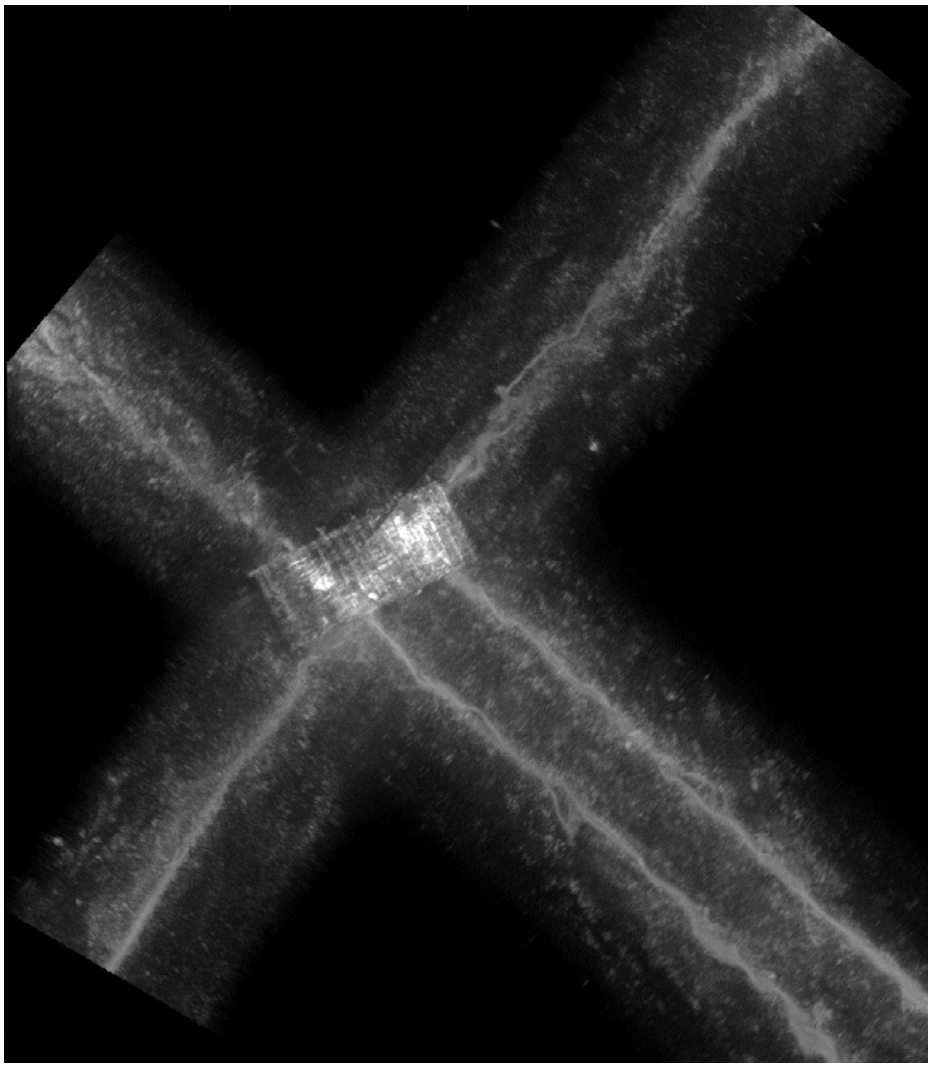

Figure 253: A composite Z-MIP rendering of a sunken barge in Lake Washington, collected during the 2021 INS integration tests.

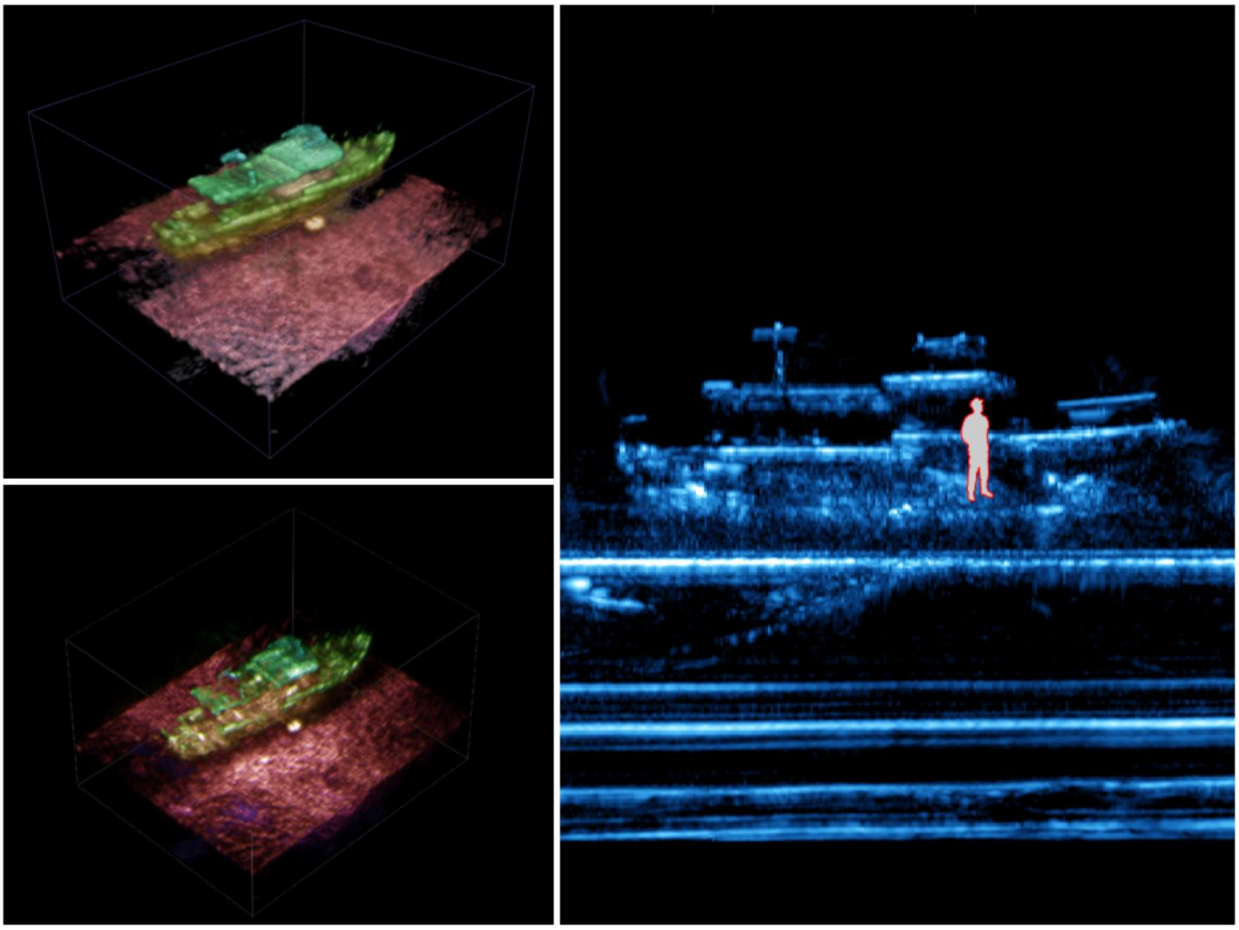

Figure 254: An Alpha-blended rendering (top left), and MIP rendering (bottom left), generated from a multi-scan dataset. A slice through the keel captured by one scan which travelled down the fishing boat is shown on the right along with a 6-foot tall polygon model for a man giving a size reference for the boat.

A significant discovery from this dataset was that attempts to use the feature-based mosaicking developed for Sequim Bay failed completely for the multi-layered datasets captured in Sequim Bay if no steps were taken to limit the MIP imagery used for alignment to only the region near the sediment-water interface. Limiting the detected features to the water / sediment interface reduces the impact of refraction, which may possibly be an explanation for why this step was necessary. The apparent sensitivity of the algorithm to environment provides a motivation for additional development, including multi-sensor feature-based navigation.

IX. Field test results: 2021 Sequim Bay UXO Testbed Survey

Overview

In Sept. 2021 a UXO testbed was again emplaced by PNNL, however a major difference differentiated the 2021 survey from the 2020 survey: the number of UXO placed in the testbed was now an unknown. In contrast, in the 2020 test the number of clutter items and exact number of UXO, including how many of each type that were placed, were given to APL-UW and the only unknown was the position.

Additional differences included the usage of an improved INS, experience on the part of the MuST Towbody team and an efficient concept of survey operations developed from the previous year. Furthermore the MuST post-processing toolkit – everything excepting the PMA tool – was in a nearly current development state.

Significantly, development of the MuST Toolbox was sufficient that in the evening, following data acquisition, the entire set of testbed surveys captured by the BOSS could be batch processed, mosaicked, and markers corresponding to targets of interest were detected and stored in a marker list (e.g. Fig. 59). The following morning, prior to survey operations, the list and mosaic were reviewed to inform survey decisions. The process was so efficient that despite an unanticipated afternoon standing down due to inclement weather the intended survey lines were completed a day early and a cursory detection map was available at the end of the survey.

This section will detail some of the day-by-day developments of the field operation including mosaic development and daily additions, the final mosaics as rendered in the PMA tool along with several of the target regions-of-interest viewed via the tools available in the ROI review tool and look-editor, and give a brief overview of the grading provided by IDA for the official APL-UW detection list. Additionally, the D. Williams 3D Mondrian detector was integrated into the beamformer following acquisition of this dataset and a review of the results of this detector will be given.

Intra-survey BOSS data products

This section will describe some of the data products produced with the MuST BOSS toolbox during the actual 2021 Sequim Bay testbed survey, including the mosaic and detection map updates. Prior to showing data products created via post-processing, however, it is worthwhile to note some of the data-products that were created live during data acquisition. Figure 255, for example, shows the navigation screen containing the positions of the targets tagged during acquisition via MuSTLive's target-tagging feature (see *Materials and Methods: MustLive*, Fig.'s 60 and 61).

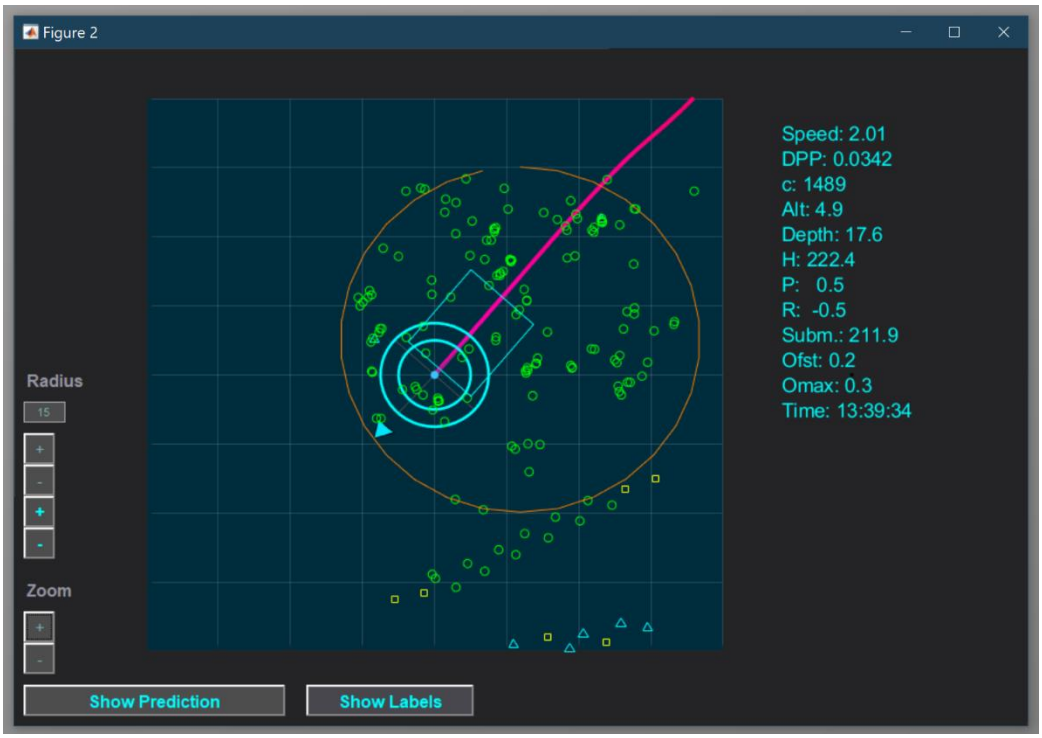

Figure 255: The navigation screen of MuSTLive showing the large numbers of objects detected after the first day of operation. Each of the green circles inside of the testbed region, demarcated by the orange circular line, is a detection noticed by the MuSTLive operator and tagged during data acquisition.

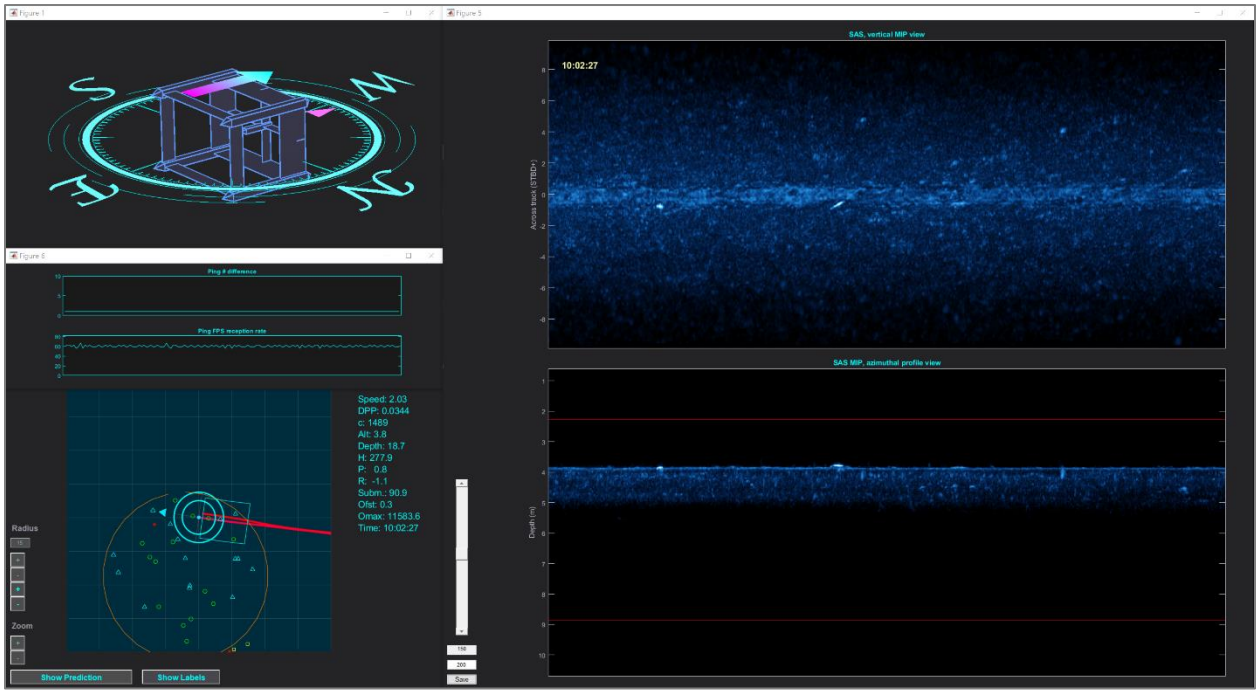

Figure 256: A snapshot of MuSTLive as the vehicle traverses the northern portion of the testbed region. Several prominent scatterers are visible in the SAS MIP images on the right, one of which (near the center) is clearly cylindrical and has a tapered end. This was the fourth day of testing and the detections created from

the mosaics of the previous days' data are visible in the navigation map screen on the bottom left of the figure.

Shifting focus to post-processing data-products created from data acquired during the day, Fig. 257 shows the mosaic generated from scans performed on 9/27/2021, the first day of the survey operations. Features are well aligned and several prominent scatterers are visible in the target region in the north. Two different scan patterns are also visible in the mosaic.

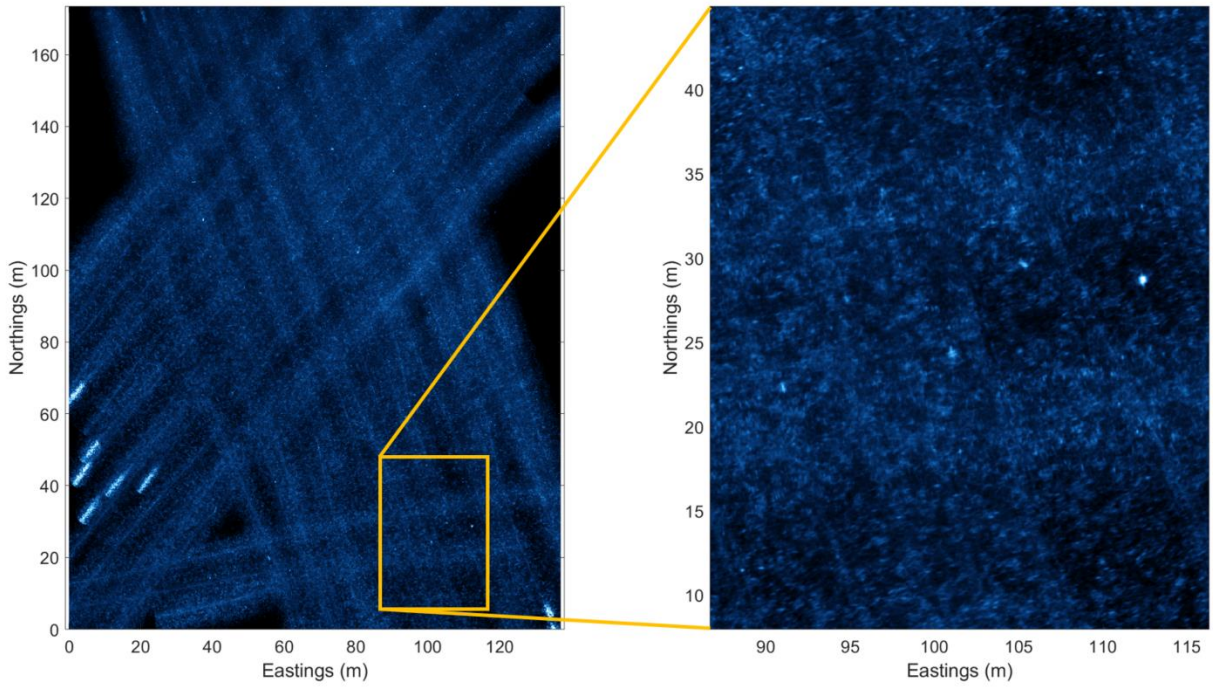

Figure 257: The day-1 mosaic for the target field and calibration line (left) and a zoomed inset showing the targets on the eastern end of the calibration line.

Figure 258 shows the detections that were generated from this set of mosaicked data. Blue triangles correspond to objects that look likely to be UXO, yellow squares are objects that are likely clutter, and green circles represent objects that are difficult to classify and need additional views. This was generated on, and available for use on day 2 of the Sequim Bay survey. The background image, created later on the third day of experiments, is a composite generated from the first two days of scanning activity.

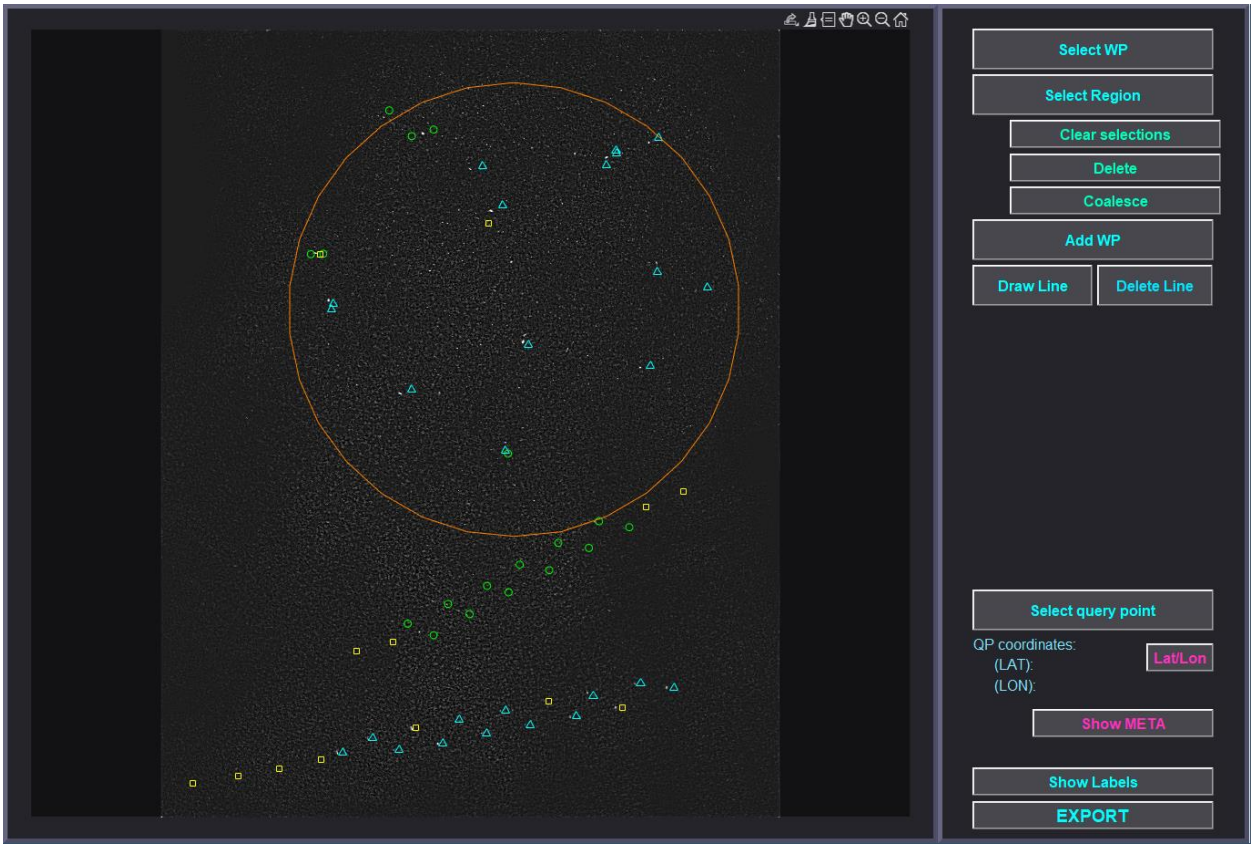

Figure 258: The detection list created from a cursory examination of the day 1 mosaic (Fig. 257). The target field detections are shown in the orange circle. The particular view tool, ("Waypoint_edit") was developed just prior to the experiment to help view marker data and view metadata associated with scan detections.

Figure 259 shows the mosaic created from the scans collected on the second day of experiments. Note that significantly fewer scans were collected due to inclement weather. The performance of feature-based renavigation improves as the number of looks at the target field increases and diversifies, however despite the relatively small number of looks the data appears well registered. The zoomed plot on the left shows the targets at the western end of the calibration line.

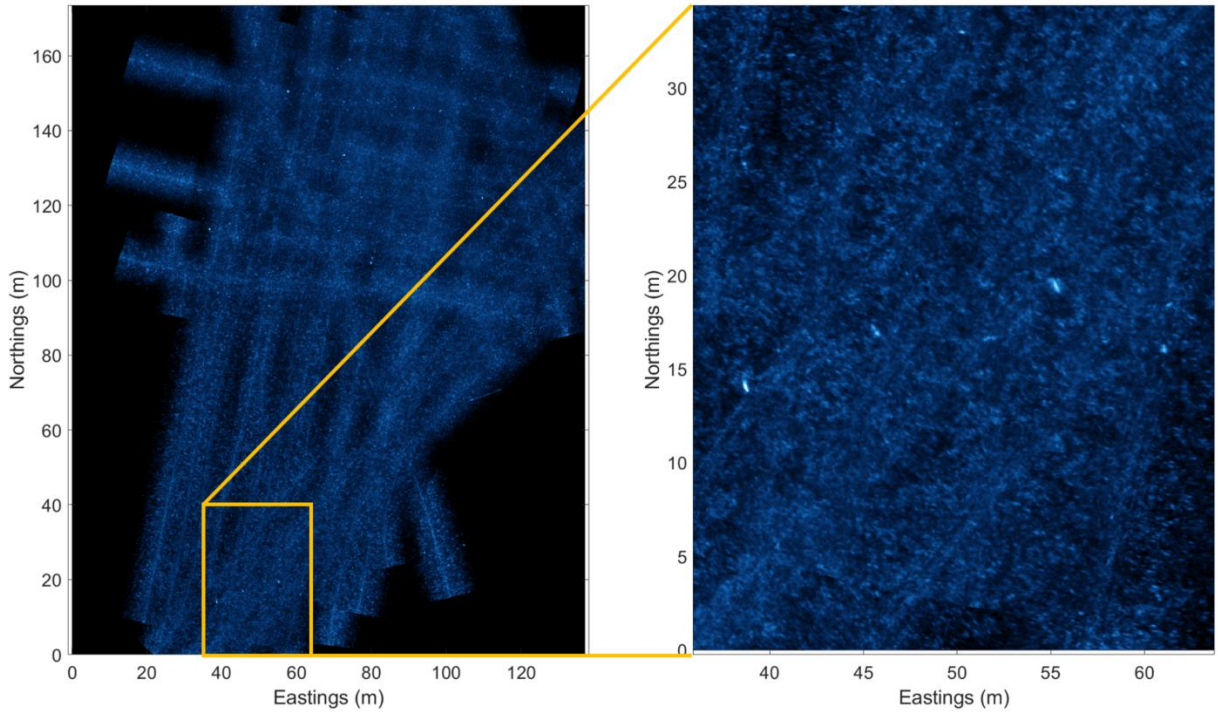

Figure 259: The mosaic generated from scans collected on Day 2 of the Sequim Bay 2021 survey. The zoomed plot on the right shows the targets on the western edge of the calibration line.

On the third day of operations, rather than generating new mosaic data from the third day of operation, effort was put forward to resolving some bugs encountered during the day and in some of the renavigation software. Efforts were re-focused on generating an improved combination of the day 1&2 feature-aligned mosaics rather than batch processing the data from day 3. The resulting combination, which performed some normalization and de-trending before computing an average, was used as the background for the "Waypoint_edit" view tool shown in Fig. 258. The aggregate is shown in 260 with a better dynamic range thresholding and various targets zoomed to show details.

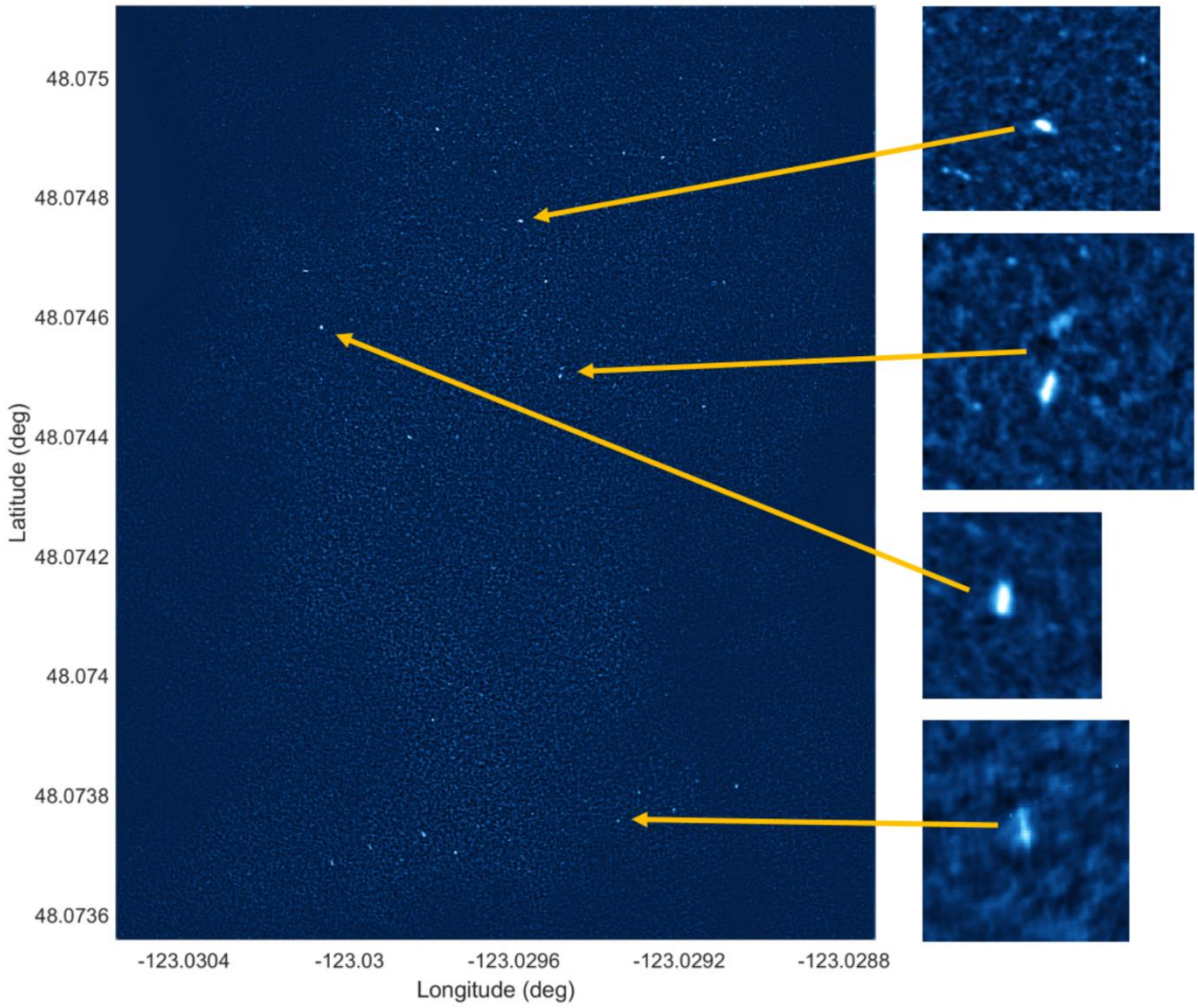

Figure 260: A combined sonar mosaic generated from data acquired in days 1 and 2 of the Sequim Bay 2021 testbed survey. Several prominent objects in the test bed and calibration lines are indicated with the zoomed snippets on the right and arrows pointing to their locations in the mosaic.

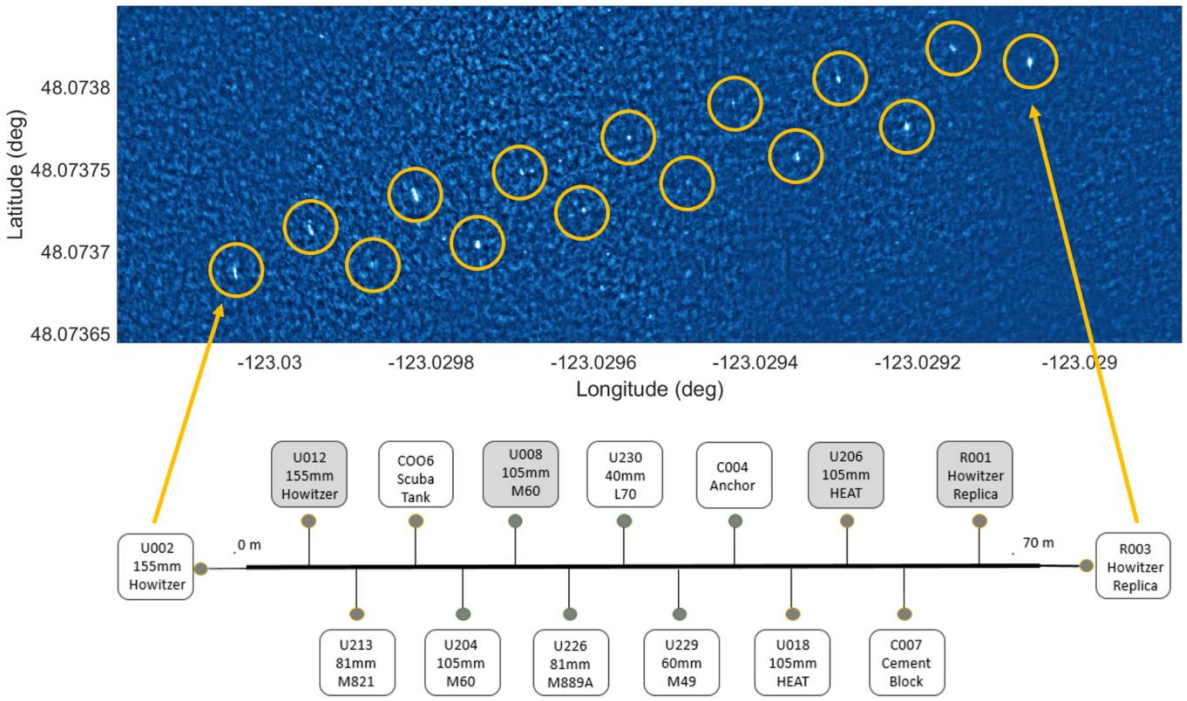

Figure 261: A view of the calibration line using the day 1 and 2 combination mosaic. The calibration line targets are circled and the calibration line key is shown on the bottom. Greyed targets in the calibration line key indicate that the targets were buried by divers.

Note that in Fig. 261 each of the targets on the calibration line is visible, however some are significantly more prominent relative to the background than others. For example, the Howitzer shells are trivial to detect, however U229 and U213, 60mm and 81mm shells, show only faint signatures. Unusually, the 40 mm target had a reasonably strong signature (U230), indicating that, at least from some aspects, targets of this size are capable of being detected by the BOSS.

Prior to the end of the experiment on day 4 no further feature-based mosaicking was performed (during the experiment), however detection lists were continually being refined via MuSTLive and Fig. 262 shows the final waypoint list assembled on day 4 (9/30/2021).

