Seabed Remote Sensing by Single-Beam Echosounder: Models, Methods and Applications

by

Benjamin R. Biffard
B.Sc., University of Victoria, 2003

A Dissertation Submitted in Partial Fulfillment
of the Requirements for the Degree of

DOCTOR OF PHILOSOPHY

in the School of Earth and Ocean Sciences

Supervisory Committee

Seabed Remote Sensing by Single-Beam Echosounder: Models, Methods and Applications

by

Benjamin R. Biffard
B.Sc. (Honours), University of Victoria, 2003

Supervisory Committee

Dr. N. Ross Chapman, School of Earth and Ocean Sciences
Supervisor

Dr. Stan E. Dosso, School of Earth and Ocean Sciences
Departmental Member

Dr. Michael J. Wilmut, School of Earth and Ocean Sciences
Departmental Member

Dr. Jon M. Preston, School of Earth and Ocean Sciences
Departmental Member

Dr. Colin J. Bradley, Department of Mechanical Engineering
Outside Member

Abstract

Supervisory Committee
Dr. N. Ross Chapman, School of Earth and Ocean Science
Supervisor
Dr. Stan E. Dosso, School of Earth and Ocean Science
Departmental Member
Dr. Michael J. Wilmut, School of Earth and Ocean Science
Departmental Member
Dr. Jon M. Preston, School of Earth and Ocean Science
Departmental Member
Dr. Colin J. Bradley, Department of Mechanical Engineering
Outside Member

Single-beam echosounders are an inexpensive, practical and non-invasive means of remote sensing the seabed. Ideally, the common single-beam echosounder should be able to tell fishers, navigators, engineers and scientists what the seabed consists of in addition to water depth. Low-frequency underwater acoustic systems (<10 kHz) can do this in some circumstances, but are expensive, offer limited resolution and potentially hazardous to marine mammals. High-frequency systems, such as single and multi-beam echosounders, are very effective at mapping bathymetry, but do not characterize the seabed directly. Instead, these systems divide the seabed into self-similar segments or classes, and then rely on ground-truth data (usually sediment grab samples) to assign seabed-type labels such as sand, etc., to the classes. However, inadequate and inaccurate ground-truth is a major problem. Single-beam seabed classification methods also suffer from a lack of discriminatory power and from artefacts such as water depth and seabed slope. The cause of these problems is that the methods lack a basis in physics and are mainly statistical. Then, the central objective in this dissertation is to develop physics-based methods to improve classification and to address the problem of ground-truth by inferring seabed characteristics directly from the acoustics.

An overview of current methods is presented along with case studies of single-beam surveys to introduce the current seabed classification method called QTC VIEW™ and to identify specific problems. A physical basis is established in scattering and geometrical theories and observations of field and model data. This leads to new classification and characterization methods that overcome the shortcomings of current seabed classification methods. Advancements also include new physical models of echosounding. The new methods are presented, implemented and evaluated.

Highlights of experimental results include a new testbed located in Patricia Bay, British Columbia. The testbed consists of exhaustive ground-truth, surveys and novel controlled experiments with various single-beam echosounders, ranging in frequency from 12 to 200 kHz. Simulated echo time series data from the numerical BORIS model and a new analytic model are used to augment the testbed. Evaluation of experimental results shows the new physics-based methodology improves seabed classification significantly and enables seabed characterization by an uncalibrated single-beam echosounder.

Table of Contents

Supervisory Committee .... **ii**

Abstract .... **iii**

Table of Contents .... **v**

List of Tables .... **ix**

List of Figures .... **xi**

Dissertation Conventions .... **xvi**

Acknowledgments .... **xix**

Dedication .... **xx**

Chapter 1. Introduction **1**

1.1 Objectives and Hypothesis *9*

1.2 Research Field Overview *11*

1.3 Research Facility *12*

Chapter 2. Current Seabed Classification and Characterization Methods using Single-Beam Echosounders **13**

2.1 Fundamentals of Single-Beam Echosounders *13*

2.2 Seabed Classification by Statistical Segmentation – the QTC Method for Time Series Data *20*
- 2.2.1 Depth Compensation 22
- 2.2.2 Stacking and Features 28
- 2.2.3 Data Reduction and Clustering 31
- 2.2.4 Interpolation and Interpretation 35
- 2.2.5 Supervised Classification 36

2.3 An Overview of All Seabed Classification and Characterization Methods *37*

2.4 Similarities and Differences among Current Seabed Classification and Characterization Methods – Potential Avenues for Improvement *40*

Chapter 3. Applications of Current Methods in Single-Beam Seabed Classification **44**

3.1 The Kitimat Log Boom Debris Survey and the Problem of Seabed Slope *45*
- 3.1.1 Implications of Seabed Slope on SBES Seabed Classification 50

3.2 Discussion and Definition of the Research Problem *52*

Chapter 4. Existing Theory and Models **53**

4.1 Intuitive Overview *53*

4.2 Geoacoustic Parameters of the Seabed *58*

4.3 Scattering Theory *62*

4.4 Heald's Model of E1 and E2 .... *63*

4.5 The BORIS Numerical Echo Time Series Model *66*

4.6 A Comparison of the Echo Time Series Models *70*

4.7 Implications for Seabed Classification and Characterization Methods *71*

Chapter 5. New Theory and Models 76

5.1 The Setup and Calibration of BORIS for Modelling SBES *76*
5.1.2 Examples of BORIS Model Simulations and Simulated Surveys 77

5.2 The Analytic Echo Model *80*
5.2.1 Implications and Qualitative Verification of the Analytic Echo Model 84

5.3 The Echo Duration Model *89*
5.3.1 Initial Verification of the Echo Duration Model 95
5.3.2 Implications of the EDM for Seabed Classification on Sloping Seabeds 99
5.3.3 Implications of the EDM for Bathymetry on Sloping Seabeds 100

5.4 Modelling Studies and Demonstration of Concepts *102*
5.4.1 The Effects of Roughness 102
5.4.2 The Ensonification Regime Effect 110
5.4.3 Effective Beamwidth and Attenuation from Modelled Data – AEM and BORIS Results 118
5.4.4 Conclusion: Physics-Bases Features and Effective Depth Compensation 128

Chapter 6. Controlled Experiments in the Patricia Bay Testbed 129

6.1 The Patricia Bay Testbed *129*
6.1.1 EM3000 Seabed Classification Survey and the Selection of Experimental Sites 131
6.1.2 Experimental Site Seabed Classification Survey 136
6.1.3 Experiment Site Ground-truth survey 137

6.2 The Depth Experiment *144*
6.2.1 Experimental Setup 146
6.2.2 Data Preparation 147
6.2.3 Data Analysis Methods – Fitting the EDM to Echo Duration Data 150
6.2.4 Results from Field Data 153
6.2.5 Discussion and Comparison to Previous Results 162

6.3 The Slope Experiment *165*
6.3.1 Compensation of Slope – The SELw Method 166
6.3.2 Data Analysis 167
6.3.3 Discussion and Conclusions 174

6.4 A Discussion of Classification Errors Caused by Seabed Slope and Standard Depth Compensation Parameters *175*

Chapter 7. New and Improved Seabed Classification Methods 177

7.1 An Overview of the New Seabed Classification Methodology *178*

7.2 Analysis Methods – Tools for Evaluating the Performance of New and Improved Classification Methods *180*

7.3 The Baseline to Compare the New Methods to: QTC IMPACT (TNORM) Processing for Envelope Data *188*
7.3.1 The Patricia Bay Testbed Full SBES Survey 189
7.3.2 Standard QTC IMPACT Processing with TNORM 190
7.3.3 Evaluation of the TNORM Classification Results 192

7.4 Removal of Non-Seabed Influences: Effective Methods of Depth and Slope Compensation *198*
  7.4.1 A Synopsis of The Intermediate Depth and Slope Compensation Methods 198
  7.4.2 The Implementation of SELw Effective Depth and Slope Compensation 200
  7.4.3 Qualitative Evaluation of SELw Depth and Slope Compensation 202
  7.4.4 Quantitative Evaluation of SELw Depth and Slope Compensation 205

7.5 Improved Discrimination: Physics-Based Features *206*
  7.5.1 The New Features 207
  7.5.2 Evaluation and Reduction of the Combined Feature Set 216
  7.5.3 Classification Results and Discussion 226

7.6 Supervised Classification *233*

7.7 Conclusions for Improved Classification *236*

7.8 Three-feature Classification and Relative Characterization *237*
  7.8.1 Relative Characterization by the Interpretation of One-Colour-Per-Record Plots 238
  7.8.2 Three-feature Classification Results 244

Chapter 8. Post-Class Seabed Characterization 249

8.1 Proof-of-Concept Methods *252*
  8.1.1 Amplitude Variability Analysis 252

8.2 Mean Envelope Inversion *256*
  8.2.1 Description of the Method 256
  8.2.2 Calibration and Testing with BORIS Data 258
  8.2.3 Mean Envelope Inversion Results 261
  8.2.4 Mean Envelope Inversion Conclusion 264

8.3 Attenuation by Echo Length (ABEL) *265*
  8.3.1 Description of the ABEL Method 265
  8.3.2 ABEL Results 267
  8.3.3 ABEL Conclusion 268

8.4 Combining the Characterization Methods for a Single Result *270*

8.5 Conclusion for Characterization Methods *272*

Chapter 9. Conclusion 273

Appendix A. Picking – the Detection of the Onset and Termination of Echoes 275

A.1 Bottom Picking *275*

A.2 Tail Picking *277*

Appendix B. Data Reduction and Principal Component Analysis 282

B.1 PCA *282*

B.2 PCA Iterative Outlier Removal *290*

Appendix C. Tabulations of Geoacoustic and Geotechnical Parameters 293

Appendix D. Scattering Theory for Single-Beam Echosounders 296

D.1 Volume Scattering Theory *296*

D.2 The Helmholtz-Kirchhoff Integral Equation *299*

D.3 Rayleigh Theory *301*

D.4 Perturbation Theory *304*
D.4.1 First-Order Perturbation Theory 305
D.4.2 Second-Order Perturbation Theory 307

D.5 Kirchhoff Approximation Theory *308*

D.6 Special Tractable Cases of KA Theory *312*
D.6.1 KA Theory with Zero Roughness 312
D.6.2 The Coherent Field from General and Gaussian Roughness 312
D.6.3 The Geometrical Rayleigh Criterion, Bragg and Lambertian Scattering 315
D.6.4 The Incoherent Field from General Roughness 316
D.6.5 The Total Mean Intensity in The Case of Gaussian Roughness 317
D.6.6 KA Theory Extensions Required to Model SBES 319

Appendix E. Field Definitions 321

Appendix F. The Setup and Calibration of the BORIS Model for the Simulation of SBES (From Section 5.1) 323

F.1 Seabed and Echosounder Parameters for the BORIS Model *324*

F.2 Other BORIS Model Parameters *329*

Appendix G. Observations of Echo Characteristics and Potential New Features in a BORIS Simulated Survey (From sub-section 5.4.4) 331

Appendix H. Dissertation Outline, Summary of Novel Contributions and Publications (From Chapter 9) 341

H.1 Publications to Date *344*

References 345

List of Tables

Table 2-1. The feature set used in the QTC seabed classification method 30
Table 4-1. A summary of the effects of geoacoustic parameters on the echo components 74
Table 5-1. Measurements of effective beamwidth compared to model roughness for BORIS data....................124
Table 5-2. Measurements of effective attenuation compared to the known values for BORIS data 125
Table 6-1. A summary of ground-truth data for each multibeam class from Patricia Bay 132
Table 6-2. Grain size analysis for the Gravel site (CM3), listed by grab location within the site 139
Table 6-3. Grain size analysis for the Sand site (CM5), listed by grab location within the site. 139
Table 6-4. Table of results for the depth experiment....................155
Table 6-5. Summary of tilt limits for reliable classification 173
Table 7-1. Confusion matrix for the Airmar38 TNORM processed survey 183
Table 7-2. Square confusion matrix for the Airmar38 TNORM testbed survey 183
Table 7-3. Full confusion matrix for the Simrad200 TNORM testbed survey 183
Table 7-4. Square confusion matrix for the Simrad200 TNORM testbed survey. 184
Table 7-6. Classification performance for the TNORM results 193
Table 7-7. Regression results for effective beamwidth and attenuation without any slope information 199
Table 7-8. Measures of effective beamwidth and attenuation with slope information 200
Table 7-9. Quantitative measures of classification performance in the removal of non-seabed influences 205
Table 7-10. Average absolute correlation coefficients to depth and seabed slope 206
Table 7-11. A summary of the new feature families. 207
Table 7-12. Thresholds for the cumulative amplitude and energy threshold features. 208
Table 7-13. The relative amplitude feature family. 211
Table 7-14. The measures of time feature family. 211
Table 7-15. The features of the three feature families that measure aspects of variability. 213
Table 7-16. The AEM cross-correlation feature family 215
Table 7-17. The spectral moments feature family 215
Table 7-18. Feature statistics from the Airmar38 testbed survey and the reduced feature set. 220
Table 7-19. Feature statistics from the Simrad200 testbed survey and the reduced feature set. 222
Table 7-20. AMI and OA statistics comparing the baseline classification result to the final classification 232
Table 7-21. AMI and OA statistics comparing the reduced feature set to the less-reduced feature set. 233
Table 7-22. Relative strengths of the major echo attributes and expected three-space colours 239
Table 7-23. The features used for three-feature classification. 240
Table 7-24. Colour interpretation table for the Airmar38 three-feature classification of Figure 7-21. 243
Table 7-25. AMI, Kappa and OA statistics comparing TNORM to SELw3feat. 248

Table 8-1. RMS roughness height estimated by amplitude variability analysis 253

Table 8-2. Model parameters and bounds for mean envelope inversion 257

Table 8-3. The calculation of correlation lengths from roughness ratio estimates from MEI 259

Table 8-4. MEI class characterization results for the Airmar38 testbed survey.. 263

Table 8-5. Free regression ABEL results for the Airmar38 SELwNF classification..................................269

Table 8-6. Fixed effective beamwidth ABEL results for the Airmar38 SELwNF classification. 270

Table 8-7. Feature interpretation table for the main classes of the Airmar38 SELwNF classification 271

Table C-1. SBES parameters used throughout this dissertation. 294

Table C-2. Sediment parameters from the APL-UW models, [APL-UW, 1994].......................................294

Table C-3. A list of the parameter values used to simulate seabeds for BORIS simulations. 295

List of Figures

Figure 1-1. Diagram of side-scan sonar deployment geometry 5

Figure 1-2. Multibeam echosounder geometry showing the combined beams.. 6

Figure 1-3. A depiction of SBES operation with a conical beampattern, hypothetical seabeds and the corresponding received echo signals 8

Figure 2-1. Schematic of a standard SBES survey setup.. 15

Figure 2-2. Beampattern cross-sections plotted in *angular* co-ordinates: (*left*) a generalized transducer with high amplitude side lobes, (*right*) the corresponding conical aperture approximation................................. 16

Figure 2-3. A conical beam incident on a flat seabed with the definitions of the seabed normal, the transducer axis and the incident angle 16

Figure 2-4. Flowchart of the QTC VIEW seabed classification methodology. 23

Figure 2-5. A cartoon plot of echo duration, measured in samples after reference depth compensation, for two seabed types.. 25

Figure 2-6. A random assortment of 25 boxes to be sorted into groups of similar boxes as an analogy for the process of statistical segmentation employed in the QTC seabed classification method. 29

Figure 2-7. A diagram showing the calculation of the cumulative integral shape features, adapted from [Lurton and Pouliquen, 1992; Pouliquen and Lurton, 1992], by J. Preston. 30

Figure 2-8. Four groupings of 25 boxes as an analogy to the data reduction and clustering steps in the QTC seabed classification methodology. 31

Figure 2-9. An example of Q-values plotted and clustered in three-dimensional Q-space, as seen in the QTC ACE software 33

Figure 2-10. The ACE result for the Comox area, Pacific Sandlance survey [Biffard et al., 2009].............. 35

Figure 3-1. Location map for all seabed surveys and experiment sites for this dissertation 45

Figure 3-2. An example raw SSS mosaic from Clio Bay. 46

Figure 3-3. SBES seabed classification for Eagle Bay overlaid on to seabed slope..................................... 48

Figure 3-4. SBES seabed classification for Clio Bay overlaid on (*lower*) sunken log density and (*upper*) seabed slope calculated from survey bathymetry 49

Figure 4-1. The formation of a SBES seabed echo shown by backscattering components. 57

Figure 4-2. Geometry for the derivation of Heald's model of E1 64

Figure 4-3. Examples of simulated surface heights and volume inhomogeneities. 69

Figure 4-4. The fractional contribution of the coherent field to the total field: RMS roughness set at 1 cm, frequency fixed at 50 kHz 72

Figure 4-5. A depiction of the four major components of a typical echo 74

Figure 5-1. Example BORIS echo time series.. 77

Figure 5-2. BORIS simulated echo amplitude time series as an example of a controlled experiment for varying water depths 78

Figure 5-3. An echogram of a simulated survey of the Suzuki50 SBES on the SandAG seabed.................. 80

Figure 5-4. A comparison of beampattern models 82

Figure 5-5. AEM simulated echo times series for a number of depths..85

Figure 5-6. AEM simulated time series for a number of pulse durations..85

Figure 5-7. AEM simulated echo time series for a number beamwidths...86

Figure 5-8. AEM simulated echo time series for a number of roughness ratios...86

Figure 5-9. Echo duration from AEM echoes for varying values of correlation length and RMS roughness height..88

Figure 5-10. Total echo energy measured from AEM simulations for varying roughness height and beamwidth. 89

Figure 5-11. Ray-trace diagrams for case I and II...91

Figure 5-12. Plots of echo duration as a function of depth and slope, and footprint major semi-axis length as a function of slope for typical SBES. 93

Figure 5-13. Measured and modelled echo duration for the 24 kHz survey of Eagle Bay............................96

Figure 5-14. Measured echo duration from the vicinity of Patey Rock, B.C 98

Figure 5-15. The correction of measured bathymetry for seabed slope based on Equation 5-4 for several typical SBES frequencies and beamwidths 101

Figure 5-16. An echogram of increasing roughness from BORIS simulations of the circ50_20 virtual echosounder on the HardFlat seabed...103

Figure 5-17. Average coherent, incoherent and total amplitudes for three domains of relative roughness.105

Figure 5-18. Echo duration vs. depth from BORIS simulations of the circ50_20 virtual echosounder on the HardFlat seabed...108

Figure 5-19. AEM simulations of echo duration as a function of roughness. 110

Figure 5-20. The four stages of ensonification visualized with a conical beam...111

Figure 5-21. Qualitative representation of ensonification area and a typical echo shape as a function of time for full ensonification and annular ensonification...112

Figure 5-22. AEM and BORIS simulated echo time series at several depths around the ensonification critical depth..114

Figure 5-23. Echo envelope time series at three depths from the Odom24 echosounder at the sand site. ..115

Figure 5-24. AEM simulated echoes with the pulse duration adjustment..116

Figure 5-25. AEM simulated echoes with and withoutpulse length adjustment 117

Figure 5-26. Echo duration as function of depth for AEM simulated echoes for a series of different beamwidths 118

Figure 5-27. Examples of linear regressions of echo duration to water depth from BORIS simulations....122

Figure 6-1. The MBES seabed classification map of Patricia Bay, B.C. 133

Figure 6-2. Seabed slope, as extracted from EM3000 multibeam bathymetry for Patricia Bay..................134

Figure 6-3. Sun-illuminated bathymetry from the EM3000 multibeam survey of Patricia Bay..................135

Figure 6-4. Seabed classification result for the Odom24 survey of the experiment sites...........................136

Figure 6-5. An enlargement of the Patricia Bay multibeam seabed classification map for the sand - gravel transition site (CM2) 142

Figure 6-6. Grain size distributions for the Sand - Gravel Transition Site (CM2). 142

Figure 6-7. The STING sounding result at location CMARS 2-5. 143

Figure 6-8. An enlargement of the Patricia Bay multibeam seabed classification map for the complex sloping site (CM1). 144

Figure 6-9. Apparatus setup aboard the CGR research boat for the depth dependence experiment. 147

Figure 6-10. An example trace with multi-path echoes. 149

Figure 6-11. Ray diagram for multi-path echoes in the depth experiment 149

Figure 6-12. Linear regressions of echo duration to depth for the depth experiment 154

Figure 6-13. Histograms of echo duration residuals. 157

Figure 6-14. Echo duration residuals plotted against altitude and auto-correlation analysis for the Odom24 over the sand site. 159

Figure 6-15. Beampattern profiles for the Odom dual-frequency echosounder 163

Figure 6-16. Three features calculated from Odom24 kHz echoes collected over the gravel site without slope compensation 170

Figure 6-17.The same three features calculated from the same echoes as in Figure 6-16 except with tilts compensated by using the transducer tilt values as the seabed slope in Equation 6-1. 171

Figure 6-18. The same three features as before calculated from echoes collected over the sand site (Odom24). 172

Figure 6-19. The same three features calculated from the same echoes as in Figure 6-18 (i.e. Odom24 over the sand site), except with tilts compensated for by using the transducer tilt values as seabed slope 173

Figure 6-20. An example showing the calculation of percent misassignment 174

Figure 6-21. Ratio of echo duration calculated with effective parameters to echo duration calculated from the standard or default parameters 176

Figure 7-1. The new seabed classification methodology, including SELw depth and slope compensation, physics-based features, three-feature classification and post-class characterization. 178

Figure 7-2. A summary of the methods used to evaluate the new and improved classification methods....180

Figure 7-3. The relationships between the three statistical measures of classification performance 185

Figure 7-4. Airmar38 and Simrad200 seabed classification from the standard processing method overlaid on the EM3000 multibeam seabed map. 191

Figure 7-5. A comparison of the TNORM classification map for the Airmar38 testbed survey to the multibeam derived ground-truth class map 194

Figure 7-6. A comparison of the TNORM classification map for the Simrad200 testbed survey to the multibeam ground-truth class map.. 195

Figure 7-7. Maps of correct and incorrect classification for TNORM processed survey data from the Airmar38 and the Simrad200 testbed surveys 197

Figure 7-8. Echo duration residual maps generated with slope information. 201

Figure 7-9. A comparison of the SELw classification map for the Airmar38 testbed survey to the multibeam derived ground-truth class map. 203

Figure 7-10. A comparison of the SELw classification map for the Simrad200 testbed survey to the multibeam derived ground-truth class map 204

Figure 7-11. TVG inaccuracy at various depth and echosounder frequencies for a two degree difference in temperature, based on Equation 7-3 and the Francois-Garrison model as presented in [Lurton, 2002]. The null at 97 kHz arises because absorption is a complicated function of temperature and frequency. 210

Figure 7-12. TVG inaccuracy at various depths and frequencies 210

Figure 7-13. Feature maps for the Airmar38 224

Figure 7-14. The correlation matrix for all available features from the Airmar38 survey as processed by SELwNF 225

Figure 7-15. The correlation matrix for all available features from the Simrad200 survey as processed by SELwNF 225

Figure 7-16. Airmar38 and Simrad200 SELwNF seabed classification overlaid on the multibeam seabed map. 228

Figure 7-17. A comparison of the SELwNF classification map for the Airmar38 testbed survey to the multibeam seabed map 229

Figure 7-18. A comparison of the SELwNF classification and OCPR maps for the Simrad200 testbed survey to the multibeam ground-truth class map 230

Figure 7-19. S-QC classification map for the Airmar38 survey, based on analysis of all features produced by the SELwNF method 235

Figure 7-20. Coloured three dimensional feature space plots for the Airmar38 and Simrad200 testbed surveys 240

Figure 7-21. The OCPR three-feature map for the Airmar38 testbed survey compared to the multibeam seabed map 241

Figure 7-22. The OCPR three-feature map for the Simrad200 testbed survey compared to the multibeam seabed map 242

Figure 7-23. 2-D projections of the Airmar38 3-D feature space showing the relationships between the features and the seabed types based on theory, models and observations. 245

Figure 7-24. The three-feature classification map for the Airmar38 testbed survey to be compared to the multibeam seabed map 246

Figure 7-25. The three-feature classification map for the Simrad200 testbed survey to be compared to the multibeam seabed map 247

Figure 8-1. A depiction of the bounds of applicability of the characterization methods relative to the major seabed parameters which are empirical functions of grain size 251

Figure 8-2. Geoacoustic parameters to be measured by the characterization methods of this chapter, plotted as functions of mean grain size calculated for the Airmar38 SBES 251

Figure 8-3. Echo peak amplitude histograms for the three major Airmar38 SELwNF classes 255

Figure 8-4. Marginal probability distributions of the AEM model parameters as found by Metropolis-Hastings sampling (part of MEI) 260

Figure 8-5. Mean envelopes for the three main classes of the Aimar38 SELwNF classification result. 262

Figure 8-6. Mean envelope for class 6 (gravel) as seen in Figure 8-5, overlaid by best-fit AEM echoes ...263

Figure 8-7. Marginal probability distributions for MEI on the mean envelope of class 6 (gravel) 264

Figure 8-8. Linear regression of slope-removed echo duration to depth for the echoes of class 1 (Airmar38 SELwNF). 269

Figure A-1: The QTC IMPACT bottom picking algorithm 276

Figure A-2. A BORIS-generated echo amplitude trace to illustrate the results of the various picking methods 279

Figure B-1. An example result of IMPACT PCA on Patricia Bay TNORM-processed 38kHz data 286

Figure B-2. An example scree plot and component loadings for standard covariance PCA as calculated from TNORM-compensated 38 kHz Patricia Bay data 287

Figure B-3. Example scree plot and weighted component loadings for MULTIVIEW type correlation PCA as calculated from TNORM-compensated 38 kHz data from Patricia Bay. 288

Figure B-4. Example scree plot and component loadings for standard correlation PCA. As calculated from TNORM-compensated 38 kHz data from Patricia Bay. 288

Figure B-5. Q-space plot of the Q-values of the Patricia Bay 38 kHz survey processed with the SELwNF method with correlation PCA before and after iterative outlier removal 291

Figure D-1. 10 dB skin or penetration depths for the seabeds of the APL-UW handbook (Table C-2), based on the skin depth model (Equation D-1), for typical SBES frequencies. 297

Figure D-2. Diagram for the derivation of Rayleigh scattering from a periodic rough surface...................302

Figure F-1. Transmit pulse and analysis for the Odom200 SBES 325

Figure F-2. The beampattern for the Odom200 SBES as constructed from two-sided beampattern profiles provided by the manufacturer...325

Figure F-3. Simulated beampattern for the virtual Circ50_20 SBES 326

Figure F-4. BORIS parameter test results for the surface roughness high-pass frequency. 328

Figure G-1. Measures of time for the example simulated survey with the Suzuki50, SandAG seabed 333

Figure G-2. Ratios of the measures of time for the example simulated survey 334

Figure G-3. Measures of amplitude and energy for the example simulated survey 334

Figure G-4. Energy ratios for the example simulated survey. 335

Figure G-5. The within-echo cumulative amplitude threshold features for the example simulated survey 336

Figure G-6. Power spectral density for the example simulated survey 337

Figure G-7. Measures of the centre frequency for the example simulated survey 338

Figure G-8. Spectral moments for the SandAG and HRHB simulated surveys. 338

Figure G-9. Measures of bandwidth as applied to FWF data from the SandAG seabed and the HRHB seabed 340

Dissertation Conventions

References to figures, equations and tables follow the following formats: Figure *C*-_, Equation *C*-_, Table *C*-_, [*author*, *date*], where *C* stands for the chapter number or appendix letter. In other to make cross-referencing easier, literature references are *always* in a format consisting of [*author*, *date*]. If a researcher's name is a subject in a sentence, the reference will follow. This referencing style is the same as Sternlicht's thesis and the recent book by Jackson and Richardson [Sternlicht, 1999; Jackson and Richardson, 2007]. This square-bracket style is preferred over the standard style of parenthetical referencing (author, date) as it is more consistent, not easily confused with other parentheses and readily searchable in electronic form (adobe PDF or MS WORD). Footnotes[1] interject information at the bottom of the page without interrupting the flow of the text. Appendices are used when the information required is lengthy and not directly necessary to all readers, or when the information is background information required in multiple places throughout the document, Appendices A, B and C for example.