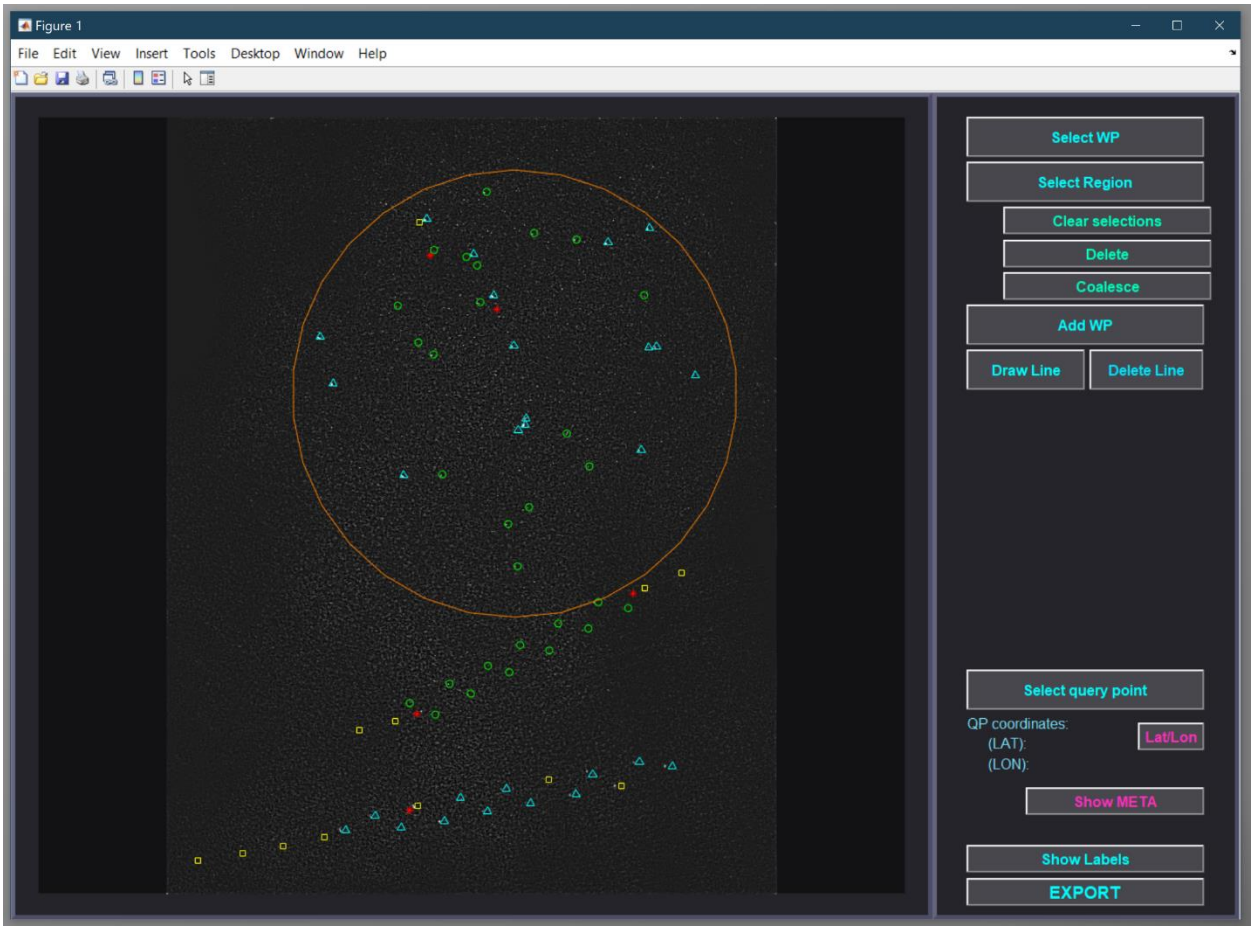

Figure 262: The Sequim Bay day 4 detection list viewed in Waypoint_edit.

Mosaic results

Following the testbed survey the entire set of scans collected in Sept. 2021 in Sequim Bay were SAS batch-processed with the renavigation solution generated via the MuST Toolbox RENAV tool (e.g. Fig. 67 – 71), and 40cm Z-MIP layers with centers: $Z_{MIP_centers}$ ={-0.5, -0.3, -0.1, 0.1, 0.3, and 0.5}. In addition to these layers, a seventh "full sediment" Z-MIP is also stored. Each of these layers can be viewed using any of the mosaic combination methods, e.g. mean, max, contraharmonic / Lehmer mean, in the PMA tool. This section will show the resulting mosaics produced at each of these layers and for each of these combination types including several regions of interest.

Figures 263 – 269 show side-by-side comparisons of layers 1 – 6 and the full MIP, zoomed out to show the whole target field. When zoomed out it is nearly impossible to see any small scale target features, however it is useful for comparing the relative types of information available at scale in each projection type. For example, the average has the most scan-related intensity artifacts and modulations, however it also gives the best depiction of local scan coverage, i.e. the number of looks acquired at a given location. Clearly the target field (north) and the calibration line to the south have both been heavily scanned.

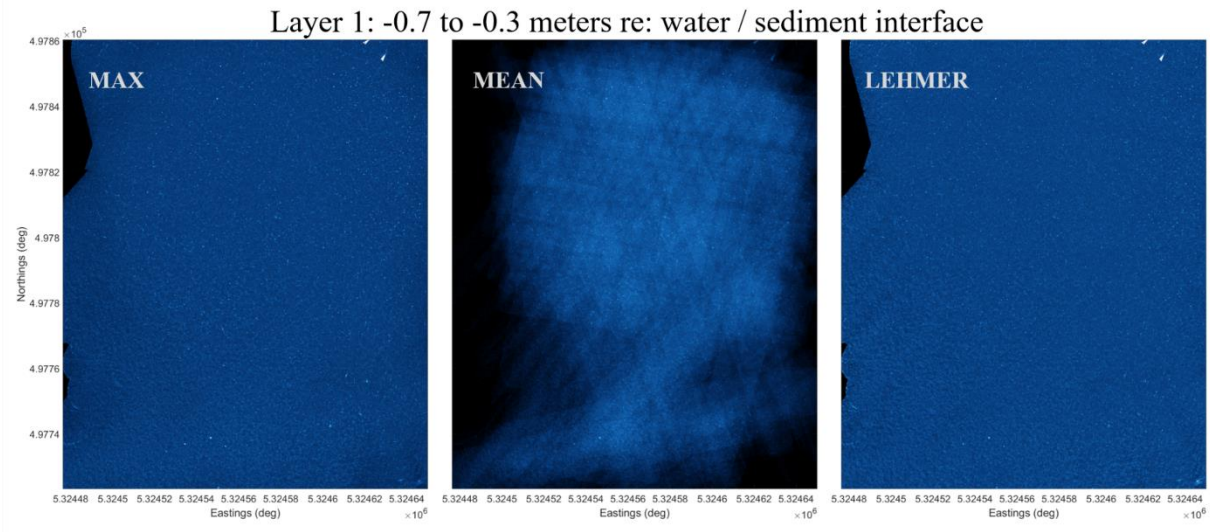

Figure 263: Sequim Bay 2021 test region mosaic set showing the max, mean, and contraharmonic (Lehmer, p = 2) mean for the lowest layer. Note the MAX mosaic is least impacted by scan beam pattern artifacts, however it gives you no indication of scan density. The mean is the most impacted, but is also useful for determining the regions of high coverage in surveyed area.

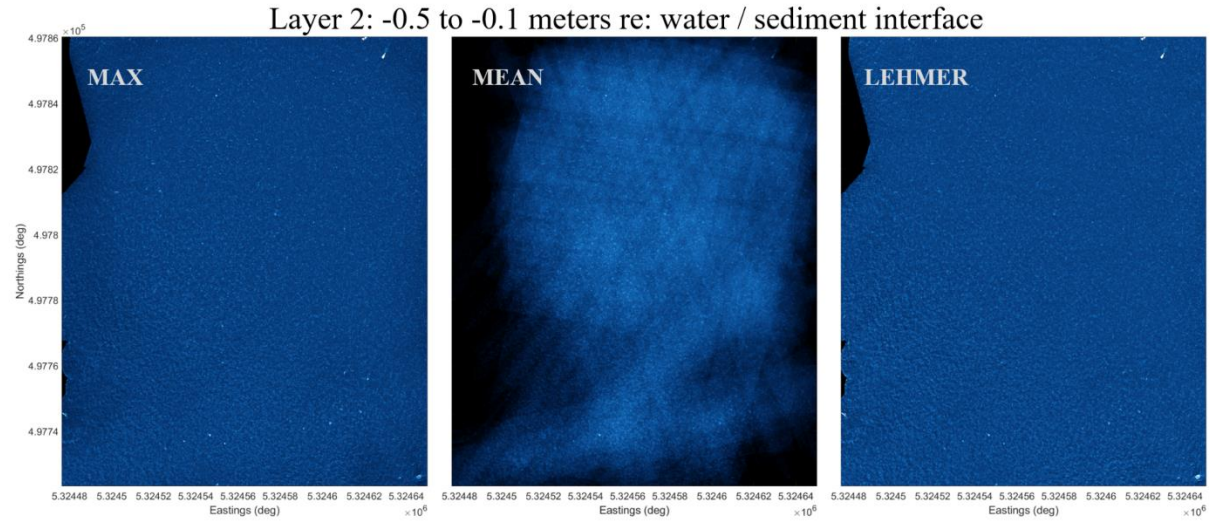

Figure 264: Sequim Bay 2021 test region mosaic set showing the max, mean, and contraharmonic (Lehmer, p = 2) mean for the second to lowest layer.

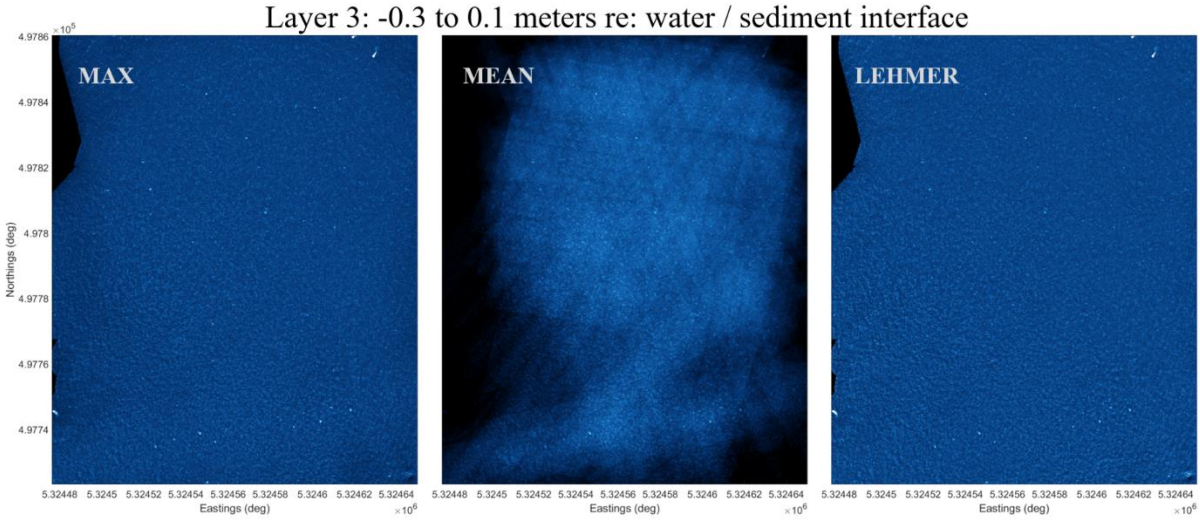

Figure 265: Sequim Bay 2021 test region mosaic set showing the max, mean, and contraharmonic (Lehmer, p = 2) mean for a layer that includes 30 centimeters below the sediment and 10 cm above it.

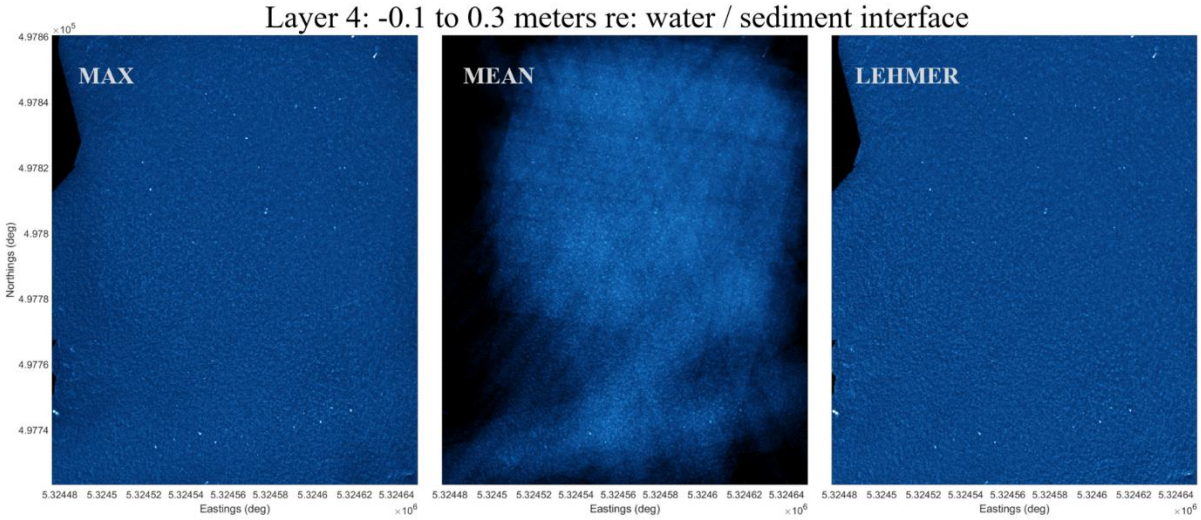

Figure 266: Sequim Bay 2021 test region mosaic set showing the max, mean, and contraharmonic (Lehmer, p = 2) mean for a very shallow layer of sediment.

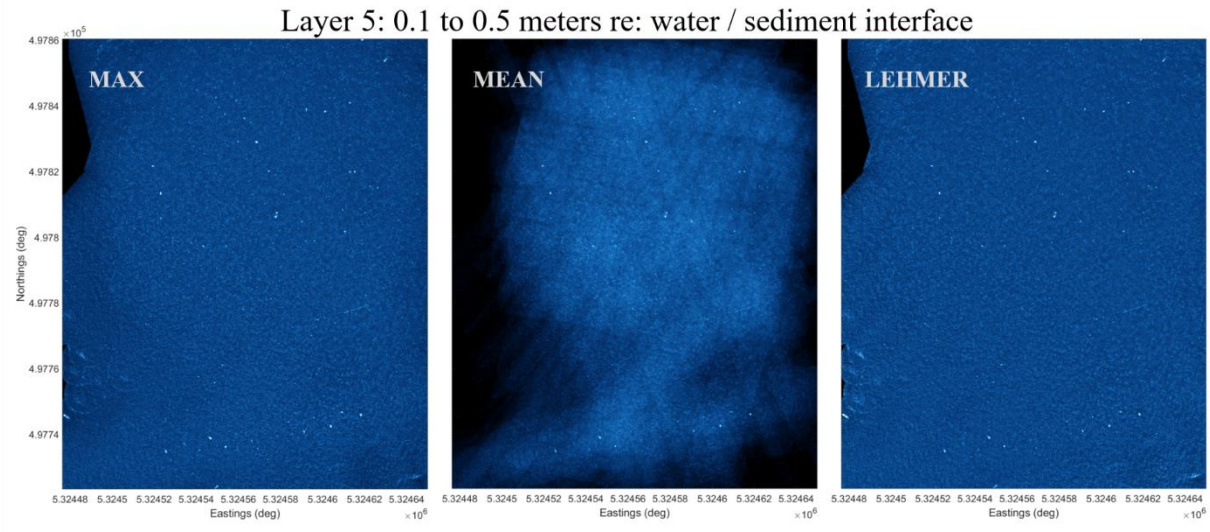

Figure 267: Sequim Bay 2021 test region mosaic set showing the max, mean, and contraharmonic (Lehmer, p = 2) mean for a region that is nominally above the sediment. This area should only contain features above the sediment, however it clearly contains sediment features as well. This is due to the fact that the point-spread function of the system has a slant in the vertical direction for any points other than directly at nadir. The result is that there is a projection of intensity above the interface which shows up in this layer.

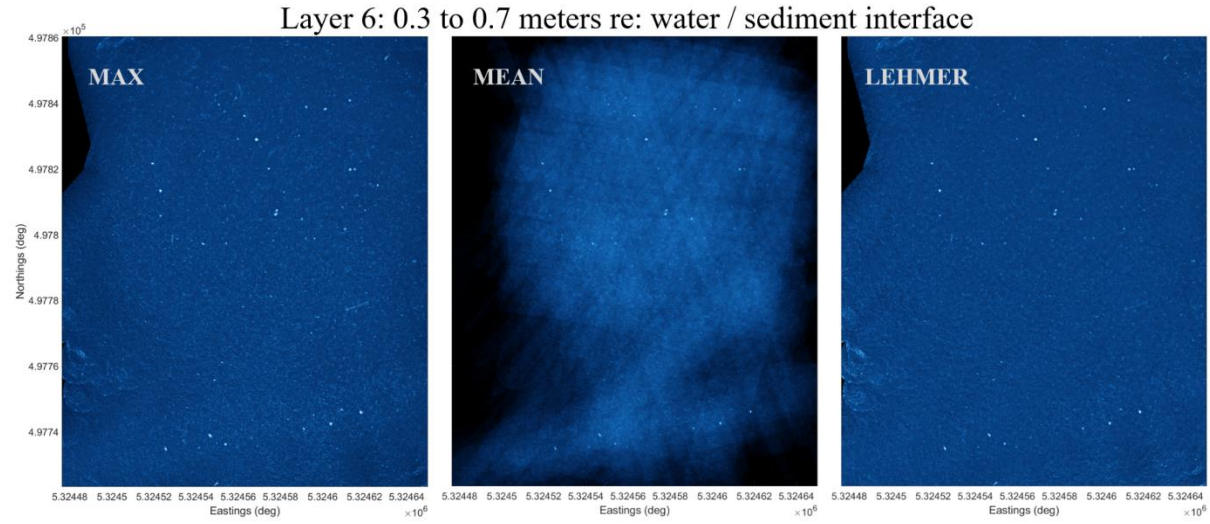

Figure 268: Sequim Bay 2021 test region mosaic set showing the max, mean, and contraharmonic (Lehmer, p = 2) mean for a layer that nominally has 30cm of clearance relative to the water / sediment interface. Note that the prominence of proud scatterers relative to the sediment is significantly higher in this layer vs. the previous layer (layer 5). At this level of height above the sediment most of the scattering due to the tilt of the point-spread-functions and aliased energy is significantly lower.

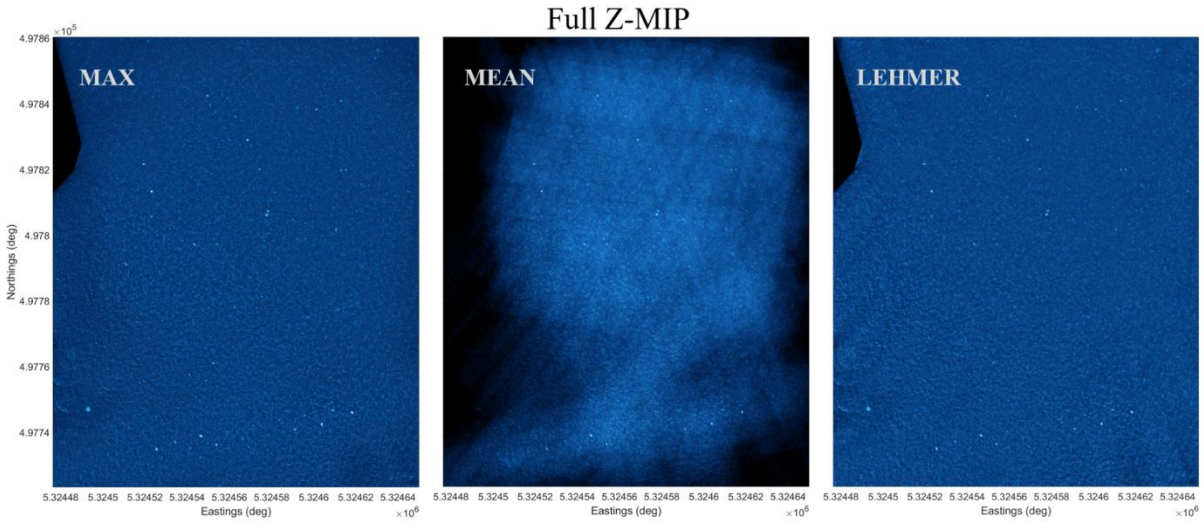

Figure 269: Sequim Bay 2021 test region mosaic set showing the max, mean, and contraharmonic (Lehmer, p = 2) mean for the full sediment MIP.

From the previous figures it is very clear that the performance of the Lehmer mean, from a scan path artifact standpoint, is much closer to the max than to the mean operation. The next figures show the same layers and combinations, however they are zoomed to show a portion of the calibration line. This allows the effects of the projection methods at target and sediment texture scales to be evaluated.

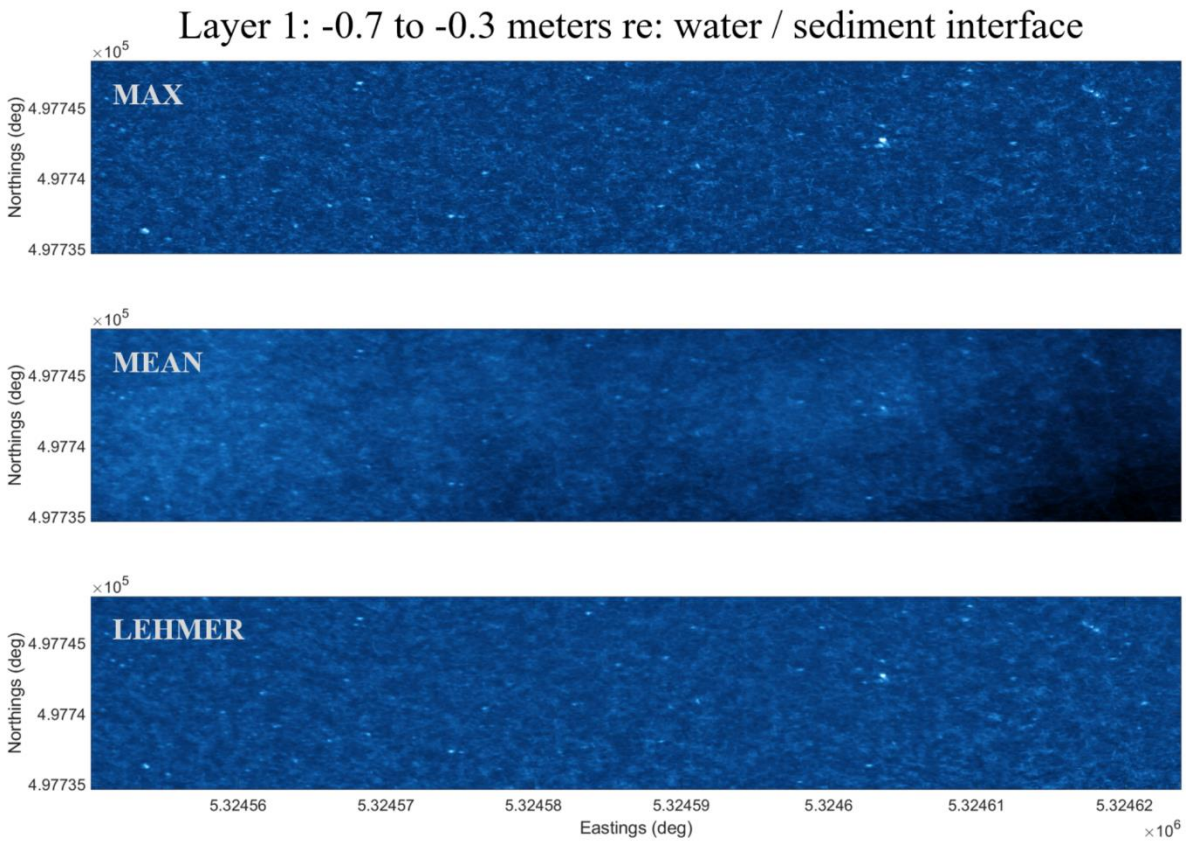

Figure 270: Sequim Bay 2021 layer 1 mosaic, zoomed in on part of the calibration line, set showing the max (top), mean (middle), and contraharmonic (Lehmer, p = 2) mean (bottom).

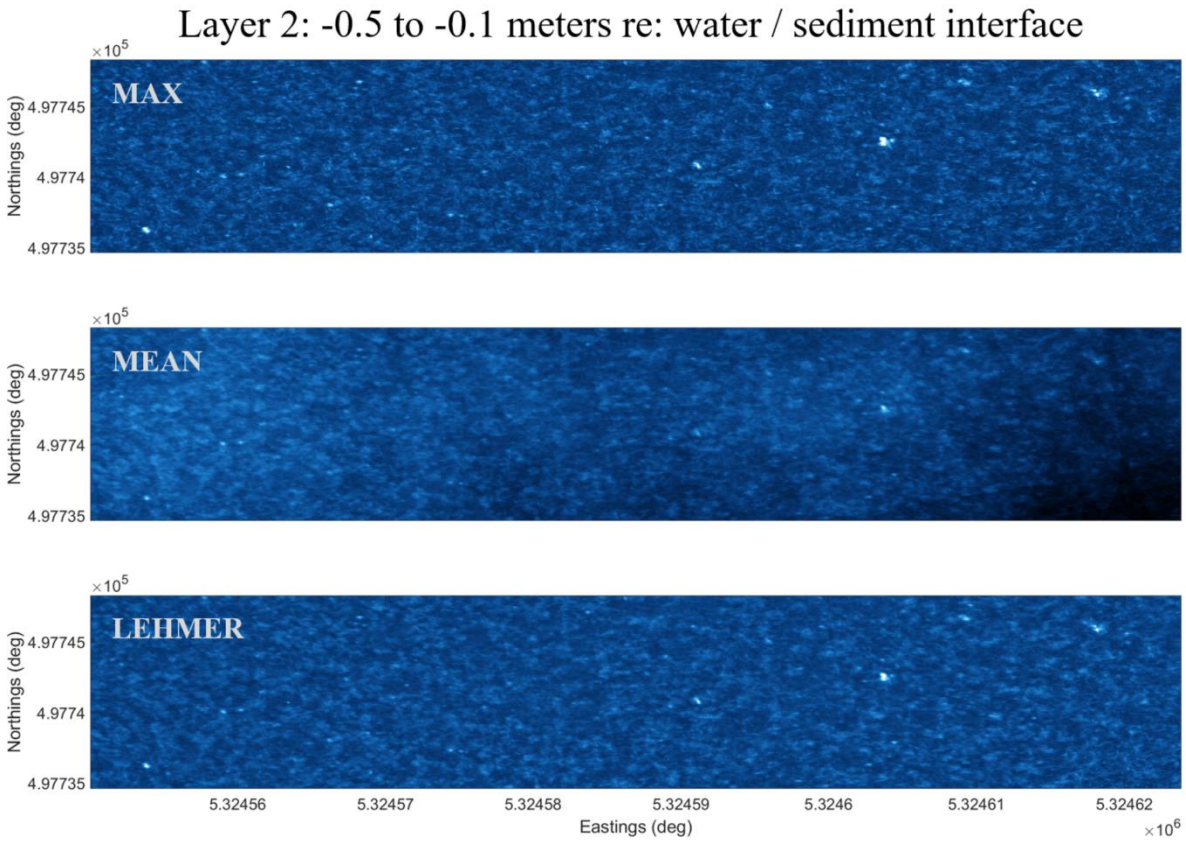

Figure 271: Sequim Bay 2021 layer 2 mosaic, zoomed in on part of the calibration line, set showing the max (top), mean (middle), and contraharmonic (Lehmer, p = 2) mean (bottom).

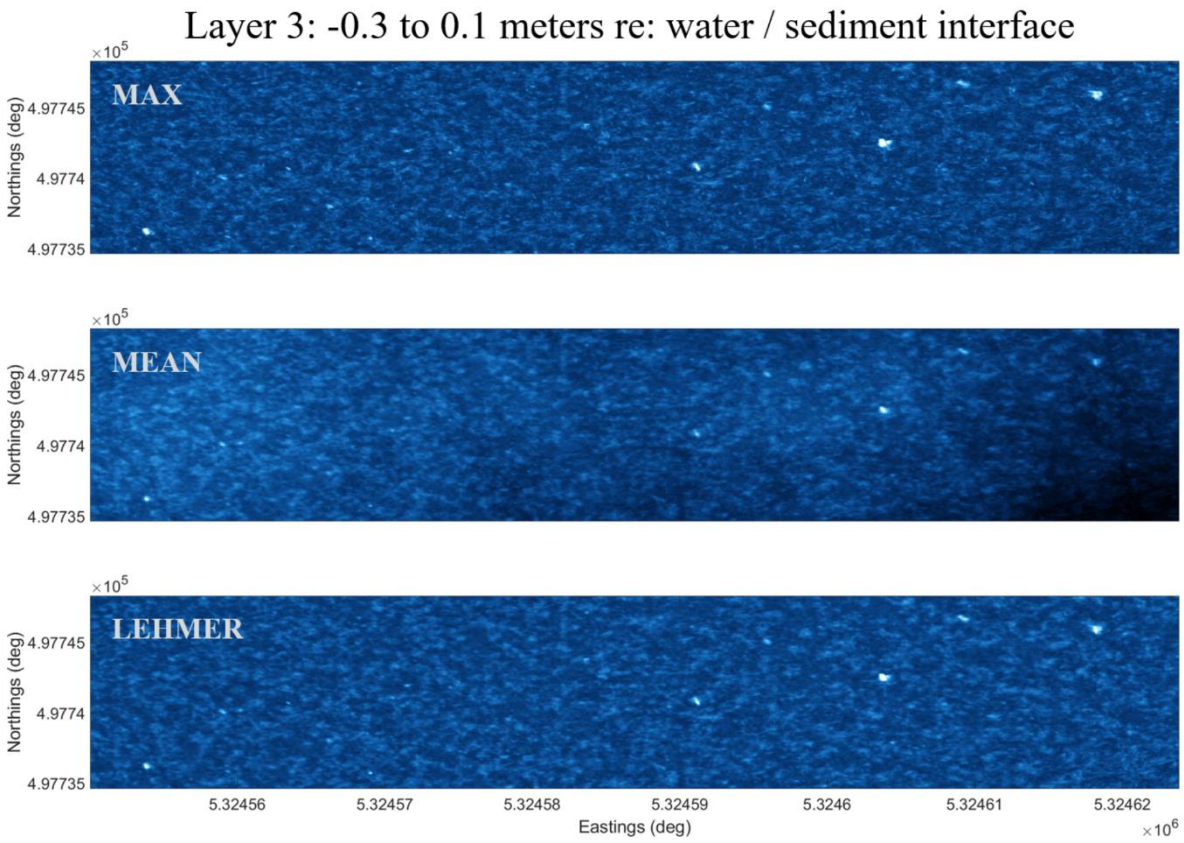

Figure 272: Sequim Bay 2021 layer 3 mosaic, zoomed in on part of the calibration line, set showing the max (top), mean (middle), and contraharmonic (Lehmer, p = 2) mean (bottom).

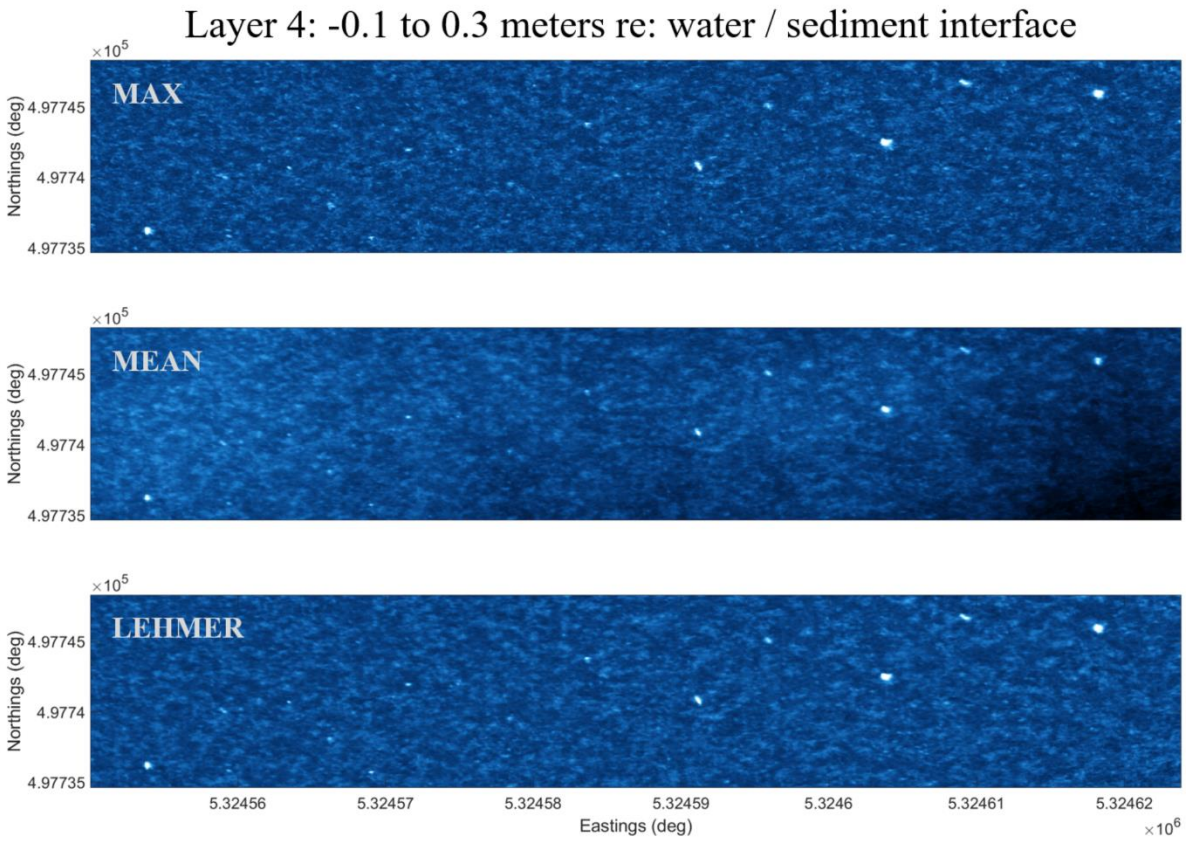

Figure 273: Sequim Bay 2021 layer 4 mosaic, zoomed in on part of the calibration line, set showing the max (top), mean (middle), and contraharmonic (Lehmer, p = 2) mean (bottom).

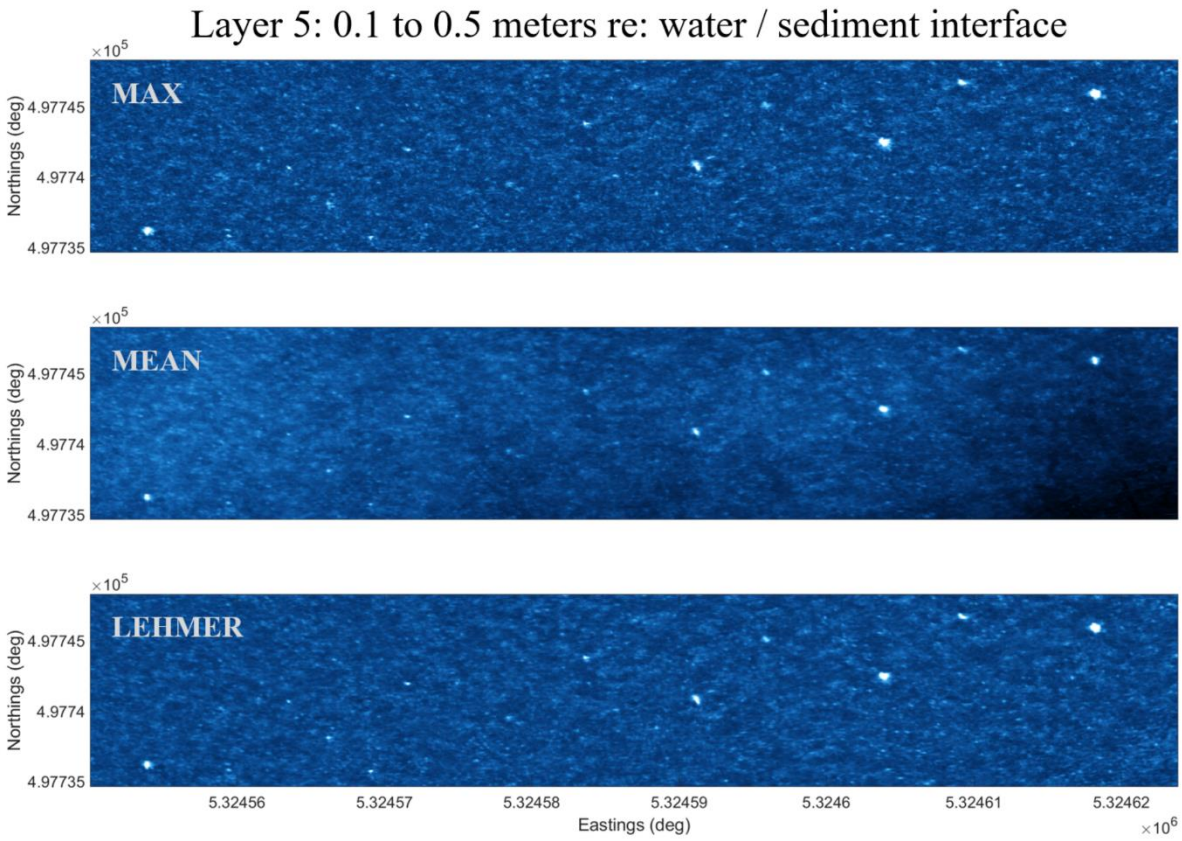

Figure 274: Sequim Bay 2021 layer 5 mosaic, zoomed in on part of the calibration line, set showing the max (top), mean (middle), and contraharmonic (Lehmer, p = 2) mean (bottom).

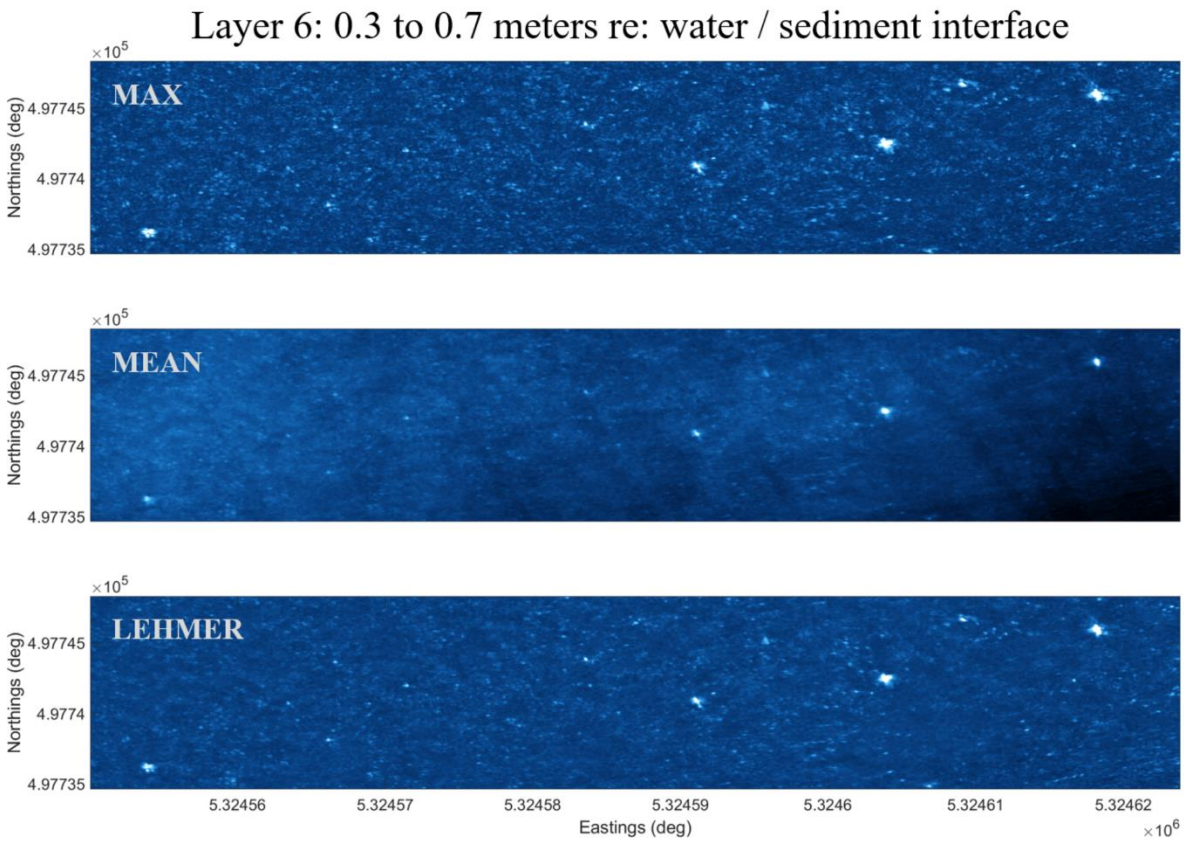

Figure 275: Sequim Bay 2021 layer 6 mosaic, zoomed in on part of the calibration line, set showing the max (top), mean (middle), and contraharmonic (Lehmer, p = 2) mean (bottom).

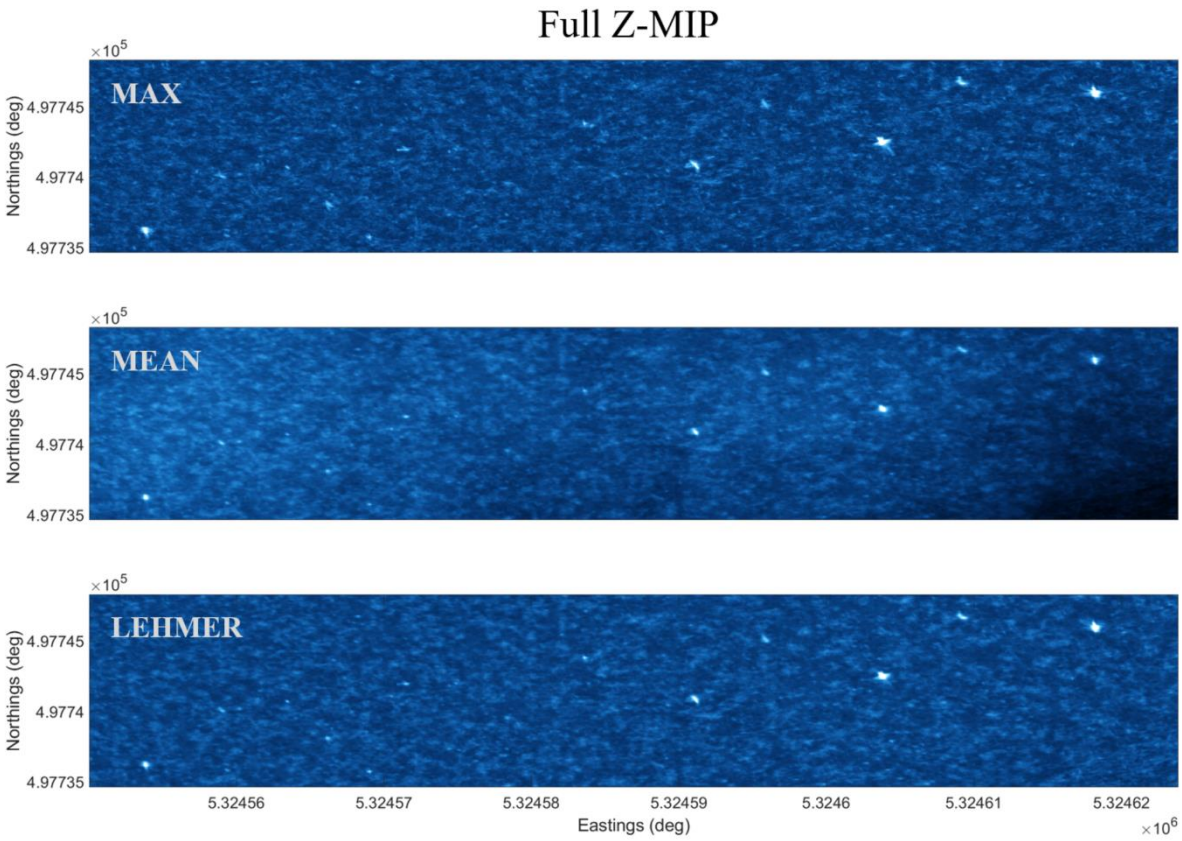

Figure 276: Sequim Bay 2021 full depth MIP mosaic, zoomed in on part of the calibration line, set showing the max (top), mean (middle), and contraharmonic (Lehmer, p = 2) mean (bottom).

Qualitatively the level of de-speckling provided by the mean is highest, the maximum is lowest, and the Lehmer mean lies in between the two. Quantifying the de-speckling properties of the Lehmer mean would be a useful follow-up effort, however it clearly provides more despeckling than the maximum operator, but is much more similar to the maximum operator in terms of how beam trend artifacts are minimized on the large-scale mosaic.

Target-of-interest ROI processing results

Pacific Northwest National Laboratories (PNNL) positioned 23 targets of interest in the testbed region. The set was comprised of 2x 155mm Howitzer shells (one of which was a facsimile), 3x 105mm HEAT rounds, 4x 105mm M60 projectiles, 5x 81mm M821 finned mortar shells, 7x M889A1 mortar shells, and 2x 60mm M49 shells for a total of 23 targets of interest. Following the official submission and grading of a target list created by APL-UW, the official answer key, which providing the target type and coordinates within the test-bed, was released to APL-UW. A marker-list was generated from this set and overlaid on the mosaics. All targets were associated with a nearby anomalously bright (relative to the mosaic background) feature excepting the 60mm shells. Views through the color-by-depth multi-aspect composite images of all 23 targets are shown in Fig.'s 277 – 279. The targets are labelled in the images and ranked based on size (top row = largest, bottom row = smallest, moving left to right).

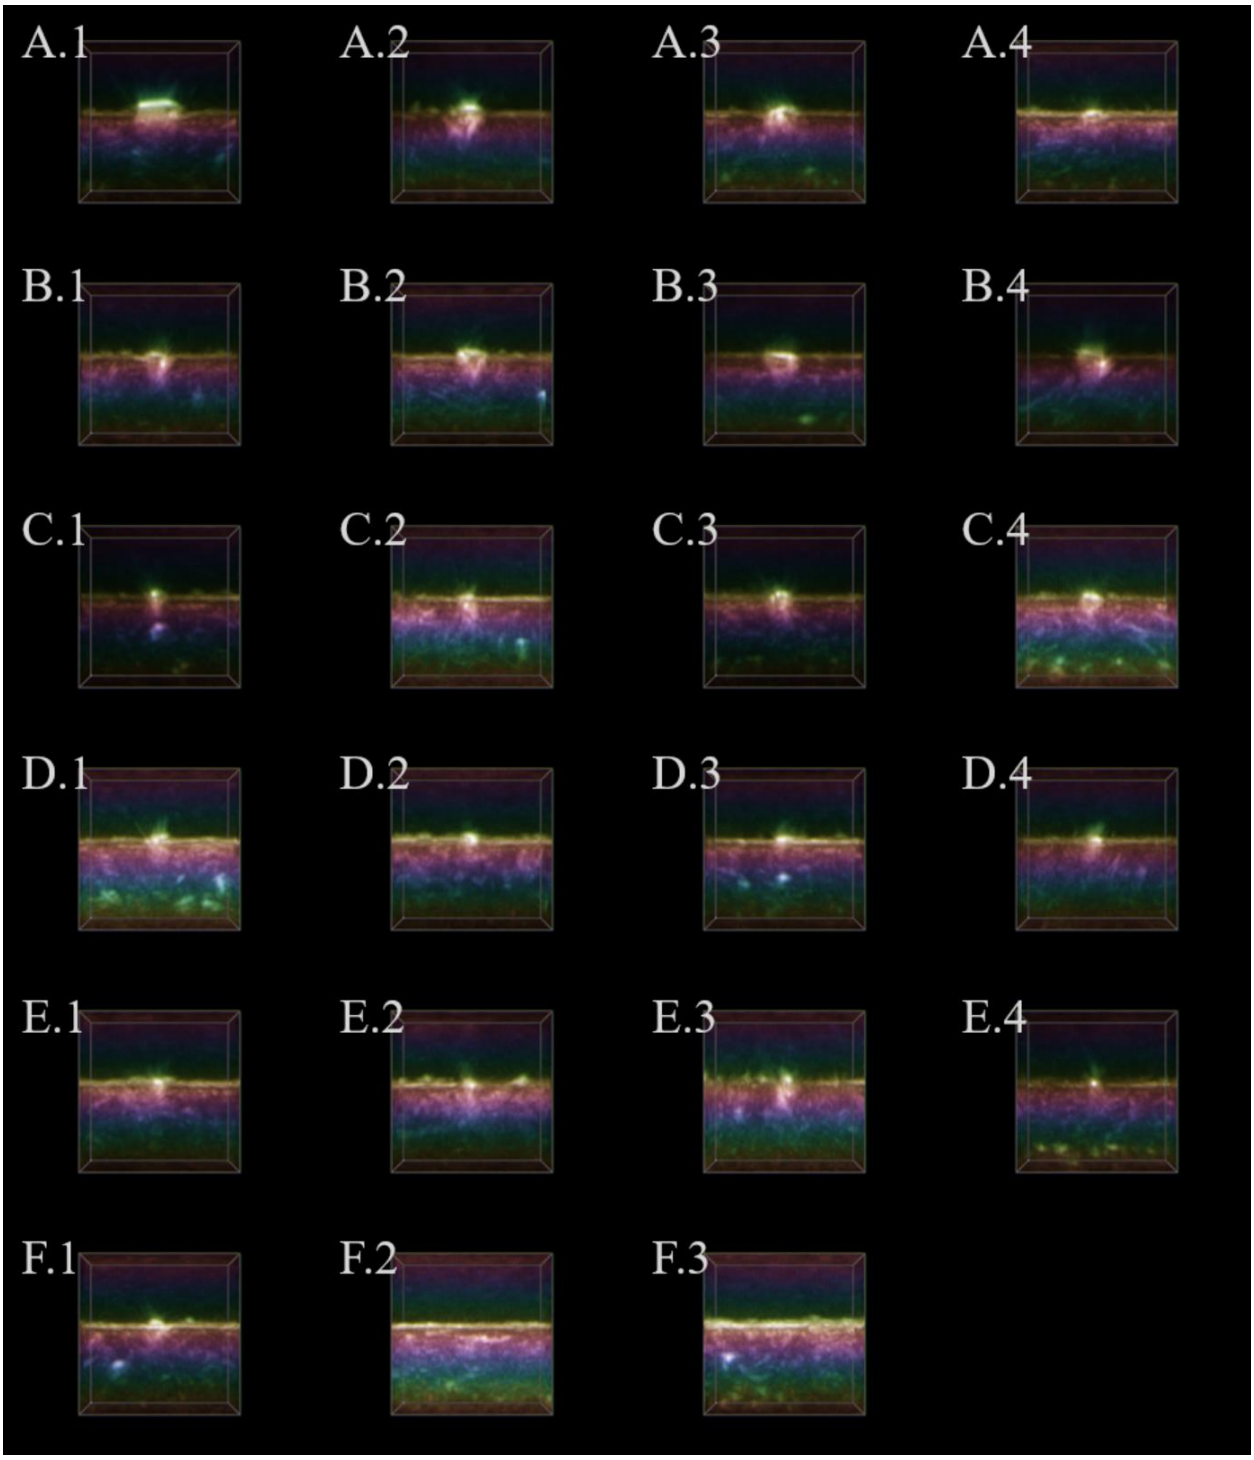

Figure 277: Color-by-depth MIP's (side view) through target ROI's for the Sequim Bay UXO testbed. A.1 – A.2: 155mm Howitzer shells, A.3 – B.1: 105mm HEAT rounds. B.2 – C.1: 105mm M60 shells. C.2 – D.2: 81mm M821 finned mortar rounds. D.3 – F.1: 81mm M889A1 mortar rounds. F.2 – F.3: 60mm M49 rounds.

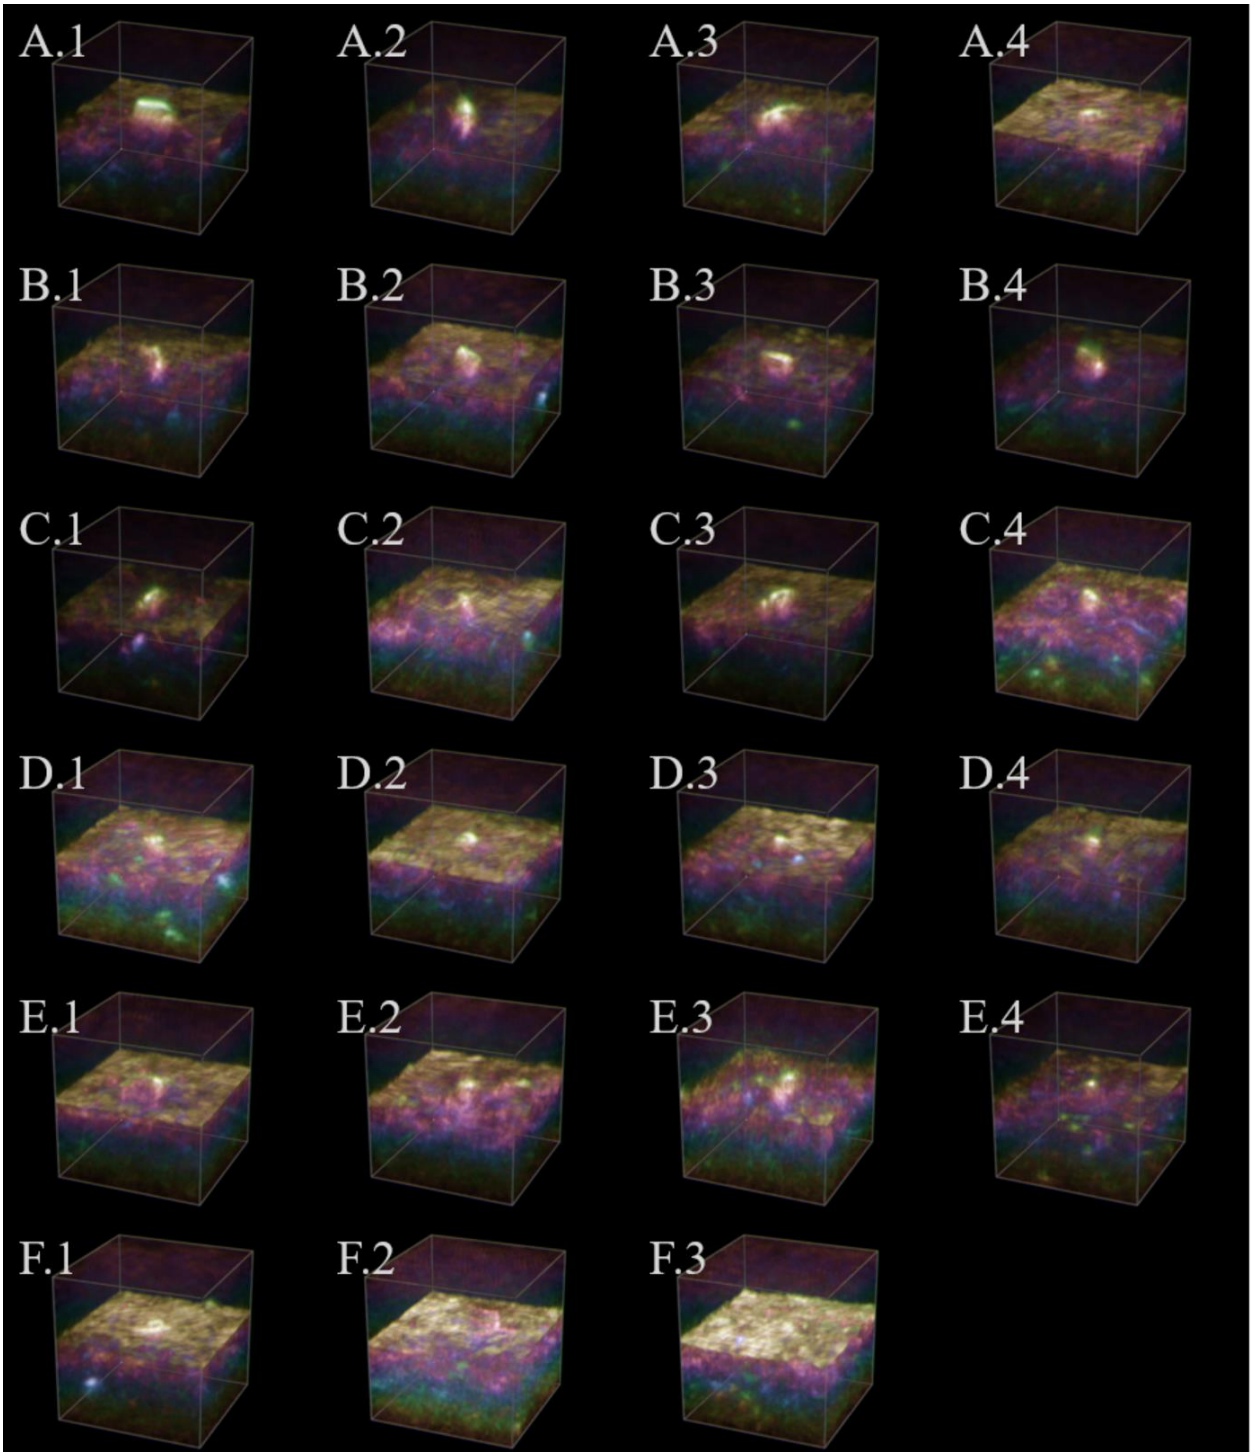

Figure 278: Color-by-depth MIP's (tilted) through target ROI's for the Sequim Bay UXO testbed. A.1 – A.2: 155mm Howitzer shells, A.3 – B.1: 105mm HEAT rounds. B.2 – C.1: 105mm M60 shells. C.2 – D.2: 81mm M821 finned mortar rounds. D.3 – F.1: 81mm M889A1 mortar rounds. F.2 – F.3: 60mm M49 rounds.

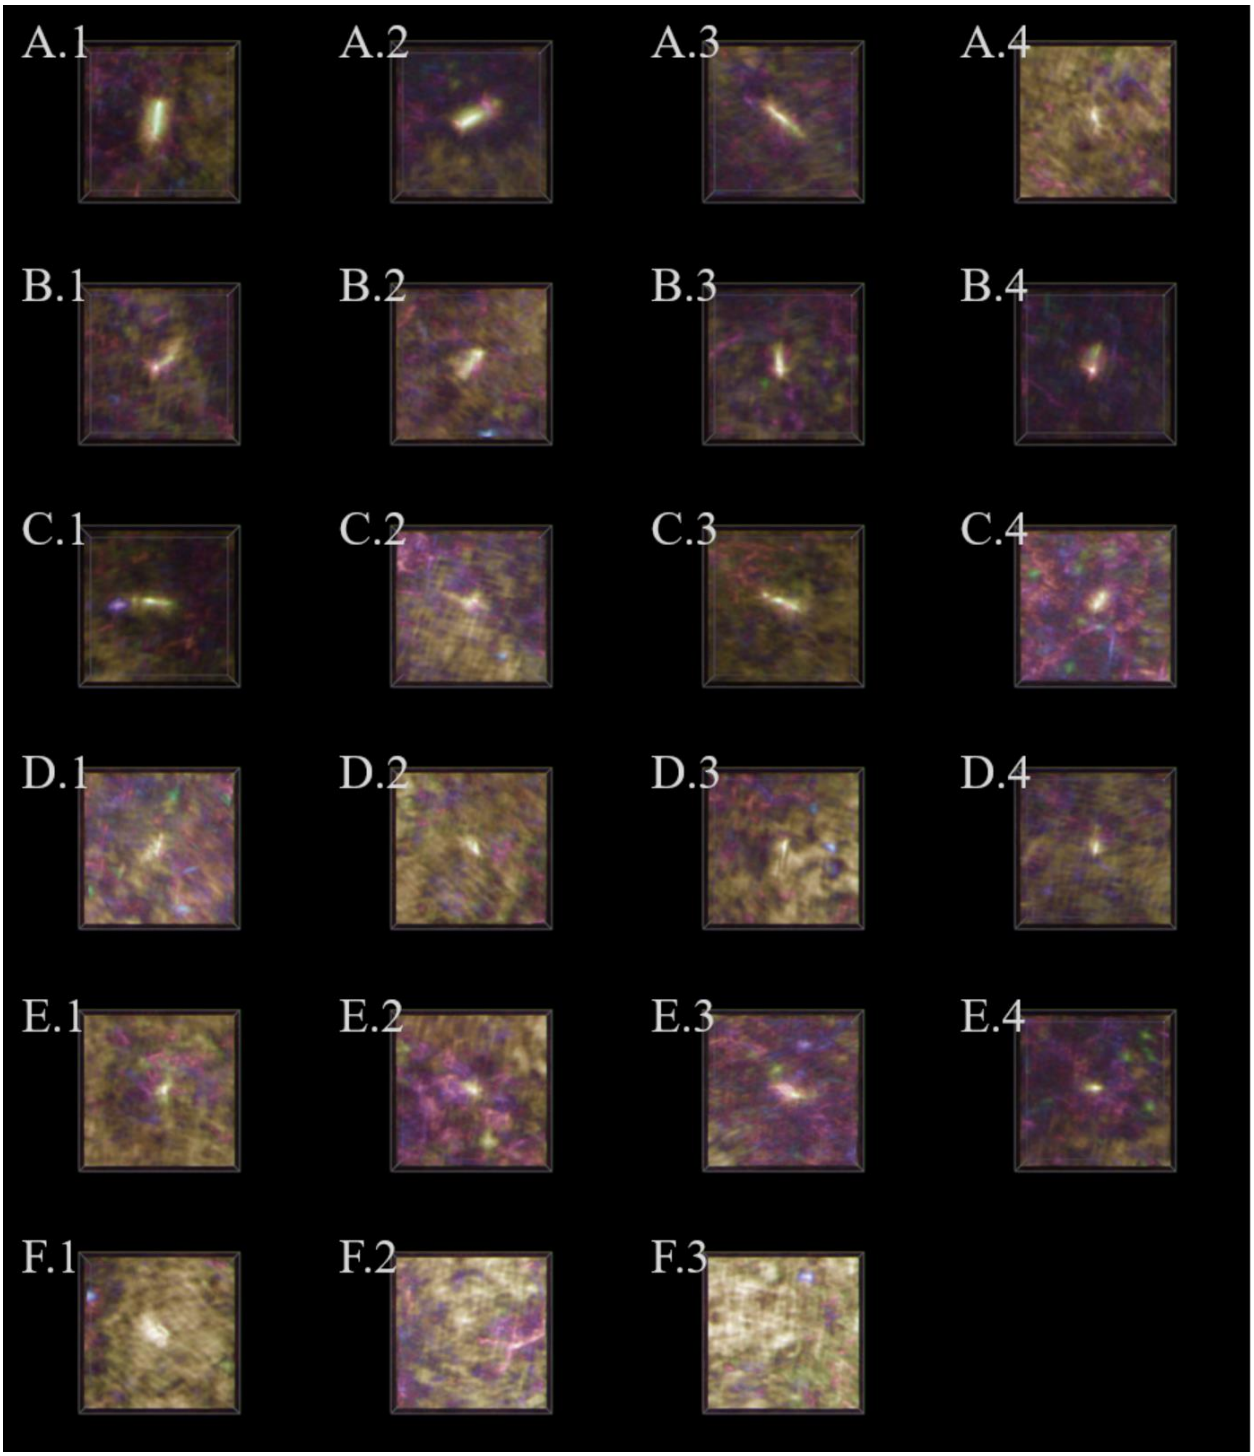

Figure 279: Color-by-depth MIP's (top) through target ROI's for the Sequim Bay UXO testbed. A.1 – A.2: 155mm Howitzer shells, A.3 – B.1: 105mm HEAT rounds. B.2 – C.1: 105mm M60 shells. C.2 – D.2: 81mm M821 finned mortar rounds. D.3 – F.1: 81mm M889A1 mortar rounds. F.2 – F.3: 60mm M49 rounds.

Note that in Fig. 279 the prominence of several of the targets (e.g. A.4, C.2, D.1) is diminished by competition in the MIP projection with nadir specular sediment glinting. This is due to the lack of any sort of glint filtering being applied to the multi-look ROI's. The following figures

show a composite of ROI views. The upper figure shows the complete set of looks through the Cartesian and cylindrical coordinate complex volume magnitudes. The bottom window of the top plot shows the unfiltered and filtered spatial (cylindrical) and acoustic color datasets. The row of plots on the bottom show rotating MIP views through the composite set of color-by-aspect (CBA), color-by-depth (CBD) and composite filtered acoustic-color spectra. This is done for each testbed target, in the same order as plotted in Fig.'s 277 – 279.

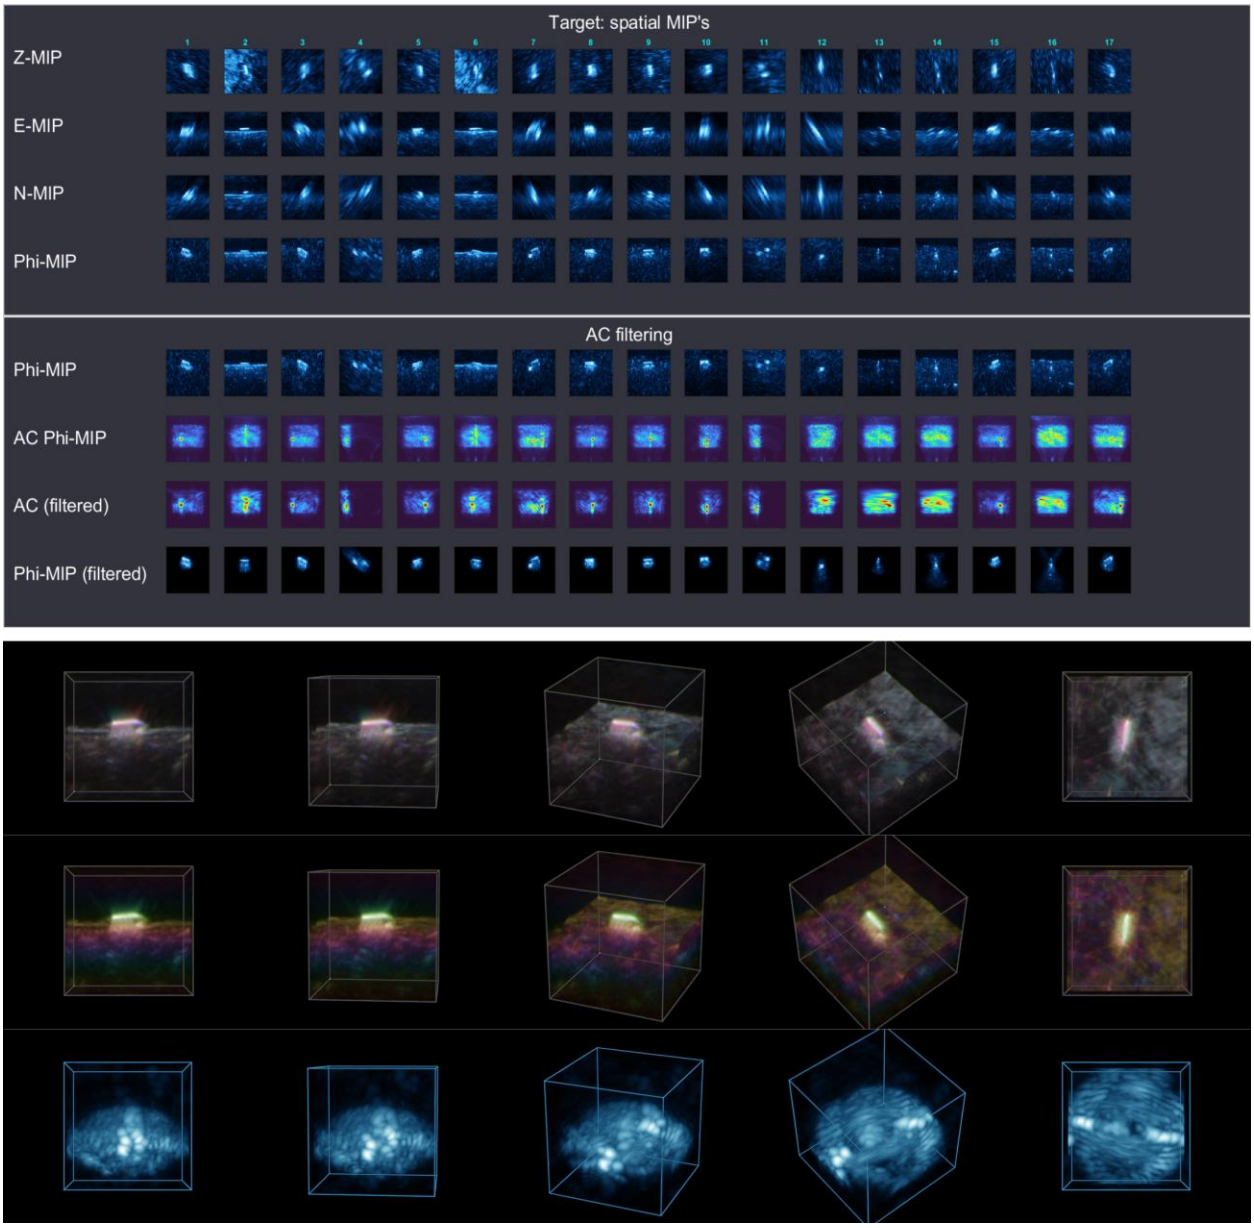

Figure 280: 155mm Howitzer shell facsimile. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum.