Chapters are generally laid out with an introduction, followed by sections and sub-sections. Chapter indexing is as follows: *chapter.section.sub-section*, i.e. "sub-section 5.4.3". Chapters generally end with their own conclusion, often including a discussion of opportunities for future research, although major sections can also have such discussions and conclusions as appropriate. Because of this layout, the concluding chapter, chapter 9, is brief.

The following is a glossary of terms, abbreviations and mathematical/physical symbols used through this document. It only defines conventions specific to this dissertation. Standard conventions, such as 'sonar' and 'dB' are not defined.

---

[1] This is an example footnote.

Terms:

Echo – this term usually refers to the part of the trace that contains the first echo from the seabed. The term may be modified by adjectives volume or surface (in reference to the processes that generate the echo), coherent or incoherent (scattering terms defined in appendix E) and full-wave form or envelope (raw amplitude or de-modulated time series).

Seabed – the interface between overlying water and the solid surface of the Earth. The term is defined here to include the bottom of lakes and rivers. It also includes the volume of material below the interface to a depth at which CW signals between 10 and 500 kHz decrease by 10 dB (corrected for spherical spreading). It also includes flora, fauna and any solid material on the interface.

Trace – a time series of samples of acoustic pressure, usually digital and usually includes a recording of the transmit pulse, water column reverberation and seabed echoes.

Abbreviations:

ABEL – **A**ttenuation **b**y **E**cho **L**ength

AEM – **A**nalytic **E**cho **M**odel

AVA – **A**mplitude **V**ariability **A**nalysis

EDM – **E**cho **D**uration **M**odel

FWF – **F**ull-**w**ave **f**orm: the amplitude (or intensity) signal including the carrier wave. Envelope time series have the carrier removed by some de-modulation process.

BORIS – **BO**ttom **R**esponse from **I**nhomogeneities and **S**urface

MBES – **M**ulti-**b**eam **e**chosounder **s**ystem

OCPR – One **C**olour **P**er **R**ecord: a technique for displaying classification data

MEI – **M**ean **E**nvelope **I**nversion

QTC – **Q**uester **T**angent **C**orporation of Sidney, B.C., Canada.

QTC4 – QTC VIEW Series 4 data acquisition system

QTC5 – QTC VIEW Series 5 data acquisition system

QTC IMPACT – QTC **I**ntegrated **M**apping **P**rocessing **A**nd **C**lassification **T**oolkit

ROV – **R**emotely-**o**perated (underwater) **v**ehicle

SBES – **S**ingle-**b**eam **e**chosounder **s**ystem

SEL – **S**tandard **e**cho **l**ength depth compensation

SELw – Effective slope and depth compensation with variable length windows

SELwNF – SELw effective depth and slope compensation with the new reduced features

SELwNF+ – SELwNF with all non-biased features

SSA – **S**mall **s**lope approximation (part of the BORIS model)

SSS – **S**ide-**s**can **s**onar

<u>Symbols:</u>

d – Water depth as determined by the echosounder, which is the shortest distance between the echosounder and the seabed along a line defined by the seabed normal (chapter 2)

D – True water depth as defined by the distance from the water surface (maybe defined relatively) to the seabed along a vertical line (chapter 5)

θ – Defined as the effective beamwidth, which is a beamwidth characteristic of SBES that determines how echo duration increases with depth, effectively representing the beamwidth from which there is significant backscattering. This parameter is defined in relation to the EDM, see section 5.3.

θ_{-3dB} – The standard value beamwidth is defined by the point at which the beampattern's amplitude sensitivity decreases by 50%.

Acknowledgments

I would like to thank my supervisor Dr. Ross Chapman. Dr. Chapman provided critical guidance and an excellent working environment. This work benefited from his experience, patience and understanding. Dr. Chapman established collaborations that were vital in this work.

I would like to thank the members of our research group; particularly Dr. Steve Bloomer. Steve assisted with much of this work. I would like to thank Dr. Jon Preston for his guidance. There are fewer than five scientists with his knowledge in this very specific field. I would also like to thank Dr. Stan Dosso for contributing to the inversion section and Dr. Mike Wilmut for contributing to the evaluation of my results by quantitative statistical methods. I would like to thank the members of my dedicated supervisory and examination committees, including Dr. Colin Bradley and Dr. Philippe Blondel, the external examiner.

I would like to thank our collaborators. Foremost is Quester Tangent Corporation. Quester Tangent contributed funding support via a National Sciences and Engineering Research Councils' industrial post-graduate scholarship, an Idea to Innovation project and a Collaborative Research and Development project. Quester Tangent provided research support from their marine division, particularly efforts of Jon Preston, Karl Rhynas, Bill Collins and Tom Younger. Our other major collaborator is the Canadian Hydrographic Service – particularly Jim Galloway and the sonar systems research group. I would like to thank Gaetano Canepa and SACLANTCEN (NURC) for providing the BORIS model. I would like to thank Dr. Cliff Robinson of Parks Canada for leading the Pacific Sand Lance habitat project. I would like to acknowledge our collaboration with the NEPTUNE and VENUS projects and ROPOS for the opportunities to collect ROV-mounted echosounder data. And finally Rick Linden (Coastal Geoscience Research Corp.), our boat operator for the slope and depth experiments.

Additional funding support is from the Natural Sciences and Engineering Research Council, the Canada Foundation for Innovation and the BC Advanced Systems Institute.

And finally, I would like to acknowledge the support of my wife Maura, our little dude Jonathan, my family, particularly my mom and dad, and my friends.

Dedication

To the pursuit of Science.

May the knowledge gained through this work be used to
explore the ocean and preserve nature.

Chapter 1. Introduction

There is tremendous interest in mapping the seabed, not only in knowing the water depth above the seabed, but also in knowing the composition of the seabed. However, relatively little of the world’s seabed has been mapped even though the seabed makes up 70.8% of the crust of planet Earth. And yet so much hinges on knowledge of the seabed. Many groups seek seabed maps and information. Fishers, fisheries management and researchers seek habitat-type information for both exploitation and conservation. Many seabed-mapping technologies were initially developed for military purposes. A particular motivation is in mine counter-measures; the seabed was mapped to locate areas of low bearing strength where anti-shipping mines could be buried to avoid detection. The nature of the seabed is also of interest for geologists. Navigation information is an extremely important application of seabed mapping for recreational boaters, commercial and military shipping. Seabed information may be used to set anchors. Other uses of seabed information include cable route planning, installation of offshore infrastructure, resource extraction (oil, gas, aggregates, minerals, etc.), seabed hazards (unstable slopes, gas hydrates and gas pockmarks) and underwater archaeology. These are just a few of the many applications of seabed information.

Early efforts in seabed mapping were called 'sounding' and the term is still accurate today. In general, this refers to probing the environment by sending out a stimulus. From ancient times until the sinking of the Titanic in 1912, sounding in the ocean exclusively meant feeding out a weighted rope until it went slack; the length of the rope was the water depth. Affixing a sticky substance to the end of the rope to sample the seabed was the next innovation. This is called lead lining and much of the seabed type information on contemporary marine charts is from lead line data. Nowadays, the term sounding is even more relevant as the probing stimulus is, ironically, sound. With the sinking of the Titanic in 1912, acoustic means were developed to detect and avoid icebergs. English meteorologist Lewis Richardson obtained the first patent for an underwater acoustic echo ranging device in 1912, and a German physicist Alexander Behm obtained a patent for an echo sounder in 1913. Canadian Reginald Fessenden built the Fessenden oscillator in 1912 and demonstrated it could detect icebergs up to 3 km away. Acoustic remote sensing systems were then developed for anti-submarine warfare through World War I and II, and concurrently used to sound for water depth for navigation. The technology was dubbed *sonar*, an acronym for **so**und **n**avigation **a**nd **r**anging. The electromagnetic equivalent of sonar, *radar*, shares many properties with sonar, including that both systems send out soundings or *transmit pulses* and detect the echoes from their target. Remote sensing of seabed topography (or bathymetry) is called hydrography. Today, hydrographers and hydrographic agencies are charged with mapping the seabed using the most advanced techniques available.

Contemporary seabed sensing techniques include underwater acoustics, physical methods, video, lidar and electromagnetic techniques. Modern physical methods also include sampling methods and various in situ measurements such as penetrometers. There are two types of penetrometers: those that are driven into the seabed or those that free-fall to impact with the seabed. Upon impact, the deceleration curve is measured and used to infer seabed characteristics – free-fall penetrometers will be applied in this dissertation. There are numerous sampling methods. The most common is the sediment grab sampler. It takes a scoop of the upper ~20 cm of the seabed – this technique is applied extensively in this dissertation. Grab samplers are subject to several problems: fine sediments can be washed out during retrieval and processing, cobbles may jam open the grab or be left behind, any layering in the sediments is destroyed and most importantly, surface roughness can only be inferred from empirical relationships to mean grain size. In general, physical sampling methods lack spatial coverage and may drift from deployment location. Other sampling methods include push cores, gravity cores, freeze cores and alike. Towed video is a technique that is limited in field of view and is labour-intensive; there are very few automated and non-subjective approaches to analyze this type of data[2]. Also, video techniques do not offer information on the seabed volume, and thin veneers may mislead video interpretation. In this dissertation, video is primarily obtained by mounting a camera on the grab sampler; the purpose of this system to verify the type of seabed, the qualitative level of seabed roughness, and to ensure sample quality.

---

[2] The SIMS system of Coastal Ocean Resources Inc. is an example of a quantitative video interpretation system, see: http://www.coastalandoceans.com/capabilities/sims.htm

Hydrographic lidar (**li**ght **d**etection **a**nd **r**anging) systems use lasers to send out pulses of light, the light reflections or echoes are then detected for seabed range and composition; for an example see [Wang and Philpot, 2007]. However, LIDAR is limited to shallow waters not more than a few Secchi disc depths (a measurement of underwater visibility) [Galloway, 2008] as light attenuates very quickly in seawater. Electro-magnetic remote sensing techniques (also called controlled-source electromagnetic (CSEM)) generally consist of towed resistivity measurements that can determine the fraction of water to sediment in the seabed (known as porosity) [Edwards, 2005]. Other geophysical techniques such as gravity and magnetic surveys are limited in spatial resolution and do not offer useful information on the upper layers of sediment, vegetation or rock that make the portion of the seabed that is of interest here (defined as the upper five metres of seabed).

Underwater acoustics is the best means of mapping the seabed. Depending on the specific technique and frequency, acoustic waves can provide extremely high resolution bathymetry (of order 0.1 m) sensing only the seabed-water interface or, in the other extreme, acoustic waves can penetrate through the entire planet but cannot resolve bathymetry to a sufficient degree. The former is referred to as high-frequency underwater acoustics and operates at frequencies from 10 kHz to 500 kHz, while the latter is referred to as seismics and operates from a few Hertz to 5 kHz. Seismic techniques are applied to map geological formations and to find mineral and hydrocarbon resources. Passive seismic techniques listen for seismic waves from sources such as earthquakes. Active sounding can be done with impulsive sources like air guns and dynamite explosions, or more controlled sources like piezo-electric transducers. Seismic techniques also make use of different types of acoustic waves; there are four main types: compressional, shear, Rayleigh and Love waves. All of these waves are propagated within an elastic solid with elastic forces being the restoring force. Water (saline or not) and air are elastic fluids, meaning they only support compressional body waves and not shear waves as fluids have no resistance to shear forces. High-frequency underwater acoustics makes use of compressional waves commonly known as sound waves.

Remote sensing with high-frequency underwater acoustics can be further subdivided into several different applications. Prominent examples are discussed here. In strict definition, the term *hydroacoustics* usually refers to the sensing of sediments, plants and animals suspended in the water volume, but can also refer to mapping bathymetry. Even by taking its strict definition, hydroacoustics is still a large field with many techniques and applications. Very briefly, normal applications entail the use of a single-beam echosounder (SBES), while more advanced techniques use multi-beam echosounders (MBES). The usual purpose of hydroacoustics is to measure fish or plankton biomass, or, as used by fishers, to detect schools of fish to facilitate capture. Back to the broader field of underwater acoustics, for applications in navigation and science, high-frequency *Doppler* sonar is used to track the movement of underwater vehicles relative to the seabed and to track the motion of water (in currents or waves) or animals. High-frequency underwater acoustics is also used for communications with modest rates of data transmission (up to several kilobits per second over several kilometres distance), or analogue voice-type transmission [Lurton, 2002]. High-frequency acoustic waves also enable accurate underwater positioning via triangulation – a method called *USBL* – **u**ltra-**s**hort **b**ase**l**ine navigation. USBL makes use of a small array of receivers and a transponder on the object to be positioned – usually **r**emotely **o**perated underwater **v**ehicles (ROVs). The last application to discuss here is the subject of this dissertation – acoustic seabed remote sensing – the use of high-frequency underwater acoustics to map the seabed for bathymetry and for seabed characteristics.

Generally, there are three types of sonar used for acoustic seabed remote sensing: single-beam echosounders, multi-beam echosounders and side-scan sonar (SSS). Each type of system has various applications, advantages and disadvantages in terms of seabed remote sensing. Side-scan sonar is the least used of the three in this dissertation. The acoustic transducer used for SSS is an elongated line array, so that the beampattern is a fan projected perpendicular to the path of survey vessel or towed-body. The advantage of this geometry is that the *echogram* forms a very useful image of the seabed, as shown in Figure 1-1. The disadvantage is that the image is not spatially consistent nor can bathymetry be generated. The backscatter image can be interpreted for seabed classification by an expert user, or by statistical segmentation methods applied in products such as QTC[3] SIDEVIEW™[4] [Preston et al., 2004a] and TexAn™ [Blondel et al. 1998], to name a few. Within the last five years, bathymetric side-scan sonars have been developed. These systems determine the beam angle of seabed echoes through interferometry and multiple arrays. They produce basically the same results as a multi-beam echosounder, but have the advantage of nearly unlimited swath angle allowing bathymetry to be collected into very shallow-water broad-side of the survey vessel [Galloway, 2008]. Another recent advancement is synthetic aperture sonar; it improves the SSS image quality by combining multiple soundings (or *pings*) as the SSS traverses a target. This implementation is very similar to synthetic aperture radar. Indeed, many techniques applied to sonar have radar analogues.

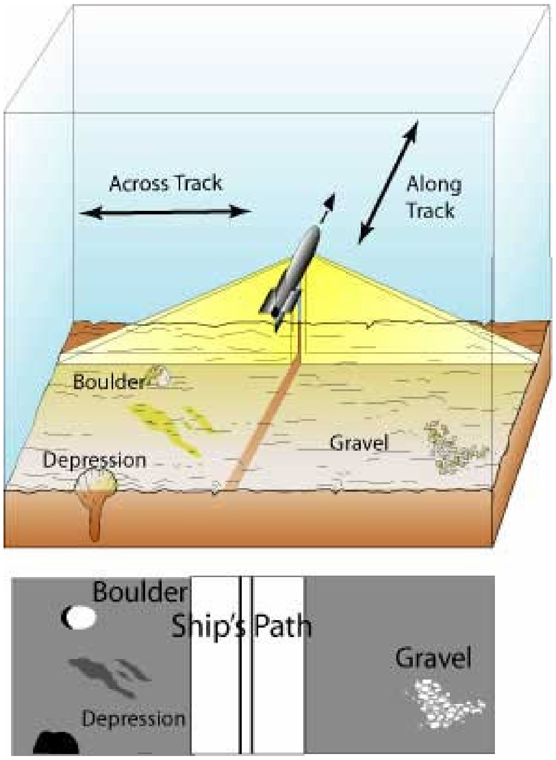

Figure 1-1. Diagram of side-scan sonar deployment geometry. (*upper*) The tow vehicle contains two SSS transducers, one for each side, projecting beams perpendicular to the survey vessel's path. (*lower*) the resulting grey-scale echogram that would be displayed on the echosounder's monitor or paper chart printout. (This is a public domain image, courtesy of the United States Geological Service.)

---

[3] QTC is an acronym for Quester Tangent Corporation, a company based in Sidney, Canada that produces and markets seabed classification technology.

[4] As of spring 2010, the QTC SideView™ software has been merged with QTC MultiView™ to become QTC SwathView™

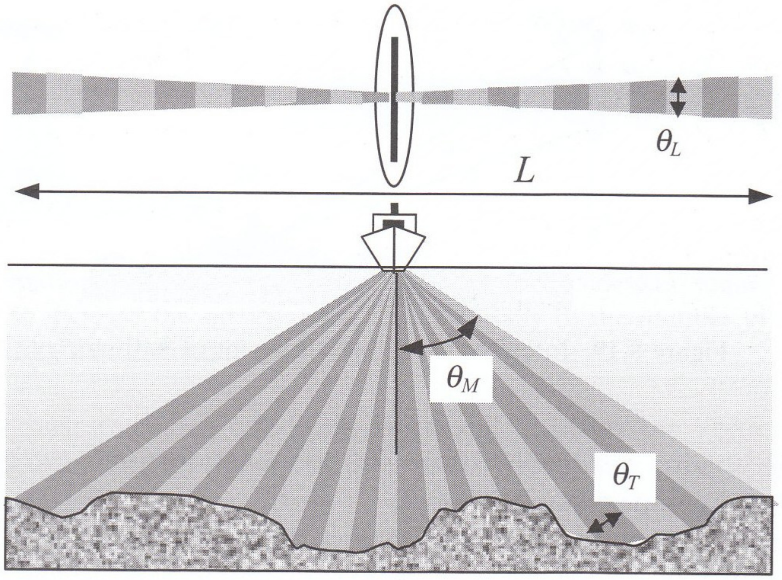

Figure 1-2. Multibeam echosounder geometry showing the combined beams. (*upper*) From above, the along-track aperture is θ_L, and the overall swath width L is a function of the water depth and across-track aperture θ_M. (*lower*) In the vertical view, individual beam aperture or beamwidth is θ_T. From [Lurton, 2002], page 268.

Multi-beam echosounders have become the instrument of choice for mapping bathymetry. They combine the areal characteristics of SSS with beam-forming to measure bathymetry across a section of seabed up to 7.5 times the water depth, as shown in Figure 1-2. The full process of beam-forming and processing the digital data for the onset of the seabed echoes (a process often known as *picking*) is complex; for a good synopsis see [Lurton, 2002]. MBES generally work by transmitting a fan-shaped pulse of sound perpendicular to the survey vessel's course (like a SSS system); a second array of transducers listens for the seabed echo signal. Beam-forming in the receive array constructs fan-shaped beams that are wide in the direction of the survey vessel track and narrow in the cross-track direction. The receive beams are perpendicular to the transmitted beam, so that the combination of the two beampatterns forms beams that are very small in angular extent as depicted in Figure 1-2; some systems achieve spatial resolutions better than 1 °. Because of these narrow beams, the system is very sensitive to vessel attitude and variations in sound speed throughout the water column. With the addition of appropriate attitude and heave sensors, sound velocity profiler systems and the requirement of large robust hull mounts, MBES become very expensive compared to

SSS and SBES. However, the quality and coverage of data produced by MBES is superior, especially for bathymetry and calibrated backscatter amplitudes. Generally, MBES do not penetrate the seabed volume to the extent that SBES can nor do they offer the same width swath imagery of SSS.

MBES are well suited to measure the backscattering strength as a function of angle. These angular response curves are characteristic for well known sediments although ambiguities arise between the contributions of roughness, impedance and volume heterogeneities [Fonseca et al., 2007]. Composite sediments (such as muddy gravel), or complications such as bioturbation, vegetation, etc., also pose problems. A prominent MBES characterization method is GeoCoder™, which is described in [Fonseca et al., 2007]. Backscatter images corrected for angle of incidence are used in statistical segmentation algorithms to separate the seabed into distinct seabed classes; this takes advantage of the large amount of data in the backscatter images. A prominent example is QTC MULTIVIEW™ [Preston, 2009; Preston et al., 2001; Preston et al., 2004b; Preston, 2004c]. There have been many studies comparing different multibeam data processing methods, e.g., [Robidoux et al., 2008], and combining multibeam classification with other data, e.g., [Pouliquen et al., 2002].

The sonar that is the subject of this research is the single-beam echosounder. SBES acoustic transducers and electronics are the simplest, cheapest, oldest and most mobile of the three major sonar types. Sub-bottom profilers are similar to SBES in that they are usually single-beam, but are distinct in that they operate in a moderate frequency range (1 to 10 kHz). Sub-bottom profilers do not resolve surface roughness features due to their long wavelengths. These systems are adept at mapping sub-bottom sediment layers down to hundreds of metres below the seabed surface – they are useful for marine geology. Sub-bottom profilers require large transducers, high power levels, often make use of wide bandwidth sound pulses (a *chirp* sonar), or combine sources of different frequencies in a *parametric* array.

SBES are the focus of this research for many reasons, including that they are the most common sonar system. SBES operate in what is generally considered the high-frequency range of underwater acoustics: 10 to 300 kHz. These systems are used as depth sounders onboard almost every ship, and also as fish-finders onboard most every fishing vessel. SBES are also used for hydrography and for habitat mapping by scientists. Most SBES, and all of the ones used in this research, use a short-duration transmit pulse of sound concentrated at one central frequency (described as CW for continuous wave). These systems project a single beam of approximately circular aperture (a conical shape) vertically downward towards the seabed and receive the seabed echo with the same transducer (i.e. a *monostatic* sonar), as shown in Figure 1-3. At these frequencies, the seabed surface is no longer mirror-like, instead it is a rough surface, causing the echo to be highly variable – referred to as incoherent or diffuse. The beam is usually sufficiently wide and the seabed sufficiently rough for backscattering from angles away from the seabed normal, and still at a low enough frequency for some backscattering from the sub-bottom material (referred to as the seabed volume) as will be shown in this dissertation. This complex interplay of surficial and volume backscattering modifies the shape of the seabed echo, as shown in Figure 1-3. The received signals then contain information on both seabed surface roughness and seabed composition, giving these systems good potential to characterize the seabed.

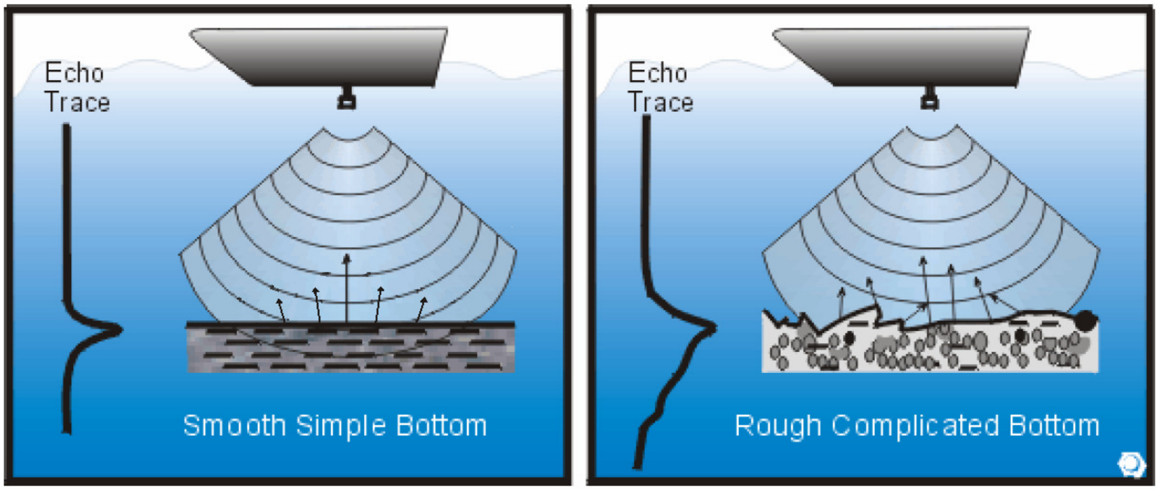

Figure 1-3. A depiction of SBES operation with a conical beampattern, hypothetical seabeds and the corresponding received echo signals (on the left of each pane). The signal of an echo is known as a *trace*. Courtesy of QTC, from [QTC, 2004a], modified to correct the surface backscatter (represented by arrows) in the smooth case.

As alluded to above, no acoustic seabed characterization method has been completely successful and there is much room for improvement in this area. Seabed classification methods for high-frequency sonar have been much more successful than their characterization counterparts have been, and with relatively simple methods as well. The reason for the difference in success is the difference in the aims of the two approaches. *Seabed characterization* is the direct determination of seabed characteristics without prior or additional knowledge of seabed properties. *Seabed classification* aims to draw borders around similar types of seabeds, grouping seabed unit areas (i.e. each sounding) together into classes. These classes are then identified with any of the aforementioned ground-truth techniques such as grab sampling or towed video (the two most common). Current seabed classification methods are *phenomenological* – they seek a qualitative description of the data without considering the cause of the phenomena. The methods are also mainly statistical, aiming to group or segment the data – therefore seabed classification is also described as *statistical segmentation*. Classification is a relative process, while characterization results in the measurement of one or several absolute quantities such as mean grain size. Characterization is much more difficult to achieve than classification, as it requires a full understanding of all of the physical processes that influence echoes. The relevant literature for current characterization and classification methods will be reviewed in detail in the relevant chapters later in this dissertation.

1.1 Objectives and Hypothesis

Given that current seabed classification methods are phenomenological and that characterization methods are not successful, the clear avenue for further research is to add a basis in physics to classification methods (effectively removing the phenomenological tag) and to use that basis to develop seabed characterization methods as well.