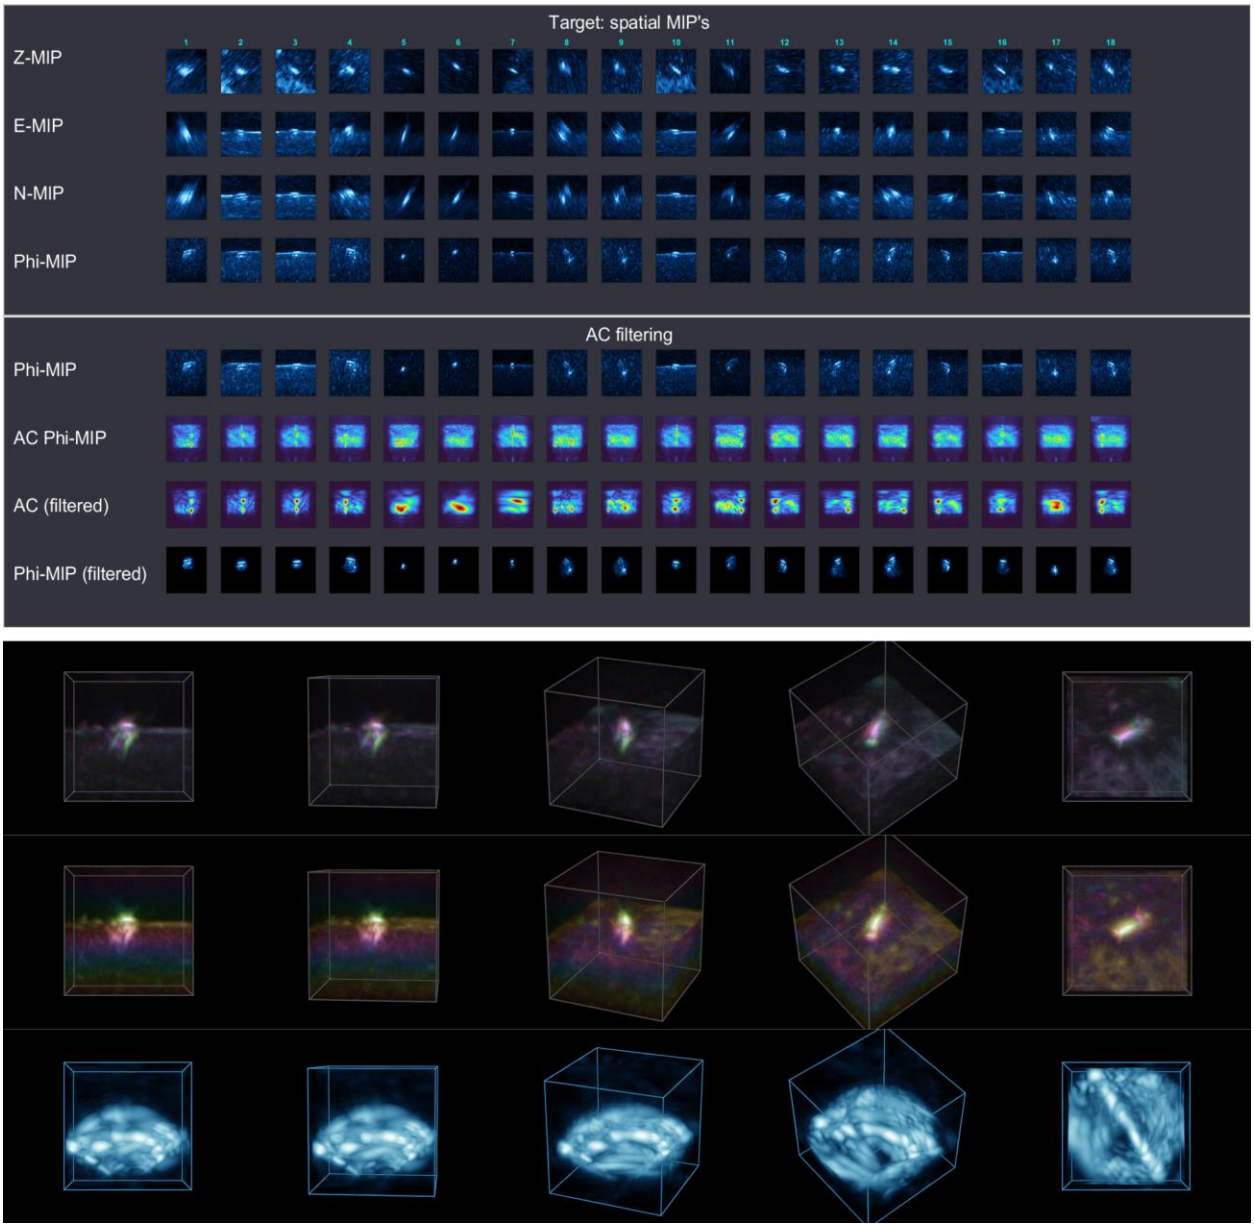

Figure 281: 155mm Howitzer shell. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum.

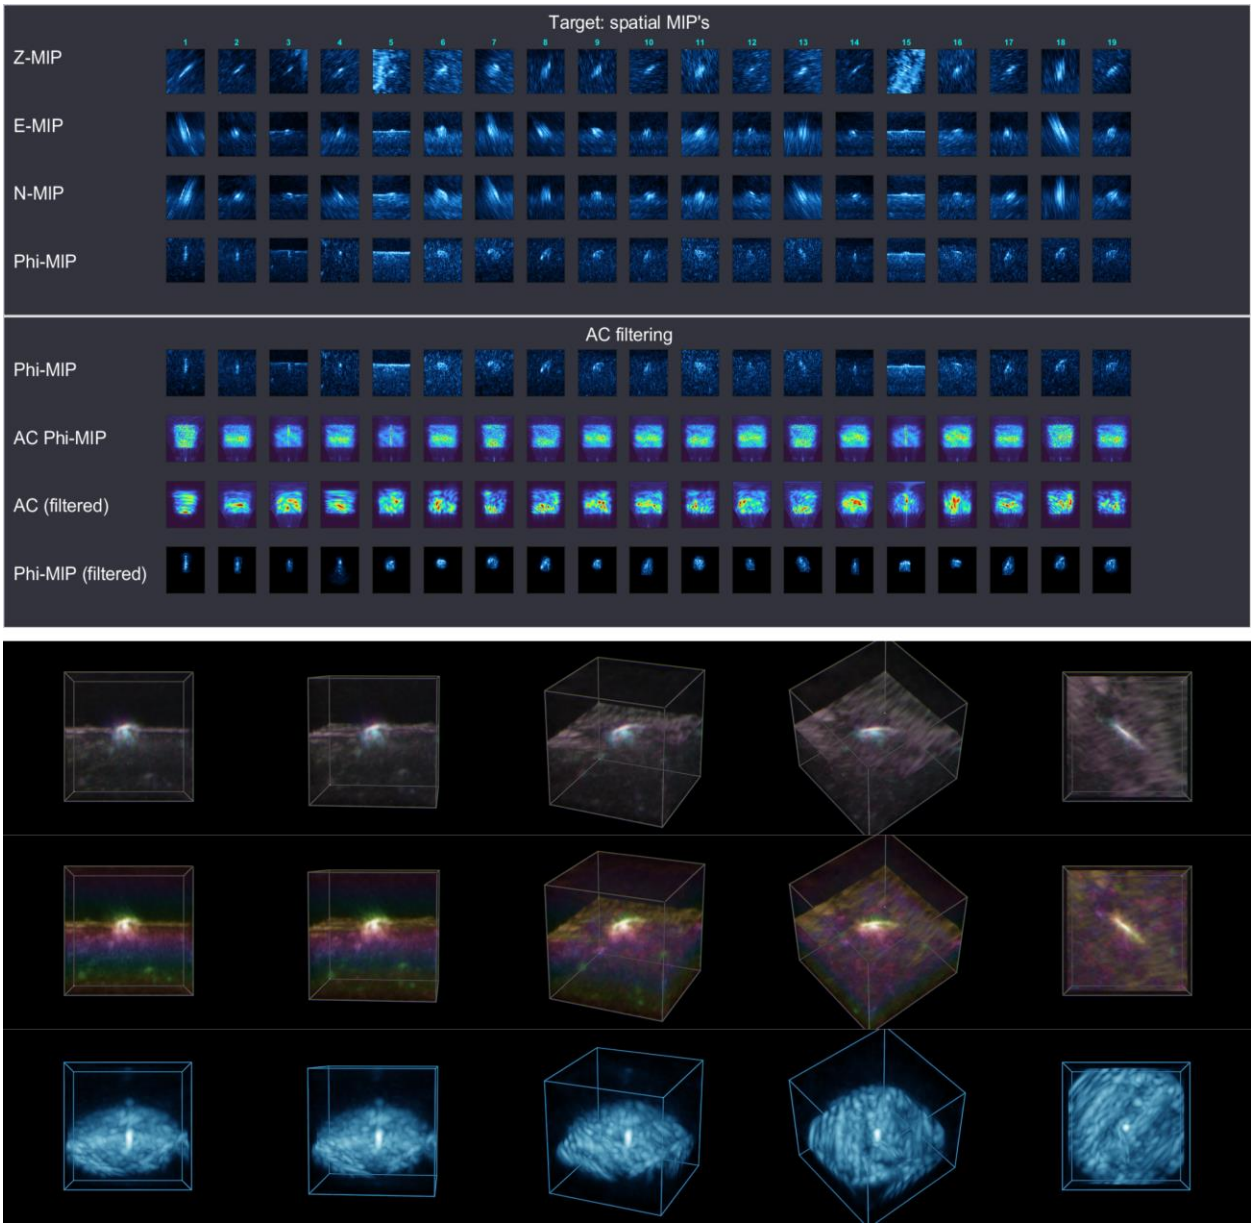

Figure 282: 105mm HEAT round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum.

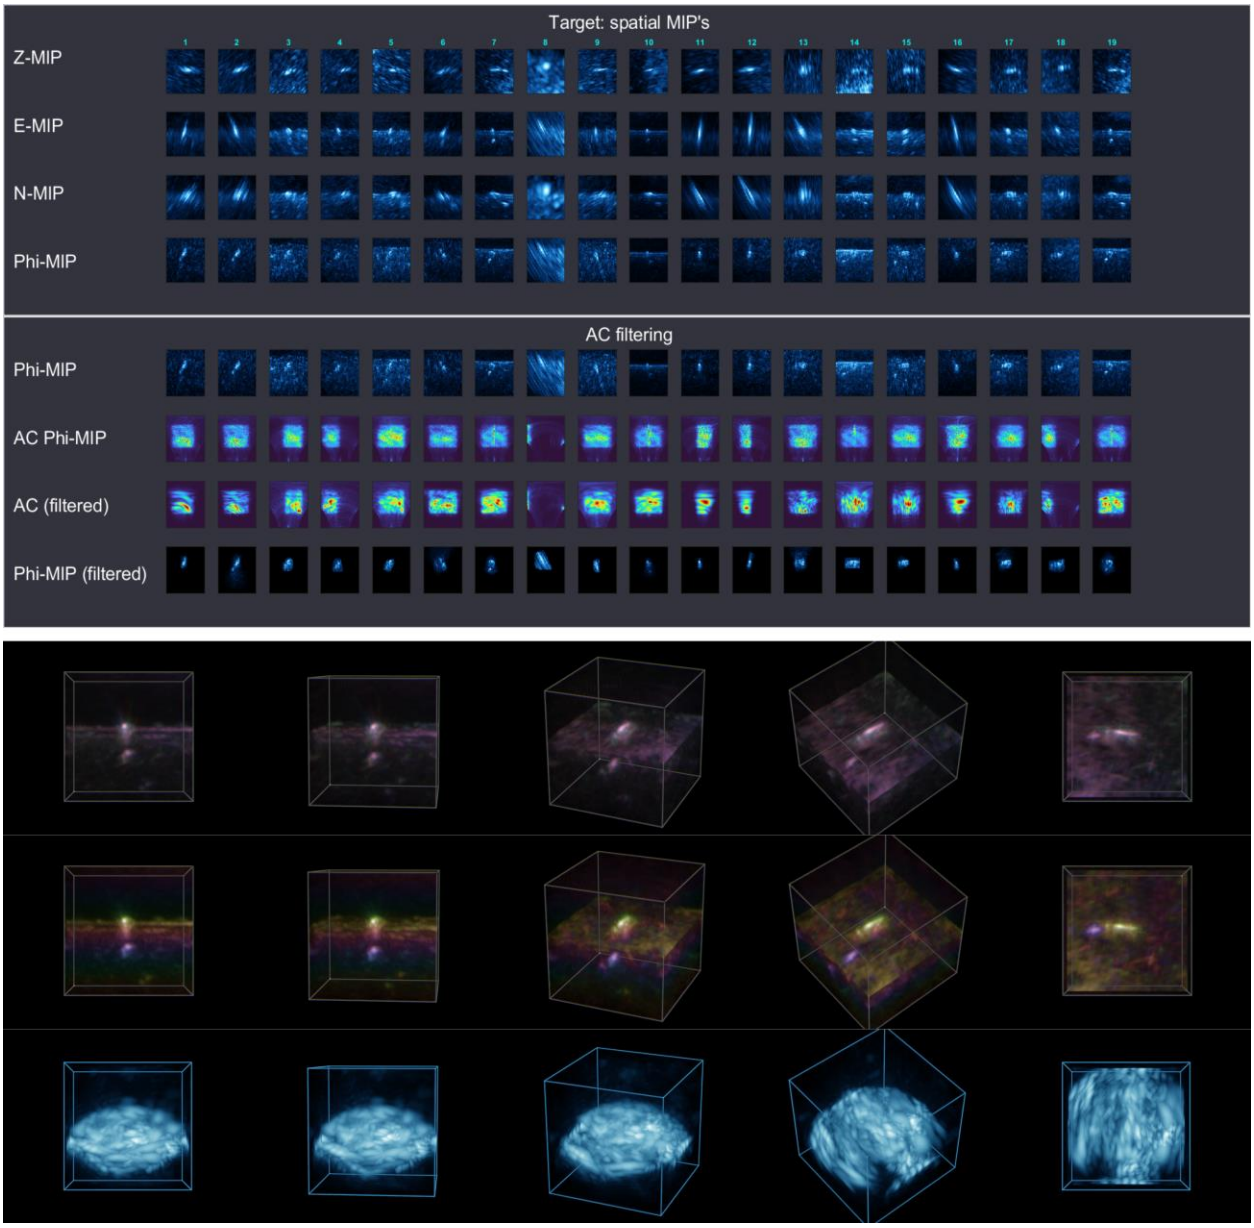

Figure 283: 105mm HEAT round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum.

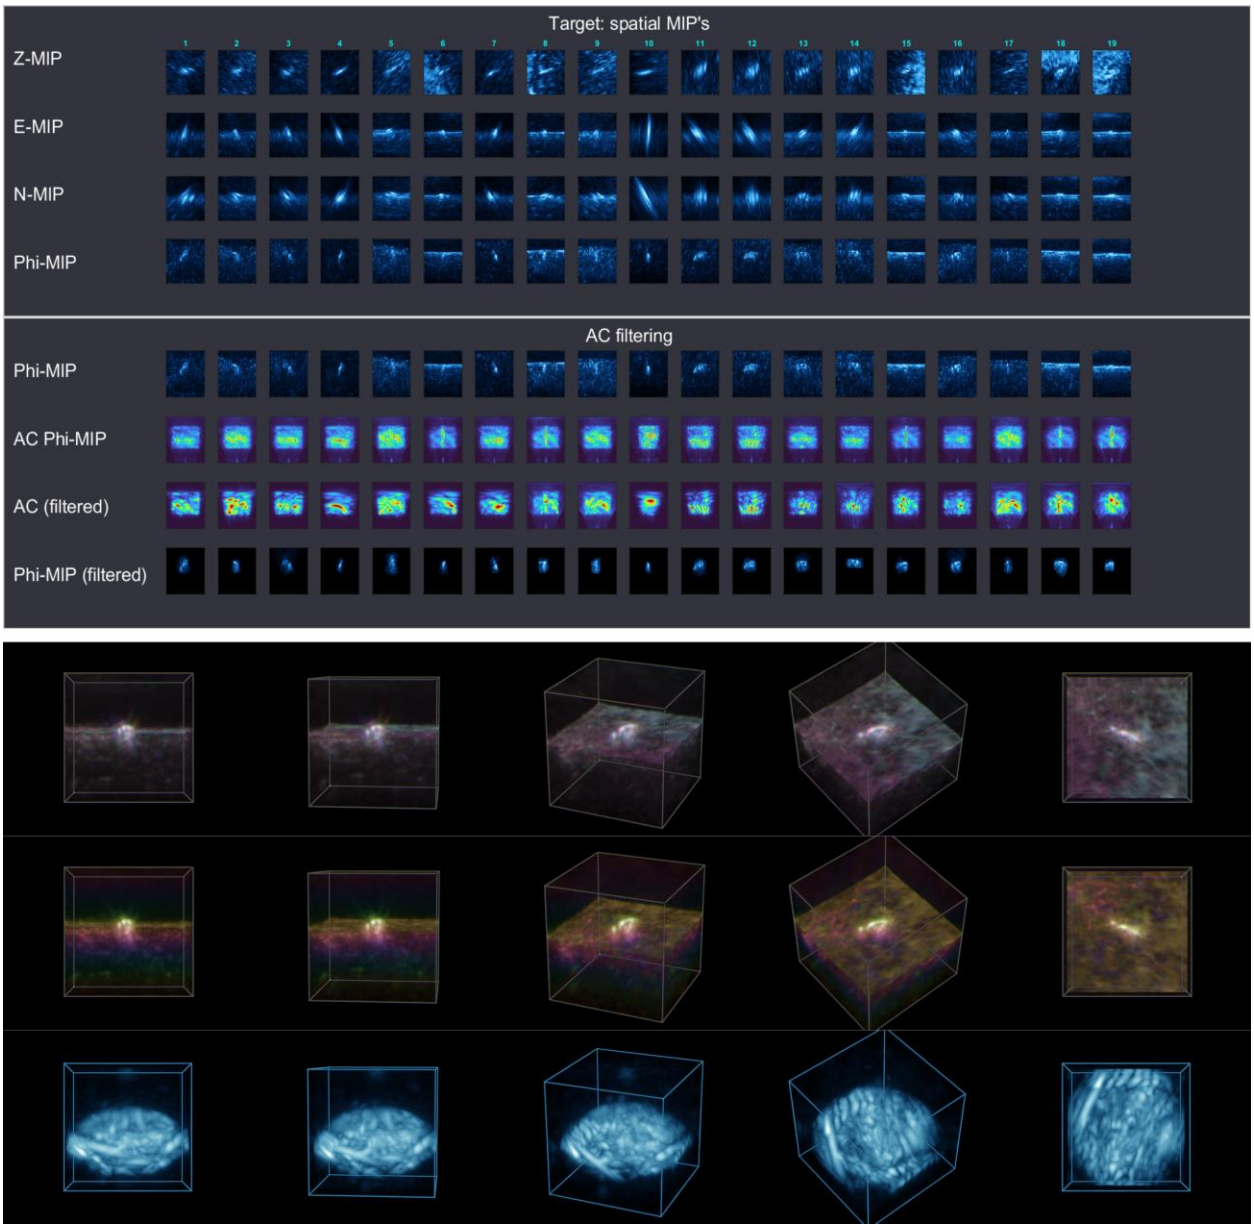

Figure 284: 105mm HEAT round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum.

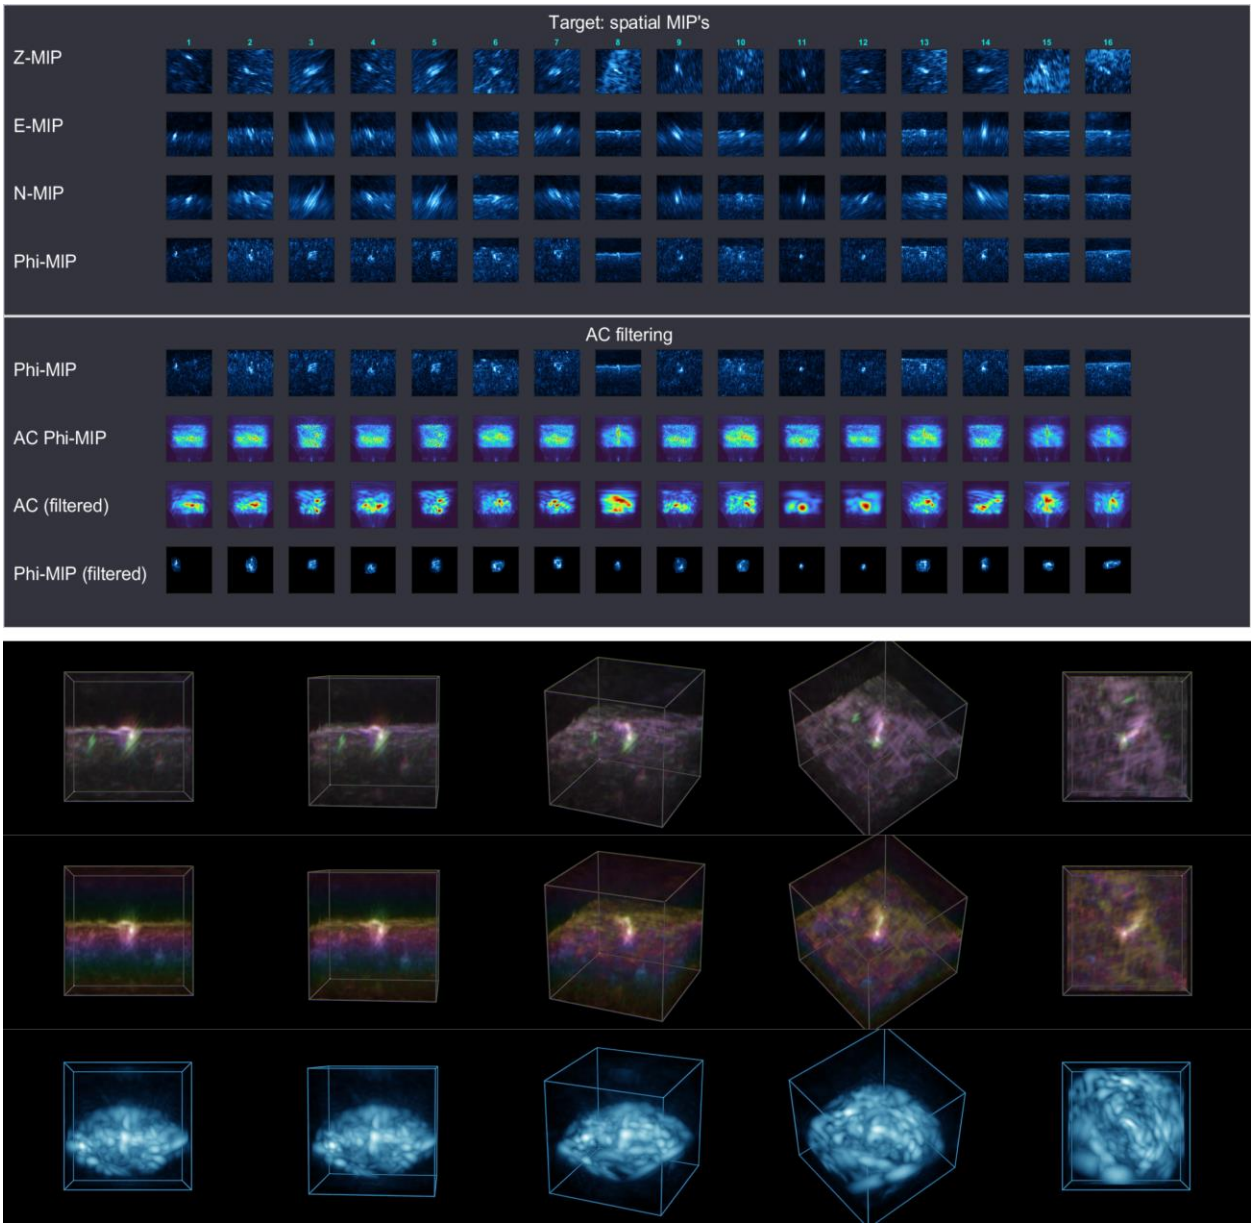

Figure 285: 105mm M60 shell. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum.

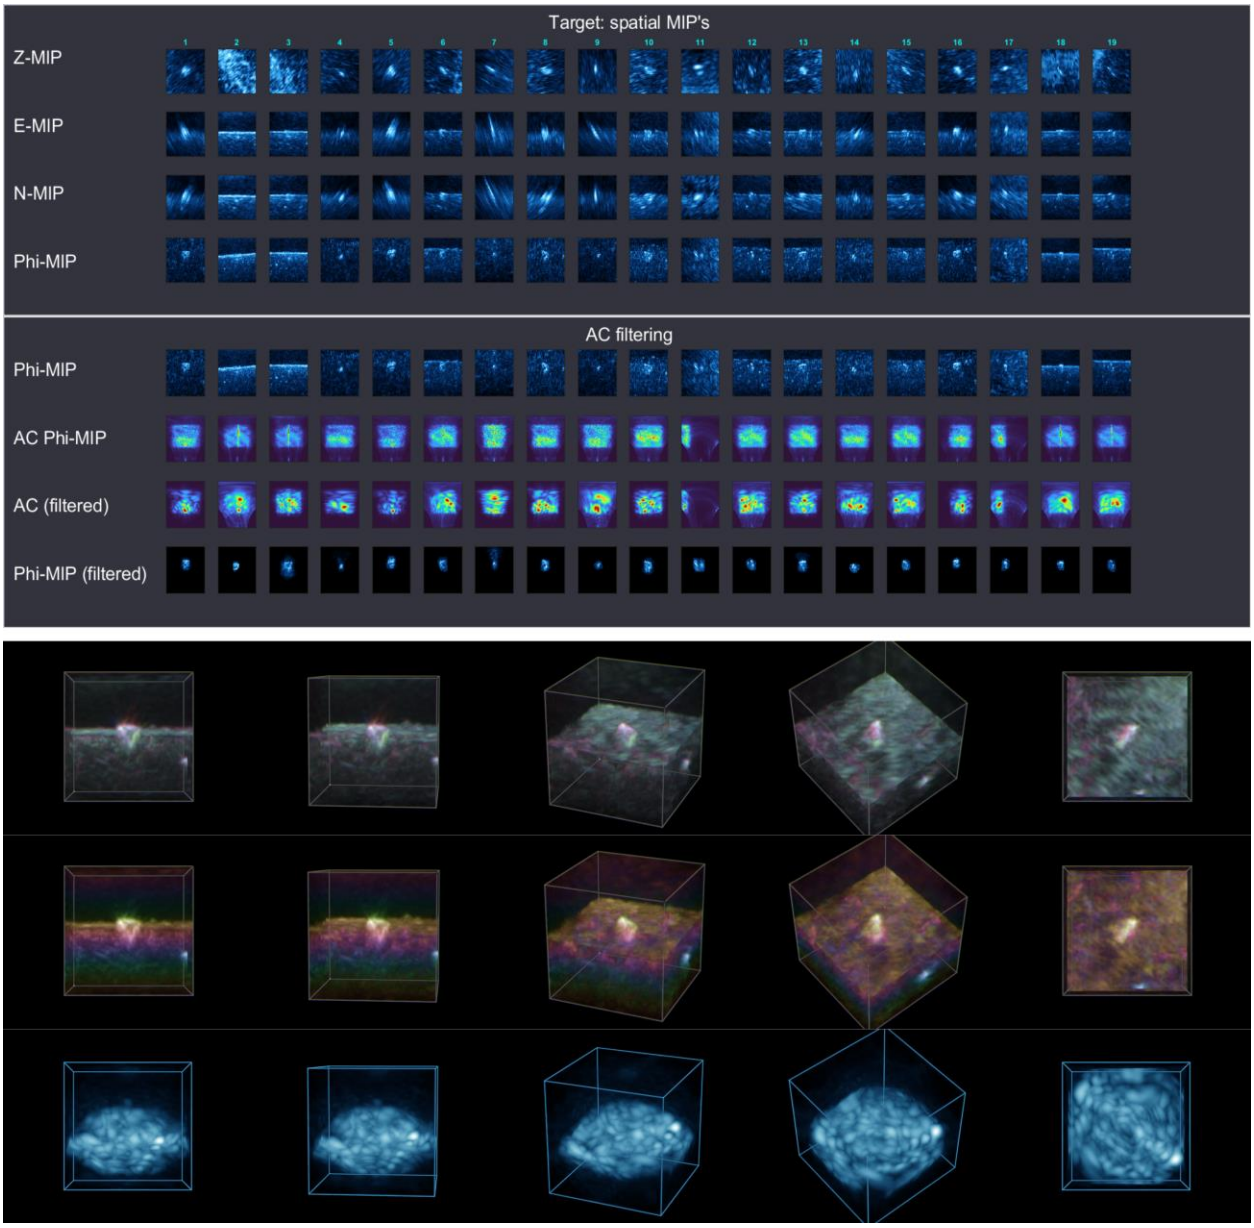

Figure 286: 105mm M60 shell. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum.

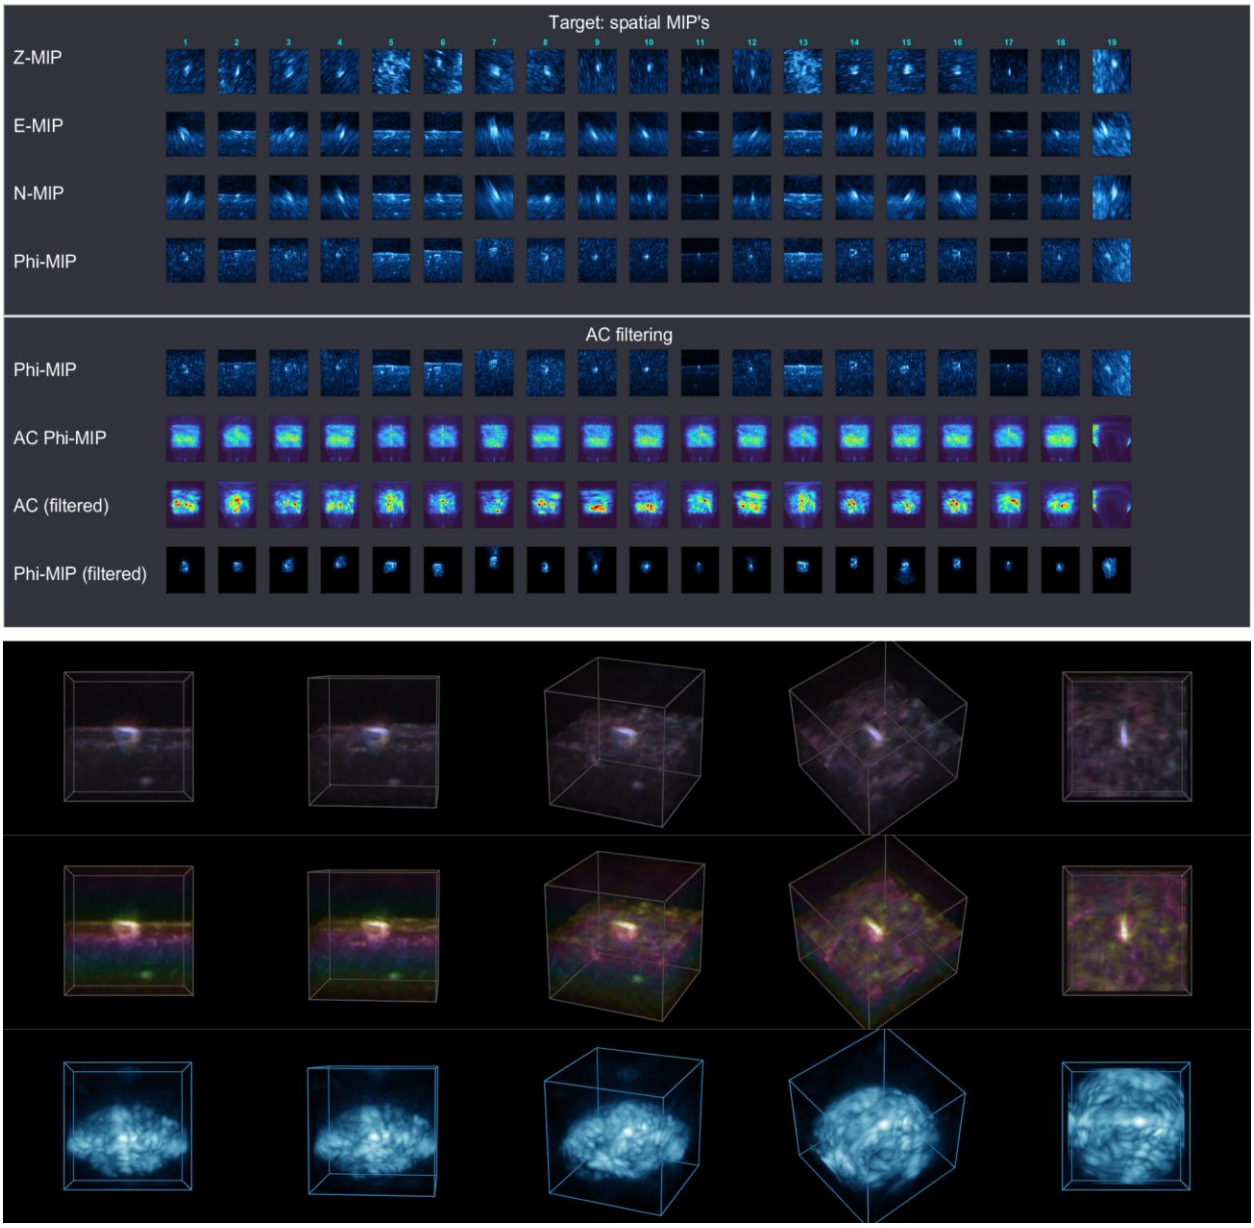

Figure 287: 105mm M60 shell. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum.