This dissertation focuses on SBES simply because SBES is the most common echosounder. Imagine the benefits if standard fish-finding and depth–sounding sonars could also indicate seabed type. Some of the many examples are shown below[5], but most

---

[5] Here are some practical examples of the benefits of standard SBES being made capable of discerning seabed type: protection of sensitive habitat areas from fishers, especially trawlers, improved anchor setting, decreased fishing effort, decreased loss of fishing gear. Seabed maps in general are very beneficial as

poignantly, shipwrecks due to anchoring mishaps could be avoided, saving lives. The benefits to science are immense, as discussed earlier. The two original SBES classification methods, QTC VIEW and RoxAnn, have facilitated much research; the new methods of this dissertation will hopefully facilitate many more discoveries.

There are many additional benefits to working on SBES: availability, cost, simplicity and ease of deployment, and more importantly, SBES have excellent potential for improved classification and characterization. MBES are becoming much more common, especially for seabed classification. However, the methods and the basis of these methods developed here for SBES could be extended to MBES to improve MBES classification and characterization methods.

In this dissertation, it will be shown that current seabed classification methods for SBES have problems with non-seabed influences on echo signals (causing errors) and can lack discriminatory ability (causing diminished precision). These problems will be explored, and solutions. This is the first of two major objectives of this dissertation - more specifically stated, the improvement of SBES seabed classification methods. Many of the advancements will be applicable to other methods, but the methodology this dissertation focuses on is the QTC VIEW™ approach[6]. QTC VIEW™ was chosen as it is widely used by scientists and is more robust to survey conditions and seabed types than its leading competitor RoxAnn™. In a recent dissertation, [Heald, 2000], a physical basis was established for the RoxAnn™ seabed classification method. This work will aim to do the same (and more) for the QTC methodology.

In another recent dissertation, a physical model was applied for SBES characterization by inversion technique [Sternlicht, 1999]. Although the techniques applied significantly advanced the field, only one method, inversion, was used. The second objective of this dissertation is to apply and combine several physics-based

---

mentioned earlier in chapter 1. For instance, a complete seabed map of coastal areas could be used to improve a number of models – current models, sound propagation models, and many different ecology models based on habitat areas.

[6] QTC View was the name of the combined hardware and software product that first employed the statistically-based seabed classification method to SBES data. The general approach is used in all QTC products. The current SBES product is called QTC IMPACT which is software only, while QTC VIEW now refers to the hardware only. The term QTC VIEW is still used to refer to the QTC methodology in the literature, but that practice will not be continued in this dissertation.

characterization methods, including inversion, to create a successful seabed characterization scheme.

The two main objectives may seem disjoint, but there are two strong links between them. First, ground-truth information is required to label seabed classes from seabed classification methods. However, the ground-truth can often be erroneous, lack spatial coverage or measure characteristics on different spatial scales than the acoustic techniques[7]. A successful seabed characterization scheme could be used as ground-truth information for labelling of seabed classes, effectively replacing standard ground-truth data sources such as grab sampling. SBES characterization results could be used to ground-truth spatially overlapping MBES classifications. The other link is that several characterization methods require (and can be improved by) large ensembles of data that are of the same seabed type. Creating large ensembles of data belonging to the same seabed type is exactly what seabed classification methods do. The classification method then effectively provides the characterization method with the type of data required for the latter to be successful. Therefore, the characterization scheme in this dissertation is called *post-class characterization*. It aims to characterize seabed classes generated by the improved physics-based classification method developed here. As a whole, this dissertation will aim to explore, develop and evaluate a comprehensive physics-based remote sensing methodology for single-beam echosounders.

1.2 Research Field Overview

The research field of seabed classification and characterization with SBES is somewhat limited. Some of the most comprehensive references are three Ph.D. dissertations: D. Caughey implemented the original QTC VIEW method [Caughey, 1996], D. Sternlicht focused on inversion of calibrated SBES data [Sternlicht, 1999] and G. Heald developed a physical basis for the RoxAnn method [Heald, 2001]. There are

---

[7] In chapter 6, an example of the limited spatial resolution of grab samplers is shown – the STING penetrometer shows a significant sub-bottom layer. Grab samplers often only sample the upper 15 cm of the seabed, while a 24 kHz SBES may penetrate many metres. [Kenny et al., 2003] has a table comparing the spatial resolutions of the various methods. [Galloway, 2004] identifies this problem. The shortcomings of standard ground-truth techniques will be explored in more detail throughout this dissertation.

few groups active in SBES research, while there is much more active research in MBES classification and characterization. However, MBES research not normally directly relevant to SBES research. Papers that are directly relevant to SBES are mostly applications or evaluations of the methods, the majority of which are published in conference proceedings. Papers applying and evaluating the methods tend to do so qualitatively. Quantitative, objective comparisons to ground-truth are rare. Rigorous application of statistics and physics is also rare.

Because of the limited nature of the field of SBES seabed classification, most readers of this dissertation are not experts in this specific field. Therefore, this dissertation provides a thorough background via chapters 1 through 4. The length of this dissertation is also due to the breadth of topics that are brought together to make a complete, inter-disciplinary, study of single-beam seabed remote sensing. This dissertation will provide quantitative evaluations were possible, and apply rigor via statistics and physics. Although there are drawbacks to working in such a field, including the need for a lengthy dissertation, there are benefits, as there is much scope to advance the field.

1.3 Research Facility

The research was conducted through the C-MARS facility, School of Earth and Ocean Sciences at the University of Victoria, Canada. The purpose of C-MARS (**C**anadian facility for **m**arine **a**coustic **r**emote **s**ensing) is to facilitate inter-disciplinary research in aquatic science by bringing together users of seabed information (biologists, geologists, etc.) to the equipment and expertise (geophysics) required to do the research. C-MARS also facilitates collaboration with industry. The collaboration with QTC has resulted in projects to develop new science and methodology to improve seabed remote sensing (chapters 5 through 8). In keeping with the C-MARS mandate, it is intended that the improved seabed remote sensing science and methods stemming from this research will be used to offer future users better seabed information to improve their research.

Chapter 2. Current Seabed Classification and Characterization Methods using Single-Beam Echosounders

There are a few good overviews of current seabed classification science and technology in the literature. The International Council for Exploration of the Sea (ICES) cooperative research report on acoustic seabed classification is probably the premier example [Anderson et al., 2007]. It has a thorough background plus literature reviews, summaries of methods and applications with recommendations for future directions. Anderson wrote a follow-up paper that focuses on comparing the different sonar platforms, survey design, different methods and future directions [Anderson et al., 2008]. ICES contributor Hamilton published an exhaustive list of acoustic seabed remote sensing literature [Hamilton, 2005]. A recent Ph.D. dissertation by Gleason has a thorough list of published applications of SBES seabed classification [Gleason, 2009]. Since the focus of this dissertation is on methodology, refer to the aforementioned literature for overviews of the numerous applications of seabed remote sensing. References pertaining to the methods presented will be shown where appropriate.

The purpose of this chapter is to present current seabed classification methods as a background to the applications of current methods chapter (chapter 3) and as a basis for the improved seabed classification method that appears in chapter 7. In this chapter, the basic fundamentals of SBES will be explored, and then the QTC SBES seabed classification methodology will be presented, followed by a review of the other methods, which will be compared to the QTC method to illustrate its advantages: robustness to any survey or seabed conditions, consistency of results. The shortcomings of the QTC method will be explored through applications of the QTC method in chapter 3.

2.1 Fundamentals of Single-Beam Echosounders

The fundamentals of SBES could encompass a great many topics, but instead this section focuses on fundamentals that will be built upon later in the dissertation. For more information, see a good reference on underwater acoustics such as [Lurton, 2002] or [Medwin and Clay, 1998].

A standard seabed classification survey apparatus consists of a number of components (Figure 2-1), including a differential global position system (d-GPS), a navigation computer, a data computer, a data acquisition unit and the SBES system. The functions of the navigation computer and the data computer may be combined, but usually a survey is run with at least two personnel – one monitoring the data quality and the other piloting the survey vessel along a survey grid using the navigation computer. The survey grid design is an important aspect of seabed surveys, with many factors influencing the design, including echosounder spatial resolution, time available and area to cover. Survey grids are orthogonal lines with even spacing. SBES surveys can be run from any vessel from a large ocean going ship or a small skiff with just one operator; the power source could be ship power or a pair of car batteries; both situations are represented in this dissertation.

A SBES system consists of an acoustic *transducer* and a head unit. The transducer converts electrical signals into sound pressure and vice versa, while the head unit produces the sounding pulse, then receives and processes the seabed echo signal for display and, in some cases, recording. Transducers are mounted permanently through the hull or temporarily deployed on poles; either is sufficient as long as the transducer is deep enough to prevent cavitation and be below air bubbles created by waves and the survey vessel's bow wake.

There are several hardware variations shown in Figure 2-1, the variations accommodate the varying capabilities of the echosounder head unit and data acquisition unit. For example, QTC VIEW data acquisition units digitize the acoustic data by eavesdropping on the analogue signals traveling between the head unit and transducer via a cable splice. QTC VIEW units then feed the acoustic data, which may include GPS data, to a data computer for quality assurance, storage and post-survey processing. Some echosounder head units (a Kongsberg Simrad EA 600, for example) are capable of digitizing and storing acoustic data themselves. The data are later transferred to a data processing computer for post-survey analysis. Such systems are called digital echosounders. Older, paper-chart type echosounders are called analogue echosounders. A hybrid digital echosounder system can digitize the data, but not store it, so the head unit feeds the digital data to the data computer for storage (e.g. a Knudsen 320M or Biosonics DT-X).

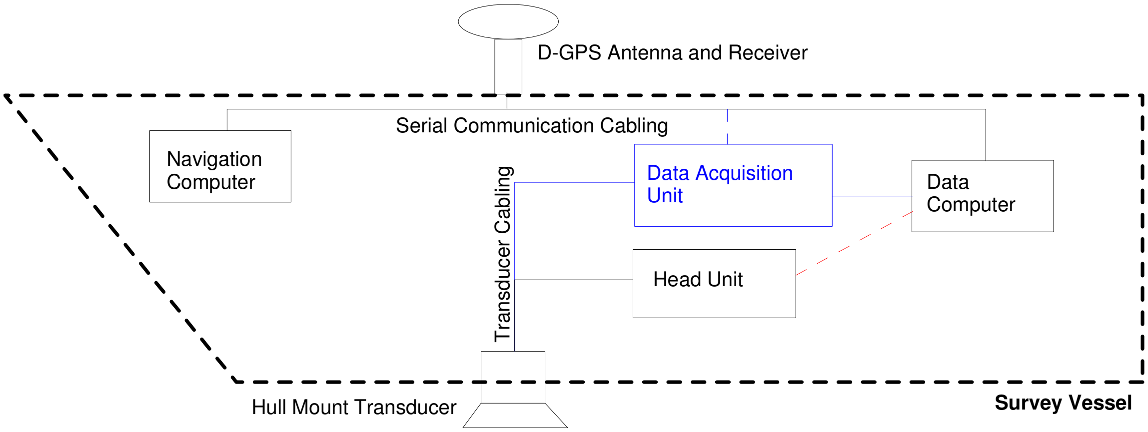

Figure 2-1. Schematic of a standard SBES survey setup. Blue indicates equipment used with a data acquisition unit such as a QTC VIEW Series 5. Without a data acquisition unit, the red connection is required. Dashed-lines are optional connections; the thick dashed-line outlines the hull of the survey vessel.

The properties of the transducer are very important as they determine many of the capabilities of the survey and the character of the data. The centre frequency of the sounding/transmit pulse, and hence the seabed echo, is matched to the resonant frequency of the transducer for optimal performance. The size, shape and material properties of the active element of the transducer determine the resonant frequency and the *beampattern*. The beampattern of the transducer is the same for transmit and receive. It is extremely important in determining the characteristics of the echo – the echo shape and duration in particular. (Echo shape and duration are key echo attributes exploited by classification methods as will be discussed later in this chapter.) For a disc-shaped transducer, the beampattern is a Bessel function of the first-kind, with secondary maxima away from the main lobe that are called *side-lobes*, see Figure 2-2 (*left*). A rectangular transducer produces a similar pattern, with the shorter dimension producing a wider beam, and the longer dimension producing a narrower beam – these transducers are not axially symmetric. Rectangular transducers have the disadvantage of transducer orientation influencing the signals – changing the heading of the survey vessel while remaining stationary would change the echoes for a rectangular beam, whereas echoes from disc-shaped, axially symmetric, transducer would be unchanged (assuming that it is directed perfectly vertical).

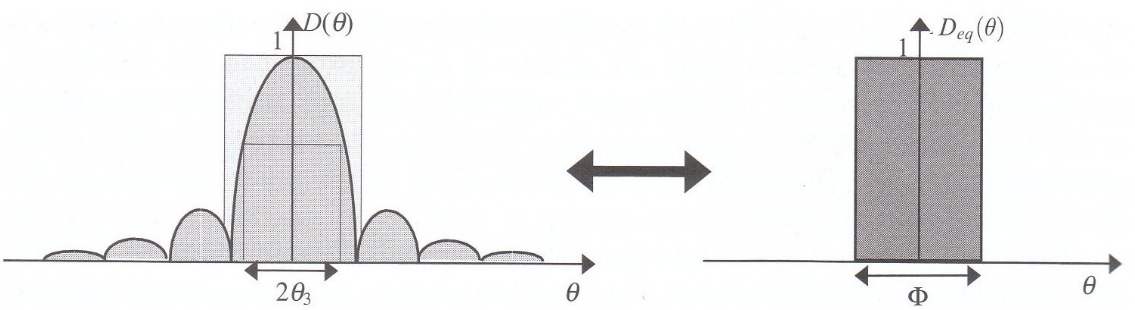

Figure 2-2. Beampattern cross-sections plotted in *angular* co-ordinates: (*left*) a generalized transducer with high amplitude side lobes (for demonstration only), (*right*) the corresponding conical aperture approximation defined by Φ, such that the sum of the amplitudes of the left and right are equal. $2\theta_3$ is the -3 dB beamwidth. From [Lurton, 2002].

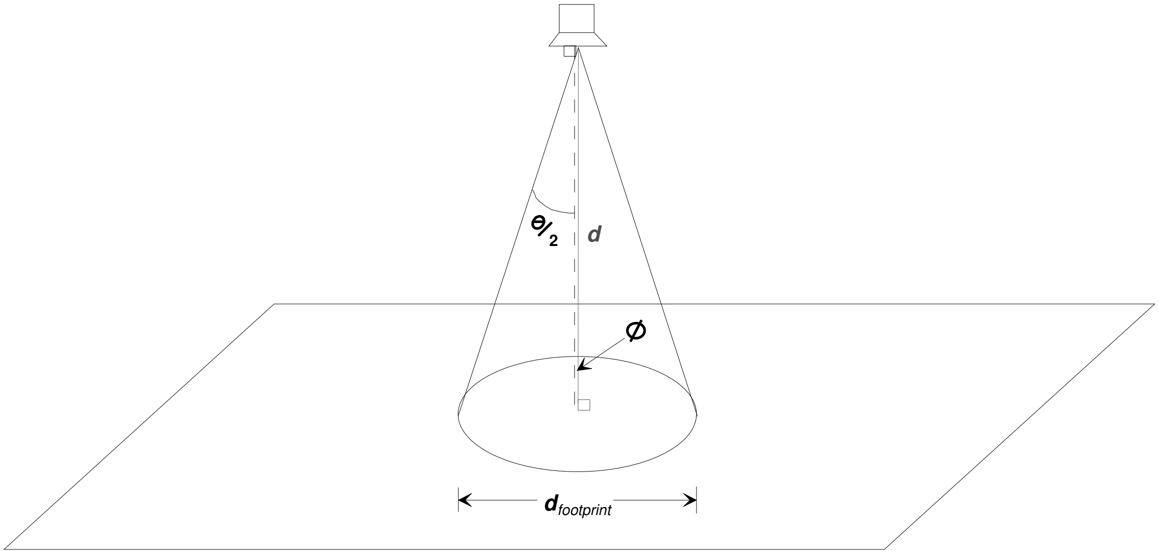

Figure 2-3. A conical beam incident on a flat seabed with the definitions of the seabed normal (grey line, normal to the seabed), the transducer axis (dashed line, normal to the transducer face) and the incident angle, φ. The incident angle is the angle between the transducer axis and the seabed normal, which is the sum of transducer tilt from vertical and seabed slope (seabed slope is defined over an area larger than the acoustic footprint).

For ease of modelling and to distinguish transducers, the beampatterns of SBES transducers are approximated by a conical aperture (like that of a non-collimated flashlight beam). The simplified beampattern is then a cone with constant amplitude throughout and no amplitude outside, see Figure 2-2 (*right*). It is defined by a single value, the *beamwidth*, which may be determined in several ways. Lurton [2002] defines it so that the conical beam has the same total amplitude as the actual beampattern as shown in Figure 2-2 (Φ). For most transducers, this is approximately the -3 dB point on the main lobe ($2\theta_3$), which is also used as the angular resolution of the beam. Most manufacturers of SBES quote the -3 dB point as the beamwidth. Some theoretical work uses the $1/e$ or the -4.34 dB beamwidth [Brekhovskikh and Lysanov, 1988]. The -3 dB beamwidth (measured in degrees) can also be calculated from the diameter of an ideal circular/disc transducer (measured in metres, d_{tx}), the centre frequency (f_c) and sound speed (c_w) using a numerical approximation: $\theta_{-3dB} = 58.9\frac{c_w f_c}{d_{tx}}$ [Lurton, 2002]. The beamwidth is important for depth compensation routines; the effectiveness of the routines and the appropriate value of beamwidth to use will be explored and defined in greater detail in chapters 5 and 6. For consistency, that appropriate value will be used now. It is called the *effective beamwidth*, denoted by θ. In general, the effective beamwidth is more similar to the $1/e$ beamwidth than the -3 dB beamwidth, but is also a function of the survey scenario. A conical beam defined by the effective beamwidth is shown in Figure 2-3.

Echosounders measure water depth by sounding with a transmit pulse confined by the beampattern. In the simplest case, the water depth is:

$$d = \frac{c_w}{2} t_{bp} \tag{2-1}$$

where t_{bp} is the time from the start of the transmit pulse to the onset of the first seabed echo. The onset of the echo is marked by a *bottom pick*, which is found by any number of bottom picking algorithms, some of which are described in detail in appendix A. Bottom picks are almost always found by detecting when the amplitude crosses a threshold, i.e. threshold detection.

The area of the seabed that intercepts the beampattern is commonly called the *footprint*. Within the footprint, the first point of contact between the seabed and the incident transmit pulse is the location of a seabed normal ray that connects to the transducer. This is true for all seabed slopes and transducer tilts and is shown as the grey line in Figure 2-3. The surface area and volume of the seabed within the footprint affected by sound is called the *ensonified* seabed, which varies over time. The time-varying area and volume of the ensonified seabed strongly influences echo shape. In the models discussed in chapters 4 and 5, the directivity of the transducer and the resulting footprint and ensonified seabed will be handled numerically and by Green's functions, respectively. For now, it is quite useful to explore the concept of the footprint using a geometrical approach. For a normally incident beam on a flat seabed, the footprint is a circle with a diameter, d_{fp}:

$$d_{fp} = 2d\tan\left(\theta/2\right) \tag{2 - 2}$$

The geometry of the footprint is shown in Figure 2-3. The footprint diameter is generally regarded as the spatial resolution of the echosounder.

As the transmit pulse propagates through the water column toward the seabed within the beam, it spreads out, losing intensity. **T**ime **v**arying **g**ain (TVG) is the gain applied to correct for losses due to spherical spreading for both the outgoing pulse of sound and the incoming echo. It removes the depth dependence from echo amplitude and, when added to the base gain, keeps the amplitudes at a recordable voltage for the analogue to digital converter in the data acquisition unit. The base gain is a function of the seabed and system characteristics and can be automatically adjusted by the system – called **a**utomatic **g**ain **c**ontrol (AGC). The simplest form of TVG corrects for spherical spreading loss and water column attenuation, α. In decibel units, the TVG is:

$$TVG = 20\log_{10} r + 2\alpha r \tag{2 - 3}$$

where α is the water column absorption coefficient in units of decibels per metre (see [Lurton, 2002] for more information). r is the range to target which is related to time from the broadcast of the sounding pulse to the target and back (t_r), by $r = \frac{ct_r}{2}$. In more complex models, the TVG depends on the type of target – a point like a single fish, or a distributed target like a school of fish or an extended target greater than the footprint like the seabed surface or a surface that does not fill the footprint. For seabed bathymetry and classification applications, Equation 2-3 is generally used for most SBES, even though it is not strictly correct. More complex and new models of TVG are discussed in sub-section 7.3.1.

TVG corrects the amplitudes of echoes for propagation losses. Normally TVG is combined with measurements of receive and transmit sensitivities so that the backscattering strength of the sonar target can be measured. However, the calibration of transmit and receive sensitivities are complicated and may not be entirely reliable. The more acute problem is that of varying water column attenuation (caused by temperature and salinity variations) that makes the standard TVG correction erroneous. When the calibration and TVG are accurate, values of the seabed backscatter strength can be used to infer seabed properties [Sternlicht, 1999; Walter et al., 1997]. However, this approach is not used here because it does not offer unique characterization of the seabed, nor is it feasible as calibrated sonars are not readily available or reliable. More reliable relative backscatter strengths are measured by a new approach introduced in chapter 7.

Another important consideration for the fidelity of the recorded acoustic signals is the sampling frequency of the analogue to digital converters employed in the data acquisition system. By the Nyquist sampling theorem, the sampling frequency must be twice the highest frequency of the signal to be sampled. The process of backscattering that generates the echo is a linear system with no emission of acoustic energy into frequencies outside of the bandwidth of the transmit pulse. Then the spectral characteristics of the transmit pulse will determine the sampling frequency of the data acquisition system. For a gated or 'burst' CW pulse, the spectrum is a sinc function. A bell or Gaussian shaped CW pulse has a Gaussian spectrum which is more compact than a gated-pulse, however the gated pulse is easier to generate and so is more common. The -3 dB bandwidth of a gated CW transmit pulse is $\sim {}^{0.886}\!/_{\tau}$ where τ is the pulse duration. This is approximated as ${}^{1}\!/_{\tau}$, which is generally taken to be the full bandwidth (as a simplifying approximation similar in rationalization to the aforementioned conical beam approximation). The highest or cut-off frequency depends on what signal form is being sampled: *envelope* or *full waveform*. A full waveform (FWF) echo signal is the raw electric signal that is proportional to the actual pressure in the water, characterized by a central frequency (the frequency of the echosounder f_c, also called the carrier frequency) modulated by the response of the seabed to the transmit pulse. The envelope signal is the amplitude of the full waveform signal found by de-modulation by either analogue circuitry or by digital signal processing, which uses the Hilbert transform. In cases, de-modulation and base-banding folds the two sides of the echo spectrum together, so that the highest frequency is half the bandwidth. In the case of a FWF signal, the highest frequency is the centre frequency plus the half bandwidth. Bandpass filters can reduce these upper limits to the FWF and envelope signals, while error in producing the transmit pulse can raise these limits. The central frequency of most echosounders has a significant uncertainty as well, which is particularly hazardous to analogue de-modulation. For these reasons and to account for the 'sidelobes' of the transmit spectrum, a safety factor of two is usually applied to the bandwidth when considering the Nyquist sampling theorem. The minimum sample frequency for both envelope (ENV) and FWF data types are summarized below:

$$\begin{aligned} f_{\min Samp}^{ENV} &= 2/\tau \\ f_{\min Samp}^{FWF} &= 2f_c + 2/\tau \end{aligned} \tag{2-4}$$

Based on these guidelines, the minimum sample frequency for recording FWF data, generalized for SBES, is approximately 625 kHz (for a highest centre frequency of 300 kHz). Until the advent of the QTC5 data acquisition system, systems of this speed had not been applied to SBES. For envelope signals, the requirements are much lower: 6.7 kHz for the usual 0.3 ms pulse duration[8], which is much easier to achieve. Because the envelope is much easier to sample, all SBES systems prior to 1998 sampled the envelope signal, not the FWF signal.

2.2 Seabed Classification by Statistical Segmentation – the QTC Method for Time Series Data

It is important to note that no single comprehensive description of the QTC seabed classification method has yet been published. This hinders communication of

---

[8] this example uses $2/\tau$ for calculating the Nyquist frequency

results to those who have not personally used the system. A symptom of this was an exchange in the Canadian Journal of Fisheries and Aquatic Sciences that was due to a lack of published information on the methodology [Legendre et al., 2002; Preston and Kirlin, 2003; Legendre, 2003]. This section presents a single comprehensive description of the QTC statistical segmentation methodology for time series data from SBES. The image classification method (for SSS and MBES) is similar except in the data preparation and feature generation stages.

QTC produces and markets several acoustic seabed classification systems, including both hardware and software. The only hardware produced is for SBES, of which there are two: the QTC VIEW Series 4[9] and Series 5, denoted QTC4 and QTC5 for brevity. QTC4 uses analogue circuitry for amplification, filtering and demodulation of the signal and then digitizes the signal with an 8 bit analogue to digital converter with a sample frequency of 20 kHz. This hardware was first described in [Prager et al., 1995], and then in greater detail in [Caughey, 1996]. The newer QTC5 systems apply anti-aliasing and bandpass filters and then sample the FWF signal with sample rates between 625 kHz and 5 MHz at 12 bit resolution [Preston and Collins, 2000]. QTC5 also employs AGC, while the QTC4's operator had to monitor the gain and adjust accordingly. Even with all its disadvantages, the QTC4 hardware is effective at recording quality echo time series data, although the quality of the data set is much more operator dependent and needs to be scrutinized more closely than QTC5 data during data analysis. The details of the hardware and associated signal processing are important for the design of the data analysis methods of this dissertation because the data analysis methods have to consider the quality and information content of the data provided.

QTC produces two software products for acoustic seabed classification: QTC SWATHVIEW™ and QTC IMPACT. QTC SWATHVIEW combines the former MBES and SSS data processing products QTC MULTIVIEW™ and QTC SIDEVIEW for processing image-based data, while QTC IMPACT is for processing SBES time series data. QTC SWATHVIEW and QTC IMPACT share the same general multivariate statistical approach to create phenomenological segmentation that offers no information

---

[9] As of September 2010, QTC no longer supports the QTC4 system. However, many units are still in operation and the data collected is still the subject of some research.

on what the classes are, just that they are distinct from one another. The major differences are in data conditioning and feature generation phases as shown as steps A and B in Figure 2-4.

Figure 2-4 shows the four major processing steps of QTC IMPACT in a flowchart. At several points along the way, especially step (A), quality assurance measures are taken to filter out poor data. Some of the sources of poor data are bad bottom picks, long time between successive echoes, sea state [Gleason et al. 2009], transmit ring down, multipath echoes and general noise. A complete description of bottom picking is given in appendix A-1.

2.2.1 Depth Compensation

In Figure 2-4 step (B), depth compensation is shown. Varying water depth throughout a survey affects the echoes strongly by altering their durations and their amplitudes. TVG accounts for spherical spreading losses. It is applied by amplifiers prior to digitization together with the base gain. The effect of depth on echo duration and shape is accounted for in post-processing and is described here. There are two approaches, and a third will be introduced in this dissertation. The industry standard method of depth compensation is generally called reference depth resampling. QTC calls their implementation 'time normalization' or TNORM. The concept is based on linear scaling of the geometry of Figure 2-3 with depth: if the depth is doubled the footprint is doubled, and so the time it takes for the leading edge of the transmit pulse to sweep across the seabed must be doubled as well. This time is referred to as the spreading or curvature duration in later discussions. In a more rigorous derivation, a model of angular dependence of the backscatter as a function of time is used to show the horizontal speed[10]

---

[10] A geometrical model of the horizontal speed finds it to be infinite at the time of echo onset and then decreases from there, as one would expect. Plots of the progression of the transmit pulse across the seabed are shown in [Sternlicht, 1999; Caughey, 1996]. All depth compensation methods scale the time axis evenly. They do not attempt to normalize the echo for instantaneous horizontal speed as this would be technically difficult, but instead normalize echoes to the same average horizontal spreading speed. By doing so, the horizontal speeds of all echoes progress in the same way and so will be depth invariant.

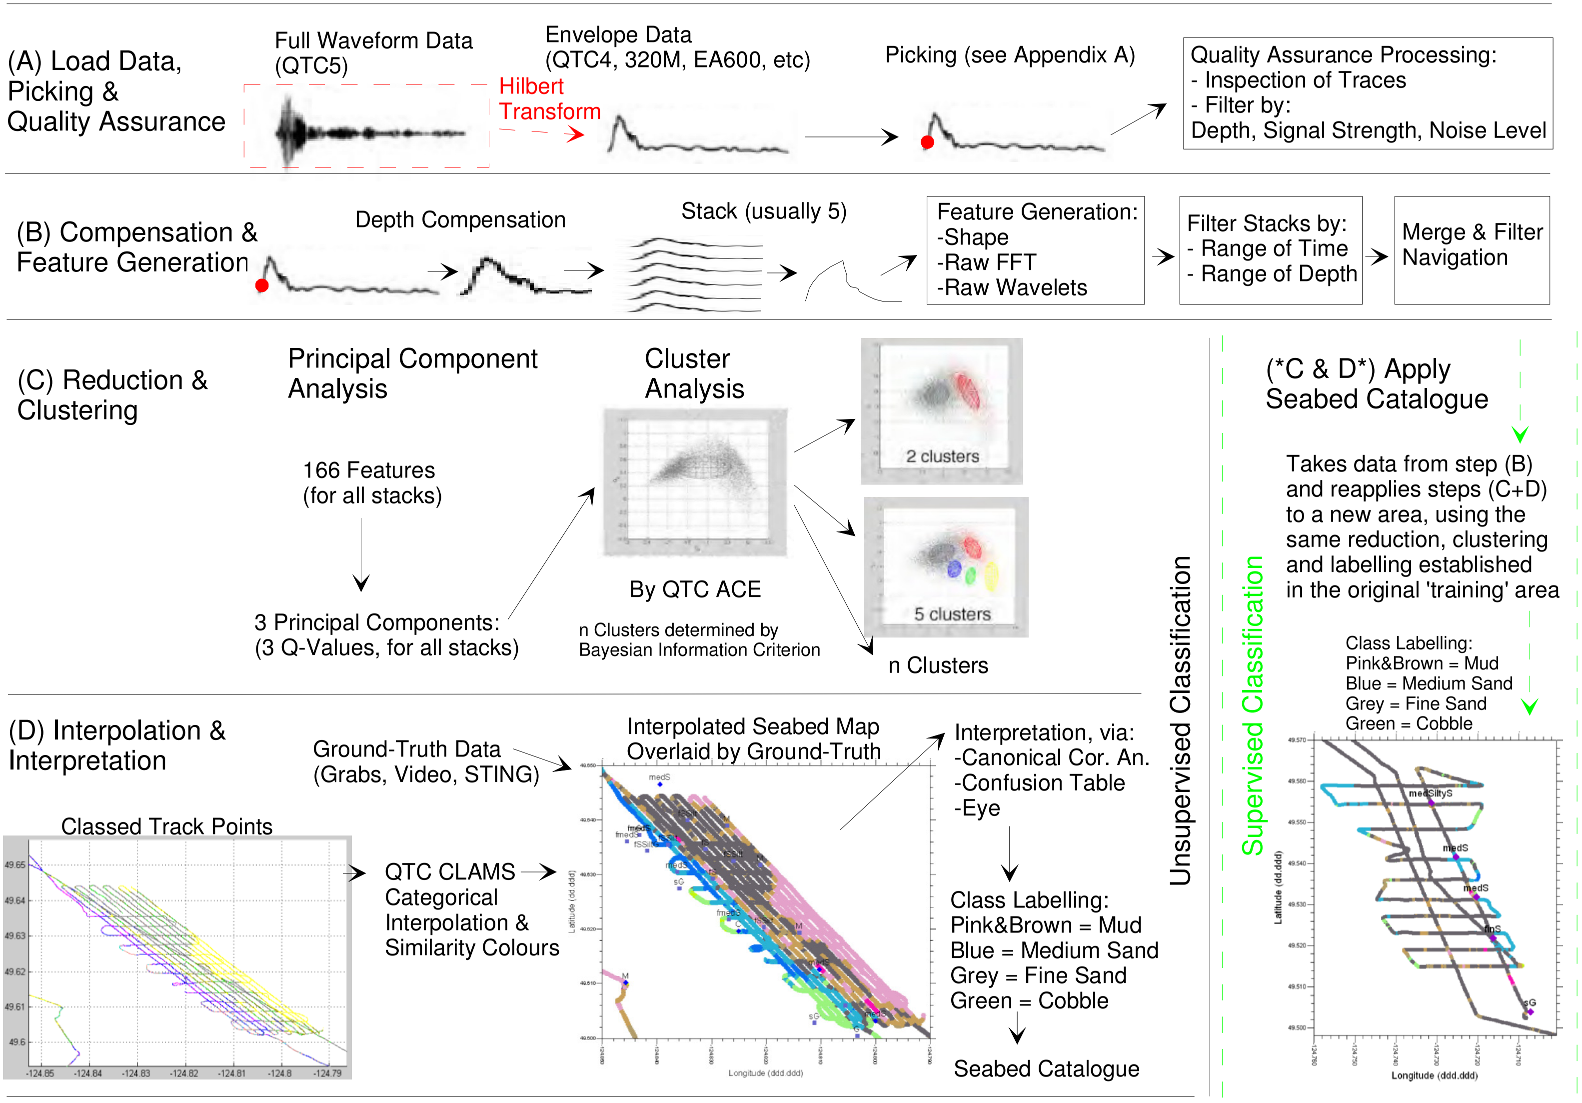

Figure 2-4. Flowchart of the QTC VIEW seabed classification methodology. Dashed-lines are alternative or optional processing.

of the leading edge of the transmit pulse across the seabed is invariant under multiplicative geometric scaling [Pouliquen, 2004; Caughey, 1996]. TNORM aims to make the spreading duration invariant of depth by scaling the time axis of the echo time series in proportion to the water depth. It normalizes the time axis, thus the name TNORM. An important aspect of TNORM is that a reference depth is chosen for all echoes to be scaled to. The scaling is the resampling of the echo time series so that:

$$t = t'\frac{d}{d'} \tag{2-5}$$

where t is the modified sample interval in the echo, t' is the unmodified sample interval, d' is the reference depth and d is the depth of the echo time series being resampled.

The problem with this method is clear, what happens at zero depth? The model implies the duration of the expected echoes goes to zero. Realistically, the duration of the echo should never be less than that of its transmit pulse unless the seabed is perfectly anechoic. Consider a seabed at a water depth approaching zero (ignore near-field effects for now) consisting of just one infinitesimally weak point scatterer: that scatterer will still scatter sound for the entire pulse duration creating an echo as long as the initial transmit pulse. Here is another thought-experiment to demonstrate this concept. If one were to distinguish between gravel and mud by an echo duration criterion, such as gravel has echoes longer than 45 samples and mud shorter than 45 samples, the apparent durations would be quite erroneous in shallow water under reference depth resampling of Equation 2-5. This is shown graphically in Figure 2-5.

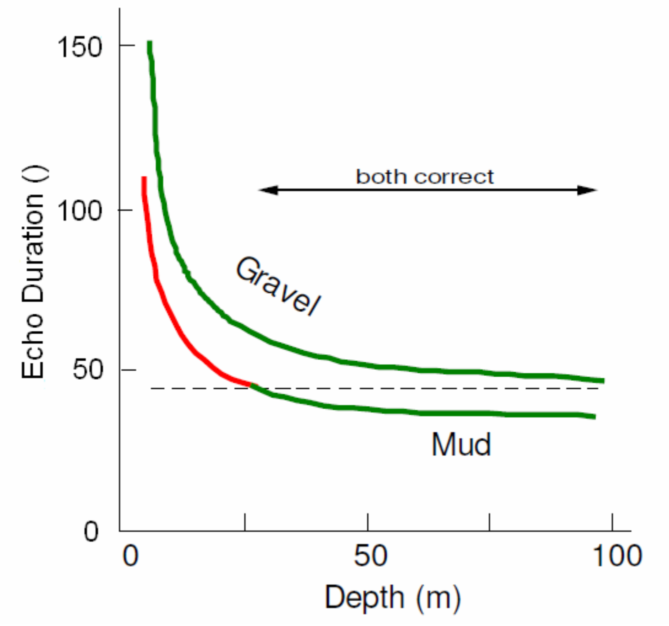

Figure 2-5. A cartoon plot of echo duration, measured in samples after reference depth compensation, for two seabed types. The red line indicates erroneous classification if one were to classify by an echo duration criterion of 45 samples (gravel above, mud below). (Courtesy of J. Preston, modified for this context)

Reference depth compensation is inadequate when the echo duration is of order of the transmit duration. This usually occurs for water depths less than approximately 20 m, but it really depends on the seabed type and probably ranges from 10 to 50 m. Another problem with reference depth compensation is the proper windowing of the echo time series. Many systems define an *analysis window* of fixed size (for QTC IMPACT it is 256 samples) and the echo must fit entirely within that window (alignment of the window is established by the bottom pick). Reference depth makes no effort to fit the echo in the analysis window – it relies on the native sample rate to determine the number of samples a detectable echo occupies within the window at the reference depth. Indeed, the choice of reference depth is very important for that reason. It is always kept higher than 20 m to ensure adequate samples (>10) within the 256 sample analysis window. Yet another problem is the downsampling of the time series for echoes at significant depth. For a reference depth at 40 m, an echo at 120 m would be downsampled by a factor of 3: a standard sample rate of 20 kHz would become 6.7 kHz, which is the lowest sample rate permissible for an echo generated by a 0.3 ms pulse (based on Equation 2-4). In all, reference depth compensation works for water depth greater than ~20 m and less than ~3 times the reference depth. This restriction is prohibitive for most surveys, especially those with large variations in depth.

An improved depth compensation and windowing method was introduced with the QTC5 hardware [Preston and Collins, 2000; Preston et al., 2000]. However, the specifics of the depth compensation used were not published until patented: [Preston, 2004]. The improved method is called *standard echo length* (SEL) depth compensation. The user defines the number of samples in the analysis window to be taken up by a normal or standard echo: n_{SEL}. The default is 128, which is half of the 256 samples in the analysis window. The duration (in time) of a standard echo is the standard echo length, or *SEL*. The SEL value in seconds is calculated based on the pulse duration, the beamwidth, a medium sand seabed and, most importantly, the water depth. The mechanism by which SEL depth compensation works is to resample the signal to a new sample rate that will sample an SEL duration in n_{SEL} samples. Accordingly, the new sample rate, f_{SEL}, is: $n_{SEL} + 1/SEL$ (the '+1' is a somewhat erroneous adjustment used in the QTC IMPACT software, replicated here for consistency). The full equation for this new sample rate that both compensates for depth variation and windows echoes properly is:

$$f_{SEL} = \frac{n_{SEL} + 1}{\frac{2d}{c}\left(\sec\left(\frac{\theta}{2}\right) - 1\right) + \frac{20}{c\kappa f_c} + \tau} \tag{2-6}$$

The term $20/c\kappa f_c$ is the extra time the transmit pulse takes reverberating within the seabed volume. It is based on the depth within the seabed at which sound penetrating the seabed decreases by 10 dB:

$$l_p = \frac{10}{\kappa f_c} \tag{2-7}$$

where f_c is the SBES centre frequency and κ is a nominal sediment sound attenuation value, which is 0.5 dB/kHz/m, a value that is typical for medium sand.

There are two issues with the SEL depth compensation method. First, the calculation of the penetration time is an approximation, based on a nominal attenuation and the water sound speed instead of the sediment sound speed. The water sound speed is used because the sediment sound speed is unknown. Many seabeds have sound speeds slightly greater than that of water so a compromise value like 1600 m/s would be better. However, this issue is somewhat inconsequential, as this term does not vary with depth. Then as long as the echoes are contained within the analysis window, the compensation will be sufficient. An improved extinction model for penetration depth is given by Sternlicht, which includes the sediment sound speed and the reflection coefficient [Sternlicht, 1999]. This issue is explored further with the full derivation of the echo duration model in chapter 5. The second issue with SEL compensation is that the new sample rate may become very low to accommodate long echoes from deep water into the SEL window, which has a fixed-length of 256 samples. If the new sample rate f_{SEL} (Equation 2-6), becomes less than the Nyquist frequency (Equation 2-4b), the bandwidth of the signal is encroached upon and aliasing errors occur. This occurs for echoes from water deeper than the SEL critical depth, which does not depend on the original sampling rate of the data acquisition system. The SEL critical depth is:

$$d_{crit} = \frac{c_w}{2\left(\sec\left(\frac{\theta}{2}\right)-1\right)}\left(\left(\frac{n_{SEL}}{s}-1\right)\tau - \frac{20}{c_w \kappa f_c}\right) \tag{2-8}$$

where s is a safety factor; its value was taken to be 2 in earlier calculations, but the QTC IMPACT calculation for QTC5 data uses a value of 12, six times the usual factor of 2. The factor of six is due to digital signal processing of the raw signal which includes elaborate decimation and band-shifting, all of this prior to envelope calculation with the Hilbert transform. Also, note that the additional sample in Equation 2-6 is gone in Equation 2-8. For the QTC5 system, the only system on which SEL compensation is currently available, the critical depth ranges between 70 and 500 m. To handle surveys that violate the SEL critical depth, anti-aliasing filters are applied to all data in the survey to make the data of uniform bandwidth. This filtering is adequate but it can remove a significant amount of high frequency information from the signal. In the extreme worse-case scenario, it can make the echoes look very smooth with few distinguishing shape characteristics.

2.2.2 Stacking and Features

After depth compensation and windowing, the echo time series are 'stacked', see the latter part of step (B) in Figure 2-4. The aim is to reduce the random, incoherent variation among individual echoes to enable better classification. Stacking is simply averaging several time series aligned by the bottom picks. After averaging, each stack is divided by its maximum to normalize the amplitude from zero to one. Aligning the time series using a threshold bottom pick for stacking has been found to be a sufficient technique for normal incidence echoes [Sternlicht, 1999]. There has been some work by others on evaluating the practice of stacking in the literature [Gleason, 2008; van Walree et al. 2005; Sternlicht, 1999], which was confirmed anecdotally by experimenting with the data shown in chapter 3. In all, it has been found that a stack of 5 to 10 echoes reduces variability while not sacrificing a significant amount of spatial resolution. For instance, in a survey running at 4 knots, with soundings being taken a rate of 3 Hz[11], the footprint of the stacked echoes would be the footprint of a single echo stretched by an additional 3.5 m, which is small considering a single echo footprint would be about 7 m for a 20° beam in 20 m water depth.

After picking, quality assurance, windowing, depth compensation and stacking, statistical measures of echo attributes, called features, are calculated. Features describe an aspect of the echo with a single value, for example, an echo's full width at half maximum. Features are then used as the input to a multivariate statistical process that groups the echoes into classes. These classes define areas of the seabed, also called seabed segments or classes, that are reasonably homogenous in acoustic character. Numerous survey results have shown that the classes differ from each other due to differences in seabed properties that are readily observable by other, non-acoustic means (ground-truth).

A suitable analogy to this process of grouping echoes by acoustic character is that of grouping boxes by attributes of size and shape. This is shown in Figure 2-6. The

---

[11] this is known as the ping rate – a ping is the transmit pulse and the frequency or interval on which these are produced should not exceed the two-way travel time in the water column – i.e. ${}^{2d}\!/_{c}$

'features' used for sorting boxes could be their length, width, thickness, length of longest diagonal and ratios of their dimensions.

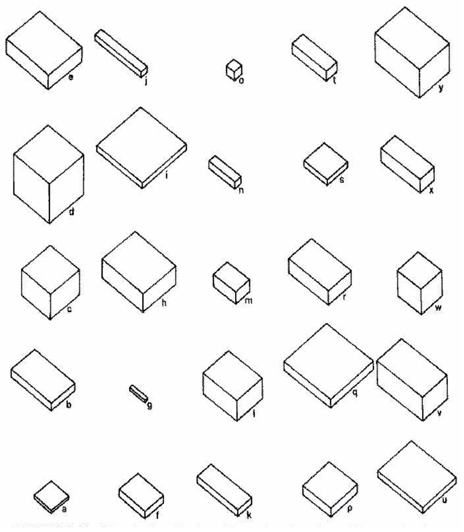

Figure 2-6. A random assortment of 25 boxes to be sorted into groups of similar boxes as an analogy for the process of statistical segmentation employed in the QTC seabed classification method.

Features used by the QTC seabed classification method fall into two categories: shape features and spectral features. The features are summarized in Table 2-1. The spectral features are the even-number components of the FFT transform of the 256 sample stacked traces (so 128 components which leads to 65 features), and the shift and scaling components from a wavelet transform with the Daubechies 4 wavelet (this wavelet somewhat resembles an echo). There are three families of shape features. The cumulative integral features are the values from the normalized cumulative sum taken at 16 points from the data. Figure 2-7 shows how this is done. An additional seven features consist of the ratios of the second cumulative integral feature to cumulative integral features numbered three through ten. The quantile features are sample values at eight evenly-spaced locations (i.e. 32, 64,…, 256) from a sorted and normalized stacked echo – the samples are sorted by amplitude from smallest to largest. The histogram features are the values of an eight-bin histogram of the data, divided by the number of samples so that the sum of the eight features is one. There are three features that have no information – they have the same value for all data. These features are the ones omitted in Table 2-1, numbers 16, 31 and 40, which are the last cumulative integral feature (the last sample of a normalized cumulative sum is always one), the last quantile feature (the last sample of a sorted normalized trace is always one), and the first wavelet feature (the shift of the first transform is always zero).

Evaluation of the performance of these features is given in appendix B-1 (as it pertains to that section) and later in chapters 3 and 7.

Table 2-1. The feature set used in the QTC seabed classification method

Feature Algorithm Name	Feature Number	Indicator of:
Cumulative Integral	1-15	General shape
-- Ratios of Cumulative Integral	17-23	General shape
Quantile	24-30	General shape
Histogram	32-39	General shape
Wavelet	41-101	Shape/scale
FFT	102-166	Shape and spectra

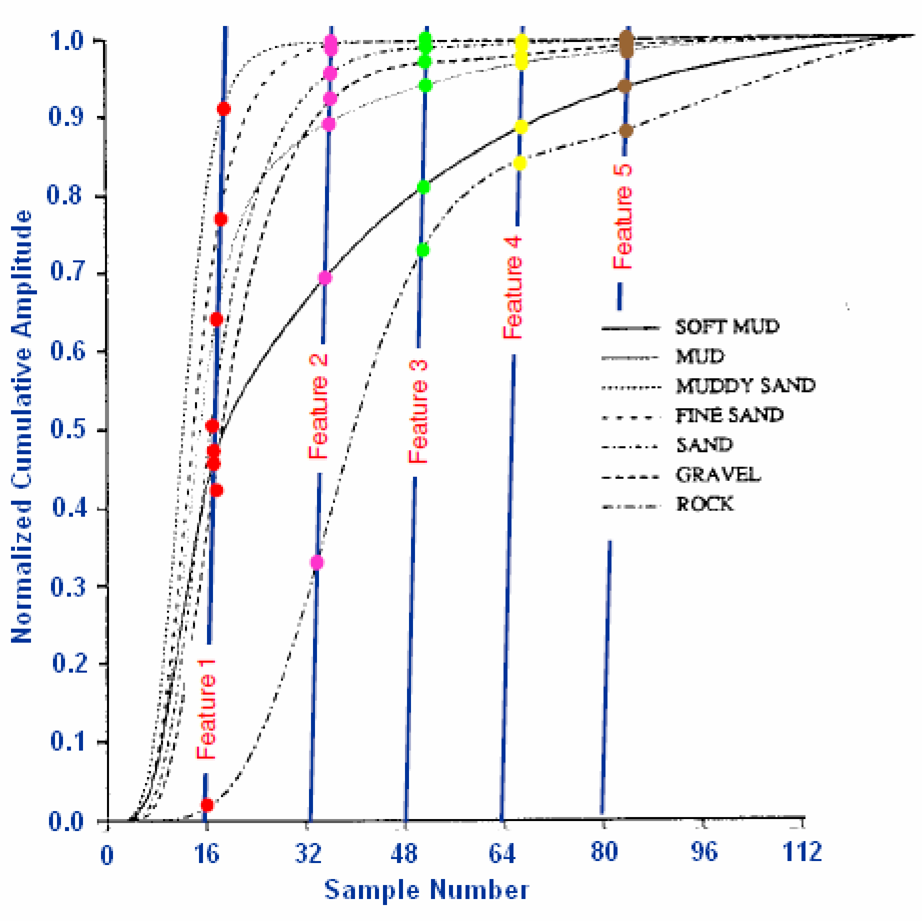

Figure 2-7. A diagram showing the calculation of the cumulative integral shape features, adapted from [Lurton and Pouliquen, 1992; Pouliquen and Lurton, 1992], by J. Preston.

2.2.3 Data Reduction and Clustering

In Figure 2-4 step C, the data reduction and clustering steps are shown. This is the aforementioned multivariate statistical process by which echoes are grouped into classes of similar acoustic response based on their feature values. Continuing with the grouping boxes analogy, this step uses the boxes' attributes to group them. The data reduction step combines the different attributes or features into two new variables, and the boxes are plotted by those co-ordinates, see Figure 2-8. The clustering step identifies groups among the boxes and then assigns individual boxes to the appropriate group.

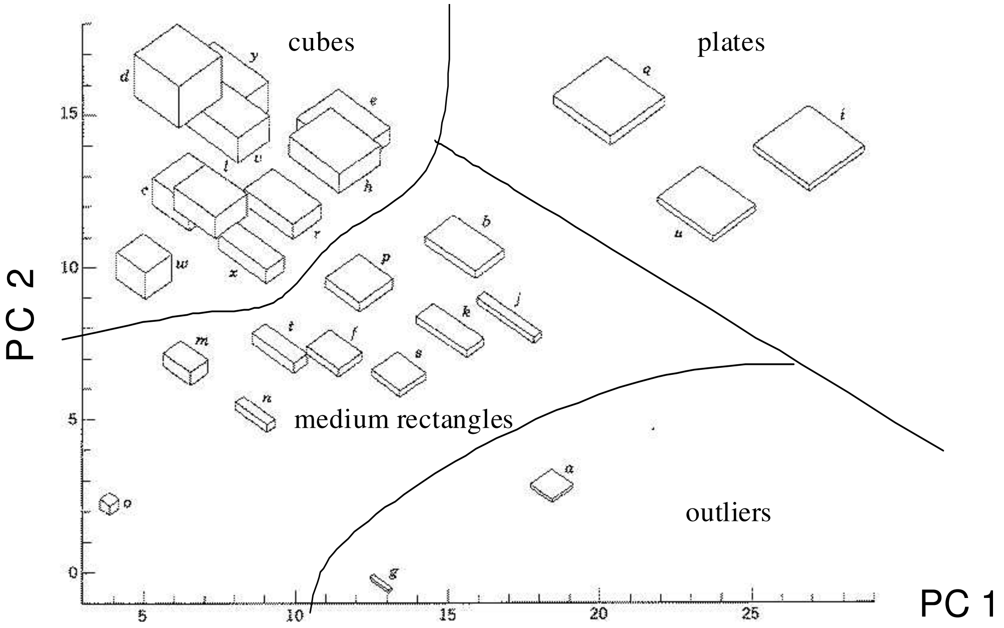

Figure 2-8. Four groupings of 25 boxes as an analogy to the data reduction and clustering steps in the QTC seabed classification methodology. PC1 and PC2 are two new variables that are combinations of the boxes' measured attributes.

The data reduction step in the QTC seabed classification methodology is handled by Principal Component Analysis (PCA). Appendix B-1 presents PCA fully, while here only a short description is made. PCA combines the echo features by linear combinations (or sums) into scores called Q-values. Only the first three Q-values are calculated as these new variables express a large fraction of the variance in the data (usually between 70 and 90%) and are conveniently viewable as a three-dimensional scatter plot. These Q-values also contain complementary information. Q-values are essentially each stack's position in a three-dimensional space defined by the three greatest directions of variance within the feature set. These directions are the principal components. Appendix B-1 presents the entire process of PCA and then uses example data from chapter 7 to evaluate the performance of PCA and the features.

An example three-dimensional scatter plot of the Q-values is visible in Figure 2-4, step C, and in Figure 2-9. The vector space defined by the three principal components in which the Q-values are plotted is called Q-space. The distribution of these points represents the 'acoustic diversity' contained in the data [Collins and Galloway, 1998; Preston et al., 2002], meaning that the relative distances between the data points represent how distinct they are from each other on the basis of the acoustics. However, the distribution of points in Q-space has no physical meaning. PCA essentially scrambles the data, combining features in an attempt to convey the most information on the diversity of the data, but losing any information on the relationships the features might have had with geophysical seabed properties. This is exactly why the process is phenomenological. The benefit of PCA is that it adjusts the content of the principal components based on the data, making the process robust to varying survey situations. It also suppresses noise (see appendix B-1), but can be subject to error if features convey erroneous data.