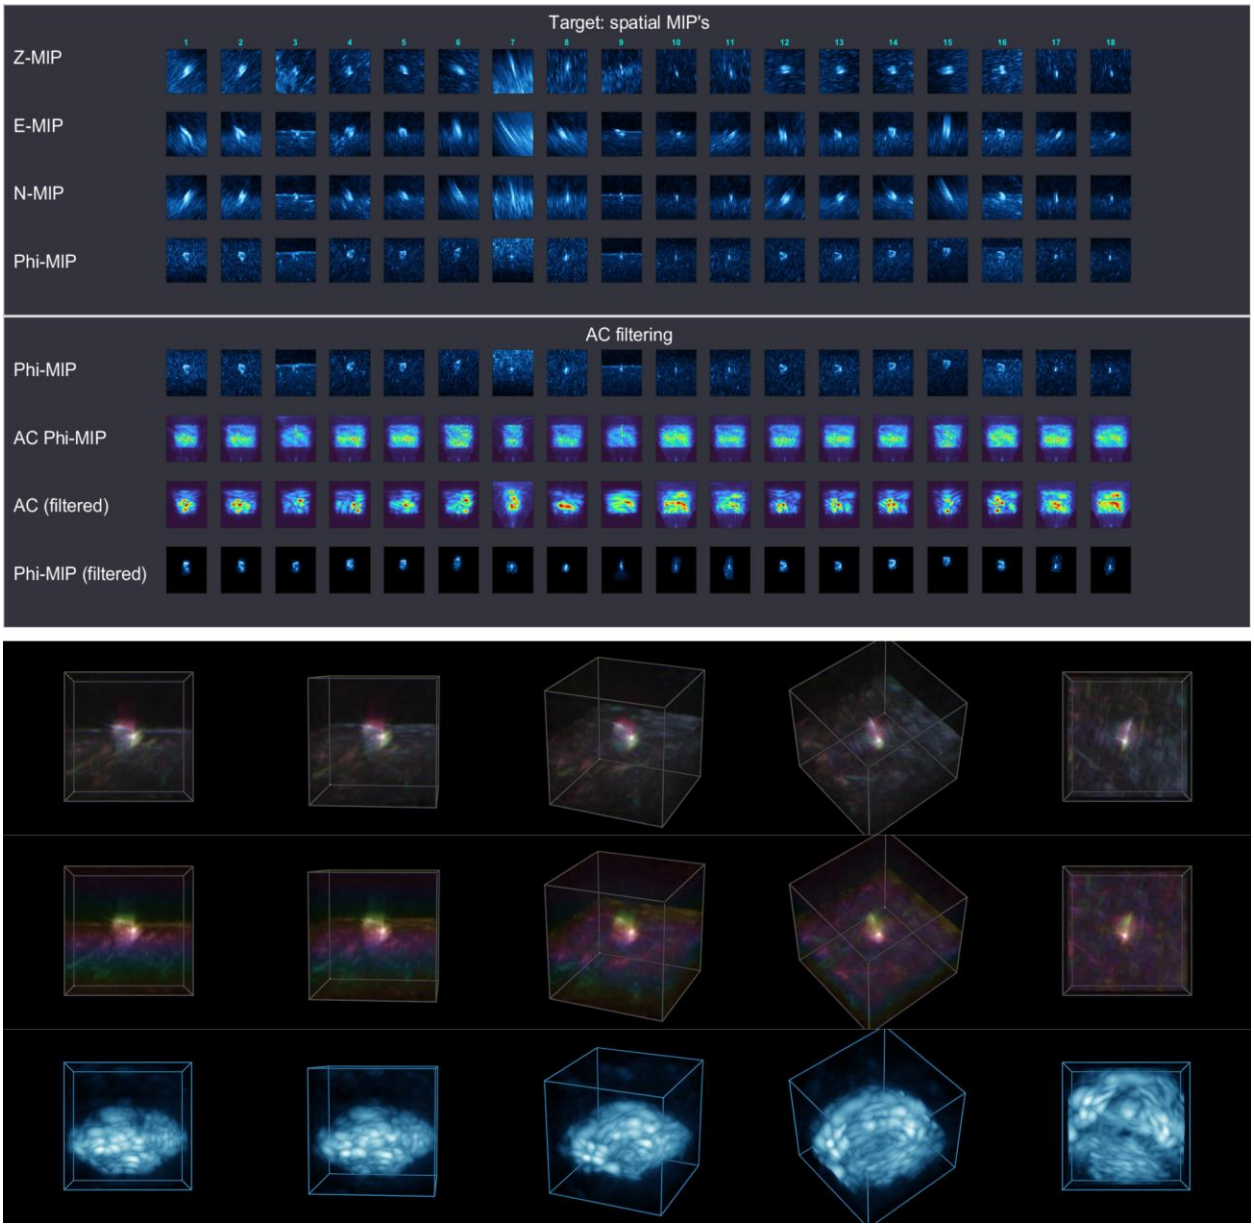

Figure 288: 105mm M60 shell. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum.

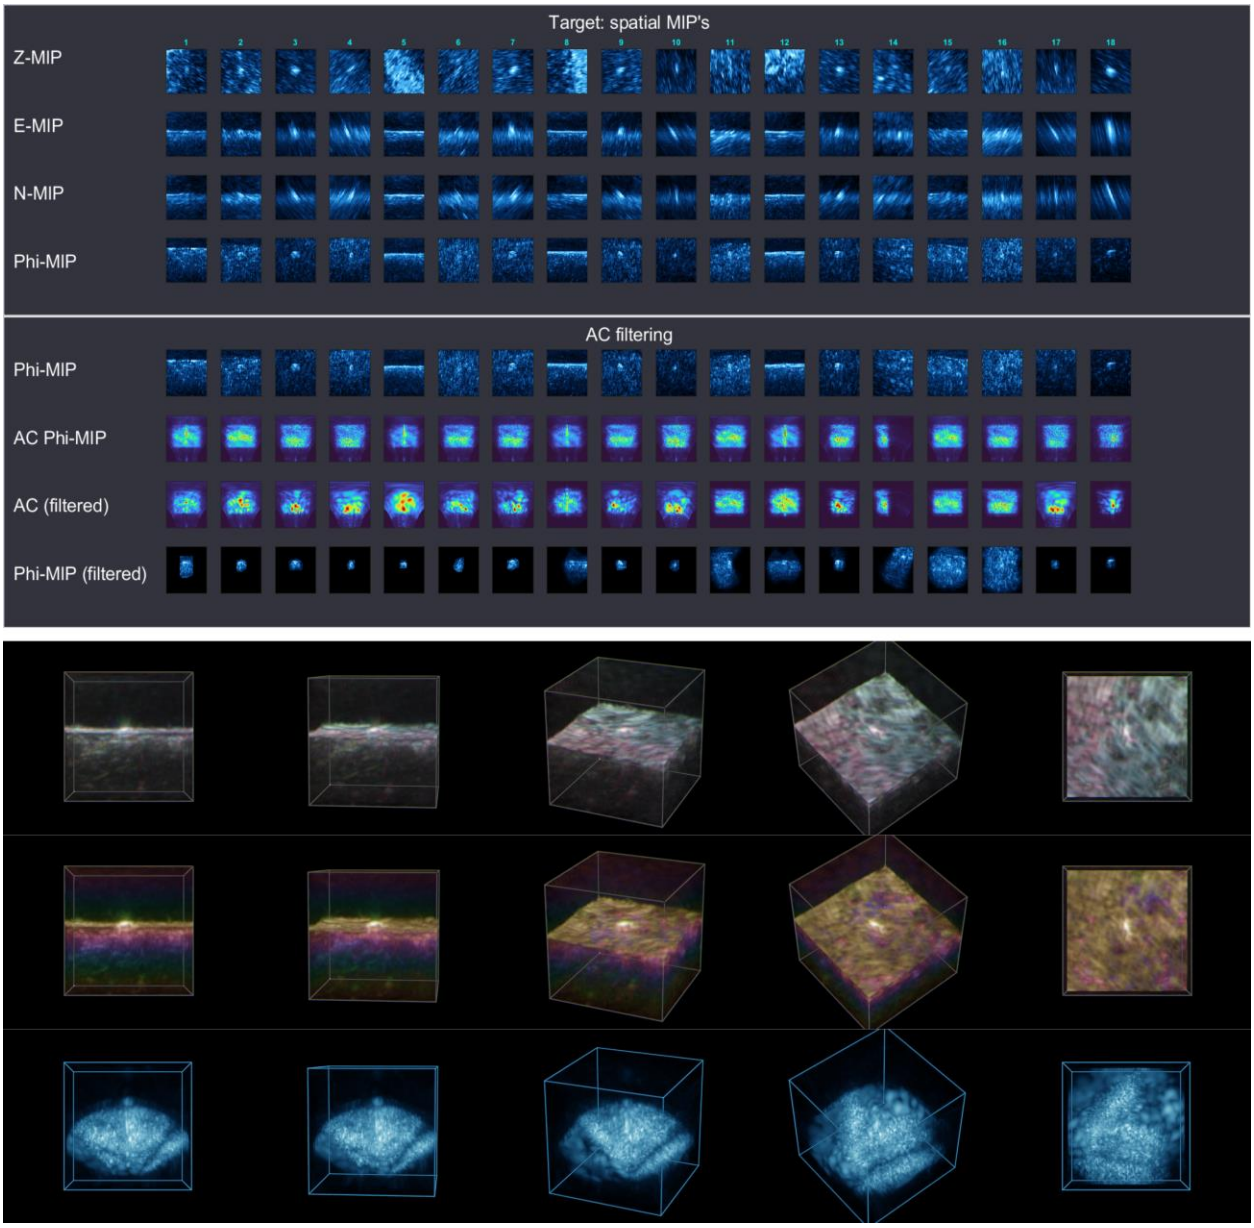

Figure 289: 81mm M821 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum.

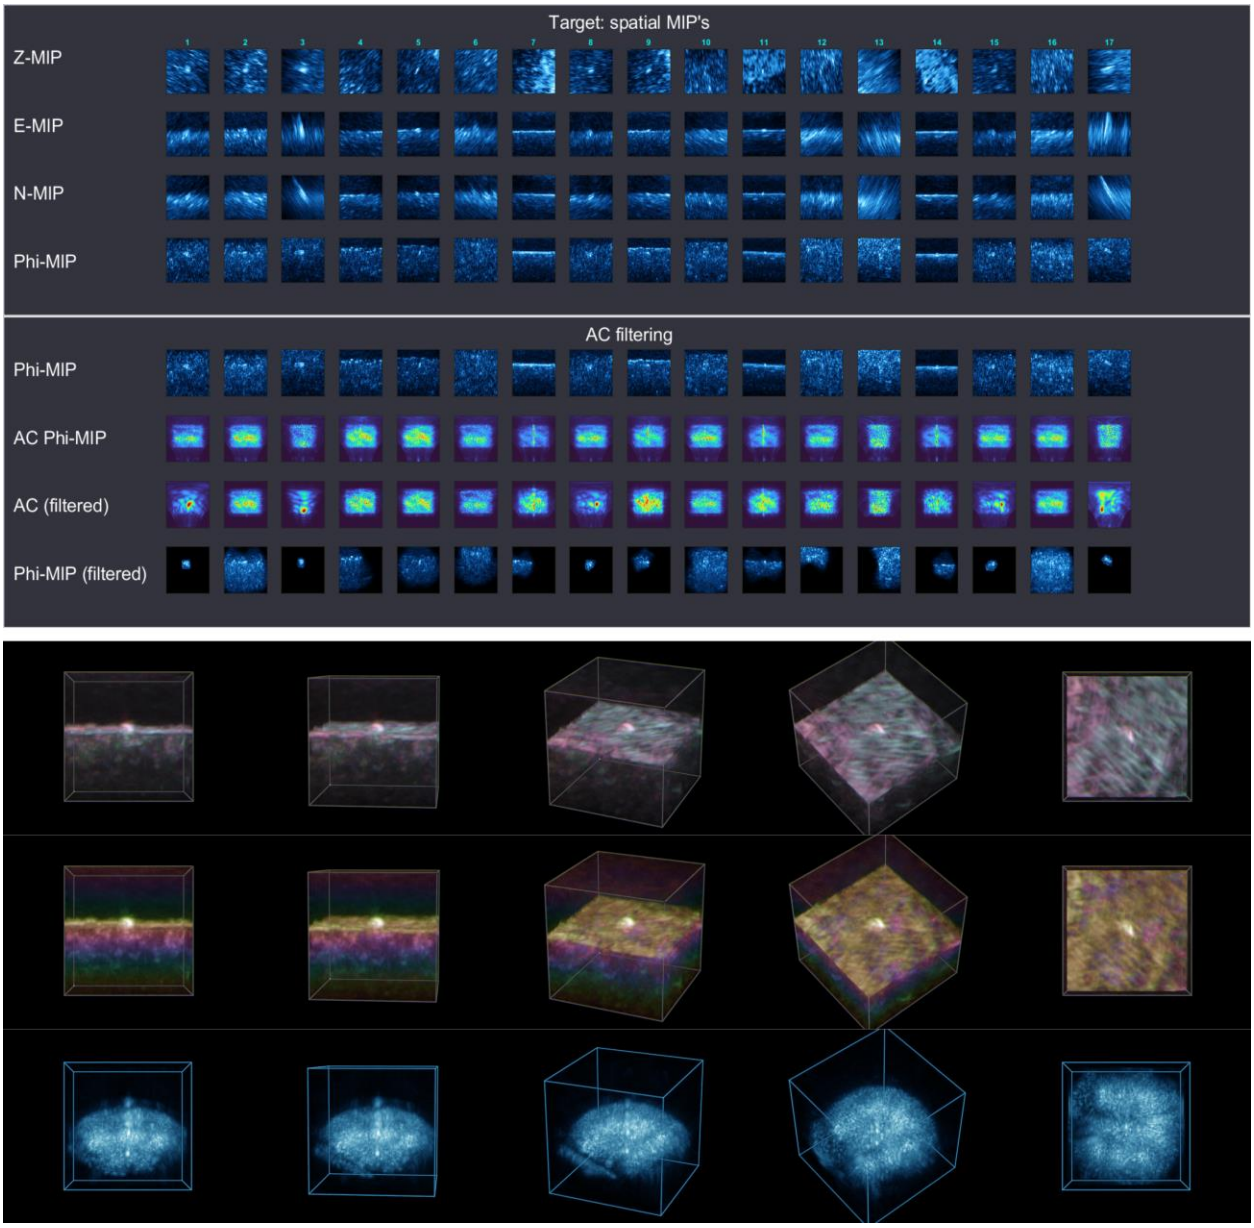

Figure 290: 81mm M821 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum.

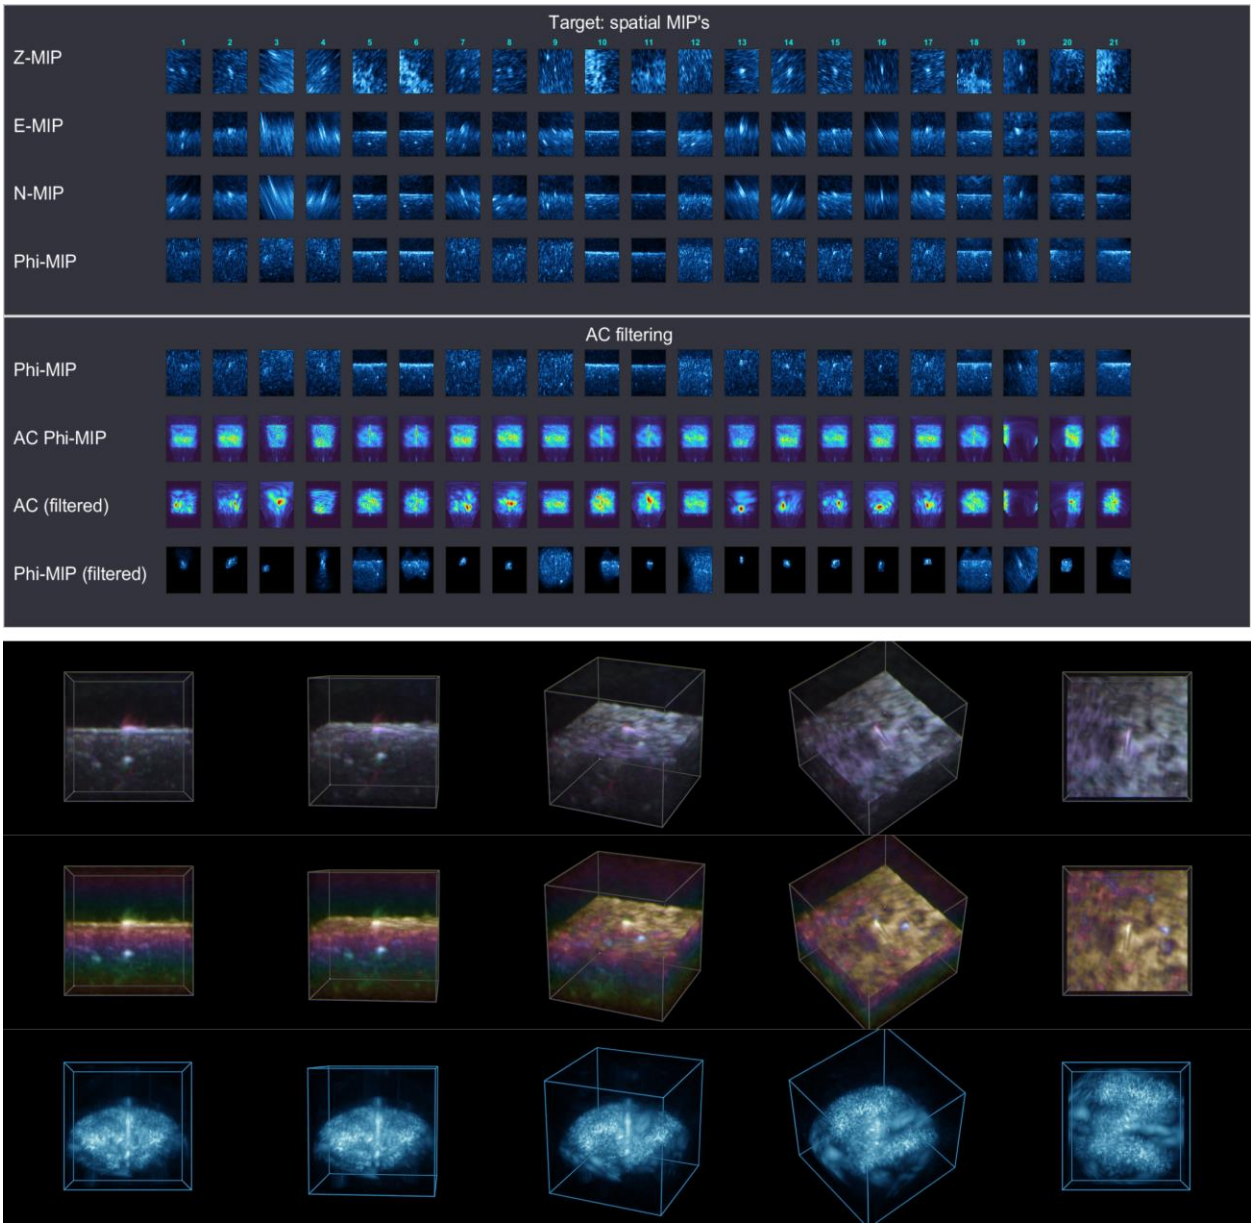

Figure 291: 81mm M821 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum.

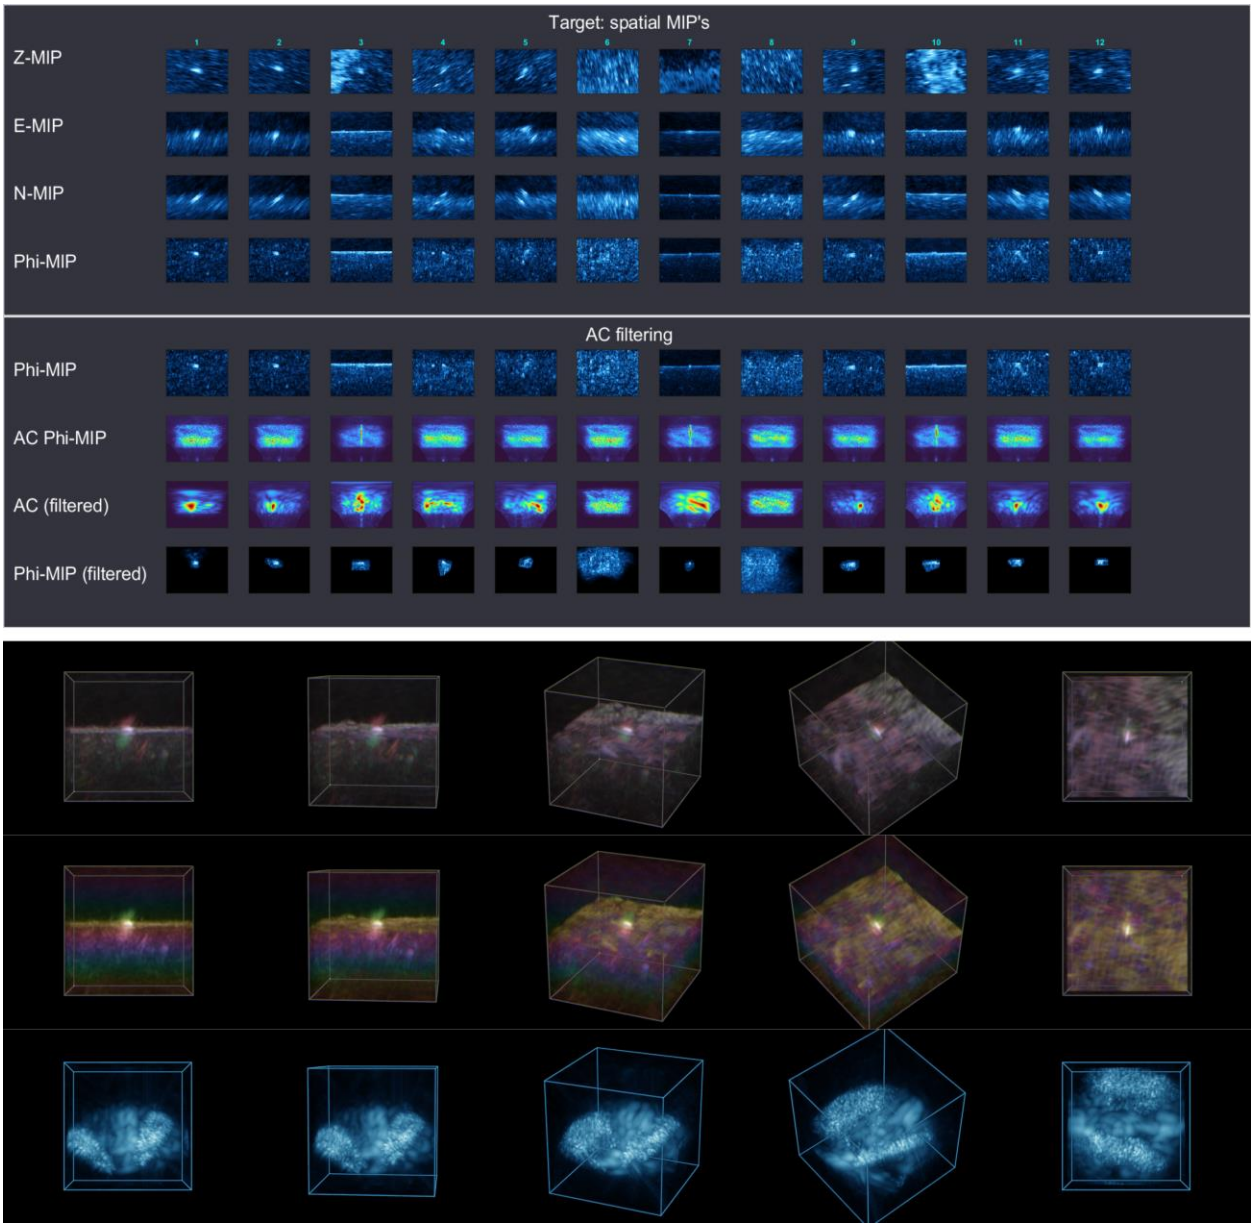

Figure 292: 81mm M821 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum.

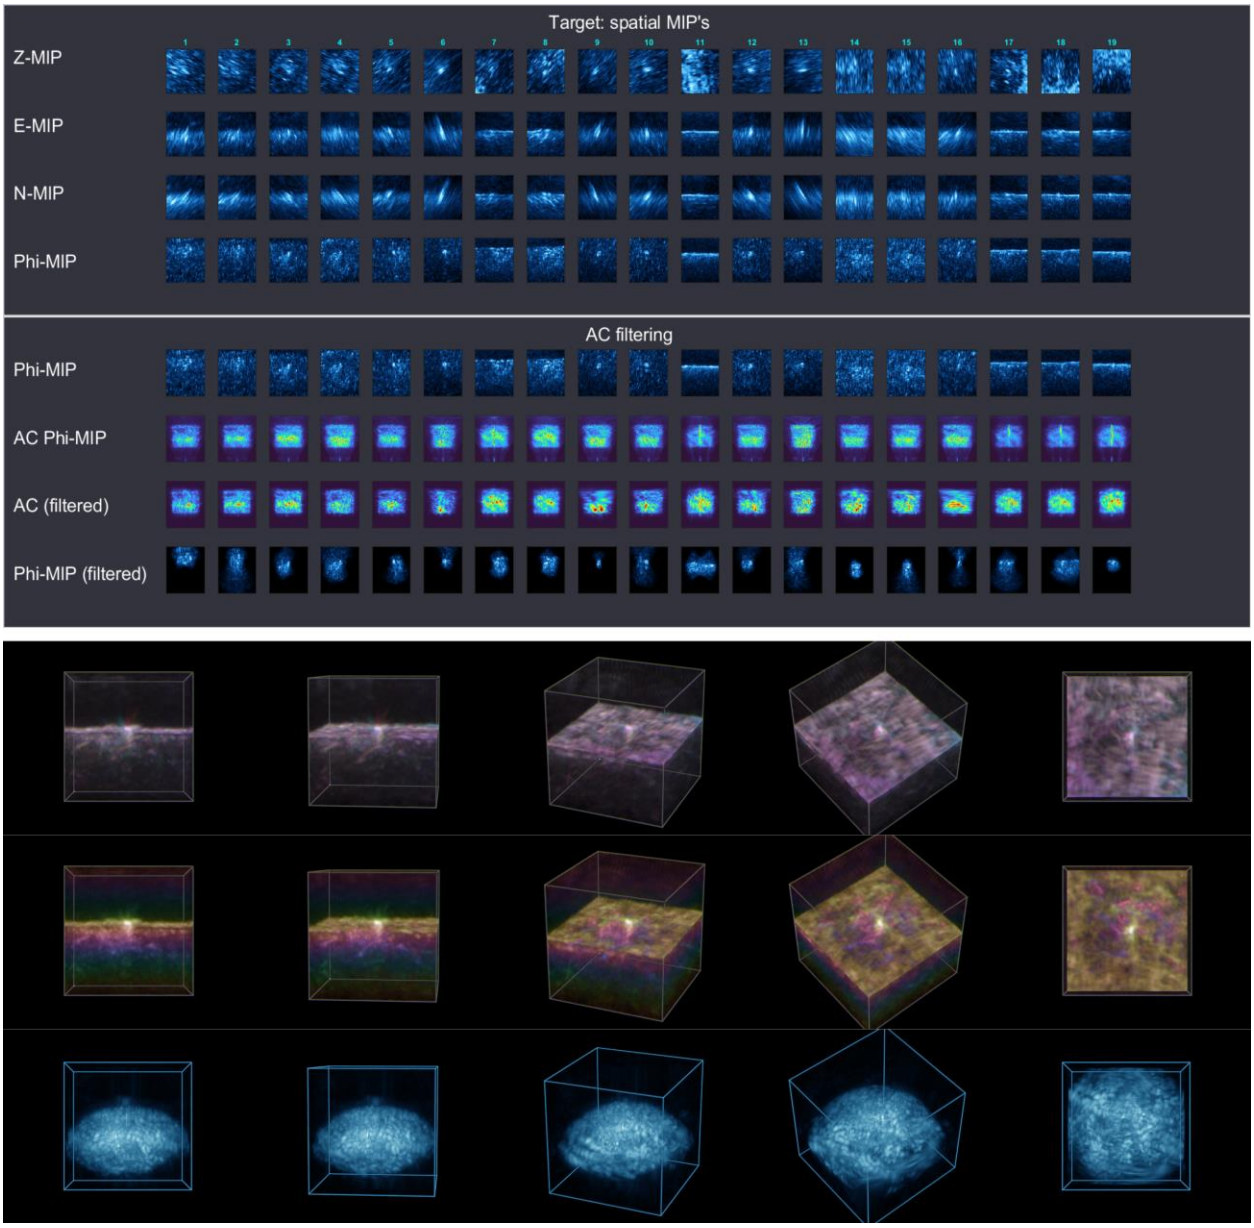

Figure 293: 81mm M821 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum.

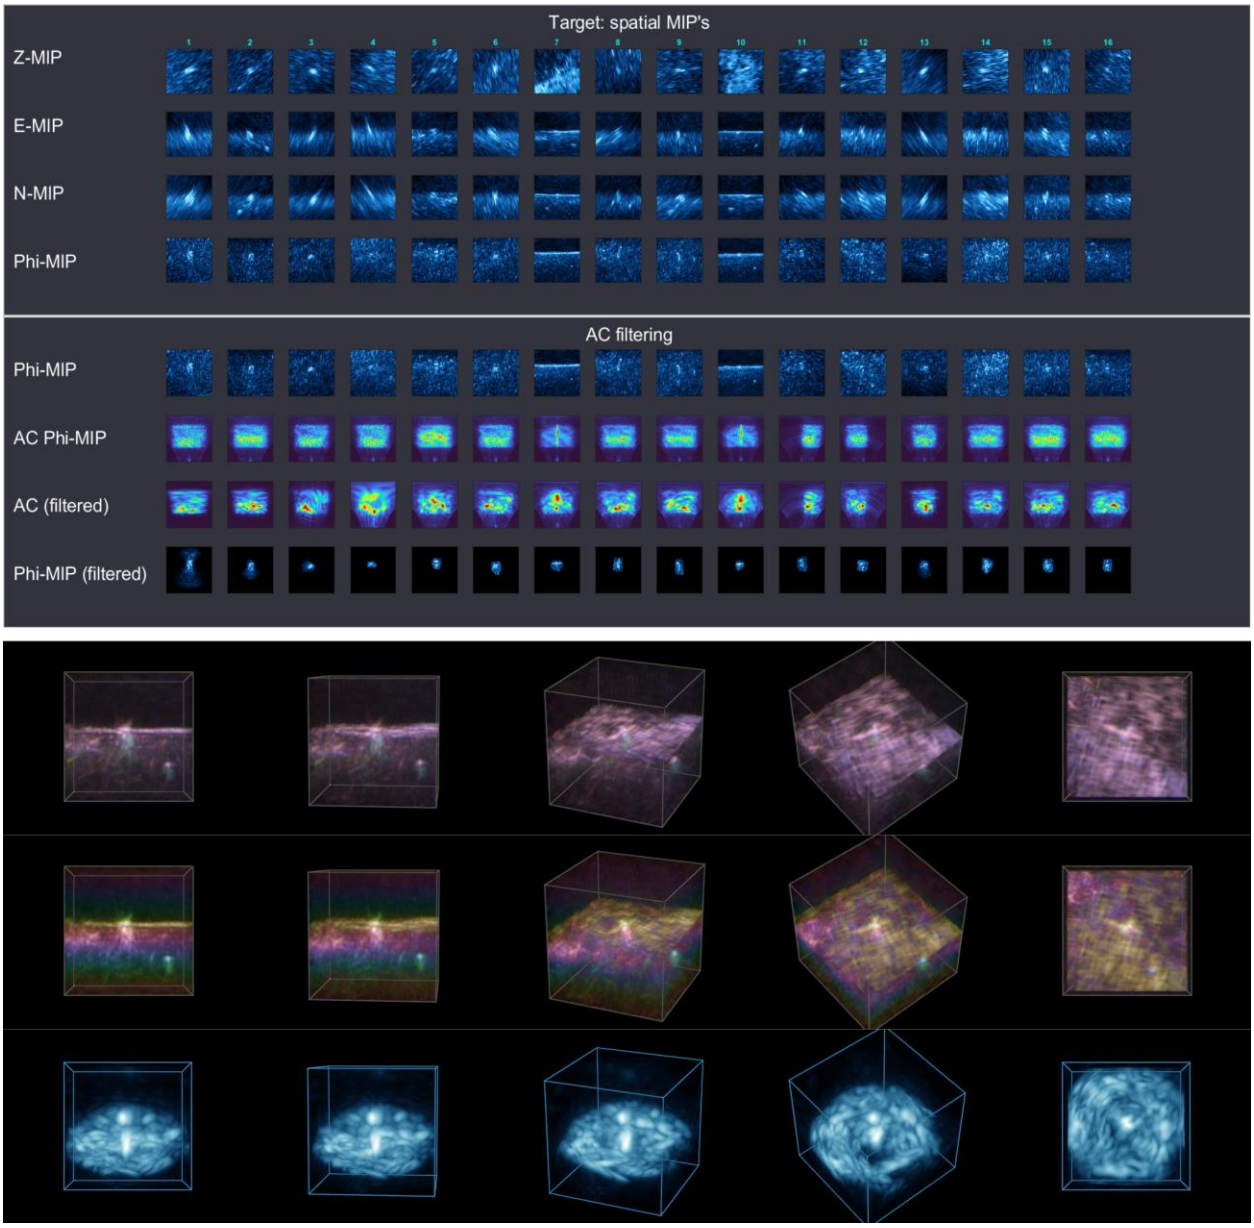

Figure 294: 81mm M889A1 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum.

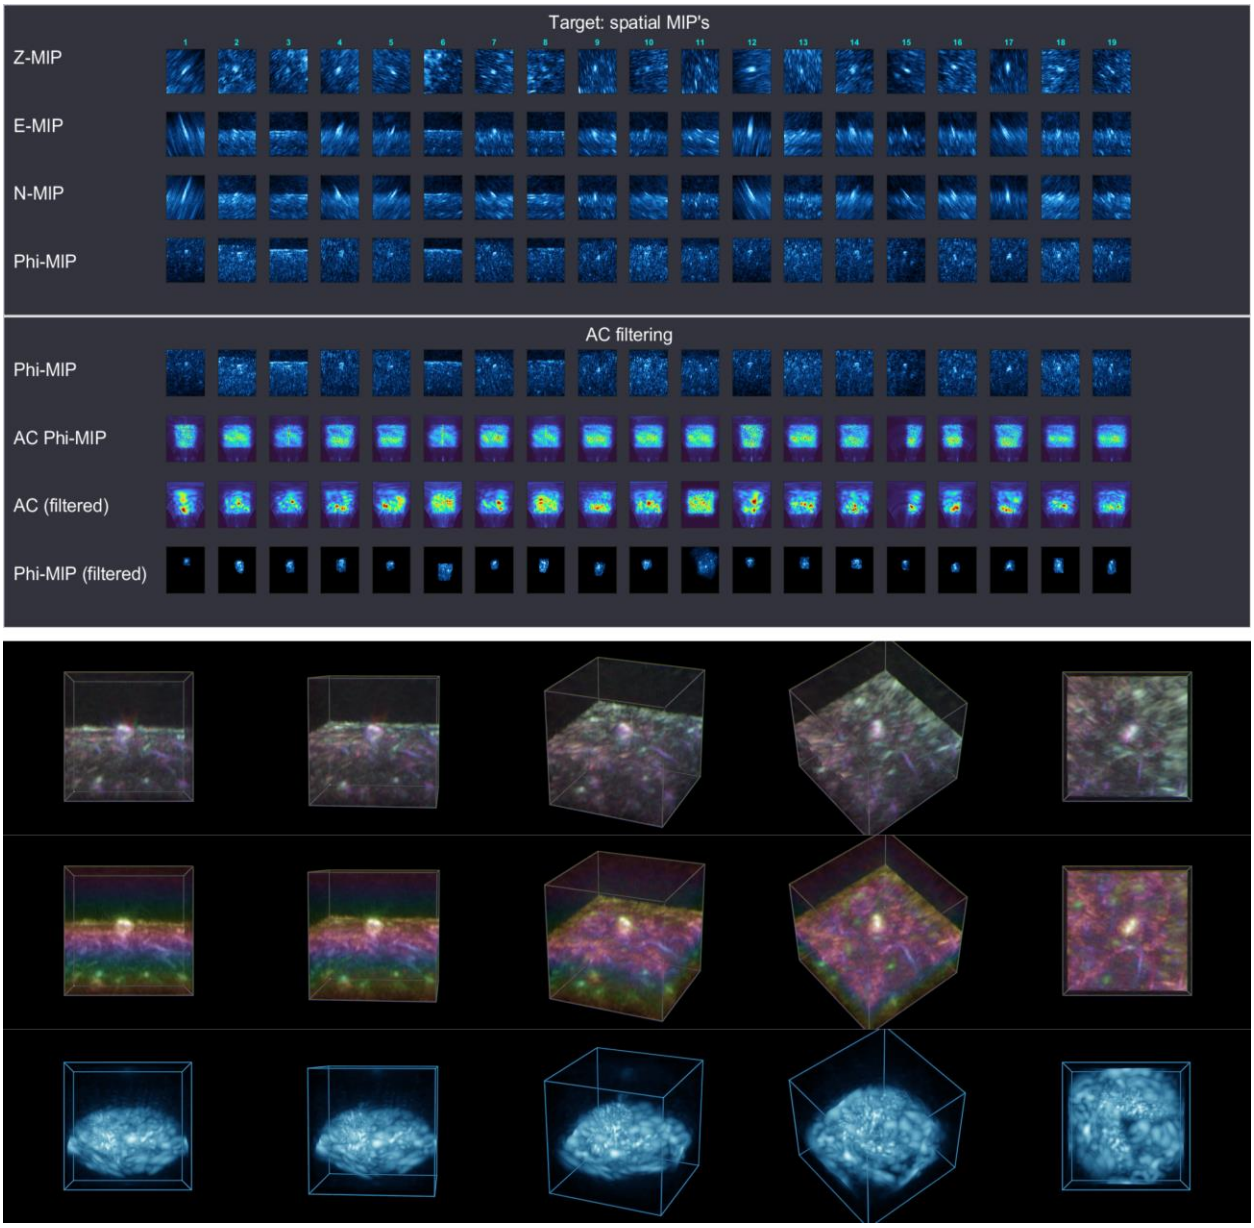

Figure 295: 81mm M889A1 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum.