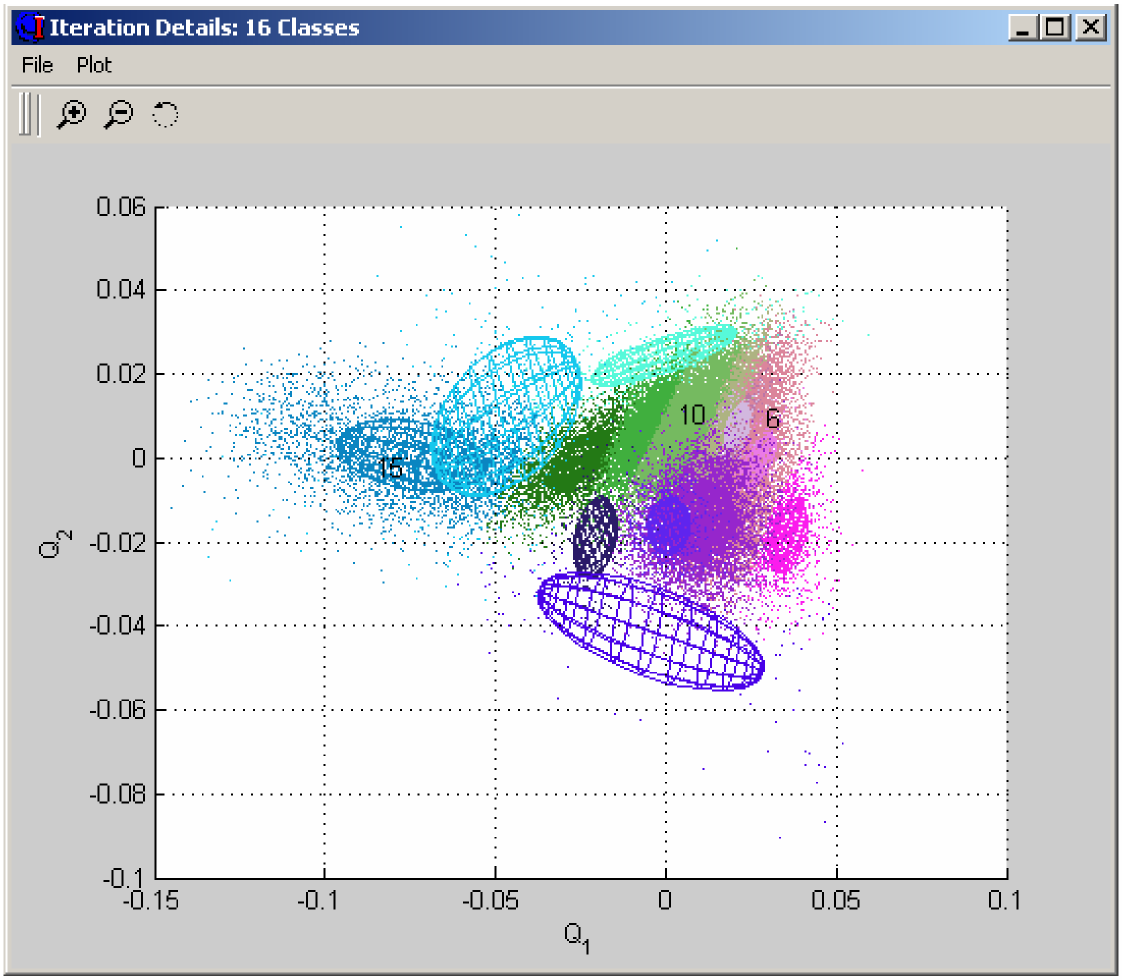

Figure 2-9. An example of Q-values plotted and clustered in three-dimensional Q-space, as seen in the QTC ACE software. The ellipses are the class covariances, which correspond to the class boundaries. The data set is the Airmar38 Patricia Bay testbed survey (as described in chapter 7), the numbers are class numbers referred to in chapter 7.

The clustering part of step C, Figure 2-4, groups data within Q-space. The groups are the seabed classes; data that are acoustically similar correspond to a consistent seabed label (such as rock, mud, tall vegetation over sand, sunken logs, etc.). This assertion is well proven by the 17 years of successful surveys. The clustering method used is an objective Bayesian k-means clustering procedure implemented in a product called ACE – for **a**uto **c**luster **e**ngine. The details of this procedure are the one area of the QTC seabed classification method not improved upon in this dissertation, so the details will not be fully explored here. There are numerous references on the subject, starting with a general reference on clustering: [Rencher, 2002; Legendre and Legendre, 1998], references on the older QTC manual clustering method: [Prager et al., 1995; Caughey, 1996; Legendre et al., 2002; Preston and Kirlin, 2003, Legendre, 2003][12], and references to papers presenting ACE, the newer objective clustering method: [Preston et al., 2004b; Preston et al., 2004c]. The advantage of ACE clustering is that it provides a means to determine, based on the **B**ayesian **i**nformation **c**riterion (BIC), the optimal number of clusters or classes for the dataset. An example ACE result is shown as Figure 2-10. The user can choose from a number of clustering results, in addition to the result with the best BIC. This is done because in some survey instances, it has been observed that the result with the best BIC will split classes beyond the ability of the ground-truth to discriminate – this is called 'over-splitting'. To avoid over-splitting, users may select a result that has fewer classes, but has a BIC within a few percent of the optimal BIC.

Similarity of clusters or classes is qualified by the ground-truth data and by a feature of QTC ACE and QTC CLAMS™[13] called similarity colours. Similarity colours works by converting cluster centres triplet positions to red/green/blue values. A good example is shown in Figure 2-9. The variability of data within a class may be quantified as well, using measures of confidence and probability[14], or a measure now called the complexity index [Gleason et al.; 2006]. All three of these statistics are available in the QTC CLAMS software. An additional way of viewing data is called 'one colour per record' or OCPR. This colours data in the same fashion as similarity colours but operates on the Q-values of individual stacks instead of the colour associated with the positions of the cluster centres. The benefit of OCPR is that it shows the data without clustering into groups. One of the drawbacks of clustering is that it forces the data into groups when the data may be a continuous variation without strict boundaries between classes. This occurs in situations like estuaries with gradual sorting from silts to coarser sands, while an example of strict boundaries could be rock outcrops in a sand field. OCPR can be a very

---

[12] An additional reference is the QTC IMPACT users manual: [QTC, 2004a].

[13] QTC '**Cl**assification **A** **M**apping **S**oftware' is a categorical interpolation utility. It takes scattered classification data and interpoleates it to a grid by taking the modal class within a search radius. In other words, it fills a spatial grid with the most populous nearby classification. It can also assign values based on the complexity and the varibility of the data, see the QTC CLAMS user manual: [QTC, 2004c].

[14] QTC has a technical note available as a reference to confidence and probability: [QTC, 2004b].

useful way of illustrating the data for a small survey with few classes. OCPR is most often used to examine the performance of clustering for over or under-splitting.

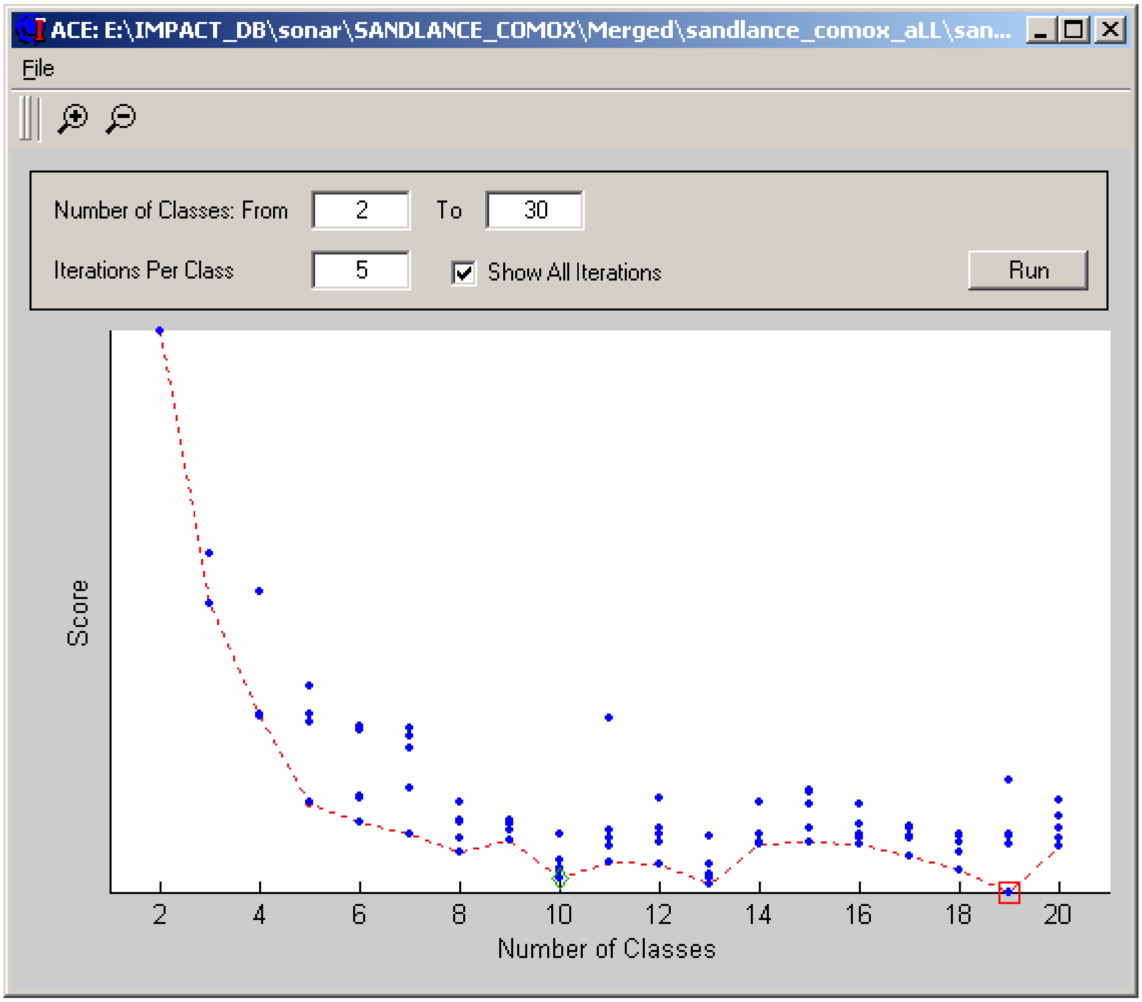

ACE: E:\IMPACT_DB\sonar\SANDLANCE_COMOX\Merged\sandlance_comox_aLL\san...

Number of Classes: From 2 To 30

Iterations Per Class 5

☑ Show All Iterations

Figure 2-10. The ACE result for the Comox area, Pacific Sandlance survey [Biffard et al., 2009]. This example was chosen to show the selection of 10 classes instead of the class model with the lowest BIC. 10 classes is preferable to 19 because of it is simpler, has nearly the same BIC as 19 or 13 classes and does not exceed the capacity of the ground-truth to distinguish the classes (i.e. 'over-splitting').

2.2.4 Interpolation and Interpretation

The final stage of the QTC seabed classification process is depicted as step (D) in Figure 2-4. Interpolation is basically the procedure of making a seabed classification map. QTC CLAMS is very useful, as software that does categorical interpolation is rare.

To display classed data points geographically, it is useful to interpolate the data to show coverage that roughly corresponds to the echosounder footprint size, and to decide which class is dominant in areas of overlap. The sign of good results is when overlapping data are consistent, such as where survey lines cross. The interpolation process is called categorical interpolation: if a map location had classes 2 and 7 within its search area (defined by a mean footprint size), the resulting class would not be class 4.5, but either one of class 2 or 7 depending on which is more frequent.

In the grouping boxes analogy, the groupings are readily identifiable with labels (Figure 2-8). For seabed classification, such labelling requires ground-truth information, which is usually plotted on the interpolated seabed class map. The seabed class labels are usually related to grain size (gravel, sand, mud), but may be modified with bedform descriptors (sand waves, pockmarks, worm-tubes), descriptions of flora and fauna (eelgrass, clams, bark debris), or anything else observable on the seabed. For simple, small area surveys, visual interpretation of the seabed class map overlaid by ground-truth is usually sufficient to generate seabed type labels for each class. This method lacks rigor. More rigorous interpretation procedures are available. Canonical correlation analysis seeks relations between the classes and the ground-truth [Preston et al., 1999; Tsehmahman et al., 1997]. Analyzing the co-occurrences of classes and ground-truth data can be accomplished by building a *confusion matrix*. The procedure is particularly well suited where there are an equal number of acoustic classes and ground-truth classes. In that case, the confusion matrix is square. Re-arranging the rows so that the largest values are on the diagonal will maximize the trace and show the best one-to-one correspondence between ground-truth and acoustic classes. A confusion matrix is employed for the interpretation of ground-truth data in chapter 6 – in that case, it is called a ground-truth interpretation table.

2.2.5 Supervised Classification

The method described in the previous two sub-sections is called unsupervised classification. It is the general approach that does not require the method to be trained or calibrated prior to surveying. Supervised classification requires training, which is basically prior knowledge. The basis of training is the PCA, clustering and class labelling information generated in a training survey that is held in a catalogue file. Then, this file contains all the information necessary to classify new data by on the basis of training survey. The process is described in Figure 2-4 and is described further in appendix B-1. Here are a couple examples of supervised classification:

- A survey is conducted in an area. The resulting catalogue may be applied to extend the survey to new areas with data gathered with the exact same procedure as the first survey.
- Training areas with well known seabed classes are surveyed. Unsupervised classification correctly separates these areas and then the resulting catalogue can be applied to all subsequent data gathered with the exact same method.

Initially, the QTC VIEW system was marketed emphasizing supervised classification. However, as was found out with use, there are numerous problems with supervised classification [Ellingsen et al., 2002]. They include variable ocean conditions altering the acoustic response [Gleason, 2009], seabed slope [Biffard, 2005; von Szalay and McConnaughey, 2002], bedforms and flora and fauna, all of which are capable of causing class confusion. An ambitious attempt to build a catalogue for supervised classification was made by J. Galloway and his group at the Canadian Hydrographic Service. The effort included 28 training areas with extensive ground-truth data, but the catalogue never produced consistent results[15]. The cause of the failure was the above problems that were eventually identified (with this authors' help) as well as difficulties with ground-truth data that were alluded to in [Galloway, 2004].

2.3 An Overview of All Seabed Classification and Characterization Methods

In the absence of one entirely successful direct characterization method, there have sprung a number of different methods. All have advantages and disadvantages. The

---

[15] Personnel communication and collaboration with J. Galloway.

QTC methodology is more often applied by the scientific community, while the RoxAnn method is more often used in commercial applications. The QTC and RoxAnn approaches both belong to the category of statistical segmentation. There are two other categories, both much less phenomenological and verging on direct characterization (listed under categories 1 and 2 below). All methods face similar problems due to the nature of seabed echoes at these frequencies as will be discussed in the following sub-chapter. Descriptions of all available methods are discussed below, broken into three categories.

1. Inversion and model-based catalogues:

A forward model of high-frequency backscatter generates a catalogue of echo time series for the operating conditions of a particular echosounder. Best matches to actual echoes are sought. Uniqueness problems limit the general applicability of catalogue searching, even with calibrated sonars. Three separate efforts in this category have been published: [Sternlicht and de Moustier, 2003] for 33 and 93 kHz echosounders; [Caiti and Zoppoli, 1998; Pouliquen et al., 1999; Bergem et al., 1999] for an 8 kHz (moderate frequency) parametric sonar; [Lurton and Pouliquen, 1992] for a 38 kHz echosounder. None of the approaches were entirely successful even with high-fidelity, calibrated SBES, primarily due to non-specificity of seabed echoes. More recent results make use of a simpler, faster forward model, which allows for inversion by optimization [Siemes et al., 2010; Simons et al., 2009]. Geoacoustic inversion techniques used on sub-bottom profiles and seismic are continuing to advance, (e.g.: [Tollefsen et al., 2006; Jiang and Chapman, 2008; Morley, 2009]). Researchers continue to attempt to adapt these low-frequency methods to high-frequency sonar: [Fonseca et al., 2007] and this dissertation.

2. Direct characterization:

There are a few methods in this category for SBES, however all require calibrated sonars or calibrated low frequency systems like sub-bottom profilers. Measures of echo energy can be used for low to moderate frequencies, where the plane wave approximation applies and the specular echo (reflection) is dominant over surficial roughness scattering. With specially calibrated sub-bottom profiles or seismics, the absolute scattering cross-section can be measured, which reliably characterizes impedances[16], and thus sediments, on both sides of the interface [Yilmaz, 2001]. Other similar approaches use multilayer theory to determine impedance. In these methods, the assumption of a normally incident plane wave precludes the effects of a conical, spherical beam interacting with a sloping seabed. Two examples are Honeywell ELAC and NRL ASCS, described in [Lambert et al., 1993; Lambert et al., 1999; Panda et al., 1994; Wood and Lindwall, 1996; Walter et al., 1997]. Characterization via impedance is the main technique; however there are other direct seabed characterization methods for low to moderate frequency acoustic system. An example is the method of spectral ratios that relies on linear dispersion within sediments to characterize the seabed [Bloomer et al., 1999; Chapman et al., 2009]. More recent work by [van Walree et al., 2006] makes use of a new theory of impedance gradient and calibrated SBES from 12 to 200 kHz to map mean grain size.

3. Statistical segmentation:

In general, there are two types of statistical segmentation methods, methods similar to the QTC approach, or those similar to the RoxAnn approach. RoxAnn is based on two features, *E1* and *E2*, which are the energy of the tail of the first echo and the energy of the second echo [Chivers et al., 1990; Burns et al., 1989; Orlowski, 1984]. 'Tail' is defined as one pulse duration after the onset, and *E1* and *E2* are interpreted as 'roughness' and 'hardness' [Schlagintweit, 1993]. Segmentation is achieved by drawing squares on a plot of *E2* versus *E1* [Burns et al., 1989]. Heald has published very innovative research, both theoretical and experimental, justifying the RoxAnn™ approach [Heald 2001, 2000; Heald and Pace, 1996]. Many researchers have duplicated both QTC and RoxAnn systems [Legendre, 2002; Siwabessy, 2001; Clarke and Hamilton, 1999]. The Kongsberg SeaBEC™ software [Lied et al., 2004] seems to have aspects that are similar to both the RoxAnn and QTC segmentation methods – it does PCA analysis directly on 128 samples of the first echo, while also having the option of

---

[16] see section 2.1 or [Lurton, 2002] for more information on the terms mentioned here.

measuring *E1* and *E2* with 'appropriate transducers'. There are several other systems that offer the RoxAnn *E1* and *E2* features. Biosonics VBT™ [Burczynski, 2001; Dommisse et al., 2005] uses *E1* and *E2* plus an energy ratio of the first echo and a fractal measure [Lubniewski and Pouliquen, 2004] with fuzzy *c*-means clustering. A newer segmentation method by van Walree et al. is somewhat in between RoxAnn and QTC [van Walree et al., 2005; Siemes et al., 2010]. It uses only six very good features with PCA and *k*-means clustering. Statistical segmentation has been applied to sub-bottom data [Bloomer et al., 2000; Lee et al., 2009]. The application of neural networks that learn to recognize echoes from different seabed types is an alternative method of segmentation. Neural networks have been applied several times with limited success. The VBT system is an example of a commercially-available neural network system [Stepnowski et al., 1996; Stepnowski et al., 2003]. Other examples include work by a Japanese group: [Motooka et al., 1995; Kitami et al, 1997; Tao and Motooka, 2006], in addition to: [Michalopoulou et al., 1995; Chakraborty et al., 2001]

2.4 Similarities and Differences among Current Seabed Classification and Characterization Methods – Potential Avenues for Improvement

All methods face similar problems due the nature of seabed echoes at these frequencies (10-300 kHz). A fundamental problem is the non-specificity of the echo data [Sternlicht, 1999], which means in many situations very similar echoes are generated from different seabeds. At these frequencies incoherent backscatter often dominates [Jackson and Briggs, 1992], making echoes highly variable, almost chaotic. This variability is called ping-to-ping variability; see [Chu and Stanton, 2010]. Some operators and researchers have even called this noise. It makes classification and characterization extremely difficult. Other problems include the repeatability [Gleason et al., 2009], accuracy, precision, interpretation of the results and the fidelity of the acoustic seabed signal. Problems with repeatability and accuracy (as from that caused by ping-to-ping variability) are known to be caused by non-seabed influences, namely depth, seabed slope and ship pitch/roll. These non-seabed influences all significantly increase echo duration and alter the scattering regime, decreasing the contribution of the specular component of the echo (this term will be defined in chapter 4 and appendix E). These artifacts almost always dominate determining the nature of echoes; the exception would be a survey under calm conditions over different, perfectly-flat seabeds. Fidelity of the acquisition system is another issue. For instance, to apply QTC SEL depth compensation, high sampling rates (625 kHz to 5 MHz) are required to conserve the bandwidth of the signal. Bandwidth can vary for numerous reasons and this must be accounted for. Time-varying gain, dynamic range, digitization and filtering are issues to be improved in current data acquisition systems. However, the most glaring issue and perhaps the single largest problem is that of the effects of depth and seabed slope. QTC SEL depth compensation is appears to perform well in surveys if the beamwidth and attenuation values are accurate; however, it is untested. Most other methods either use erroneous depth compensation (proportional or TNORM depth compensation) or none at all. Seabed slope can affect the data by lengthening echoes and reducing their amplitude, which is similar to the effect of depth, but in cases of seabed slope greater than approximately thirty degrees the effect can be even stronger than that of depth [Biffard et al., 2005; Preston and Biffard, 2007; Biffard et al., 2007]. The effect causes major misclassification for slopes greater than the ½ beamwidth [Biffard et al., 2005], and minor problems even at slopes of two to four degrees (section 6.3). Addressing problems with depth and slope artefacts is seen as a scientific advance.

A common topic of research is to compare the QTC methodology to the RoxAnn method, see [Hamilton et al., 1999] for example. The general conclusion is that they produce similar results. RoxAnn has the advantage of easier interpretation; while the QTC methodology is more robust (it does not depend on having the second echo and has many more features to draw upon). There is significant doubt about the depth invariance of the RoxAnn methodology. The literature seems to indicate that RoxAnn gets around the effects of depth variation by effectively classifying using the ratios of *E1* and *E2* [Burns et al., 1989], and that *E1* and *E2* are sums of energy so that the duration of the echoes does not really matter. Greenstreet et al. evaluated the RoxAnn approach and observed a systematic variation in classification over the same areas on separate passes [Greenstreet et al., 1997]. The variation was attributed to variation in the tilt of the transducer; an effect equivalent to slope. Heald showed that *E1* and *E2* are complex functions of water depth, but did not comment on the practical effects of depth in the context of seabed surveys or compensation of this effect [Heald, 2001]. Biosonics VBT shares the *E1* and *E2* features with RoxAnn. Dommisse found that the VBT does not employ any depth correction and cannot be used if the water depth varies [Dommisse, 2005]. Lubniewski and Pouliquen demonstrated the depth dependence of the fractal dimension measures (a feature applied in VBT), which implies the measure is also slope dependent [Lubniewski and Pouliquen, 2004]. There are some serious concerns about the non-seabed influences of slope and depth on RoxAnn and related methods.

The approach of van Walree et al. may be a good compromise between the RoxAnn and QTC methods [van Walree et al., 2005]. Their method uses a combination of fewer features with their version of PCA and cluster analysis for segmentation. This has the advantage of easier interpretation while maintaining the robustness of a first-echo only, PCA driven approach. An alternative to the RoxAnn and QTC methods is neural networks. However, training neural networks is not entirely successful for the same reasons as for supervised classification: varying survey/oceanographic conditions, unaccounted for seabed type variations, echo variability and non-uniqueness, unreliable ground-truth data and poor compensation for non-seabed influences on the data.

In comparing the segmentation methods to the inversion and direct characterization, one finds the latter methods are underdeveloped compared to the segmentation methods. The reason for this is that the inversion and characterization methods are not generally applicable. Many methods have limited bounds of applicability in terms of centre frequency and roughness height. Many methods also require calibrated, often specially designed, scientific SBES and controlled environments in order to demonstrate their ability. Some of the newer inversion efforts are more promising. One approach applies a simpler forward model with a genetic algorithm optimization [Simons et al., 2009; Siemes et al., 2010]. One of the issues with such an approach is that the inversion offers no estimate of confidence of the results. This will have to be addressed by applying inversion techniques that offer error estimates, such as the techniques applied in [Dosso, 2002]. Another issue is that the method does not consider the effect of ping-to-ping variability, which will be considered in the inversion efforts presented in chapter 8.

There are other studies that compare SBES results to MBES. Parnum et al., compare a duplicate of RoxAnn to a segmentation method applied to MBES data [Parnum et al 2009]. When comparing the acoustic results to ground-truth information, they found the MBES 90% accurate and those from SBES approximately 80% and 70% accurate at 200 kHz and 38 kHz respectively. Some studies prefer to combine SBES and MBES classification with additional data sources and this has turned out well [Pouliquen et al., 2002; Kenny et al. 2003]. Other work has used multiple SBES at different frequencies – the different frequencies cause the echoes to come from different scattering regimes so that the classification is based on different seabed characteristics. Even though more attention is moving toward MBES, there is still a clear role for SBES. This dissertation will make use of all of the above: multiple frequencies and multiple data-sources.

Chapter 3. Applications of Current Methods in Single-Beam Seabed Classification

In this chapter, four different applications of the standard QTC seabed classification method will be briefly discussed. These surveys are applications of the QTC seabed classification method for novel purposes in habitat mapping and assessment, and applications with novel experimental geometries. These include:

1. a log boom debris survey neat Kitimat, B.C.;
2. an ROV-mounted SBES survey for seabed classification of gas hydrates in Barkley Canyon;
3. an ROV-mounted SBES survey for seabed classification of siliceous sponges on Fraser Ridge;
4. a survey for habitat mapping of Pacific Sand Lance habitat along the East coast of Vancouver Island.

In demonstrating the capacities of the current technology, these surveys also present a rigorous methodology of interpreting survey data. The analysis of the surveys identifies the limitations of the current approach – non-seabed influences causing misclassification, poor discrimination and difficulties in interpreting the classification results.

The locations of the four surveys are shown in Figure 3-1. All survey and experiment sites are located within British Columbia, Canada. Surveys 2, 3 and 4 as listed above will only be referred to in this dissertation and will not be presented in full, in order to shorten this dissertation. This author's research on ROV-mounted SBES surveys can be found presented in [Bloomer et al., 2007], while the Pacific sand lance surveys were presented at two meetings of the pacific sand lance consortium [Biffard et al., 2008; Biffard et al., 2009] and will be presented at an upcoming conference[17]. These

---

[17] The citation for which will be: Biffard, B. R., Robinson, C., Chapman, N. R. and Bloomer, S. F. Acoustic Seabed Classification Investigation of Pacific Sand Lance Habitat, In *161st Meeting of the Acoustical Society of America*, May 2011.

surveys, which are side-projects to the main research of this dissertation, demonstrate the problems of slope and inadequate discrimination in SBES seabed classification. The Pacific Sand Lance survey in particular demonstrates the use of OCPR as introduced in sub-section 2.2.3 and objective interpretation of ground-truth data as introduced in sub-section 2.2.4.