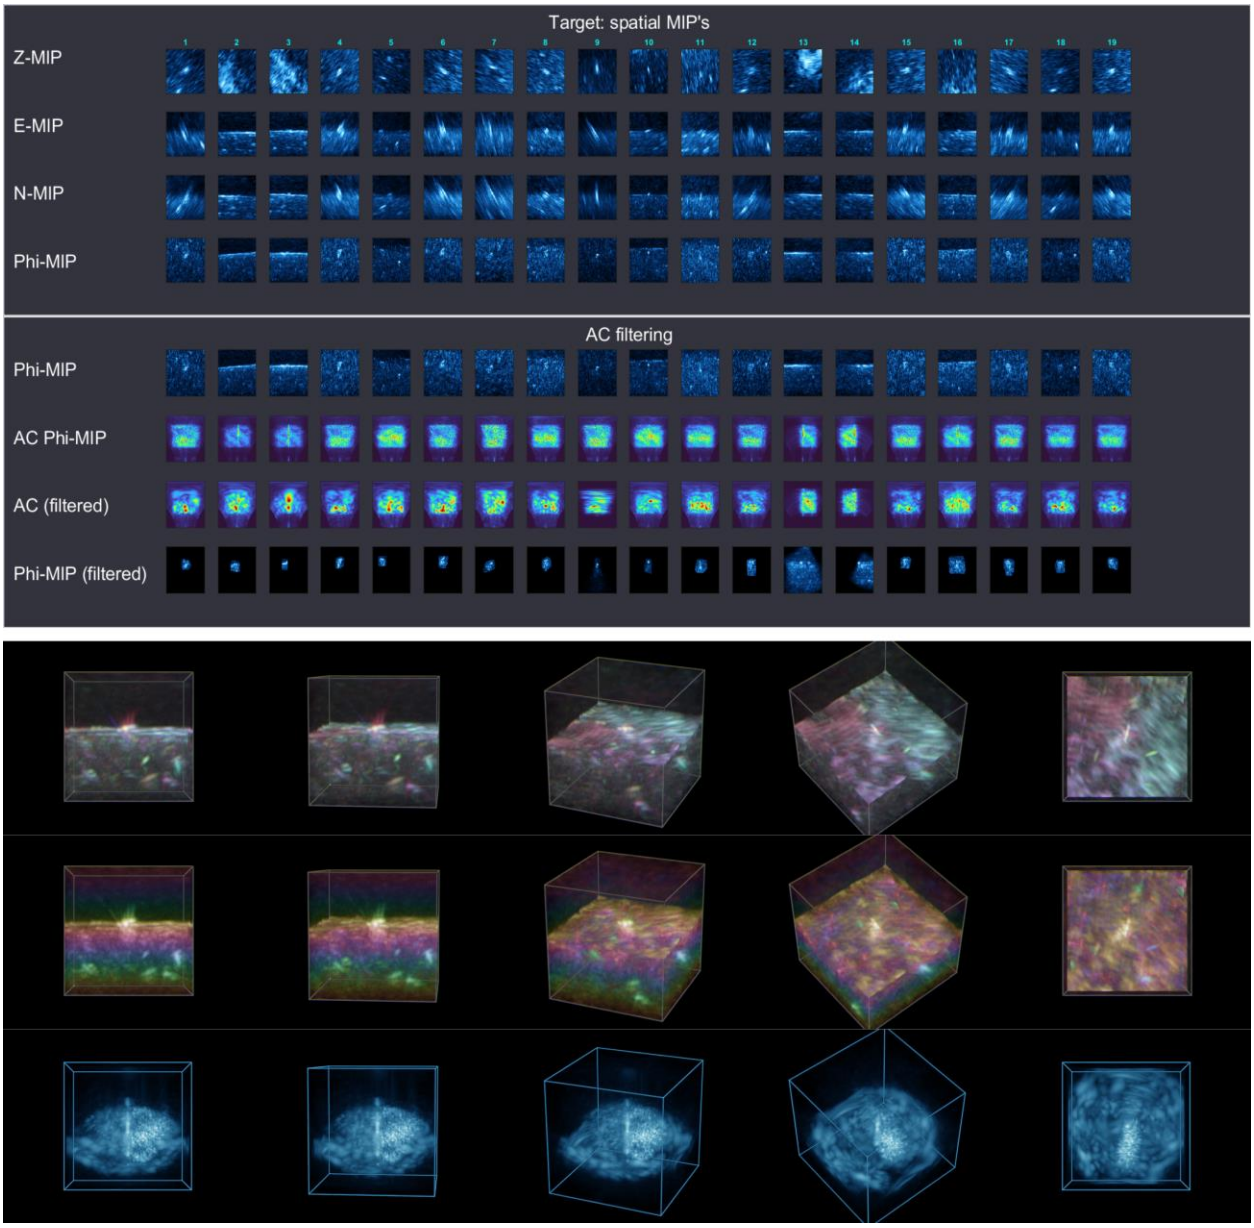

Figure 296: 81mm M889A1 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum.

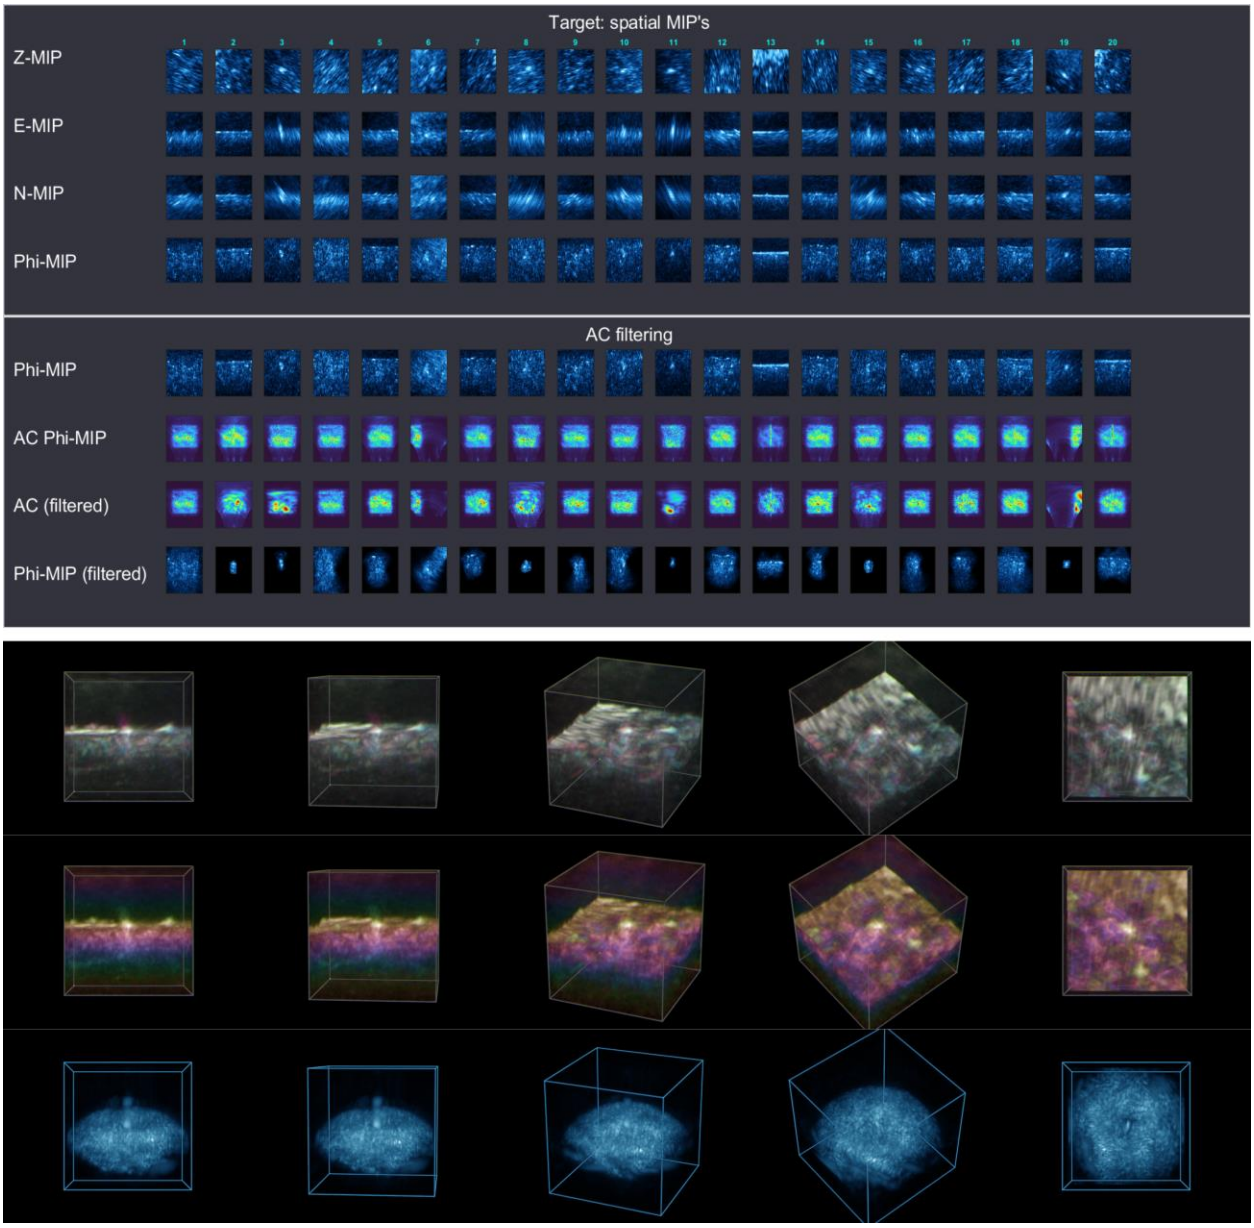

Figure 297: 81mm M889A1 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum.

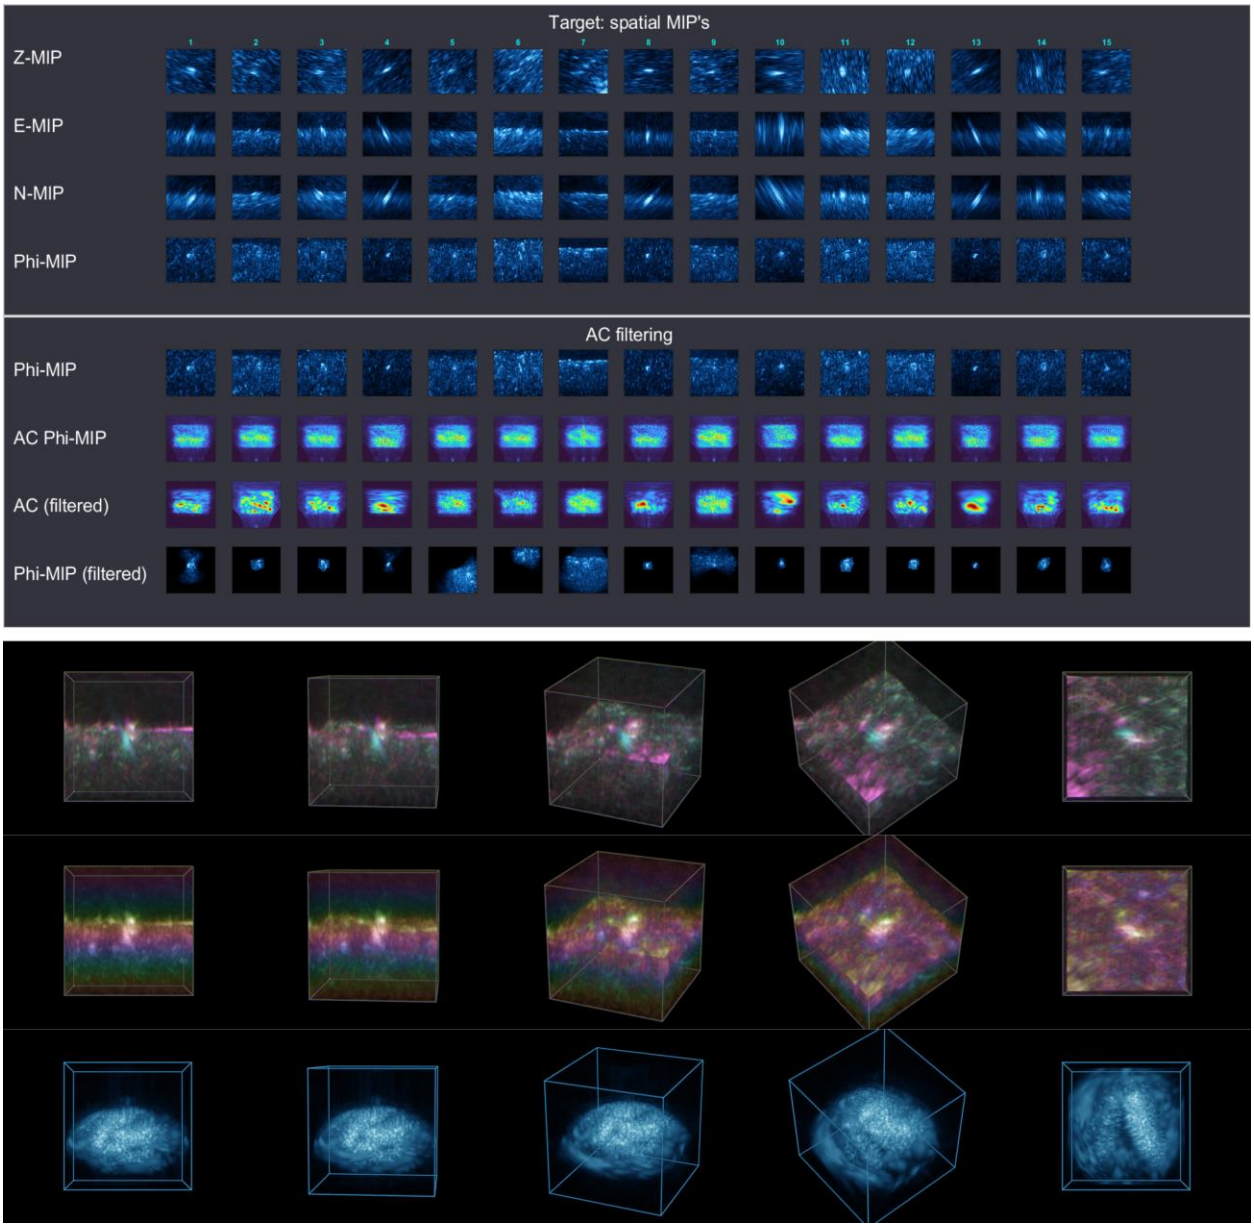

Figure 298: 81mm M889A1 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum.

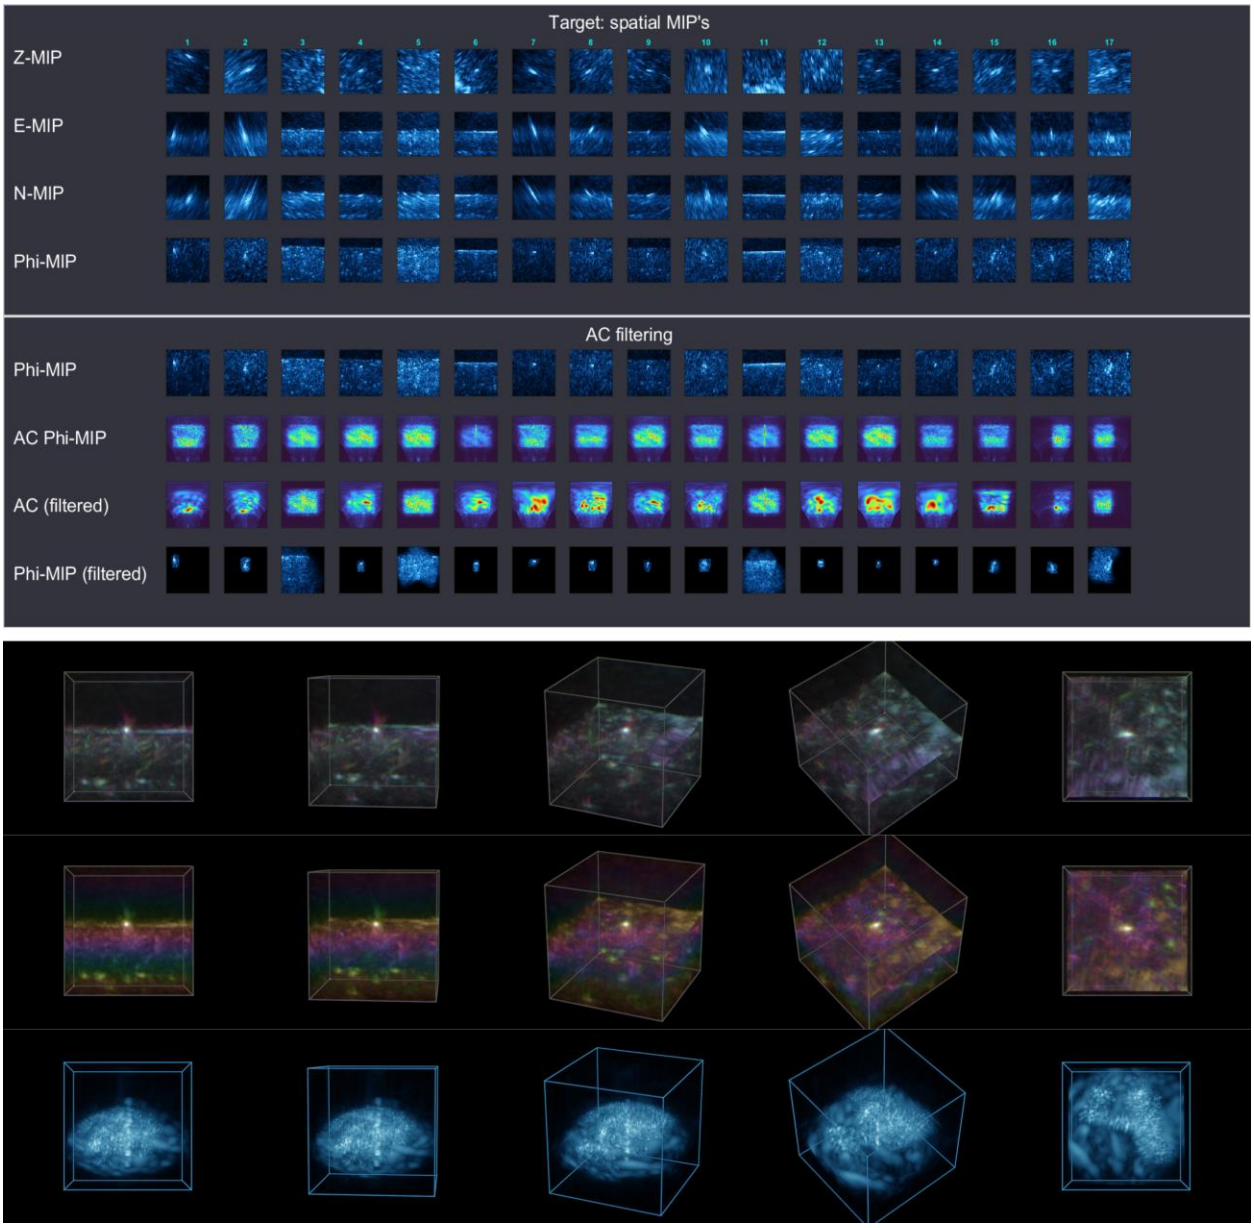

Figure 299: 81mm M889A1 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum.

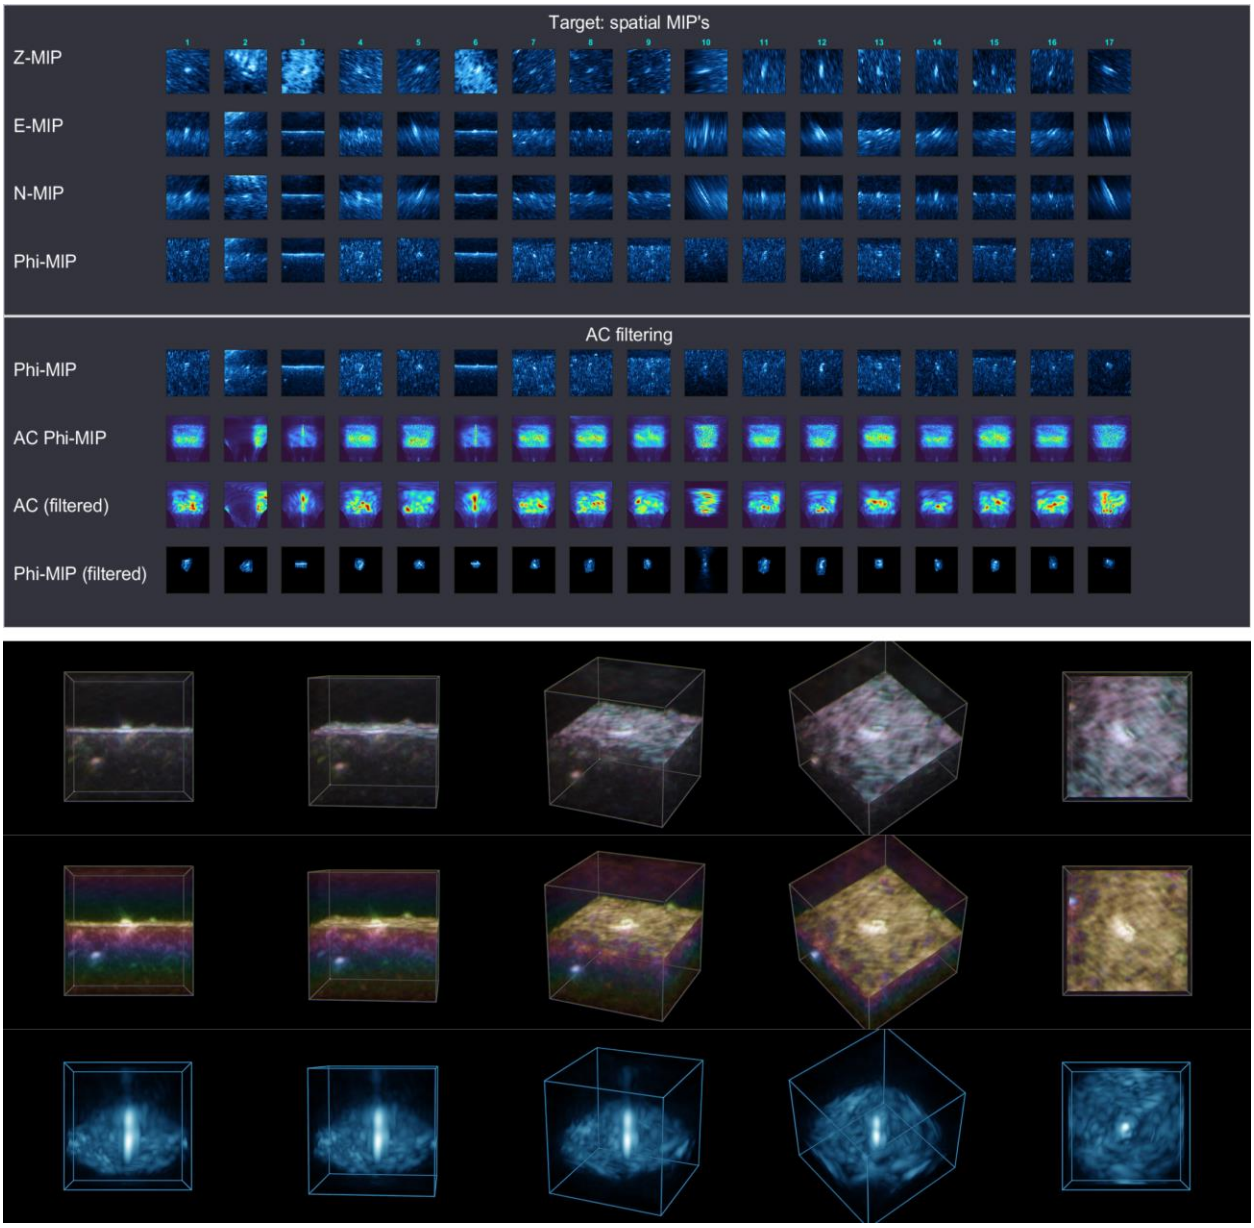

Figure 300: 81mm M889A1 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum.

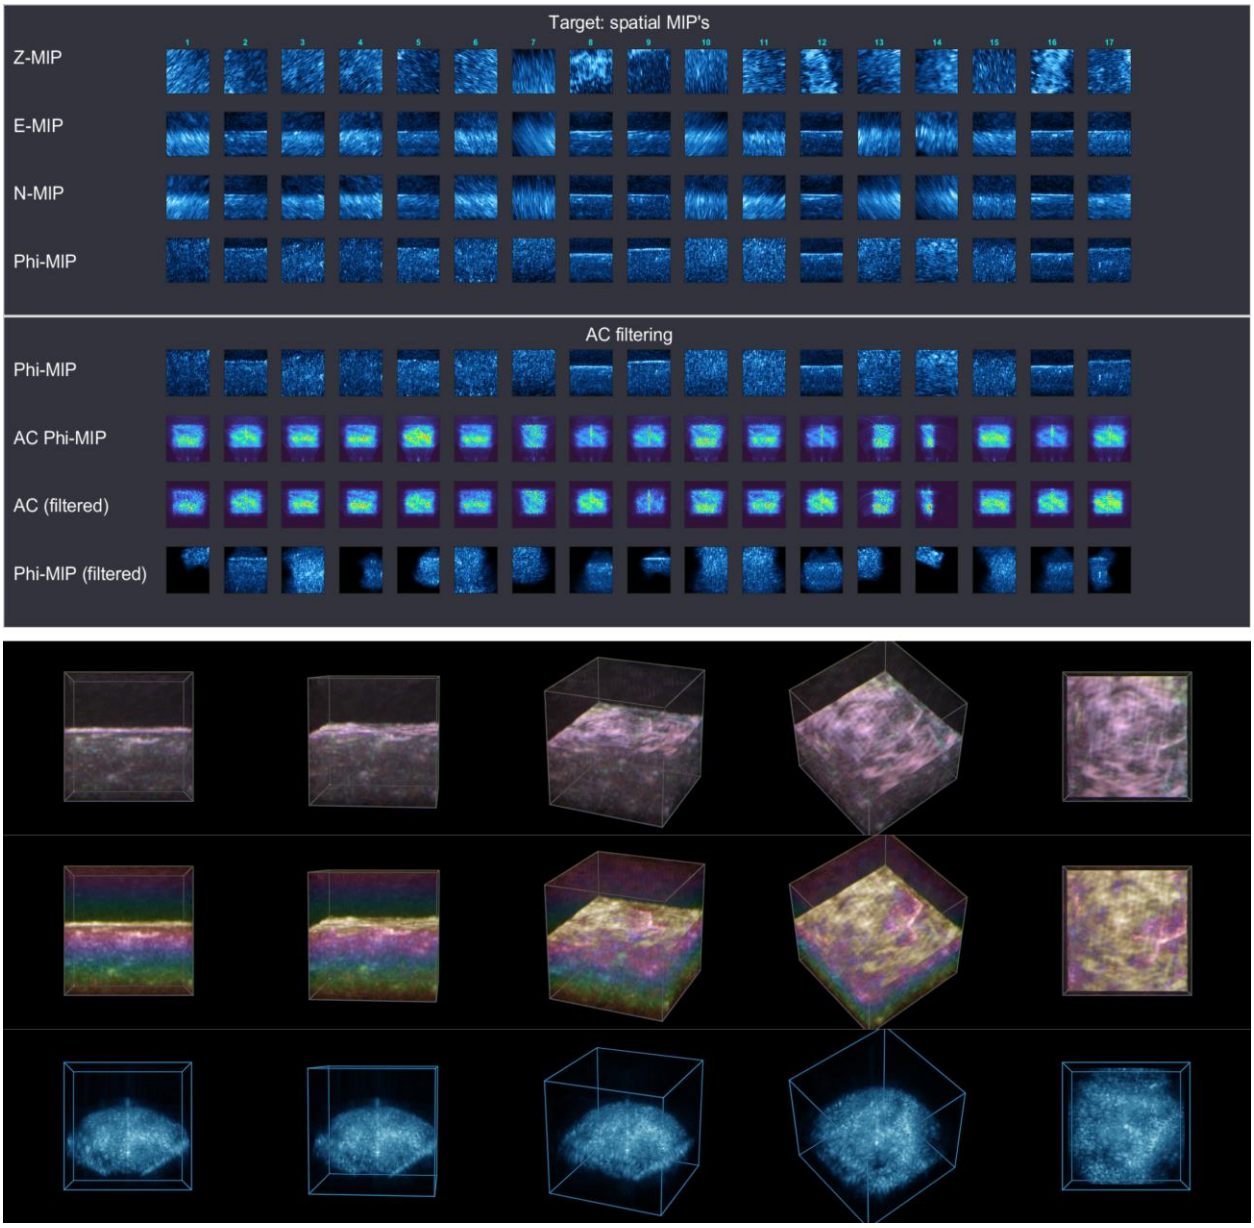

Figure 301: 60mm M49 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum.

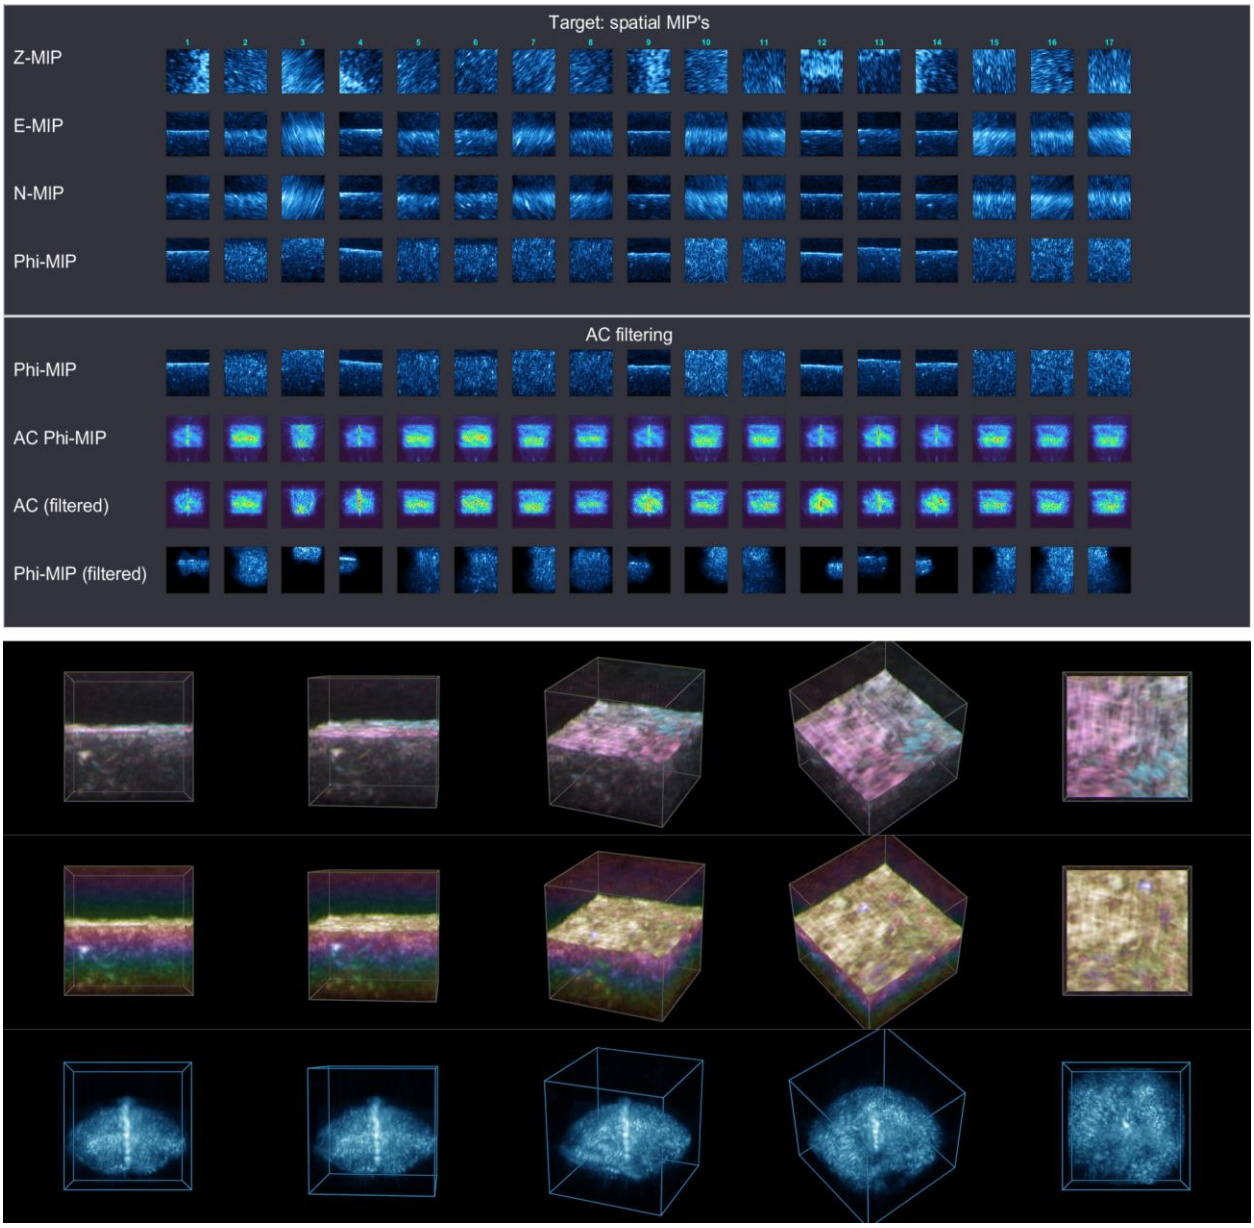

Figure 302: 60mm M49 round. Top figure: single-look projects and filtered and unfiltered acoustic-color data and imagery. Bottom figure: multiple MP views through (top row) color-by-aspect composite images, (middle row) color-by-depth composite images, and (bottom row) the composite spatial spectrum.