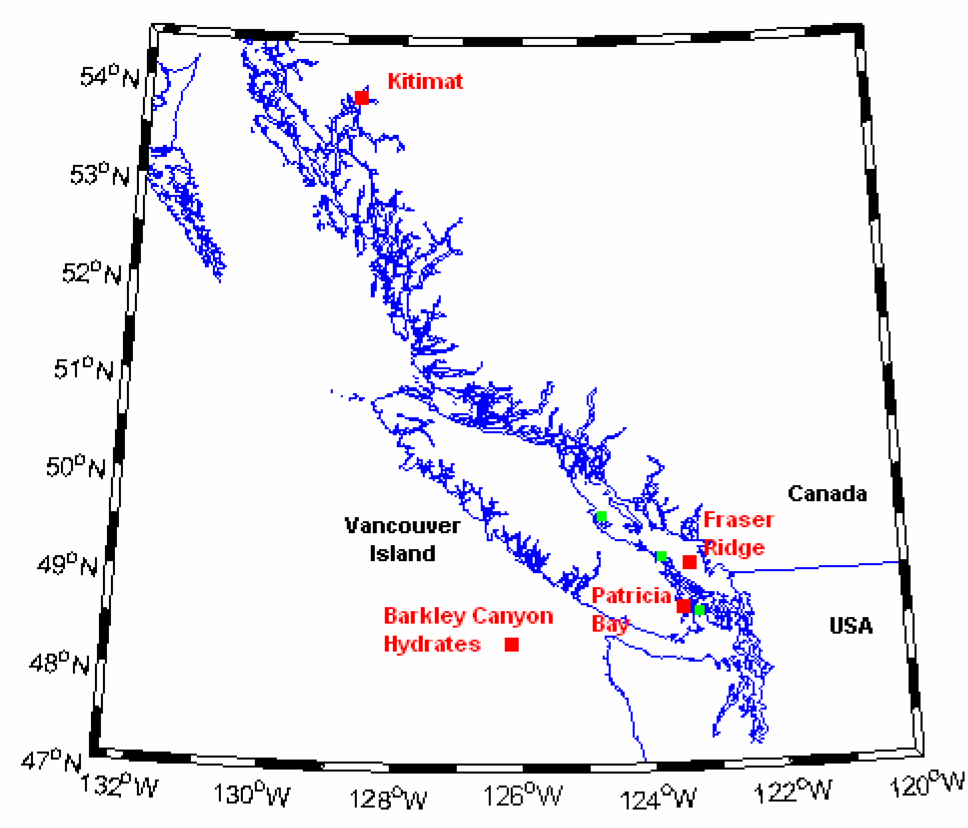

Figure 3-1. Location map for all seabed surveys and experiment sites for this dissertation. The unlabelled green squares are the three Pacific Sand Lance survey locations.

3.1 The Kitimat Log Boom Debris Survey and the Problem of Seabed Slope

In this survey, the standard classification methodology results in classes that represent degrees of seabed slope, instead of seabed type. The investigation described here shows that seabed slope completely dominates the acoustic diversity of the seabed and is the most likely cause of poor performance in seabed classification. This issue has not been addressed properly before. Compensating or addressing seabed slope in some way is vital to improve the performance of single-beam seabed classification.

A standard SBES survey was conducted May 12-13, 2003 in two small fjords (Eagle and Clio Bays) on the north coast of British Columbia, Canada, see Figure 3-1. This survey was a follow-up to an initial SBES survey with a QTC4 data acquisition system, side-scan sonar, grab-sampling and extensive towed video. The aim of the initial project was to investigate the effect of log booming activities on benthic flora and fauna of the bays. Clio and Eagle Bays are nearly identical in size and topography: both bays are steep-sided and U-shaped. These steep-sided bays are typical drowned hanging valleys created by glaciation. The singular difference between the two bays is that Clio Bay has been a log booming site for last 40 years. Sunken logs and log debris have added a large amount of organic material to the seabed so that the seabed is primarily sandy organic mud. The side-scan sonar (SSS) image in Figure 3-2 shows the large of number of sunken logs on the seabed. Eagle Bay is homogenously muddy sand (as determined by video observations and grab-sampling).

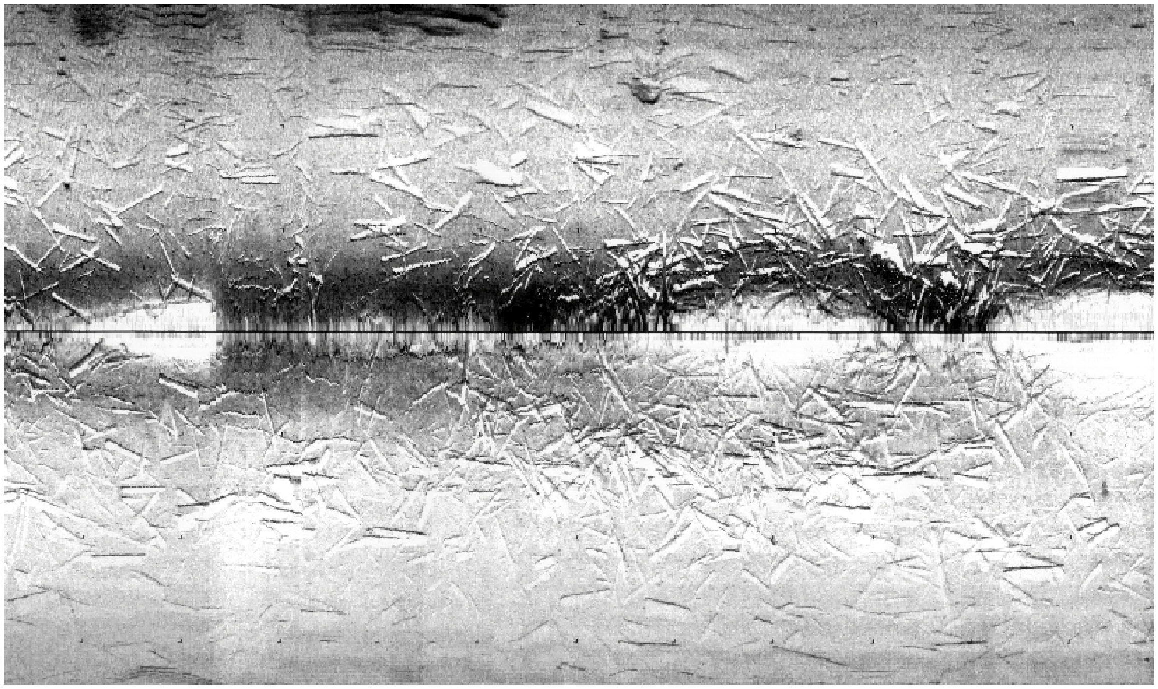

Figure 3-2. An example raw SSS mosaic from Clio Bay. The survey progressed from left to right, nadir is in the centre of the image (marked by a black horizontal line).

The aim of the surveys was to map the extent of the log booming affected seabed in Clio Bay and compare Clio Bay to the untouched Eagle Bay. However, the initial 200 kHz SBES survey failed to produce useful seabed classification. To address this, a follow-up survey was run with a dual frequency (24 and 200 kHz) Odom Mk.II SBES system with a QTC5 data acquisition unit (only the 24 kHz frequency was used). More information on this system is available in appendix C and in [Biffard et al., 2006]. The data were processed with QTC IMPACT software version 3.40 using the standard procedure and parameters as discussed earlier in chapter 2 and in appendices A and B. ACE produced five classes for the combined Eagle Bay (Figure 3-3) and Clio Bay (Figure 3-4) data set. The data sets were combined as both bays were surveyed in the same way on consecutive days.

Although the data collected in the follow-up survey was of higher quality than the initial survey (more powerful echosounder, lower frequency, wider beam), the seabed classification produced by the second survey also did not correspond or correlate well with the ground-truth information. In addition to grab sampling, the ground truth consisted of video and side-scan data interpreted into maps of seabed type and sunken log density (Figure 3-4 *lower*). Instead of representing distinct seabeds, the classifications correlate very well with bathymetric slope calculated from SBES bathymetry. In both bays, classes 3, 1 and 2 correlate with slopes greater than 5° and respectively represent increasing levels of slope. Inspection of the data shows the echo time series of classes 3, 1 and 2 increase in duration with increasing slope. Class 4 dominates Eagle Bay representing sandy mud, while classes 3, 1 and 2 are present on the edges of the bay in areas of slope. Class 4 is also present in Clio Bay representing the non-slope sandy organic mud seabed regardless of the presence of sunken logs. Class 5 occurs in the South-East corner of Clio Bay representing gravel. Classes 3, 1 and 2 dominate much of Clio Bay where there is a much larger area of sloped seabed than in Eagle Bay. Apparently, the slope of Clio Bay dominated the QTC classification precluding information about the extent of the log booming debris and sunken logs.

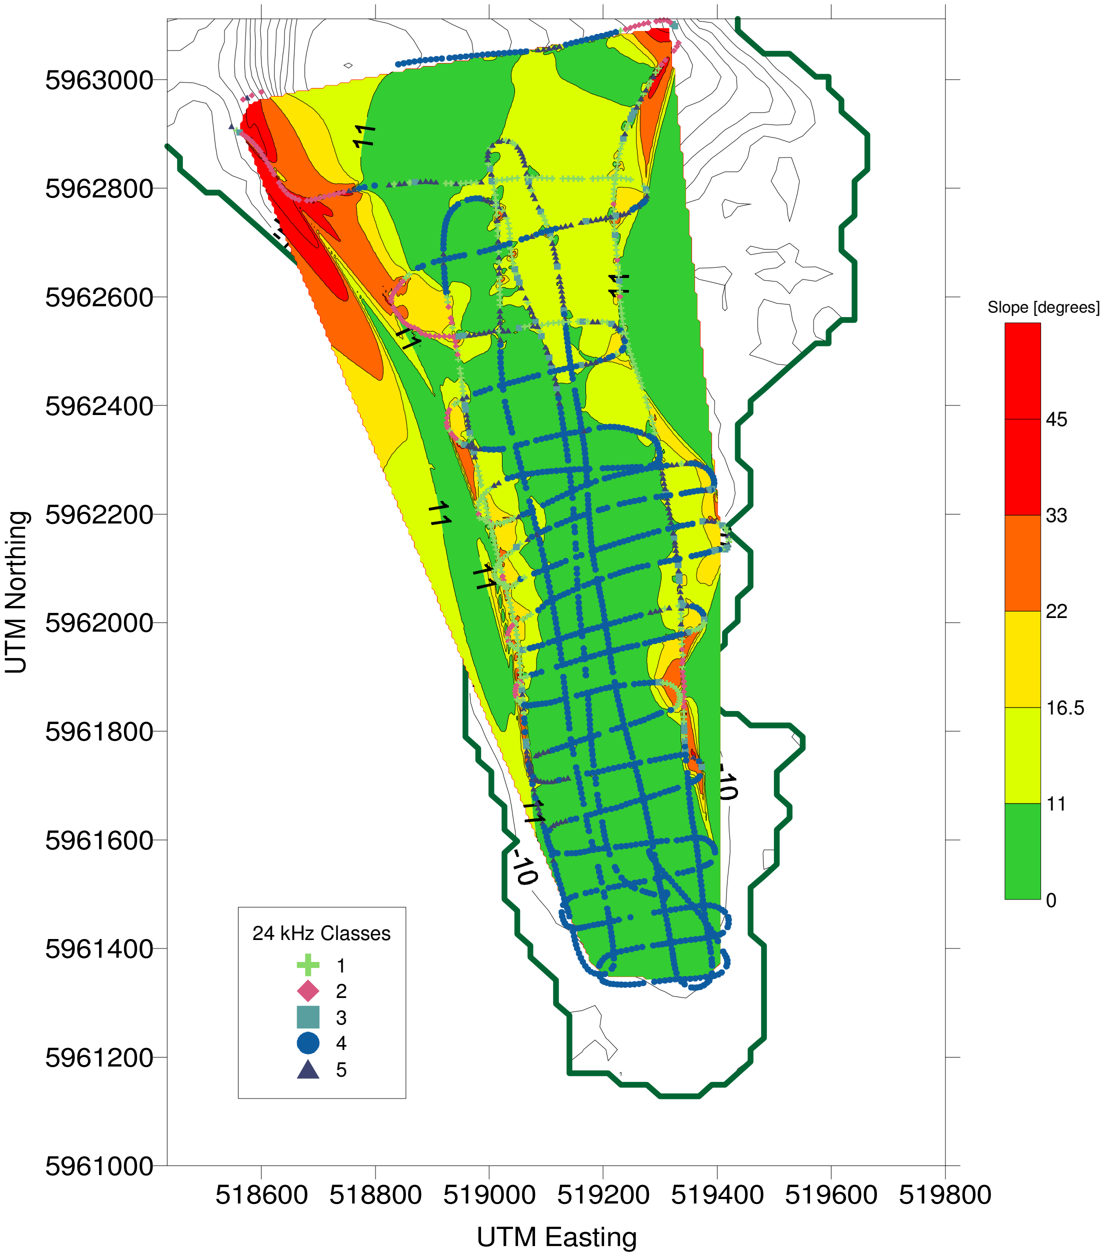

Figure 3-3. SBES seabed classification for Eagle Bay overlaid on to seabed slope. The coastline (green line) and bathymetry were generated by the initial survey.

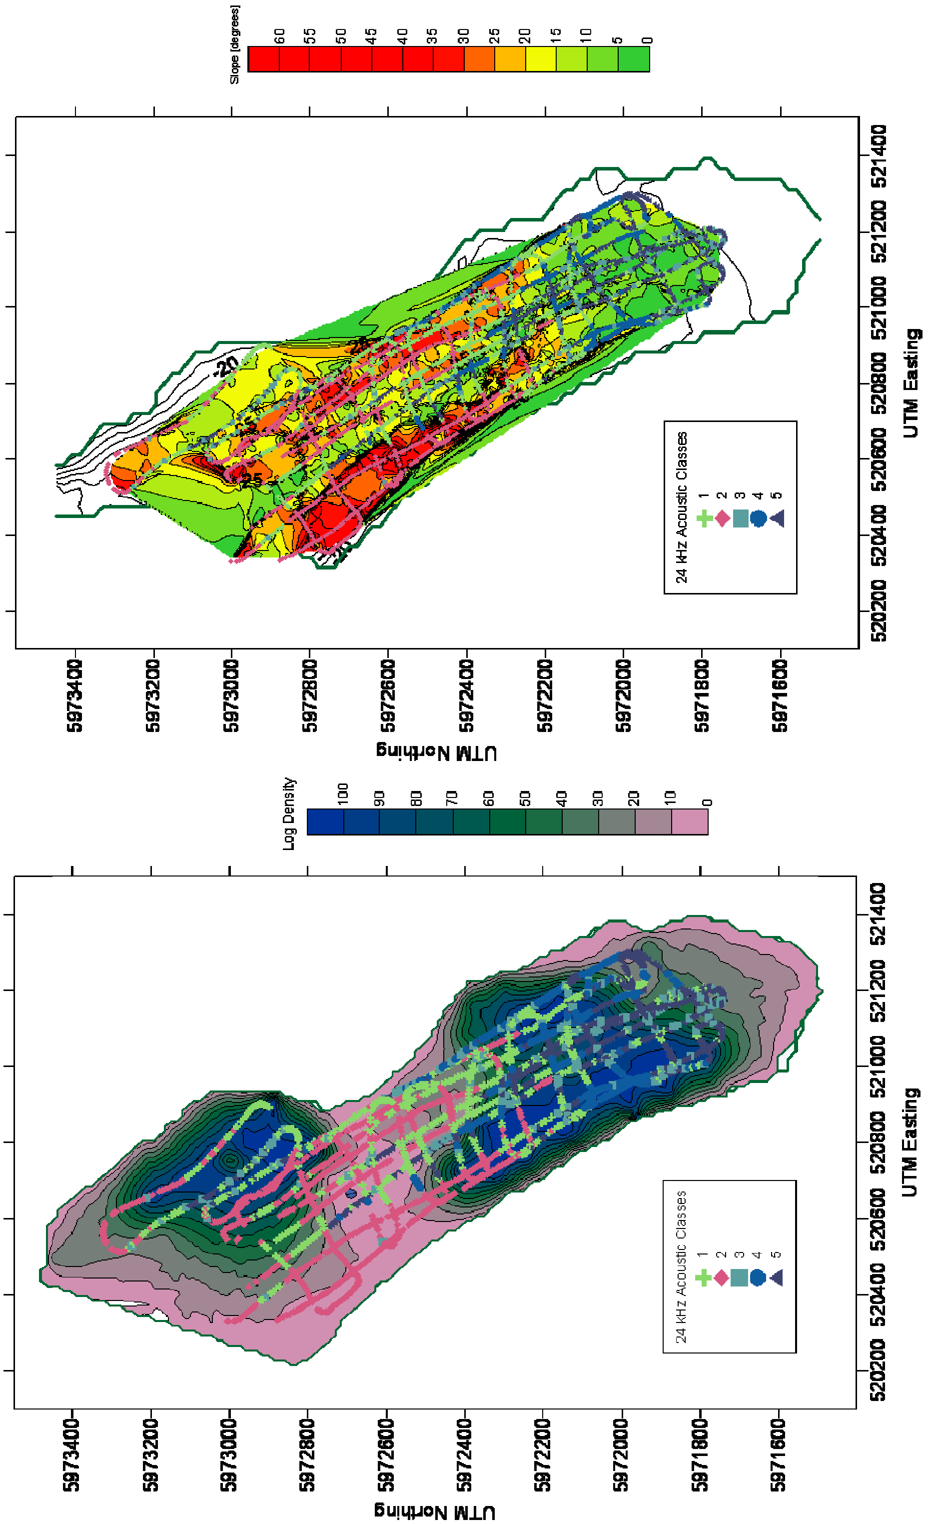

Figure 3-4. SBES seabed classification for Clio Bay overlaid on (*lower*) sunken log density and (*upper*) seabed slope calculated from survey bathymetry. The coastline (green line) and bathymetry were generated by the initial survey.

3.1.1 Implications of Seabed Slope on SBES Seabed Classification

The potency of slope in affecting SBES was clearly demonstrated by the inability of the methodology to map sunken logs, log-booming debris or seabed-type in Clio Bay. The seabed classification results correlated highly with seabed slope and only weakly to the presence of gravel in the SW corner of Clio Bay. In addition to the surveys of this dissertation, there are a number of other examples of surveys affected by seabed slope. Some authors have speculated or realized that slope is affecting their results, [Gleason, 2009; Cooke et al., 2003; Von Szalay and McConnaughey, 2002]. The magnitude of the effect of seabed slope is significant as discussed in [Biffard et al., 2005b]; that discussion is expanded here.

With similar transducer parameters as those used in this survey, modelling results published in [Sternlicht and de Moustier, 2003a] can be used to infer that the total backscatter strength drops by 15 dB for most sediments through a change in angle of incidence from 0° to 20°. No other quantitative analysis of the effect of slope exists, so some new results are presented here to aid this discussion. Later in section 5.3, a model of echo duration will be presented that is useful to quantify the effect of seabed slope. Echo duration represents overall echo shape as all shape features are correlated to echo duration as will be shown in chapter 7 (Figure 7-14, Figure 7-15). For the SBES of this survey, echo duration increases by 6.3 times through a 0° to 20° change in angle of incidence. Through 45°, the increase is 23 times. More echo duration model results are plotted in Figure 5-12. The echo duration model can also be used to show that the increase echo duration due to a change in seabed type, from gravel to clay, is roughly 4 times (for shallow water). It is apparent that modest slopes of 10° to 20° will completely dominate any change in echo character caused by a change in seabed type. In a more conservative scenario, seabed slopes as small as 4° could cause misclassification. Seabed slopes equal to or greater than the one-half beamwidth are also significant because the direct path specular reflection is lost, significantly affecting the character of the echo (as will be shown in section 5.3). Figure 2 shown in [Sternlicht and de Moustier, 2003a] confirms the strong dependence of surface scattering and also indicates that volume backscatter strength is largely independent of angle of incidence or slope. Seabed slope acts to spread volume scattering over a larger footprint, increasing the duration of echoes.

Echo envelope traces from Clio Bay and example modelled echoes found in [Sternlicht, 1999] also illustrate these changes in backscattering for increasing seabed slope. These echoes show that backscatter in the tail of the echo (due to volume scattering) peaks later in the echo and is stretched out in time, while the interface reflection diminishes to zero with increasing tilt/slope. The finding of a tolerance slope of 5-8° in [Von Szalay and McConnaughey, 2002] is accurate as this is the slope at which interface scattering is significantly decreased and modified in shape. This level of seabed slope is roughly equal to the one-half beamwidth.

As seen in the above results, slopes greater than a few degrees have an observable effect on echoes, while echoes from slopes greater than one-half beamwidth are completely distinct from echoes from the same seabed with zero slope. For slopes greater than the ½-beamwidth, QTC SBES classification clearly becomes slope-driven, overriding any signal from the seabed morphology.

Seabed slope would also undoubtedly have a large effect on the RoxAnn / echoPLUS / VBT methods. RoxAnn's *E1* feature (or *B1* in VBT) and VBT's *B3* (the energy in the rise of the first echo plotted against the energy of the tail of the first echo), would be significantly affected as the surface scattering component decreases and the echo tail is stretched with increasing slope. RoxAnn's *E2* (or *B2* in VBT) would be affected by a dramatic decrease in the energy of the second path echo, which would be out of proportion to the decrease in *E1*. It is likely that for most situations the second echo would disappear for slopes greater than one-half beamwidth, as was observed in [Hamilton et al., 1999]. With the present methods, including QTC VIEW, RoxAnn and impedance measurements, successful classification of seabeds that have seabed slopes greater than the one-half beamwidth is not possible.

A redeeming aspect is that there is significant correlation of seabed type with slope. As slope increases, the angle of repose of progressively coarser sediments is exceeded. Seabeds with slopes greater than 30° are rock or boulders, with few exceptions. However, modest slopes of 5 to 20° affect the classification results adversely, causing most of the classification errors seen throughout this chapter. This sorting of seabed type by slope is similar to the effect of depth: sediments tend to get progressively finer with depth. Because of these two effects, classification errors caused by slope and depth are not as readily apparent as one may assume. For this reason and many others, controlled experiments are required to investigate and solve this problem.

3.2 Discussion and Definition of the Research Problem

The surveys described in this chapter all show classification errors caused by seabed slope. TNORM depth compensation has a limited range of reliability, while the main issue with SEL depth compensation is in defining appropriate values of beamwidth and attenuation to be used. These problems will be addressed in chapters 5 through 7.

The surveys also have instances where clearer discrimination between seabed types is needed. The results presented in [Bloomer et al., 2007; Biffard et al., 2009] demonstrated cases of over or under-splitting which was partially handled by the new use of OCPR. Inappropriate clustering is likely due to the limited nature of the features used – only shape features are currently useful. Adding features that offer extra dimensions of information may solve this; these features should be based on characteristics of echoes that are not already utilized: echo duration, amplitude, spectral content and variability.

There is also difficulty interpreting the results of phenomenological segmentation because the only class labeling information available is from ground-truth data that are mostly localized measurements that may not characterize what the acoustics is responding to. Being able to associate feature values that have connections to the physical characteristics of the seabed to the classes would have help assign labels to classes. Even better would be to a characterization scheme to offer direct information on the seabed that could be used to aid in labeling the acoustic classes from segmentation. To do this, a basis in physics (in theory, models and experiments) is needed to develop new features, improve depth and slope compensation and to develop a successful characterization scheme. The basis will be established in chapters 4, 5 and 6, and then applied in chapters 7 and 8.

Chapter 4. Existing Theory and Models

In chapter 2, a review of current seabed classification methods found that the few methods with a basis in physics are not yet in widespread use nor have been successful in practical applications, the work of Sternlicht is an example of such a method [Sternlicht, 1999]. Instead, it is the statistical methods that are widely used. Without a solid basis in the physics of the acoustic scattering process, these statistical methods will always be phenomenological and require ground-truth data; properties that hinder their success as noted earlier. The RoxAnn method was originally based on anecdotal observations, but now has a solid basis on the work by Heald [Heald 2001, 2000; Heald and Pace, 1996], which has aided its interpretation and improved its reliability by noting conditions under which the method works better. Such a basis does not yet exist for the QTC method. The physical understanding introduced in this chapter will be built upon throughout the remainder of this dissertation as a basis for an improved QTC method. The application of this basis will be used to shed light on the nature of the problems experienced in chapter 3, and offer improvements in two main areas: removal of non-seabed influences of depth and seabed slope (chapter 6) and improved discrimination by purposely designed features that exploit relationships between echo attributes and geoacoustic properties (chapter 7).

Compiling a comprehensive theory of echosounding is not a straight-forward task. There is no one complete exact theory, instead a multitude of theoretical models have been developed, each with unique conditions of applicability. Comparing the different theories is also made difficult by the different conventions and definitions used. To avoid getting caught up in these issues of semantics, this theory chapter either unifies them or defers to the references when appropriate.

4.1 Intuitive Overview

This overview of scattering theory and results is based on the following sources: [Ogilvy, 1991; Medwin and Clay, 1998; Lurton, 2002; Biffard et al., 2007]. A recent book by Jackson and Richardson is also an excellent resource [Jackson and Richardson, 2007].

The ultimate goal for many theories is to predict a phenomenon based on the initial conditions. In this case, the problem is defined as a signal from a SBES incident on the seabed. There are many ways to simplify the definition of the problem, the most common way (as shown in chapter 2) is to assume a conical beampattern (Figure 2-3), normally incident on a flat seabed. However the simplest case of echosounding is that of a perfectly flat seabed that does not require any specification of the beampattern. In this case, the Snell-Descartes law of reflection and refraction would dictate that the only return, or *backscatter*, to the transducer would be the acoustic ray that is normally incident at the seabed – all other rays would be reflected away from the transducer (assuming far field effects). The echo would be a reflection of the transmit pulse, modified by spherical spreading and by the reflection coefficient. The normal incidence reflection coefficient is a function of the impedance contrast between the seabed and the water:

$$\Re_{\perp} = \frac{z_w - z_s}{z_w + z_s}, \quad z_w = \rho_w c_w, \quad z_s = \rho_s c_s \tag{4-1}$$

where z_w is the impedance of the water and z_w is the impedance of the seabed as determined by the appropriate density ρ and sound speed c. This is the simplest form of the reflection coefficient. The situation is more complicated if the seabed is an elastic solid and/or has absorptive properties, which is the case for sand and gravel sediments. Such sediments do support shear waves with a non-zero shear modulus. Fine sediments, such as silts and clays, are usually elastic fluids and do not support shear waves. In either case, the effects of both shear and absorption on the reflection coefficient for most seabeds are minimal, especially at normal incidence[18]. The type of echo arising from a normally incident beam on a perfectly smooth surface, as described earlier, is the *specular* echo. The specular echo is a reflection, like that of a mirror, which preserves all

---

[18] The reflection coefficient is a strong function of angle; it increases from $\Re_{\perp}$ to one (perfect reflection, no transmission), remains at one for all incident angles above a critical angle (usually 45 to 60°). For incidence at or above the critical angle, shear and absorption reduce the reflection coefficient allowing some transmission. The effects of shear and absorption on the reflection coefficient are only important for wide beams, or beam at non-normal incidence, see Figure 3.4 in [Lurton, 2002], pp. 67.

characteristics of the transmit pulse except its direction. Specular echoes are *coherent*, meaning they have definite phase relationship with the transmit pulse. The theories that deal with specular echoes are well defined – seismic systems and sub-bottom profilers deal mostly with this type of echo which allows for inversion and direct characterization.