From Fig.'s 280 – 302 it can be seen that there is a relatively significant decline in object prominence between the 105mm and 81mm UXO sizes. Unlike the ROI's for the 105mm projectiles, each of which shows a high SNR cylindrical-like target, only about half of the 81mm targets show any cylindrical features and the rest appear blob-like. Beneath 81mm the targets are difficult or impossible to see in the ROI despite many other features in the vicinity being visible, including buried clutter or sediment layering. In contrast, objects 105mm wide and larger are very prominent, and qualitative observation of UXO features in both the spatial and spectral domains is trivial for the 155mm targets.

It is highly likely that the difficulty of detecting smaller targets is due to the bandwidth of the system. Recommended improvements to the system that would likely enable detection of smaller objects include: 1) replacing the omnidirectional sources on the MuST with directional sources to reduce the strong source coloring and back-plate reflection (Fig.'s 23 – 25 in *Materials and Methods: MuST eBOSS calibration*), and 2) extending the bandwidth of the chirp on the high end of the bandwidth, preferably to 30 kHz or higher. Understandably there are limitations to the bandwidth based on the transmitter frequency response and pre-whitening can only compensate for so much before the noise becomes a problem. Additionally, adding more bandwidth at the top end will increase the spatial sampling requirements which may require a change in how the towfish is operated. It is believed, however, that these modifications will likely improve the ability to detect the smaller targets by 1) increasing the resolution of the system, and 2) increasing the *ka* of the smaller targets, improving the SNR.

Detection results

The 3D Mondrian detector created by ARL-PSU, (SERDP project MR21-3543, PI: David Williams), was applied during batch SAS processing to create a detection list which includes a prominence vs. background score *tau* (τ), as well as a 3D target location. The detection threshold τ was adjusted and plotted with different values for Day-1 data. Finally, a very low threshold was used for detection in data collected from the entire survey, allowing the threshold to be chosen in post-processing, and several observations are made in this section from these results.

Figures 303 and 304 show the detections created with the 3D Mondrian detector from Day-1 data, overlaid over a basic target-field marker set showing the target field boundary (orange line) and the calibration fields (green circles, yellow squares, and blue triangles south of the testbed region). The Mondrian detector results are color-coded based on target depth relative to the sediment. The radius of the circle changes with target prominence.

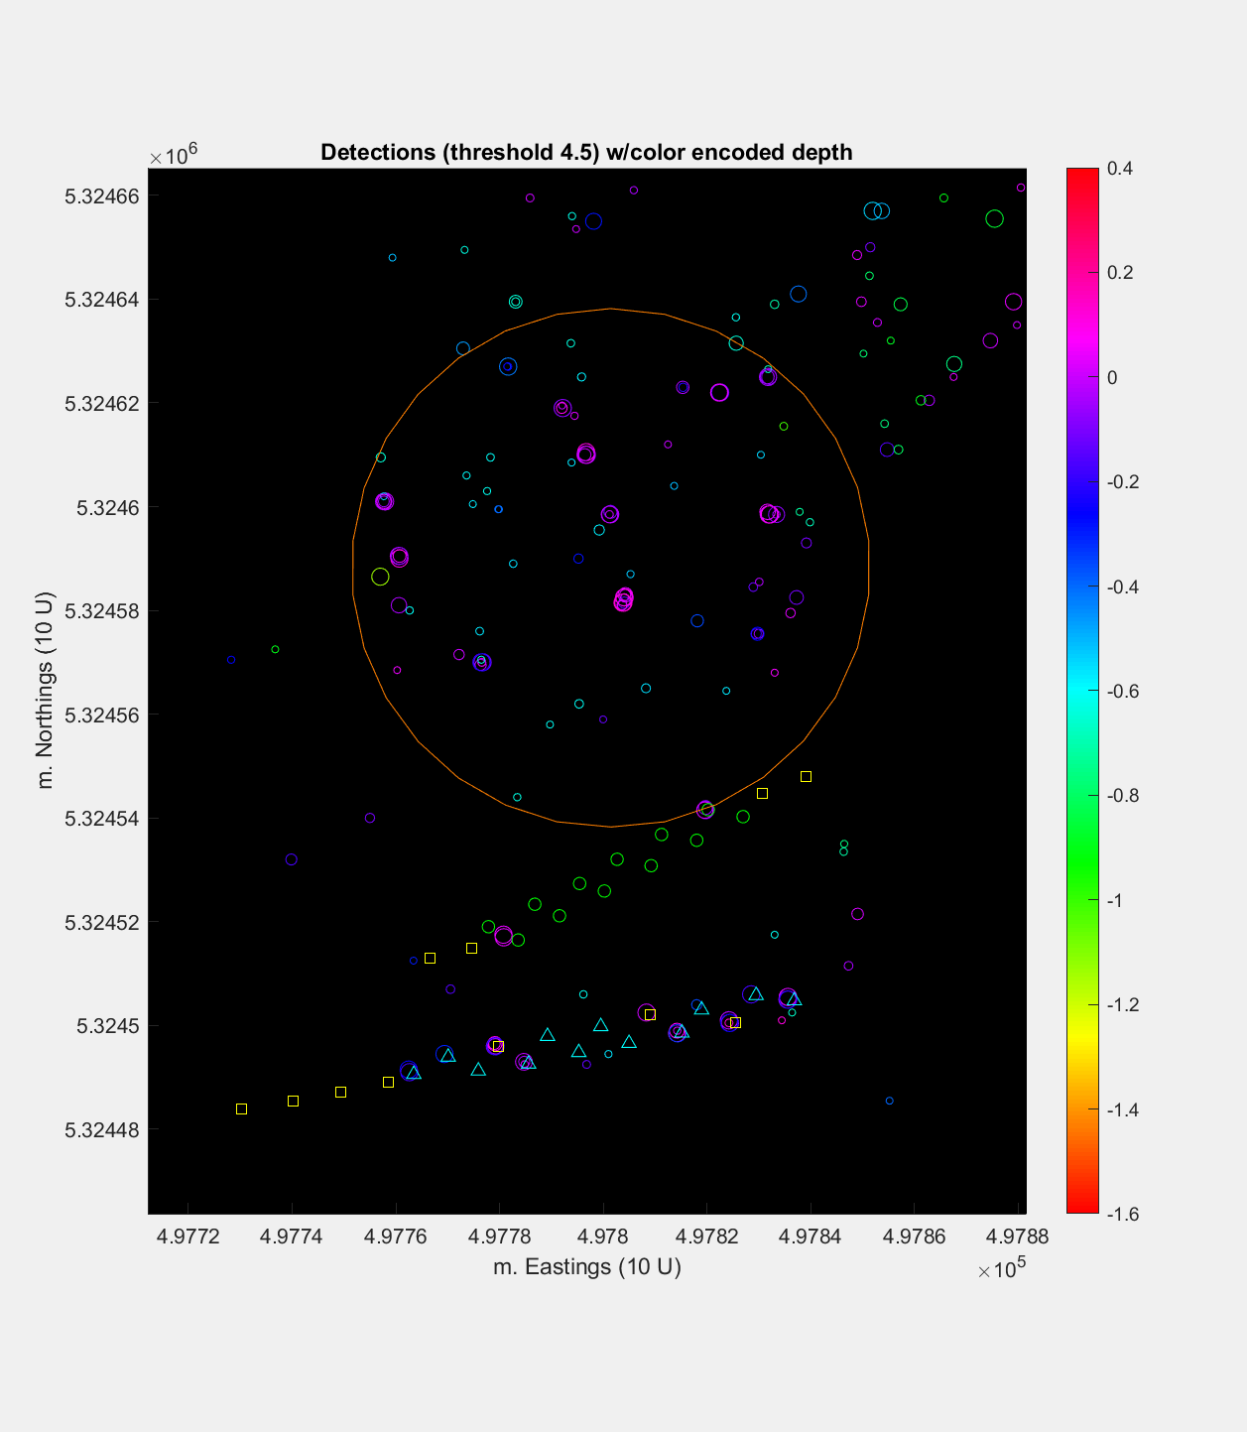

Figure 303: 3D Mondrian detector results with color-by-depth encoding using the color-scheme on the right, and radially encoded prominence (i.e. a wider radius = more prominent target). The prominence threshold value is $\tau = 4.5$.

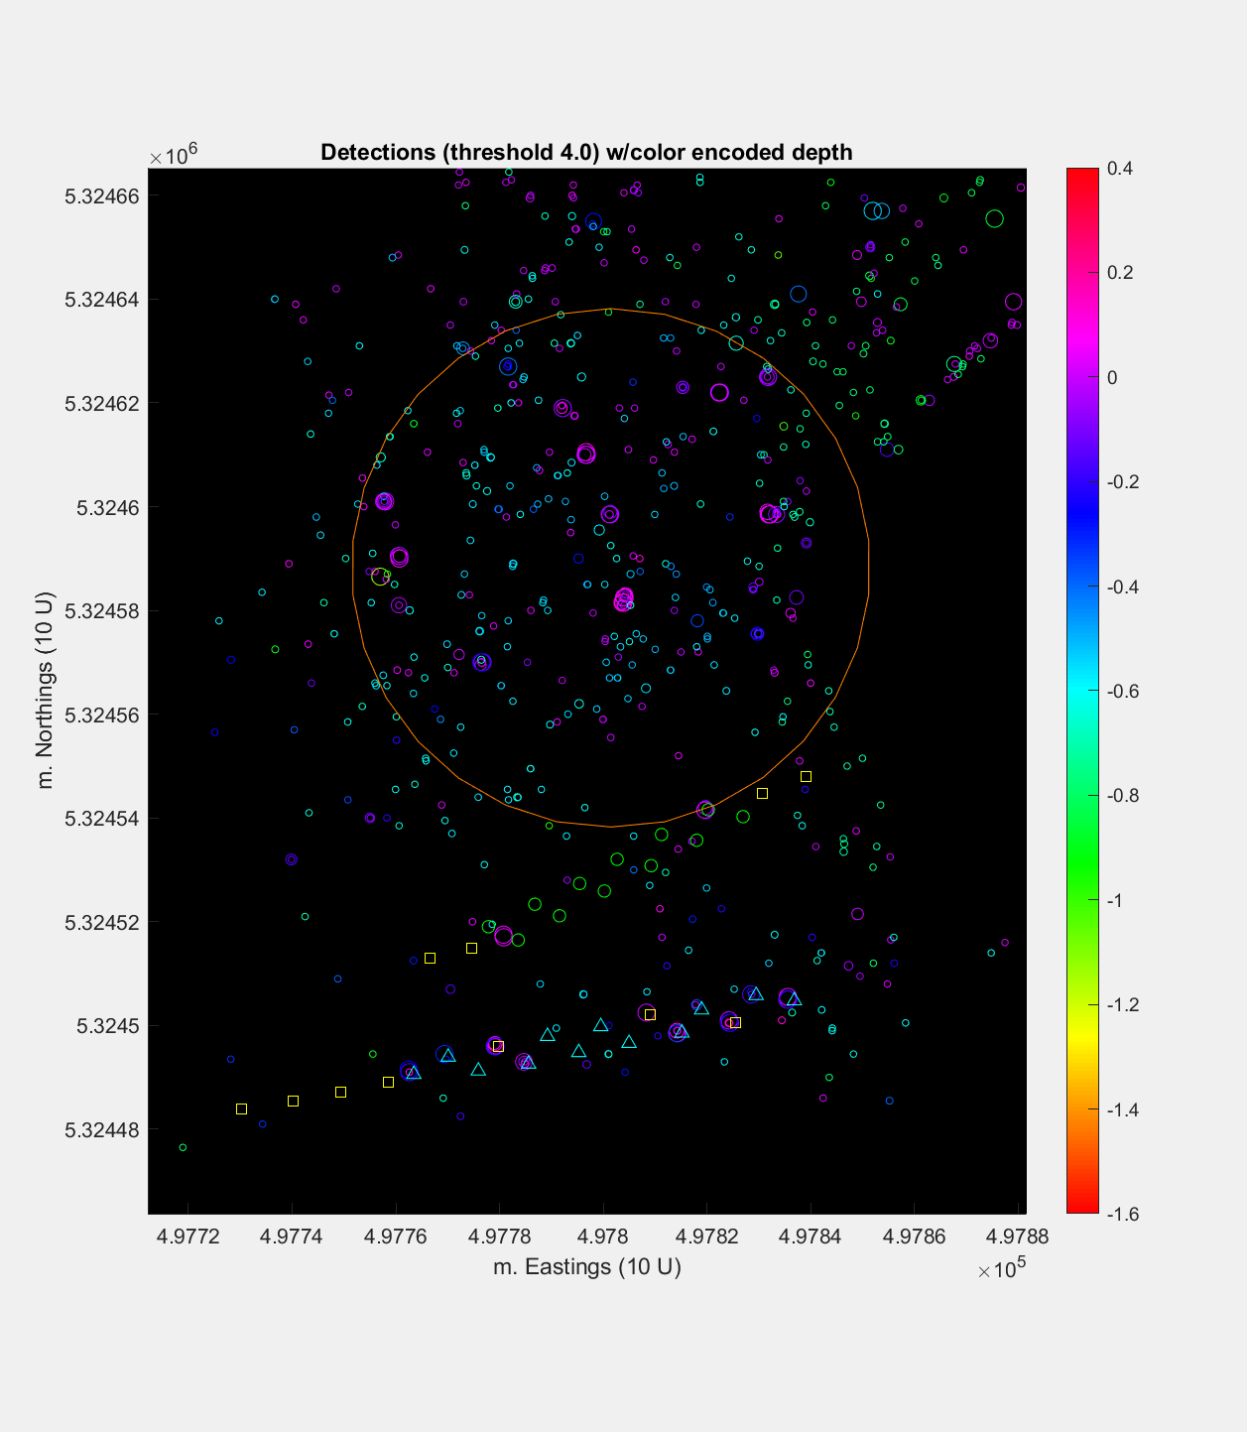

Figure 304: 3D Mondrian detector results with color-by-depth encoding using the color-scheme on the right, and radially encoded prominence (i.e. a wider radius = more prominent target). The prominence threshold value τ has been reduced to 4.0.

Notice the massive increase in the number of detections in Fig.'s 303 and 304 as τ is reduced from 4.5 to 4. Many of the new detections are less prominent scatterers that appear at depths a half meter below the sediment surface or lower. The most prominent scatterers tend to be near

the surface; this can be noted by observing that most of the large circles in Fig. 303 are a magenta shade. There are several prominent scatterers that are buried in the region; one large one is visible in the northwestern portion of the testbed and appears to be buried ~30cm.

The full set of data from days 1 – 4 was then run through the 3D Mondrian detector to create a detection list for the whole survey. The detector was run at a low threshold (τ = 3.5) to allow objects of small prominence to be detected and determine the threshold of detection necessary to observe objects in the calibration line. The number of detections ended up being very large (~3000). A three-dimensional scatterplot was made of the detections in the target field which is shown in Fig. 305. In this scatterplot the point radius is determined by detection prominence and height is color encoded. Note that the height *Z* is relative to the interface and not in absolute geo-rectified units.

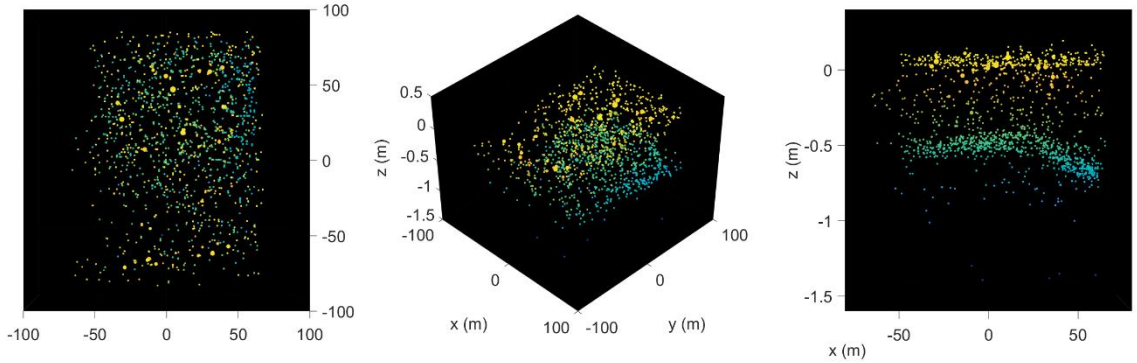

Figure 305: The entire set of 3D Mondrian detections for the Sequim Bay 2021 survey captured with τ = 3.5.

Interestingly, many buried scatterers tend to coalesce around a layer at -0.5 meters. This is most visible in the plot on the far right of Fig. 305. This lower layer curves down sharply towards the east end of the target field, indicating it is likely not a systematic error with the detector. This increased density of detections may be indicative of a change in sediment density: objects sink through the lighter top layer until they reach a density boundary that they don't sink into further. There was no extremely visible interface in the acoustic data at this depth indicating that the transition is probably gradual rather than immediate. Figure 306 is an example of an across-track MIP through a portion of the test-bed. There appears to be some moderate changes in intensity ~ 1 meter below the interface, but nothing extremely visible about half a meter below the interface.

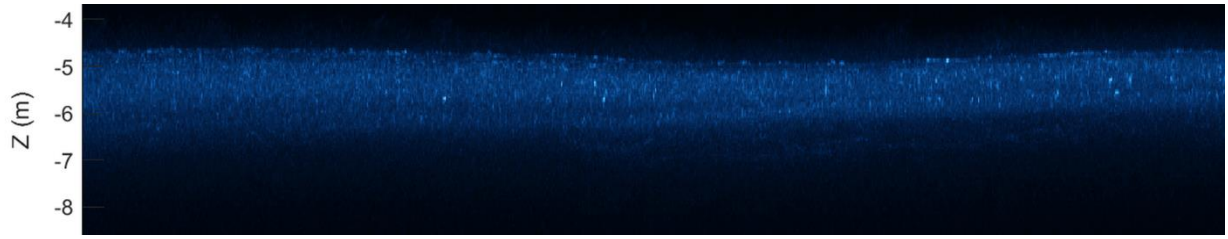

Figure 306: Example of an across-track MIP through beamformed data crossing through the testbed. Note that there doesn't seem to be a strong sediment interface ~0.5 meters below the interface, despite there being an increased number of prominent scatterers collecting at this depth as observed in Fig. 305.

The large number of detections allowed a spatial histogram to be created by counting the number of detections lying within spatial bins spaced over the testbed region. This analysis was done prior to acquisition of the testbed target location key but results were still useful for observing the performance of the renavigation solution: Figure 307 shows the spatial histogram (25cm bins) using the INS renavigation solution for assigning detection location, and Fig. 308 shows the detection histogram following the navigation update created from feature-based navigation. The detections collapse to very tight spreads upon utilization of the feature-based navigation solution.

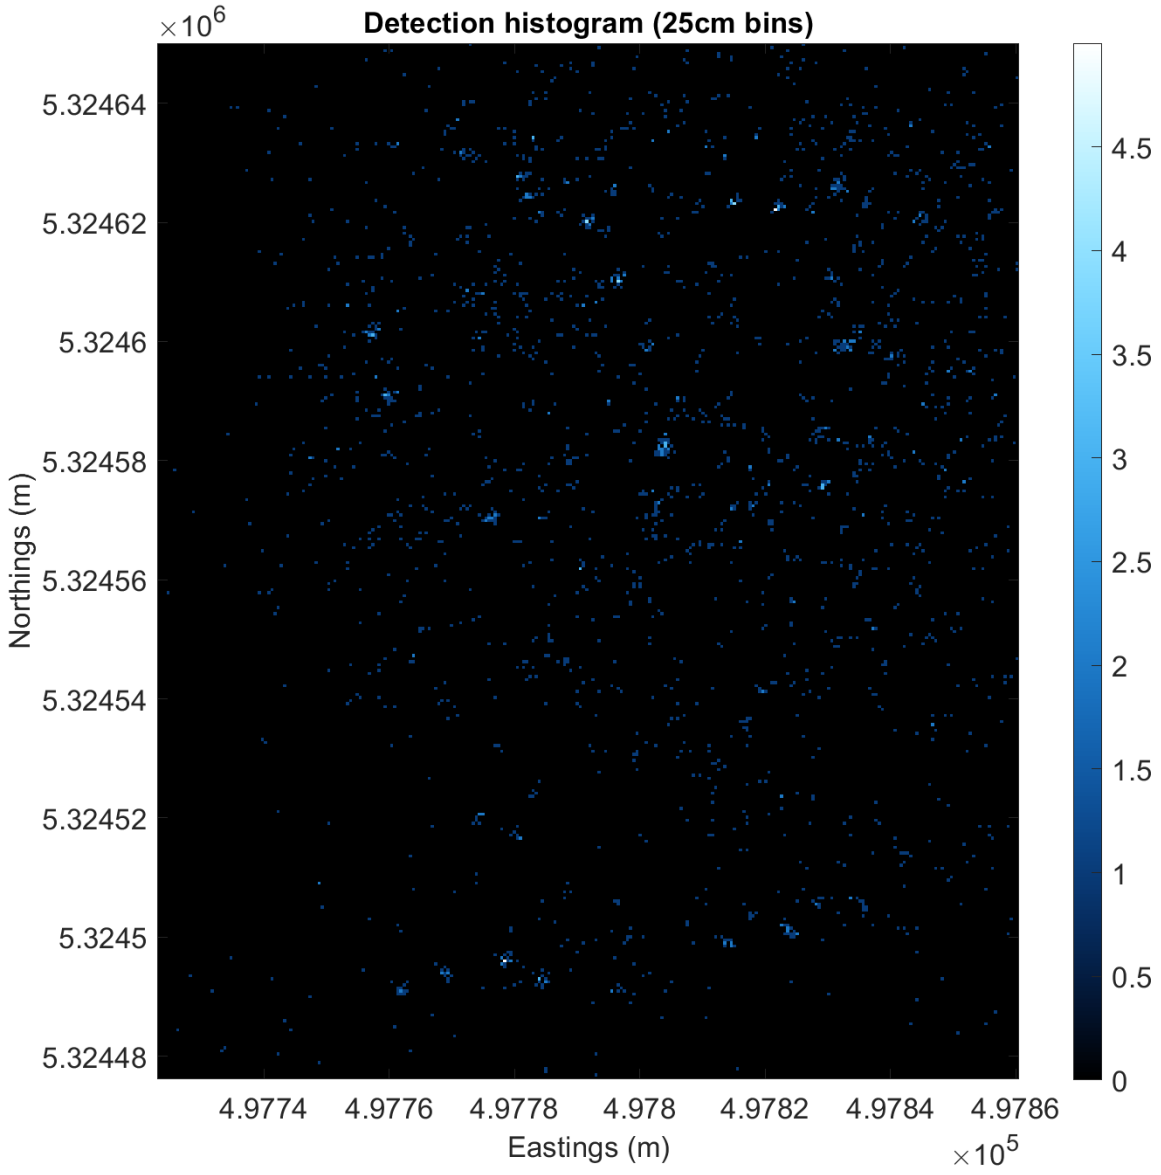

Figure 307: A spatial histogram (25 x 25 cm bins) of the 3D Mondrian detections. Clusters of detections can be seen in regions containing prominent targets. The calibration line is visible in the southern end of the displayed region.

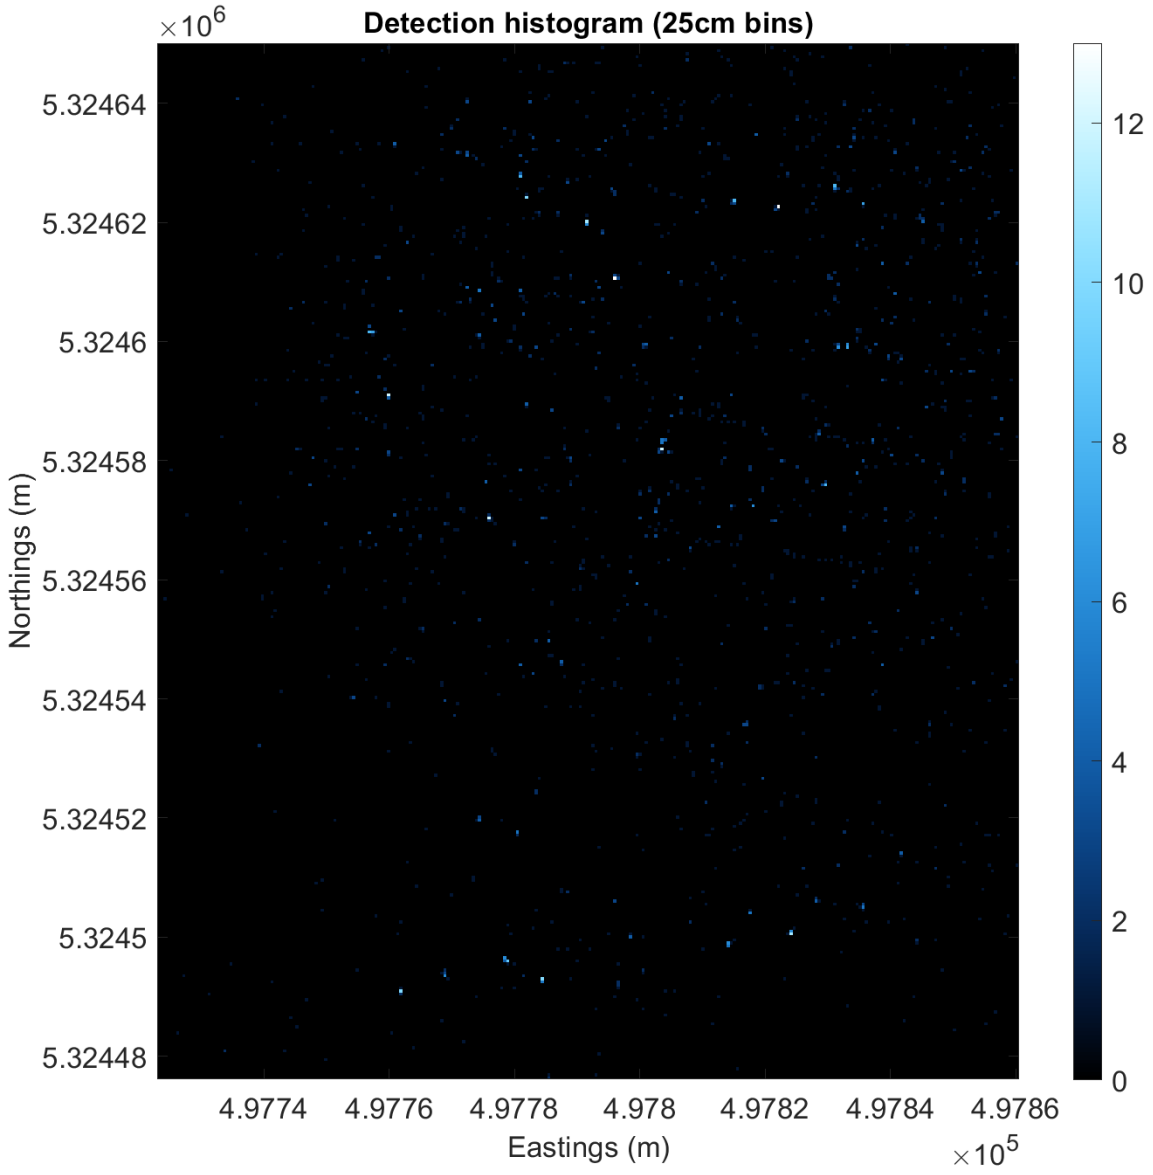

Figure 308: A spatial histogram (25 x 25 cm bins) of the 3D Mondrian detections *after a position update using feature-based navigation correction*. The detection clusters visible in Fig. 307 now collapse to very tight spreads.

Note that though many objects are detected in Fig.'s 307 and 308, several of the calibration line targets corresponding to smaller UXO are missing. (Anomalously, the 40mm is quite visible). There are a number of tunable parameters in the 3D Mondrian detector beyond the threshold of detection and no rigorous attempt was made to optimize these parameters for detection of UXO-scale objects. Optimal parameter adjustment for the present use case would be a potentially fruitful avenue of research in a follow-on effort.

Technical Evaluation

As with the 2020 Sequim Bay experiment and survey result evaluation, the technical evaluation was a part of project MR18-5004 ("Multi-Sensor Towbody (MuST) for Detection, Classification, and Geolocation of Underwater Munitions", PI: Kevin Williams [11]). The present project was integral to evaluation, however, so results are included here even though they will be discussed in more detail – including classification methods – in the technical reports associated with MR18-5004.

Following the survey operations, data were extracted for manually detected regions-of-interest using the mosaic data-products and a previous version of the post-mission analysis toolbox. These datasets were fed into a classifier developed by the PI of MR18-5004 and a list comprised of target scores and locations was provided to IDA for grading. The resulting ROC curve is shown in Fig. 309.

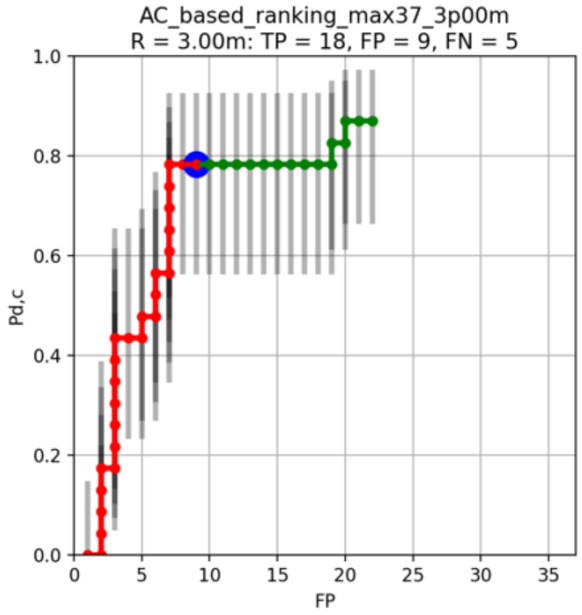

Figure 309: 2021 Sequim Bay ROC curve for the APL-UW target list.

The Pd/c value was (23 – 5)/23 = ~0.78 which was identical to the previous year. The target makeup for 2021 was significantly different that the previous year, however. In 2020, 2 of 9 targets were 81mm shells and both were undetected. In 2021, 12 of 23 (over half) were 81mm, making the 2021 set significantly more challenging from the sensor's perspective. In contrast to 2020, 11 of the 12 81mm targets were detected and 9 of them properly classified. It is likely that this is an instance in which additional training data would improve performance, and information from alternate sensors could be useful for reducing false positives. The current project supplied most of the tools used for detecting regions of interest and extracting the associated SAS-processed acoustic data, however the survey and classification effort was funded via the ESTCP MuST development project.

Conclusions and Implications for Future Research

Leveraging the development of an efficient beamformer for volumetric synthetic aperture imaging, a software framework for live and post-survey processing has been developed that simplifies analyzing sonar data captured by the Buried Object Scanning Sonar (BOSS) for UXO identification and classification. The beamformer, three-dimensional fast factorized backprojection (3D FFBP), is an extension of a well-known accelerated time-domain beamformer originally developed in the field of synthetic aperture radar [18] for the purpose of combining the flexibility and generality of time domain beamforming with the speed of Fourier beamforming. The excessively burdensome nature of volumetric beamforming ($\mathbf{O}(N^5)$, see pg.'s 33 – 35 and Fig. 35) combined with the very natural extension of factorized beamforming to the BOSS synthetic aperture array geometry resulted in an extreme acceleration factor: using basic Matlab implementations, 3D FFBP was found to be over three orders of magnitude faster than standard backprojection. Combined with GPU acceleration, real-time fully sampled beamforming of the entire BOSS swath is possible using the accelerated algorithm, and this has been incorporated into a program called MuSTLive, which is now used routinely with the Multi-Sensor Towbody to monitor data being acquired by the BOSS mounted on that system.

Development of data-driven autofocusing routines aimed at correcting navigation-induced focus aberrations in imagery was originally planned as part of this project, however this was soon found to be unnecessary when field data acquired by BOSS systems became available. The sufficiency of high-grade commercial off-the-shelf inertial navigation systems for achieving well-focused synthetic aperture imagery was first observed using a dataset acquired from an AUV-mounted BOSS system experiencing heavily non-linear platform motion. Well focused imagery from this dataset was attained using only the onboard inertial navigation system mounted inside the AUV. Data-driven "micro-navigation" is a standard procedure applied in ordinary synthetic aperture sonar post-processing to create well focused imagery, however the micro-navigation procedure leverages array overlap between pings [38], something that is generally not possible with BOSS systems due to the nature of their array construction. The originally planned autofocusing effort was meant to address the issue of aberration correction in the absence of the micro-navigation capabilities found in ordinary SAS systems, a very challenging task, and it was therefore a significant discovery to learn that creating well-focused imagery was not contingent on data driven routines which would have required extensive testing and evaluation in numerous environment types. In hindsight, the likely explanation for this simplification of the imaging process is a combination of the BOSS system's long wavelengths, which are much longer than those typically used by SAS imaging systems, and the significantly shorter aperture lengths resulting from imaging at closer ranges.

This lack of necessity for developing data-driven autofocusing routines opened up further avenues of research that included development of a feature-based "renavigation" algorithm for aligning features between geo-rectified scans, enabling reliable multi-aspect data fusion and improving target detection and classification, and the development of an extensive software framework (the "MuST Toolbox") providing an intuitive set of user interfaces, data visualization tools, and post-processing capabilities to the end user.