If the seabed surface varies in height on lateral spatial scales from one-tenth of the wavelength up to the diameter of the footprint, it has *roughness*. Roughness of a larger scale is seabed slope. Smooth seabeds are like mirrors, they only return acoustic energy to a monostatic sonar transducer on normally orientated rays, while rough seabeds backscatter non-normal acoustic rays from many directions. Backscatter from rough surfaces can have a mirror-like specular or coherent component[19] but is usually dominated by diffuse or incoherent backscatter. Incoherent echoes are characterized by chaotic fluctuations within the echo time series and between adjacent echoes of a survey – i.e. these echoes have high ping-to-ping variability. These chaotic fluctuations are due to interference among various backscattered acoustic rays. Surficial scattering is controlled by the distribution of roughness heights. Most theories parameterize the distribution with a measure of the amplitude of the height variation (e.g. root-mean-square or RMS roughness height) and a measure of the shape of the distribution. Seabed surface scattering theories will be discussed at length later in this chapter because of the dominant contribution of surface scattering on seabed echoes.

Backscattering from the seabed volume is often neglected when considering the formation of SBES echoes. There are several factors that control volume scattering. First, how much of the transmit pulse penetrates the seabed and then how much is backscattered. The amount penetrating the seabed is determined by the transmission coefficient, the magnitude of which is $1-|\Re(\theta)|$, where θ is evaluated as the angle of incidence on the small local surface that is directly affected by seabed roughness. In the idealized case of sound waves normally incident on a smooth, homogenous seabed that does not support shear waves, the transmission coefficient is frequency independent. Sound transmitted into the seabed volume is extinguished 0.1 to 5 m below the seabed surface as controlled by the attenuation coefficient. Above a few kilohertz, attenuation in

---

[19] The terms specular, mirror-like and coherent all refer to the reflection from a smooth surface.

marine sediments increases linearly with frequency, so higher-frequency sound (e.g. 300 kHz) passing through the seabed surface is quickly attenuated, while lower-frequency sound (12 kHz) penetrates to a depth of several metres, generating significant volume backscatter along the way. The strength of the volume backscatter component is usually modelled by a volume backscattering cross-section[20] or a measure of the density of volume scatterers. Volume scattering is usually assumed to be isotropic, single scattering and frequency independent, although different approaches exist. Volume scattering is certainly incoherent due to the random positions of the scatterers [Jackson et al., 1986b]. The scatters, or inhomogeneities, consist of sediment laminae (i.e. layers), large distinct grains (e.g. stones within a sand seabed), or features caused by bioturbation: bubbles, burrows, wood debris, plants, animals, shells, etc.

Overall, seabed echoes are comprised of four components: the specular echo, coherent surface backscatter, incoherent surface backscatter and volume backscatter. The relative contribution of those four components (and coherent vs. incoherent), the beampattern, water depth, seabed slope and nature of the transmit pulse are the factors that shape seabed echoes. All of these factors are summarized in Figure 4-1.

---

[20] The exact physical definition of a scattering or backscattering cross-section depends on the situation, including such factors as the proportion of the incident wave intercepted, the area or volume of the scatterer (be it a target, swarm of targets, a surface or a volume). See [Lurton, 2002] for details. In general, scattering cross-sections express the ratio of the incident intensity to the scattered intensity.

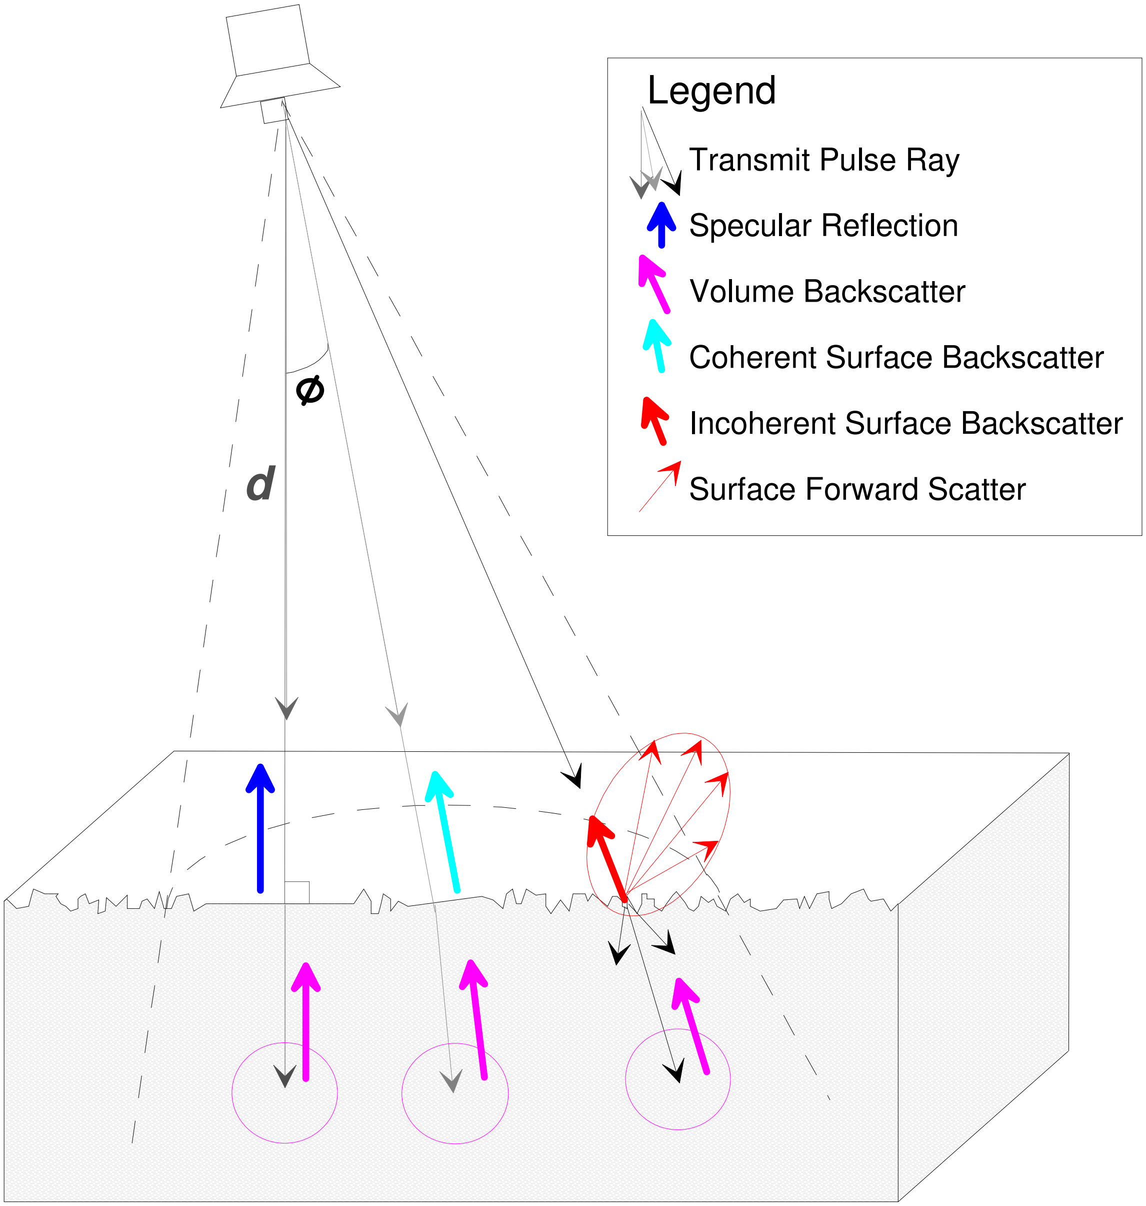

Figure 4-1. The formation of a SBES seabed echo shown by backscattering components. The down-going transmit pulse is contained within the beampattern and footprint (dashed lines), the axis of which is at an incidence angle *Φ* to the seabed (this is equivalent to seabed slope for a vertical transducer). In cases of low roughness, the specular reflection is the first arrival, indicating water depth *d*. Volume scattering is usually isotropic (shown by the magenta circles), while surface scattering is not (shown as a red ellipse) – forward scatter does not contribute to the received echo. Coherent surface backscatter is consistent echo to echo and includes the specular reflection plus scattering from an ensemble of facets that are also consistent.

4.2 Geoacoustic Parameters of the Seabed

The various theories and models pertaining to echosounding all require geoacoustic and geotechnical parameters. The relevant geotechnical parameters were defined in chapter 2: water depth, angle of incidence (including seabed slope), beamwidth/beampattern, centre frequency and transmit pulse duration. Geoacoustic parameters are characteristics of the seabed that influence the acoustic response of the seabed, such as sound speed, density, attenuation, surface roughness characteristics and volume scattering characteristics.

There are several tabulations of geoacoustic seabed parameter values, which fall into two groups due to differences in models of seabed roughness. The BORIS numerical model of echo time series[21] and associated theory [Pouliquen et al., 1999] specifies the roughness height distribution using correlation functions. Berntsen published a tabulation developed at SACLANTCEN[22] for use as input for the BORIS model to generate seabeds with realistic properties [Berntsen, 2001]. There are several example sets of parameters included with the BORIS manual and software [Canepa et al., 2004]. The other group of tabulations is based on models compiled in the APL-UW handbook [APL-UW, 1994] which originates from [Jackson et al., 1986], improved in [Mourad and Jackson, 1989] and evaluated in [Jackson and Briggs, 1992]. This approach specifies the roughness height distribution using coefficients of a power law spectrum. The time-domain model of SBES echoes in [Sternlicht, 1999] uses the geoacoustic parameters of the APL-UW handbook [APL-UW, 1994]. Both sets of parameters have values consistent with the values predicted by theory [Brekhovskikh and Lysanov, 1991; Biot, 1956a; Biot, 1956b; Buckingham, 2005] and measured in situ [Hamilton and Bachman, 1982; Richardson and Briggs, 1993; Briggs, 1989].

In this dissertation *both* sets of parameters are used. The seabed characterization methods in chapter 8 make use of the APL-UW handbook geoacoustic parameter definitions and values, while BORIS modelling efforts later in this chapter and in

---

[21] The BORIS model will be described in detail in section 4.5. It is used extensively later in this disseration.

[22] SACLANTCEN is now called NURC – **N**ATO **U**ndersea **R**esearch **C**entre

chapters 6 and 7 use BORIS geoacoustic parameters, the values of which are tuned to the situation based on information from the literature. One of the difficulties of this is comparing the parameters. The APL-UW handbook does have a formula that relates RMS roughness height, h_{RMS}, to the parameter it uses called spectral strength, w:

$$w = 0.00207 {h_{RMS}}^2 h_0^2 \tag{4-2}$$

where h_0^2 is a reference length of 1 cm; all units are in centimetres. This equation assumes a roughness spectral power law exponent of $v' = 3.25$ (the power law will be presented later). A table of geoacoustic parameters from the APL-UW handbook can be found in appendix C, Table C-2 [APL-UW, 1994]. The parameters that are common to both sets are the sediment sound speed c_s, the sediment density, ρ_S, and the attenuation κ_s.

RMS roughness height is a more generic measure of the roughness distribution than is spectral strength. It is applied in several models, so it is useful to explore its definition here. If the seabed surface heights, h, are measured precisely over a surface s, the RMS roughness height is:

$$h_{RMS} = \sqrt{\left\langle h^2 \right\rangle_S} \tag{4-3}$$

(angle brackets indicate a spatial average). The height distribution for most seabeds is considered to be Gaussian [Berntsen, 2001], which makes sense considering the central limit theorem[23]. For a Gaussian height distribution, the probability of a given height, $P(h)$ is:

$$P(h) = \frac{1}{h_{RMS}\sqrt{2\pi}} \exp\left(-\frac{h^2}{2{h_{RMS}}^2} \right) \tag{4-4}$$

The Fourier transform of the probability density function ($P(h)$ in Equation 4-4), results in the characteristic function. It is not possible to reconstruct a seabed directly from the characteristic function, nor is the characteristic function something that can be readily measured. The characteristic function is important as it appears in several scattering theory results representing phase shifts due to the seabed undulations.

---

[23] For a definition of the central limit theorem, see any statistics reference, e.g. [Rencher, 2002]. In this case the seabed heights are the sum of many random processes – the heights should be normally distributed about their mean.

The power spectrum of the seabed surface heights can be used to construct a seabed surface stochastically through a process called Fourier synthesis [Bergem et al., 1999]. This process is important for numerical scattering models, such as the BORIS model. For general scattering theory, the power spectrum is also important as the results of the theory are often functions of the power spectral density or a correlation function. Correlation functions can be measured directly or are Fourier transforms of the power spectra. The relationships between correlation functions and the power spectra can be expressed as follows:

$$W(\mathbf{k}) = \frac{{h_{RMS}}^2}{(2\pi)^2} \int_{-\infty}^{\infty} C(\mathbf{R}) e^{i\mathbf{k}\bullet\mathbf{R}} d\mathbf{R} \tag{4-5}$$

$$C(\mathbf{R}) = \frac{\langle h(\mathbf{r}) h(\mathbf{r}+\mathbf{R}) \rangle_S}{{h_{RMS}}^2} \tag{4-6}$$

where **r** and **R** are location vectors on the seabed surface and **k** is the acoustic wavenumber vector, with magnitude: $k = {}^{2\pi}\!/_{\lambda}$. Equation 4-5 is the relationship between the correlation function $C(\mathbf{R})$ and the power spectrum $W(\mathbf{k})$, Equation 4-6 is the definition of the correlation from measured roughness heights. Model correlation functions are usually either Gaussian or exponential and are functions of the correlation length, λ_o, which is a term that will become important later. The correlation length represents the scale over which the surface heights are significantly correlated, which is the point at which the autocorrelation of the surface heights decreases by a factor of ${}^{1}\!/_{e}$.

Roughness heights can be measured using methods such stereo photography [Briggs, 1989]. From roughness height data, (i.e. $h(r)$), the power spectrum, the correlation function, the height distribution and then the characteristic function can be obtained. In general practice, a model of the height distribution or power spectrum is assumed and then parameters of the model are measured from the data and tabulated. Most power spectrum models result in Gaussian height distributions.

A common model for the surface heights power spectrum is that of a power-law. The power law spectrum is based on fractal theory and is becoming widely accepted [Briggs, 1989]. The seabed surface is modelled as a fractal surface of sine waves. The spatial wavelengths and amplitudes of these waves are related such that the smaller waves are reduced sized versions of the larger ones and vice versa. This is a definition of a fractal, the power spectrum of which is a power law:

$$W(k) = \frac{w}{(h_o k)^{v'}} \qquad (4\text{-}7)$$

where w is the spectral strength as mentioned in Equation 4-2, v' is the spectral exponent, and h_o is a scale factor, usually 1 cm. The spectral strength controls the vertical size of the surface, while the spectral exponent controls relative fractions of the spatial components, which is analogous to the role of the correlation length. However, there is no published relationship between the power-law exponent and the correlation length. The power-law spectral exponent is related to the fractal dimension of the surface.

The two modelling approaches (the APL-UW handbook and BORIS) both apply the power-law spectrum for their models of the seabed surface, although the BORIS model has the option of other spectra, including a Gaussian spectrum for the purposes of modelling the air-sea surface. Another difference is that the BORIS model specifies the spectral strength by way of a conversion from the RMS height, which is presumably of the form of Equation 4-2. An unbounded power law is also unrealistic; spatial wavelengths cannot be smaller than the sediment grain size, while seabed height variation larger than the acoustic footprint is considered bathymetry and seabed slope. The BORIS model also invokes filtering to limit the power law to reasonable bounds. The approach of the APL-UW handbook and Sternlicht [APL-UW, 1994; Sternlicht, 1999] is that the spectral exponent has a constant value of 3.25, while the BORIS model allows values from 3 to 4, while real values may range from 2 to 5 for some extreme cases. Jackson et al. found that values of the spectral exponent range from 3 to 3.5, with a Gaussian height distribution [Jackson et al., 1986]. The consensus from measurements is that the power law spectrum is the best overall model for the seabed, although the negative slope of both the Gaussian and exponential correlation functions does match the trend – these models may be applicable in special cases or as approximations when the theory cannot be solved with power law spectrum.

The way the models handle volume scattering is also different. BORIS uses an RMS measure of the sum of relative sound speed and density variations[24] (and is thus like an RMS 'height' of the volume inhomogeneities), while the APL-UW handbook and more recent work by Sternlicht use two sets of parameters that are directly related to the volume scattering cross-section [APL-UW, 1994, Sternlicht, 1999]. In the older BORIS-3D model [Berntsen, 2001; Caiti and Zoppoli, 1998], the volume model also appears to use a volume scattering cross-section, the values of which are consistent with those found in [Sternlicht, 1999]. The newer BORIS volume model also requires vertical and horizontal correlation lengths that help represent layering – this leads to anisotropic scattering which seems more realistic. Unlike surface scattering, there does not appear to be any formal conversion between the two different approaches, so comparisons will have to rely on measured values. An important geoacoustic parameter common to both approaches is attenuation. It is discussed at length in the next section.

4.3 Scattering Theory

Scattering theory is important to the goal of understanding the physics of echosounding and will facilitate many of the advances of this dissertation. However, a complete review is not strictly necessary to the main text of this dissertation, so the reader is referred to appendix D: A review of scattering theory for single-beam echosounders for the full treatment. Otherwise, the reader can proceed to the specific scattering theories pertaining to the BORIS and Heald models (sections 4.4 and 4.5) and the summary sections (4.6 and 4.7).

---

[24] The relative measure of volume inhomogeneity used by BORIS is: $\mu = \frac{c_V - c_V'}{c_V} + \frac{\rho_V - \rho_V'}{\rho_V}$

4.4 Heald's Model of *E1* and *E2*

G.J. Heald applied variations of the propagation model method presented in the previous section to derive an expression for the energy in the tail of the first direct path echo, called *E1*, and the total energy in the second multipath echo, called *E2* [Heald, 2000, Heald, 2001]. *E1* and *E2* are the echo features for the RoxAnn seabed classification method as described in chapter 2.3. This section will describe the details of Heald's derivation of *E1* as it is the basis for the derivation of an analytic time series model for the first echo that is shown in Section 5.2. The derivation of *E2* is similar to *E1*, the details of which are not important for this dissertation.

Integrating the backscatter intensity from an area element *dA* over the elevation angle (which is a function of time) leads to an analytic expression for the time-dependent backscatter intensity, see Figure 4-2 for the geometry required. The elemental backscatter intensity (not shown) is a function of the scattering coefficient, directivity function and an angular-dependent range similar to Equation D-29. The time-dependent backscatter intensity is:

$$I_{totH}(t) = I_0 \int_{\theta_{trail}}^{\theta_{lead}} \frac{S(\theta_t)\,\beta^2(\theta_t)}{(R'_0)^4}\, dA \tag{4-8}$$

$$dA = 2\pi\, R'^2_0 \tan\theta_t\, d\theta_t \tag{4-9}$$

$$R'_0 = D\sqrt{(1+\tan^2\theta_t)} \tag{4-10}$$

where θ_t is the angle of incidence of the area element. It is an integration variable that is replaced by θ_{lead} and θ_{trail}, so that the integration area is defined as the area between the leading and trailing edges of the transmit pulse as it traverses the footprint. I_0 is the source intensity. The source and the receiver both have the same Gaussian beampattern function $\beta(\theta_t)$, which is also dependent on the $1/e$ beamwidth, $\theta_{-4.3dB}$:

$$\beta^2(\theta_t) = \beta_t\beta_r = \exp\left(-2\left(\frac{\tan^2\theta_1}{\tan^2\theta_{-4.3dB}}\right)\right) \tag{4-11}$$

The backscattering coefficient $S(\theta_t)$ resulting from the high frequency limit solution to the Helmholtz-Kirchhoff integral (assuming Gaussian roughness) is given in [Brekhovskikh and Lysanov, 1988] and repeated in [Pace, 1990] and may be written as:

$$S(\theta_t) = \frac{\Re_{\perp}{}^2 \lambda_0{}^2}{32\pi h_{RMS}{}^2} \left(\frac{\exp\left(-\left(\frac{\lambda_0 \tan\theta_t}{2h_{RMS}} \right)^2 \right)}{\cos^4 \theta_t} \right) \tag{4-12}$$

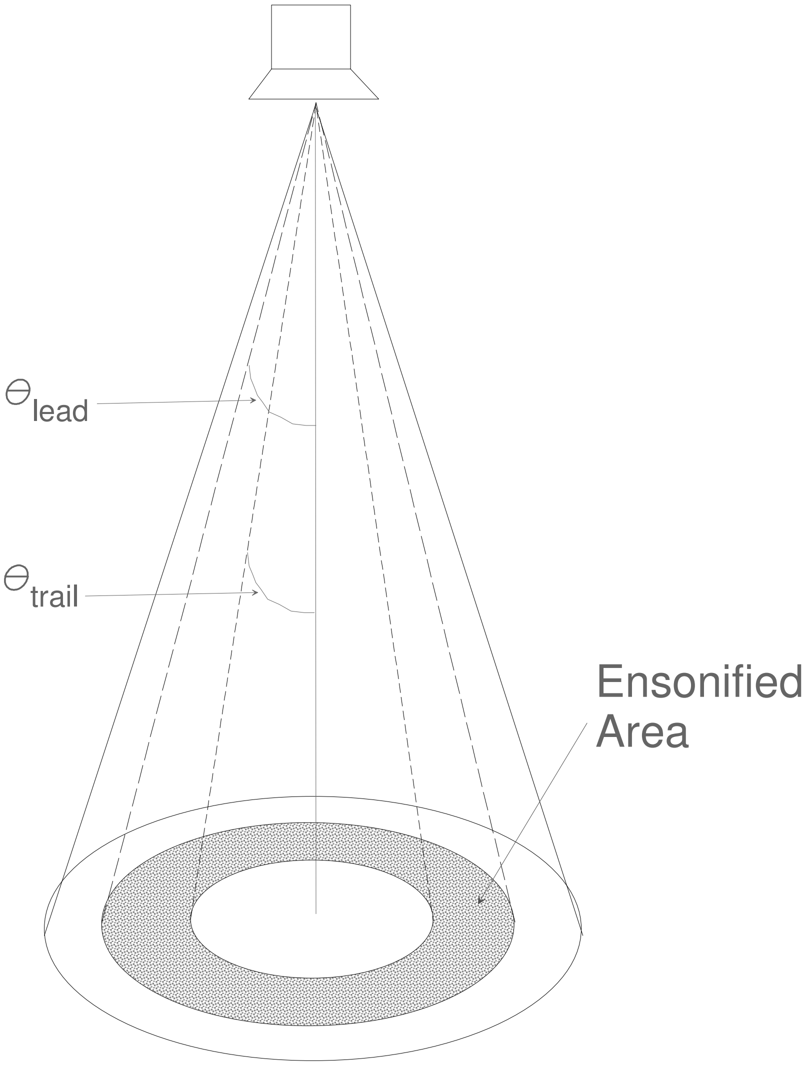

Figure 4-2. Geometry for the derivation of Heald's model of E1. The vertical solid grey line is the seabed normal and transducer axis (for normal incidence), the length of which is proportional to the true water depth, *D*. The ensonified area shown should not be confused with the areal differential, *dA*, which would be an infinitesimal ring defined by the angle θ_t.

The scattering coefficient is dependent on the RMS slope of the seabed, as it contains the ratio of RMS height and correlation length:

$$s_{RMS} = \frac{h_{RMS}\sqrt{2}}{\lambda_0} \tag{4-13}$$

RMS slope is defined for roughness of spatial scale less than the footprint size, so for clarity, it is called RMS *micro*-slope later in this dissertation. It expresses the mean slope of the scattering facets.

The scattering coefficient is also independent of frequency – it cancels out in the derivation. This is expected as the intensity calculated here is the total average intensity. The scattering coefficient does not handle the directional preferences of the coherent and incoherent fields – it is the simplest solution to the HK equation possible generated by taking the 'high-frequency limit', see [Brekhovskikh and Lysanov, 1988].

To solve Equation 4-8, the angular terms must be resolved. They are:

$$\begin{aligned} \theta_{lead} &= \sqrt{\frac{ct}{D}} \\ \theta_{trail} &= \sqrt{\frac{c(t-\tau)}{D}} \end{aligned} \tag{4-14}$$

The angular terms are valid for $t > \tau$ with *t=0* at the onset of the echo; this means the model is only defined for the tail of the echo (i.e. one pulse duration after onset). Section 5.3 will derive new angular terms and a new solution for Equation 4-8 so that the model applies to the entire first echo. The solution to Equation to 4-8 using Equation (4-14) is:

$$\begin{aligned} I_{totH}(t) &= \frac{M}{2\beta^2}\exp\left(\frac{-\beta^2 ct}{D}\right)\left(\exp\left(\frac{-\beta^2 c\tau}{D}\right)-1\right) \\ \text{where: } M &= \frac{I_0\Re^2T^2}{16h_{RMS}{}^2D^2}, \beta^2 = \frac{T^2}{4h_{RMS}{}^2} + \frac{2}{\tan^2\theta_{-4.3dB}} \end{aligned} \tag{4-15}$$

To find *E1*, $I_{totH}(t)$ is integrated over time, which is straightforward but not shown. Equation 4-15 is the only known analytic expression of a high-frequency echo time series.

4.5 The BORIS Numerical Echo Time Series Model

The BORIS model is an accurate and widely-accepted model that generates realistic echo time series [Canepa et al., 2004; Pouliquen et al., 1999; Bergem et al., 1999, Canepa et al., 2008]. BORIS is a time domain model similar to the models already described: Sternlicht's model [Sternlicht, 1999; Sternlicht and de Moustier, 2003a; Sternlicht and de Moustier, 2003b] and Heald's model [Heald, 2000; Heald, 2001]. The version of BORIS used in this dissertation is BORIS-SSA version 1.0, 2006. BORIS is an acronym for **BO**ttom **R**esponse from **I**nhomogeneities and **S**urface augmented with the **S**mall **S**lope **A**pproximation. The only aspects of echosounding in situ that BORIS does not replicate are the input and output transfer functions of the transducer and the water column effects of attenuation and ray-bending. Both sets of effects are reasonably assumed negligible for the simulation scenarios of this dissertation. As will be demonstrated later in chapter 5, BORIS can be used to simulate the raw signal acquired by an echosounder. In order to use the model effectively to create realistic simulations, the theoretical background must be understood. The BORIS documentation, [Canepa et al., 2004], states:

> "It is not possible to run BORIS-SSA intuitively: the meaning of all the options and the correlations between them have to be well comprehended before attempting a simulation with personal parameters."