With regard to the development of the various capabilities for SAS beamforming, data visualization, and post-processing, significant effort was focused toward creating simple and intuitive interfaces that lead the user through the necessary steps required to convert raw data acquired from BOSS systems into useable data products that may be images, acoustic-color datasets, or labelled and geo-rectified target dig lists for remediation efforts. The *Materials and Methods: The MuST Toolbox* portion of this document describes in detail the software tools, including user interfaces, developed for this purpose. These tools, albeit in prior development states, were leveraged during two blind UXO testbed survey operations performed by the MuST system in Sequim Bay.

All of the tools currently in the MuST Toolbox could benefit from further development, however the post-processing step which is currently in most need of further development is feature-based renavigation. This is the one link in the chain of post-acquisition events that doesn't yet have a well-developed user interface and, due to its data-driven nature, its performance is also the most sensitive to external variables. For example, the feature-based renavigation algorithm developed as part of this project was found to be very reliable in Sequim Bay, but this performance collapsed when naively applied to data captured at Lake Washington, a very different environment. By enabling multi-aspect data fusion and creating despeckled mosaics, however, feature-based renavigation has become an invaluable performance booster for target detection and classification, and its continued development – especially the improvement of feature matching across different environments – is a highly recommended path of future research.

Place holders have also been added to the Post-Mission Analysis (PMA) tool for insertion of target-recognition algorithms. It is anticipated that various research efforts aimed at applying machine learning techniques to the UXO classification problem will create target recognition algorithms that are adaptable to working with the various data products created by the PMA tool. Eventually, it would be beneficial if these algorithms could be integrated into the PMA tool as user options for aiding in object classification.

The post-processing toolbox is comprised of a set of scripts that are accessed through Matlab's command-line interface. The necessity for having a Matlab license and some degree of programming knowledge, (specifically, reading scripts related to feature-based renavigation), are barriers-to-entry that remain, and should be removed in the future. Routes towards removing these barriers include: 1) continued development of feature-based renavigation, which needs algorithmic development for stability across environments in addition to software additions like the creation of a user interface, and 2) either obtaining the Matlab compiler to create a standalone executable that combines all of the elements in the MuST Toolbox, or migrating the MuST Toolbox to a non-Matlab software framework that doesn't require a license. Migrating away from Matlab may be the preferred rout for industry collaborators, however retaining a Matlab software framework has the advantage of making the capabilities developed as part of this project more accessible to the broader community performing research related to UXO detection and classification.

This project was not originally intended to include a direct evaluation of the UXO detection and classification capabilities of any particular BOSS system, however the developments of this

project became coupled with the developments of the MuST, a UXO detection and classification platform funded by ESTCP project MR18-5004 ("Multi-Sensor Towbody (MuST) for Detection, Classification, and Geolocation of Underwater Munitions" [11]). By incorporating the developments of this project into the data analysis procedure used with the MuST system, both projects benefitted: the performance of the MuST as a UXO detection platform could be evaluated, and this project gained a huge amount of data that could be used for testing and development. Based on the results of the MuST testbed performance evaluations, several recommendations can be made regarding design changes that could improve the performance of the UXO detection capabilities of the MuST and the quality of the target information extracted by the tools and procedures developed as part of this project. These recommendations include: 1) incorporating transmitter sources with directivity (as opposed to the omnidirectionality of the current sources), to prevent backplate reflection interference, which causes a major interference fringe to form in the transmit beam pattern (see Fig.'s 23 – 25), and 2) extending the transmit signal bandwidth. The Sequim Bay 2021 experiment clearly showed a sharp decrease in detection performance for objects around 80mm in diameter. Though it is possible to make a chance observation of small targets (e.g. the 40mm round visible in the calibration line), the stark lack of observation of any of the 60mm targets in the testbed indicates that objects of this size have passed beneath some threshold for reliable detection. Detection performance would likely improve by implementing the previously described measures: eliminating the large interference fringe from the backplate has the potential to create a higher effective bandwidth and system resolution, and extending the bandwidth will directly improve system resolution. Additional recommendations include integrating a USBL onto the towbody to bound integration drift and improve long-term position accuracy for surveys which require long periods without GPS signal acquisition, and integrating a multibeam onto the towbody for sediment parameter estimation, precise sediment/water boundary marking, and creating additional features for renavigation.

Finally, the post-processing toolbox integrates compensation for non-linear platform trajectories and soundspeed variation directly into the beamformer, but there are a number of ways that future improvements could improve performance. For example, the current beamformer implementation assumes a single soundspeed: the in-situ measurement made by the sound velocity sensor onboard the MuST. Combining a modified beamformer that compensates for varying sediment sound velocity with through-the-sensor remote sediment parameter estimation via a multibeam sonar (e.g. "Demonstration of Physics-based Inversions of Multibeam Echosounder for Sediment Properties," SERDP project MR18-1406, PI: Brian Hefner), could improve focus for buried objects and therefore improve detection performance. Also, though the MacArtney FOCUS towbody proved extremely stable, integration of motion compensation into MuSTLive would improve imagery, especially when the towbody is being maneuvered.

Literature Cited

1. S. G. Schock, T. Arnaud, J. Wulf, J. Sara, M. Ericksen, "Buried Object Scanning SONAR," IEEE Journal of Oceaning Engineering, vol. 26, no. 4, pp. 677-689, (2001)
2. S. G. Schock, J. Wulf, G. Quentin, J. Sara, "Synthetic aperture processing of buried object scanning sonar data", *Proceedings of MTS/IEEE OCEANS 2006,* pp.2236-2241 vol. 3, (2005)
3. E. Lavely, "Sensor Phenomenology and Feature Development for Improved Sonar-based Detection and Classification of Underwater UXO," Interim Report, SERDP project MM-1533 (2009)
4. J. A. Bucaro, "Structural Acoustic UXO Detection and Identification in Marine Environments – Interim Report for SERDP MR-2103 Follow-On", Interim Report, SERDP project MR 2103 (2015)
5. R. A. Leasko, "Munitions detection using unmanned underwater vehicles equipped with advanced sensors" Interim Report, SERDP project MR 201103 (2012)
6. J. R. Kennedy and T. M. Marston, "Multipass and non-concentric target CSAS," Final Report, SERDP project MR 2439 (2017)
7. S. G. Kargl, K. L. Williams, A. Espana, T. Marston, D. Plotnick, "From Basic Research to a Prototype Towbody Platform: Progress on Underwater Unexploded Ordnance Remediation," In Proceedings of UACE, Crete, Greece, 2019
8. G. Dinn, "Field Experience with a New Sub-Bottom Investigation Tool," Proceedings of IEEE OCEANS 2012, 2012, pp. 1-9, doi: 10.1109/OCEANS.2012.6405076.
9. Seachirp-3d, accessed at 4/5/2022 at http://www.soacsy.fr/seachirp-3d.html.
10. M. Gutowski, J. Malgorn and M. Vardy, "3D sub-bottom profiling — High resolution 3D imaging of shallow subsurface structures and buried objects," OCEANS 2015 - Genova, 2015, pp. 1-7, doi: 10.1109/OCEANS-Genova.2015.7271468.
11. K. L. Williams, "Multi-Sensor Towbody (MuST) for Detection, Classification, and Geolocation of Underwater Munitions," ESTCP project MR18-5004
12. Datasheet available from www.neptune-sonar.co.uk. The specific model used by the MuST is D/11/BB.
13. T. M. Marston, S. G. Kargl, D. S. Plotnick, K. L. Williams, "Improving downlooking SAS resolution with transmitter spatial diversity," in Proceedings of IEEE OCEANS conf. Seattle, WA, 2019
14. Datasheet available from http://www.teledynemarine.com/Pathfinder_DVL
15. D. C. Brown, S. F. Johnson and D. R. Olson, "A point-based scattering model for the incoherent component of the scattered field," J. Acoust. Soc. Am., vol. 141, no. 3, pp. EL210-EL215, 2017.
16. S. F. Johnson, D. C. Brown, "SAS Simulations with Procedural Texture and the Point-based Sonar Scattering Model," in Proceedings of IEEE OCEANS conf. Charleston, S.C. USA, 2018
17. D. Brown et al, "Final Report: Sediment Volume Search Sonar Development", Final Report, SERDP project MR-2545 (2021)

18. L. M. H. Ulander, H. Hellsten, and G. Stenstrom, “Synthetic-aperture radar processing using fast factorized back-projection,” IEEE Trans. Aerosp. Electron. Systems, vol. 39, no. 3, pp. 760-776, 2003.
19. J. R. Pearson Magoon, M. A. Nelson, D. D. Sternlicht, “Frequency domain beamformer for a 3-D sediment volume imaging synthetic aperture sonar,” in Proceedings of IEEE OCEANS conf. Biloxi, MS, USA, 2009
20. Rauch, H.E.; Tung, F.; STriebel, C. T. “Maximum likelihood estimates of linear dynamic systems.” AIAA Journal. 3 (8): 1445-1450. August, 1965.
21. Information about IXblue Delph INS post-processing software can be found at: https://www.ixblue.com/delph-ins, (accessed 10/30/2020).
22. Information about Sonardyne Janus INS post-processing software can be found at: https://www.sonardyne.com/product/janus-ins-post-processing-software/, (accessed 10/30/2020).
23. Stankwitz, Herbert C., Rodney J. Dallaire, and James R. Fienup. "Nonlinear apodization for sidelobe control in SAR imagery." IEEE Transactions on Aerospace and Electronic Systems 31.1 (1995): 267-279.
24. Cutrona, Louis J. "Comparison of sonar system performance achievable using synthetic-aperture techniques with the performance achievable by more conventional means." The Journal of the Acoustical Society of America 58.2 (1975): 336-348.
25. Cutrona, Louis J. "Additional characteristics of synthetic-aperture sonar systems and a further comparison with nonsynthetic-aperture sonar systems." The Journal of the Acoustical Society of America 61.5 (1977): 1213-1217.
26. Gough, Peter T., and David W. Hawkins. "Imaging algorithms for a strip-map synthetic aperture sonar: Minimizing the effects of aperture errors and aperture undersampling." IEEE Journal of Oceanic Engineering 22.1 (1997): 27-39.
27. Plotnick, Daniel S., and Timothy M. Marston. "Utilization of aspect angle information in synthetic aperture images." IEEE Transactions on Geoscience and Remote Sensing 56.9 (2018): 5424-5432.
28. Cook, D. A., J. T. Christoff, and J. E. Fernandez. "Broadbeam multi-aspect synthetic aperture sonar." MTS/IEEE Oceans 2001. An Ocean Odyssey. Conference Proceedings (IEEE Cat. No. 01CH37295). Vol. 1. IEEE, 2001.
29. Robinson, Marc, Mahmood R. Azimi-Sadjadi, and Jaime Salazar. "Multi-aspect target discrimination using hidden Markov models and neural networks." IEEE Transactions on Neural Networks 16.2 (2005): 447-459.
30. Fawcett, John, et al. "Multiaspect classification of sidescan sonar images: Four different approaches to fusing single-aspect information." IEEE Journal of Oceanic Engineering 35.4 (2010): 863-876.
31. Williams, David P., et al. "Adaptive underwater sonar surveys in the presence of strong currents." 2016 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2016.
32. Dugelay, Samantha, and Warren Fox. "Active contours for synthetic aperture sonar snippet registration." OCEANS 2015-Genova. IEEE, 2015.

33. D. G. Lowe, “Object recognition from local scale-invariant features,” *Proceedings of the International Conference on Computer Vision*. **2**. pp. 1150–1157. 1999. doi:10.1109/ICCV.1999.790410.
34. H. Bay, T. Tuytelaars, L. Van Gool, “SURF: Speeded Up Robust Features,” *Proceedings of ECCV 2006* pp. 404-417. doi: 10.1007/11744023_32
35. M. A. Fischler, R. C. Bolles, “Random Sample Consensus: a paradigm for model fitting with applications to image analysis and automated cartography.” Communications of the ACM. June, 1981. doi: 10.1145/358669.358692
36. T. M. Marston, J. L. Kennedy and P. L. Marston, “Coherent and semi-coherent processing of limited aperture circular synthetic aperture (CSAS) data,” in Proc. IEEE OCEANS conf. Waikoloa, HI, USA, Sept. 2011.
37. Fishman, Elliot K., et al. "Volume rendering versus maximum intensity projection in CT angiography: what works best, when, and why." Radiographics 26.3 (2006): 905-922.
38. Bellettini, Andrea, and Marc A. Pinto. "Theoretical accuracy of synthetic aperture sonar micronavigation using a displaced phase-center antenna." IEEE journal of oceanic engineering 27.4 (2002): 780-789.
39. https://www.boydski.com/diving/dives/PB4Y-wreck.htm, accessed 5/6/2022
40. Williams, Kevin, et al. "Results from an Informal Demonstration of a Buried-UXO Detection, Classification, Geo-Location System." UACE 2021.

Appendix A: Buried Cable Tracking

The MuST underwent a system check during the week prior to the 2021 Sequim Bay testbed survey. While conducting the system check an attempt was made to locate a communications cable stretching from north to south across the eastern portion of Lake Washington. Though not directly related to UXO remediation these results are included because of the frequent need in industry for pipe and cable tracking and burial depth monitoring. In addition to UXO surveying, the MuST is an ideal platform for performing these types of surveys, so results from this brief cable survey attempt are included in this appendix.

The NOAA Lake Washington nautical navigation chart contains a roughly ~400meter wide strip labelled as a "Cable Area" in the western portion of Lake Washington. This region is shown in Fig. A.1.

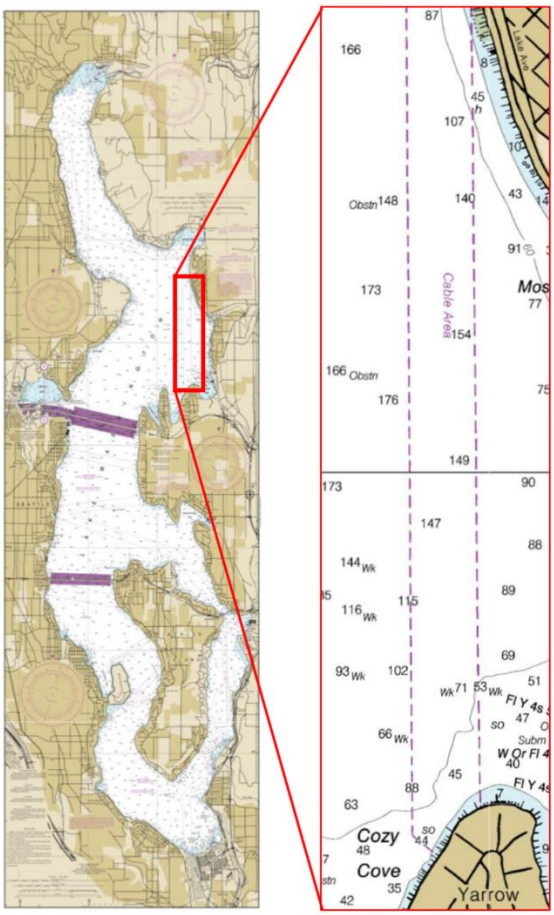

Figure A. 1: The NOAA Lake Washington navigation chart (left) and a zoomed view of the "Cable Area" demarcated by the dashed vertical purple lines. The region is roughly ~400 meters wide.

Using the "Cable Area" region labelled in the NOAA chart as a starting point, several short scans were conducted in this region to determine if the cable could be located using the MuST. This was found to be the case and, following a larger collision hazard survey using the high-frequency sidescan sonar on the MuST, a longer cable tracking survey was conducted. The towbody was moved to the northern end of the cable area, and the towfish was scanned along a southerly

course. Figure A.2 shows the navigation map display of MuSTLive as the Towbody is approaching the cable location. Note that this is an image of a replay of the sonar file, so cable marker points from previous scans are visible in the map. More specifically, two shorter scans were conducted in this northerly region, and a roughly 150 meter long segment is visible in Fig. A.2.

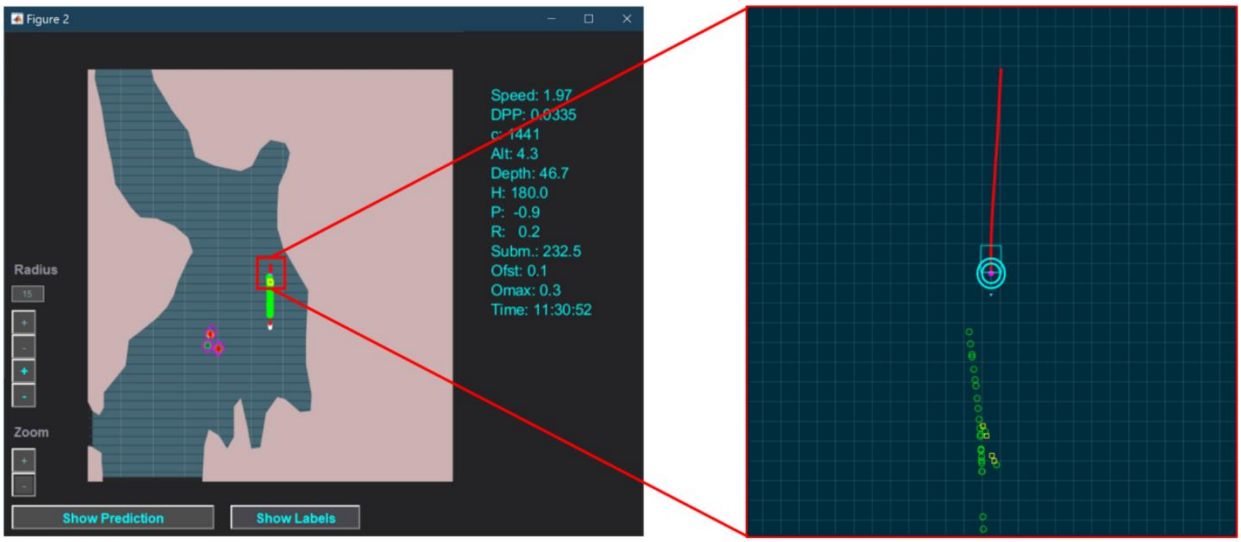

Figure A. 2: The navigation map pane zoomed out to show location context (compare to A.1), and zoomed in to show the location of the towfish and cable markers. The ~150 meter long segment of markers that the towfish is approaching was from a previous, short cable survey, and is being used as a navigation guide for the longer survey being conducted.

In this particular acquisition the towfish travelled adjacent to the location of the cable marked from two shorter surveys conducted the previous day, and for several hundred meters, the cable is not in the sonar swath (Fig. A. 3).

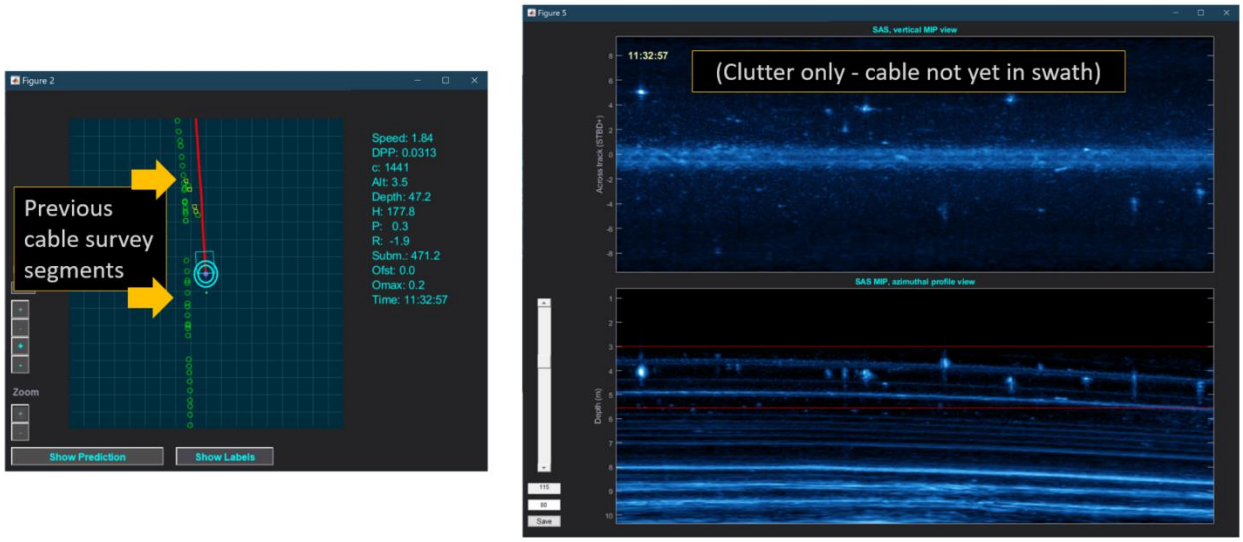

Figure A. 3: The towfish traveled south next to the location of the cable but the cable is not yet in the sonar swath. The ship line and towfish line must be adjusted to position the swath over the location of the cable.

Based on the cable markers from the previous day's surveys, the MuST was moved to starboard and crossed the cable. The cable crossing was easily visible in the MuSTLive SAS MIPs, as shown in Fig. A.4.

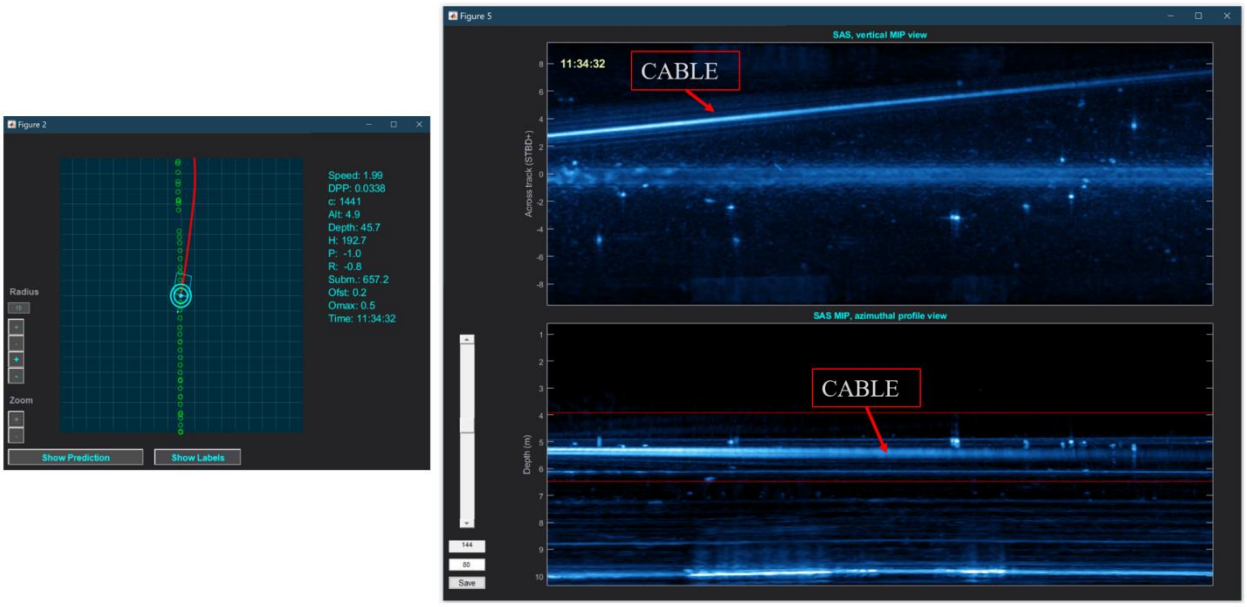

Figure A. 4: To position the cable in the swath of the sonar the MuST was steered starboard over the cable location identified in the navigation window markers positioned from previous short cable scans. The cable immediately becomes visible in the sonar imagery (right figure). Burial depth changes over the region shown in the image but ranges from ~70 centimeters to ~20 centimeters.

Feedback from MuSTLive was used to keep the cable within the field of view of the sonar. The course correction made by the towbody navigator following MuSTLive positioning feedback is shown in Fig. A.5.

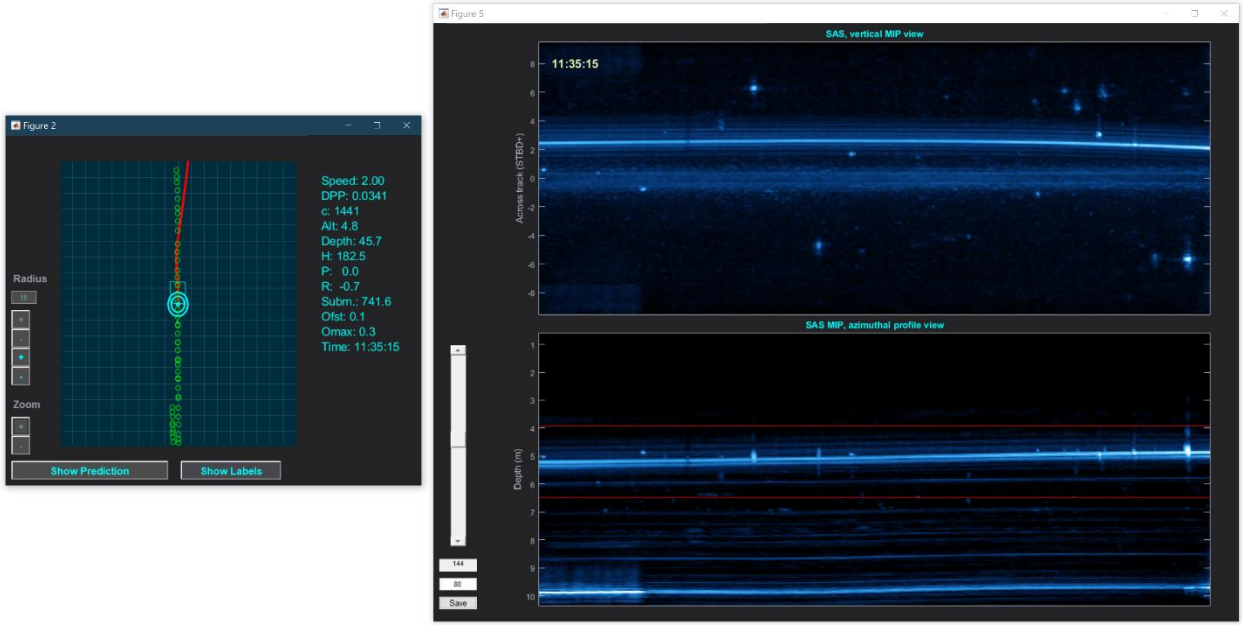

Figure A. 5: Following acquisition of the cable the towbody is straightened out and flown to keep the cable near nadir. The nadir glint of the first sediment layer (typically mud) in the Lake Washington environment is not incredibly strong and the cable is easily dominant over the nadir reflection.

Many clutter objects exist in the Lake Washington sediment. Consider the bright pipe-like object in Fig. A.6 that lies about a meter from the cable: this object is very bright, but is also buried at the same burial depth as the cable.

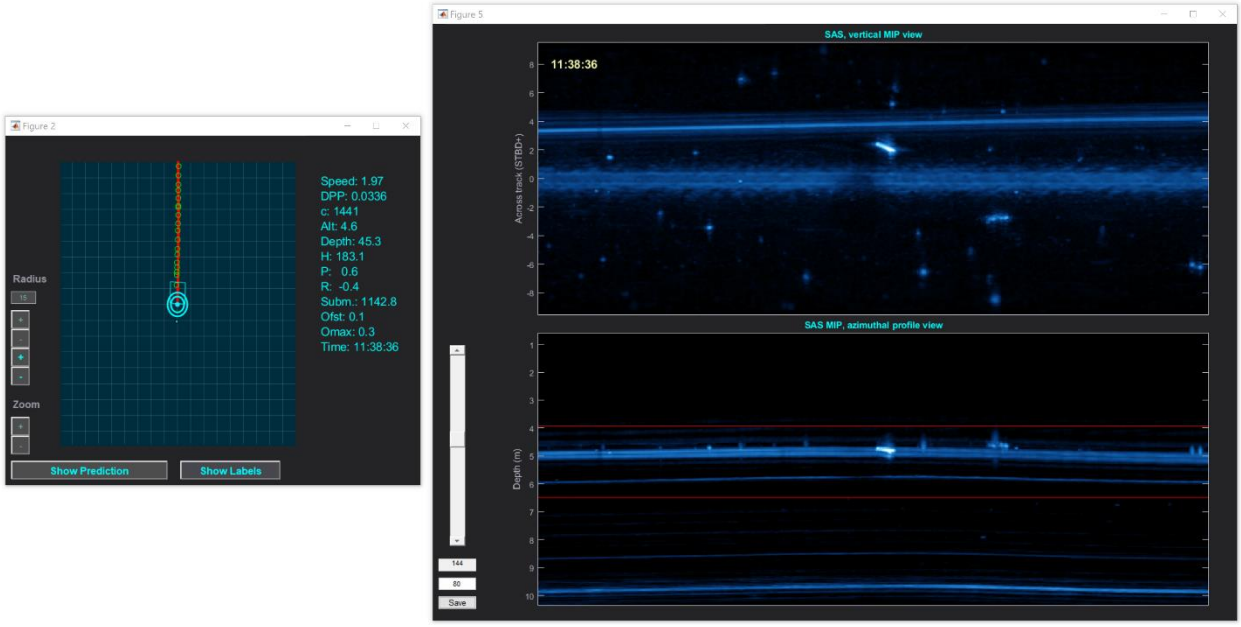

Figure A. 6: Note the pipe-like clutter item in the SAS image. The feature is very bright and sets the white-point for the color scale of the image. Lake Washington is full of man-made clutter; multiple pipes with coherent edge glints were observed during this single cable survey.

The towbody continued to track the cable until the boundaries of the region over which a high-frequency sidescan sonar safety sweep had been conducted (always done prior to scans to ensure no collision hazards exist) were reached, and the Towfish was surfaced.

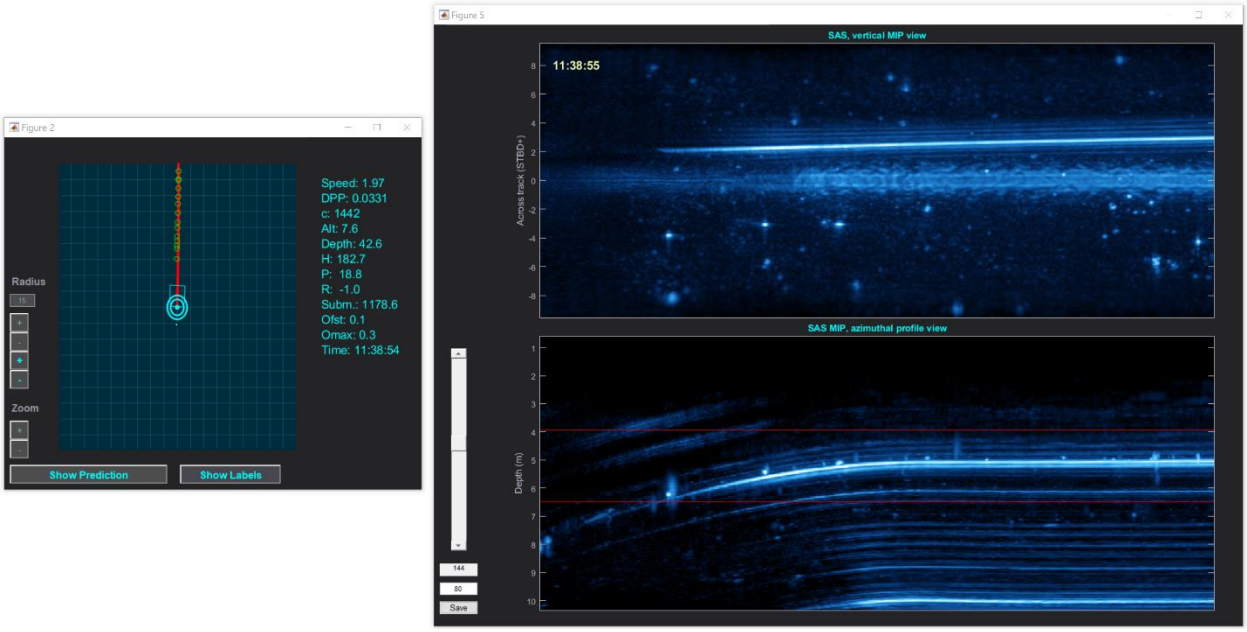

Figure A. 7: The towfish begins to surface and the cable moves out of the view of the sonar.

In total, the towfish tracked the cable for >600 meters. The initial positioning was aided by some short segment surveys done on the previous day, (the markers from which are visible in some of the previous figures), but the long cable survey was trivial to perform. The initial hazard survey conducted with the high frequency sidescan sonar mounted on the MuST revealed a smooth mud interface, however due to burial no sign of the cable was visible. This short survey demonstrated the utility of the MuST system for localizing buried cables, as well as providing immediate navigation feedback for towfish operation.

Appendix B: list of publications

Below is a list of papers published as part of this project. Note that one paper, ("Spatially Variant Autofocus for Circular Synthetic Aperture Sonar"), covers developments from a previous SERDP research project; the drafting and publishing of this paper was performed with permission by SERDP. Note that some of the topics contained in the paper are relevant to the current project and informed some aspects of the research performed.

Additionally, the 3D FFBP algorithm has only been published as an IEEE Oceans conference preceding; a larger scale submission showing results with the MuST system is planned for submission to a peer reviewed journal such as the Journal of the Acoustical Society of America or IEEE Journal of Oceanic Engineering.

Conference papers:

- T. M. Marston, K. L. Williams and D. S. Plotnick, "Sub-aperture recombination for sliding window SAS processing in 3D down-looking synthetic aperture sonar," EUSAR 2021; 13th European Conference on Synthetic Aperture Radar, 2021, pp. 1-6.
- T. M. Marston, D. S. Plotnick and K. L. Williams, "Three dimensional fast factorized back projection for sub-sediment imaging sonars," OCEANS 2019 MTS/IEEE SEATTLE, 2019, pp. 1-6, doi: 10.23919/OCEANS40490.2019.8962594.
- T. M. Marston, S. G. Kargl, D. S. Plotnick and K. L. Williams, "Improving downlooking SAS resolution with transmitter spatial diversity," OCEANS 2019 MTS/IEEE SEATTLE, 2019, pp. 1-5, doi: 10.23919/OCEANS40490.2019.8962536.
- K. L. Williams, S. G. Kargl, A. L. Espana, T. M. Marston and D. S. Plotnick, "FROM BASIC RESEARCH TO A PROTOTYPE TOWBODY PLATFORM: PROGRESS ON UNDERWATER UNEXPLODED ORDNANCE REMEDIATION." UACE2019 Conference Proceedings, 2019.

Journal articles:

- T. M. Marston and J. L. Kennedy, "Spatially variant autofocus for circular synthetic aperture sonar." J. Acoust. Soc. Am. 149, pp. 4078 – 4093, May 2021. doi: 10.1121/10.0005198.