BORIS is based on KA theory (appendix D.5) with small slope approximation (SSA) corrections from Perturbation theory (appendix D.4). The remainder of this section describes this theoretical background and its implementation. However, the literature provides little information on using BORIS to create real SBES echoes; no information on using BORIS for this purpose has been published. All published BORIS simulations have been done with either sub-profiler type sonar [Bergem et al., 1999; Berntsen et al., 1999; Caiti and Zoppoli, 1998] or bistatic scattering experiments [Isakson et al., 2005]. Significant work has gone into to the setup of a simulation environment, including tuning all the parameters for generating realistic SBES echoes. Since this is new information, it appears in the new theory and models chapter, in section 5.1 and appendix F.

BORIS calculates the sound pressure level[25] of seabed backscatter at discrete time intervals, similar to how an acoustic data acquisition system samples acoustic signals. In fact, the same rules about sampling rates and preserving the signal bandwidth as discussed in chapter 2 apply to BORIS. For each time step, BORIS calculates the surface and volume contributions to the total echo separately:

$$p(\mathbf{P},t) = p_s(\mathbf{P},t) + p_v(\mathbf{P},t) = \int_s dp_s(\mathbf{P},t) + \int_v dp_v(\mathbf{P},t) \tag{4-16}$$

Equation 4-16 expresses that the pressure field received at location **P** at time, *t*, from the seabed is the sum of the elementary backscattered pressure fields integrated over the seabed surface *S*, and the seabed volume *V*. The individual equations for the elementary fields are:

$$\begin{aligned} dp_S(\mathbf{P},t) &= -\frac{\cos(\varphi(\mathbf{R}))}{2\pi c_0 R_0^2} p_0 \beta_i \beta_r(\mathbf{R}) \Re_{01}(\mathbf{R}) e'\left(t - \frac{2R_0}{c_0}\right) d\mathbf{S} \\ dp_V(\mathbf{P},t) &= -\frac{-n_1^{\,2}(z_1)}{2\pi c_0^2 R_0^2} \mu(z_1) p_0 \beta_i \beta_r(\mathbf{R}) \mathrm{T}_{01}(\mathbf{R}) \mathrm{T}_{10}(\mathbf{R}) \\ &\qquad F^{-1}\left\{e^{-\alpha|f|R_1} f^2 E(f)\right\} e\left(t - 2\left(\frac{n_1 R_0 + R_1}{c_0}\right)\right) d\mathbf{V} \end{aligned} \tag{4-17}$$

where **R** is a time-dependent position vector relative to **P**, R_o is the water depth, p_0 is the initial amplitude (usually 1), c_0 is the average sound speed in the water, and φ is the incident angle relative to the local seabed normal. The β functions are the directivity, $\Re$ is the reflection coefficient and $e'(t)$ is the first time derivative of the transmit pulse, $e(t)$. In the volume expression, n_1 is the refractive index, z_1 is the depth below the surface, $\mu(z_1)$ is the sum of the relative fluctuations of sound speed and density from their respective bulk values, the *T* terms are the transmission coefficients, $F^{-1}\{F\}(t - t_0)$ indicates the inverse Fourier transform of $F(f)$ at the time $(t - t_0)$ and α is the sediment attenuation coefficient.

---

[25] relative to the transmitted sound pressure level measured at one metre, the standard value is unity.

Unlike the Sternlicht and Heald models, these expressions contain no scattering cross-section. Instead, BORIS applies a numerical evaluation of the integrals of Equation 4-16 that is effectively a *direct application* of the Kirchhoff Approximation: the total acoustic field is the sum of all the infinitesimal contributions (dp_S / dp_V) from every surface / volume sample. Models of the seabed volume inhomogeneities and surface heights are generated stochastically to fit the input parameters and then sampled at very high resolution (>10 times the acoustic wavelength), for example see Figure 4-3. Then the seabed is completely known, and this allows the direct calculation of the acoustic fields. A surface plane is defined for each surface sample with its neighbouring points; the scatter is determined primarily by the orientation of its normal vector. BORIS sums the field from every surface sample. The equivalent philosophy is applied to the volume component as well. As seen in the surface tile and volume cube shown in Figure 4-3, the seabed model is limited in extent to save memory, (in the case of Figure 4-3 128 MB was used). However, this model is continuous on its edges so that the surface volumes model may be tiled together to create a seabed of unlimited dimensions. An advantage of this approach is that it allows the reflection coefficient to vary; BORIS uses the full spherical reflection coefficient including the effects of shear and absorption. The application of Equation 4-17 in BORIS is slightly modified to include the capability to handle bistatic scattering. Bistatic and forward scattering experiments like those of the SAX experiments [Thorsos et al., 2005] tend to occur at high angles of incidence at which KA theory is no longer accurate. Perturbation theory enters the fray via corrective terms applied to the elementary fields of Equation 4-17 called T-matrices [Thorsos and Borschat, 1994]. T-matrices describe the coupling between incident and scattered fields correcting the faults of KA theory: sub-wavelength scale roughness, scattering at grazing incidence, multiple scattering and shadowing.

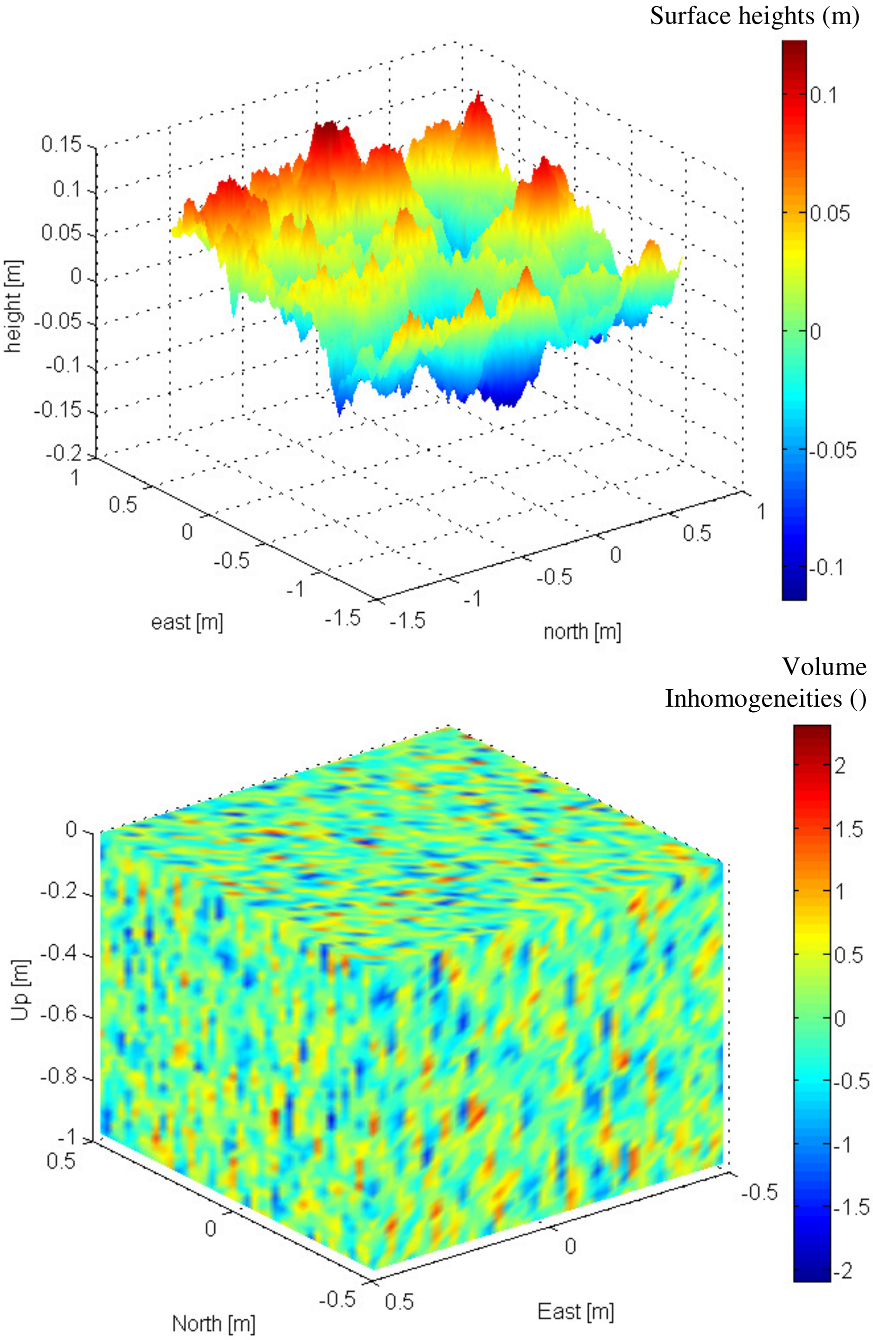

Figure 4-3. Examples of simulated surface heights (*upper*) and volume inhomogeneities (*lower*). This seabed is intended to model PSL burying habitat – coarse sand with high roughness (5 cm RMS) and vertically orientated inhomogeneities (buried PSL or their trails). A normal seabed would have horizontally orientated inhomogeneities (horizontal correlation length 3 to 5 times greater than the vertical). Volume inhomogeneities are labelled in percent difference rather than fractional difference (i.e. a factor of 100 greater than the values of appendix C).

The inclusion of perturbation theory increases the region of applicability of BORIS to grazing angles of incidence[26]. A few references demonstrate that KA theory is only applicable to modest angles (~20°) at SBES frequencies [Jackson et al., 1986b, Lurton, 2002]. The published theoretical background for the original BORIS model shows this as well [Pouliquen et al., 1999]. These limits are the same as those of general KA theory: Equations D-17 and D-19. However, the BORIS model lacks a few of the assumptions of KA theory, such as a constant reflection coefficient and has corrections from perturbation theory (the curvature limits probably no longer apply). As long as the seabed is specified and spatially sampled correctly, the limits of applicability should be much better than those of Equations D-17 and D-19. Even so, the conditions of Equations D-17 and D-19 are not violated by any simulation presented in this dissertation.

4.6 A Comparison of the Echo Time Series Models

There are two major differences between the models presented: the type of acoustic field created and the method of solution (which is analytic or numerical). The advantage of the analytic results of Heald (section 4.4) and section 5.2 is computation speed and the capability to determine analytically the sensitivity of the modelled echoes to the parameters. The acoustic field modelled by Heald and Sternlicht is the total *average* amplitude or intensity [Heald, 2000; Sternlicht, 1999], while the BORIS model creates an amplitude time series for each stochastic realization. The BORIS results will then have the inter-echo variability of real data. Also, BORIS is a full-waveform model, while the Sternlicht and Heald models produce echo intensity envelopes. Full-waveform data allow the calculation of coherent and average incoherent echoes (see appendix E). The BORIS model also outputs separate time series for the volume and surface components. Although the BORIS model has many advantages, it is extremely slow: 2 to 45 minutes to compute a single realization as will be discussed in section 5.1. BORIS

---

[26] This angular limit depends on roughness and frequency. The BORIS manual indicates that the SSA corrections were intended to facilitate simulations at grazing incidence beyond the critical angle for operational SSS and bistatic experiments [Canepa et al., 2004].

with SSA is a composite model, combining KA and perturbation theory, which provides it with the advantage of wider bounds of applicability. Several researchers apply the scattering theories in this way – combine them for a more complete model, [Jackson et al., 1986] is a prominent example. The BORIS model is used throughout the remainder of this dissertation for studying the sensitivity of echoes to seabed parameters, for testing new methods (chapter 6) for depth and slope compensation, for developing new features (chapter 7) and characterization methods (chapter 8). The model to be developed in section 5.3 is also used to the study sensitivity of echoes to seabed parameters and is used as the forward model for Bayesian inversion in chapter 8.

4.7 Implications for Seabed Classification and Characterization Methods

Overall, the scattering theories are consistent with each other and with the data (there are many examples in the literature: [Jackson et al., 1986; Michalopoulou et al., 1994] and any of SAX experiments [Thorsos, 2005]). They all show similar effects on echo shape and frequency content as discussed when the theories were introduced.

A prominent paper, [Jackson et al., 1986], concludes with a statement about the interplay of the backscatter mechanisms. The results contained in [Jackson et al., 1986] pertain to echosounders in 10-100 kHz range, and are paraphrased here. At near normal incidence for fines (silts and clays), volume backscattering is the dominant mechanism. For the remaining sediments (sand, gravels, and those with a minor fine fraction), which is the majority of the sediments observed, surface backscattering dominates, especially near normal incidence. Away from normal incidence, like the outer rays of a wide beam SBES (>30°) or a moderate beam with some seabed slope, volume backscattering is again dominant. These findings are extremely important in designing seabed classification methods.

Sediments that satisfy the Rayleigh criterion generate echoes dominated by the specular reflection. In general, smooth sands and fines meet the criterion. However, fines usually do not have a strong impedance contrast so the specular component is minimal.

Smooth sands, which have moderate impedance, are unique in that they generate SBES echoes that are near replicas of the transmit pulse. The level of coherent backscatter versus incoherent backscatter is also largely determined by the Rayleigh criterion, although as the roughness increases, the number of large facets that generate coherent surface scattering increases as will be demonstrated by BORIS results in the next chapter. Figure 4-4 contains plots of the ratio of the coherent to total amplitudes, showing that the coherent field really only dominates for smooth surfaces of RMS roughness less than 4 mm (for a 50 kHz SBES) or frequencies less than 20 kHz for 1 cm RMS roughness. This is based on the Rayleigh criterion, (Equation D-23), for normal incidence which simplifies to: $8h_{RMS} < \lambda$. The coherent field remains important outside these criteria, as the processes of scattering are ineffective at directing non-specular rays back to the receiver until there is a sufficient level of roughness. In sub-section 5.4.1, BORIS results will show the specular coherent echo is still apparent at 2 cm RMS roughness for a 50 kHz SBES.

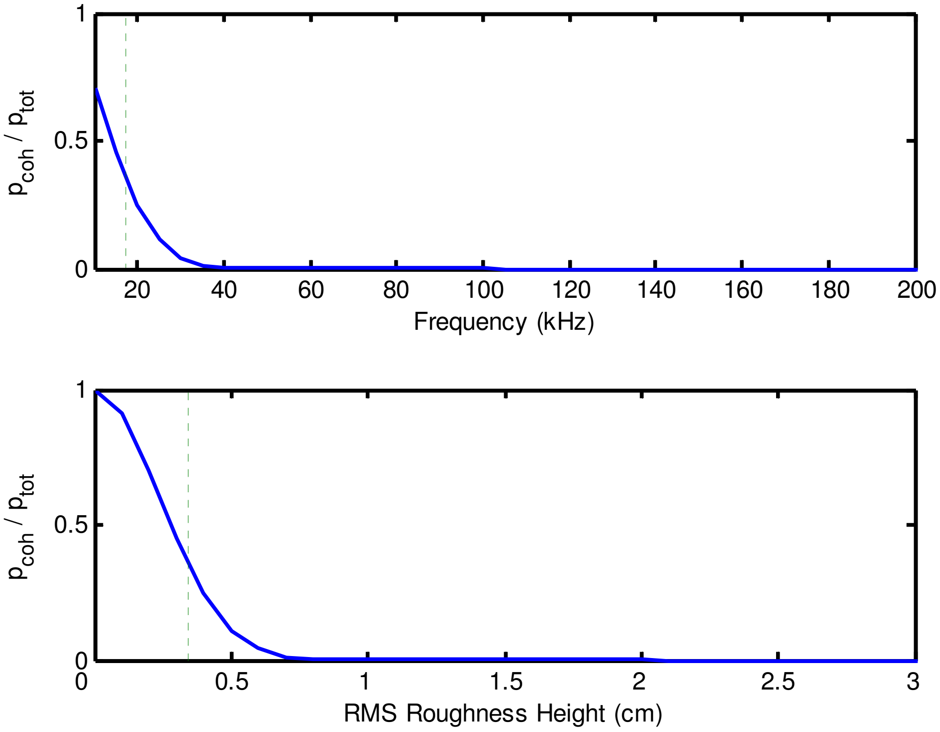

Figure 4-4. The fractional contribution of the coherent field to the total field: (*upper*) RMS roughness set at 1 cm, (*lower*) frequency fixed at 50 kHz. The green dotted lines are the location of the Rayleigh criterion.

The contribution of volume backscattering is often neglected in the literature. For instance, van Walree et al. write that volume scattering becomes important to echo shape when the penetration depth is greater or equal to the pulse duration times the sound speed (the 'spatial' extent of the pulse) [van Walree et al., 2005]. This is not strictly true – consider a low roughness fine sand with a modest reflection coefficient and a high density of volume inhomogeneities – the pulse would not penetrate very far but would offer significant backscatter, especially at off-nadir angles contributing strongly to the tail of the echo. In general, the specular reflection is the initial return, followed by surface backscatter, while volume backscatter dominates the tail of the echo, see Figure 4-5. The shapes and strengths of each echo component are important. The strengths are determined by the reflection coefficient, roughness height, volume attenuation and volume inhomogeneity. The shapes are determined by the shape and duration of the transmit pulse (the specular reflection is a replica of the transmit pulse), water depth, seabed slope and the aforementioned scattering parameters. The effects of depth and slope tend to stretch the echo in time, as will be shown in more detail in chapter 5. The coherent fields do not add echo-to-echo variability, while incoherent fields are completely variable within echoes and from echo-to-echo. The effects of seabed parameters on the echo components are summarized in Table 4-1. This summary is based on the existing theory in this chapter, plus new theory and experiments to be presented in chapters 5 and 6. Developing such a comprehensive view of echosounding has been neglected as a necessary basis for understanding and designing seabed classification methods.

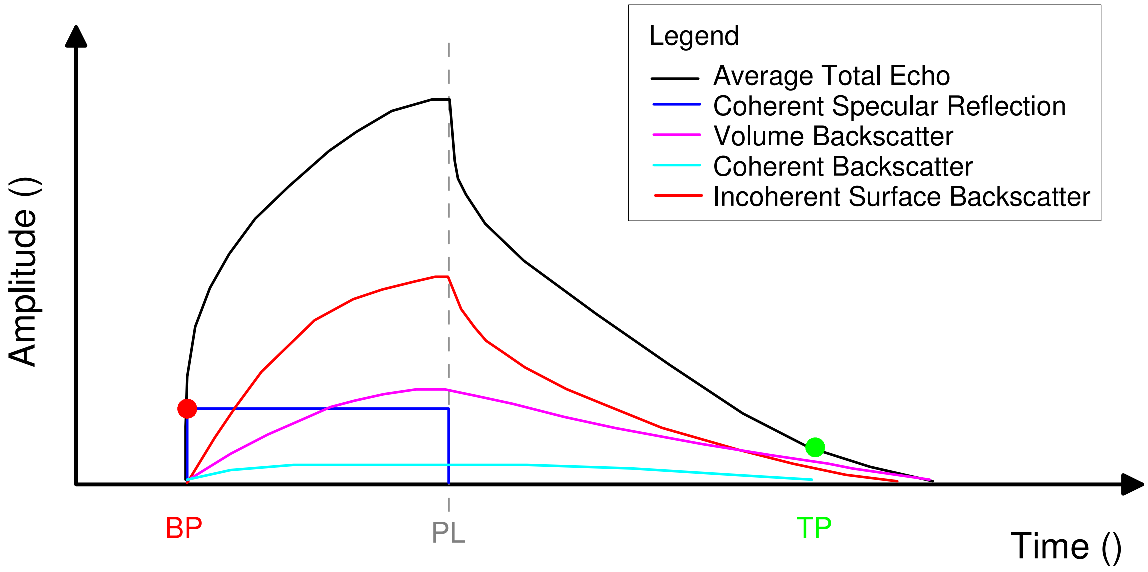

Figure 4-5. A depiction of the four major components of a typical echo (~50 kHz SBES, medium sand seabed) – for demonstration only. BP and TP are the bottom and tail picks, indicated by the large red and green dots. PL is the pulse duration, which is generally the time of peak echo amplitude. This illustration is based on theory and observations from chapters 4 and 5, plus real and simulated echo time series.

Table 4-1. A summary of the effects of geoacoustic parameters on the echo components (acoustic fields). The coherent scattered component from the seabed volume (entries indicated with a *) only exists for the case of low frequency, low relative attenuation and seabed layers that act as sub-bottom reflectors, a scenario which is rare for the SBES used in this dissertation. Also rare are roughness heights large enough to cause shadowing and multiple scattering which would diminish volume scattering.

Echo component ▶	*Coherent:*		*Incoherent:*	
Seabed parameter ▼	Specular	Backscattered	Surface	Volume
Roughness height	decrease to <10% at Rayleigh criterion	increase amplitude and duration for height > Rayleigh criterion	increase amplitude & duration	*
Impedance	scale up amplitudes	scale up amplitudes	scale up amplitudes	scale down amplitudes
Volume attenuation	-	decrease amplitude & duration*	-	decrease duration & amplitude
Volume inhomogeneity height	-	increase amplitude, decrease duration*	-	increase amplitude, decrease duration

Perturbation, Rayleigh and Bragg scattering theories predict the process of selective scattering, indicating preferential scattering in certain directions and frequencies. In KA theory, the effects of selective scattering become more continuous because of a more realistic (and continuous) model of seabed roughness. The Rayleigh criterion is also frequency dependent. The ratio of incoherent to coherent scatter determines the variability of echoes, within the echo and the repeatability of the echoes. The interplay of the mechanisms of scattering also affects the amplitude of the echoes, the shape and the duration.

In all, there are several physical based characteristics of echoes that respond, through the relative strengths of the scattering mechanisms, directly to the geoacoustic properties of the seabed. The characteristics are amplitude, shape, variability, frequency response, the ratio of incoherent to coherent amplitude and duration. These characteristics are measurable from the echo time series data. Current seabed classification methods only take advantage of shape, while there are some indirect measures of duration and amplitude. The approach presented in [van Walree et al., 2005] does improve on this with some spectral measures. However, no seabed classification methods have designed their algorithms based on the physics of scattering; that is until now. Characterization methods usually only exploit one aspect and do not attempt to understand the multi-faceted process as a whole. A clear improvement is to intentionally design echo features that directly measure the aforementioned echo characteristics, and apply a characterization method that combines a number of methods for a more complete picture.

Chapter 5. New Theory and Models

This chapter builds upon the theoretical foundation established in the previous chapter by developing two new models and calibrating the BORIS model. All of these models are used later in chapters 6, 7 and 8 to remove non-seabed influences, to form the basis of improved discrimination with new features, to test new processing methods and to enable characterization. Although the BORIS model has existed for a number of years, it is unclear, based on the documentation available, how to generate realistic simulations of SBES, because either this information has not been published or values for model parameters have not been established. Regardless, the details on how to generate realistic simulations SBES and SBES in survey conditions are given here, as these details are relevant to evaluate the validity of any BORIS results. The same level of detail is provided for the other two models presented. This chapter also aims to demonstrate the scattering processes discussed in chapter 4, and to discuss implications of the new models to seabed classification and characterization.

5.1 The Setup and Calibration of BORIS for Modelling SBES

The setup and calibration of the BORIS model is a necessary step in order generate realistic simulations. There is much necessary detail pertaining to this subject that is potentially useful to other researchers attempting to utilize the BORIS model (particularly for simulating SBES surveys). However, such detail is not strictly necessary for this chapter, so these details are left to appendix F. Some detail is necessary to understand the results to be presented in this section. Those details include an overview of the input parameters, the procedure of running simulated surveys and controlled experiments

There are three major groups of input parameters and data used by the BORIS software. First, the technical parameters of the echosounder to be simulated are established as discussed in appendix F and tabulated in Table C-1. The geoacoustic parameters describing the seabed are also discussed in appendix F and tabulated in Table

C-3. Several sets of these geoacoustic parameters were established, based on values from the literature, representing the range of seabed types. Variations from these seabeds are noted explicitly; otherwise, the seabed used is named in reference to the seabed types listed in Table C-3. The third group of parameters are the remaining parameters, which consist of administrative parameters, parameters that specify various model types or parameters that specify the location and orientation of the echosounder relative to the simulated seabed.

5.1.2 Examples of BORIS Model Simulations and Simulated Surveys

Figure 5-1 is an example of a BORIS echo time series. BORIS produces the surface and volume echoes separately, which are added together to create the total full-waveform echo from which the echo envelope can be extracted using the Hilbert transform. This example is interesting in that the volume echo clearly dominates. This is a benefit of BORIS: to be able to analyze the volume and surface echoes separately.

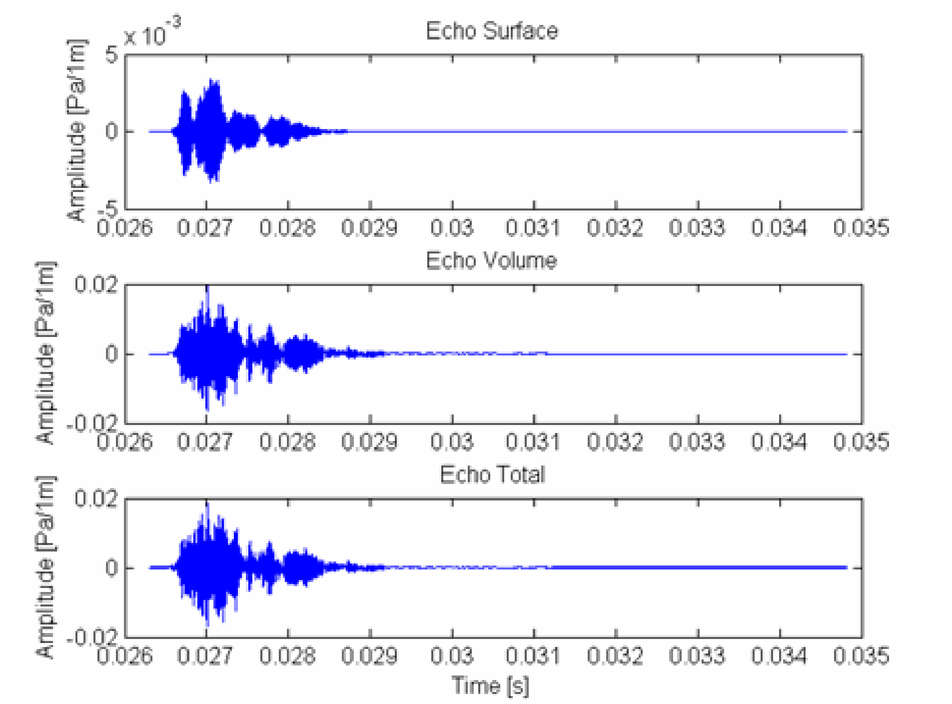

Figure 5-1. Example BORIS echo time series from a simulation of the Circ50_20 virtual SBES incident on the SiltyClay seabed at 20 m water depth, (see appendix C for echosounder and seabed parameters).

BORIS is an excellent tool for experimenting with the effects of various parameters, such as water depth, roughness, seabed slope, etc., as one can hold all other parameters constant while modifying just one parameter. Figure 5-2 shows an example of such a controlled experiment where all parameters are known and fixed, except for the one parameter that is the subject of the experiment. It is difficult to carry out a controlled experiment like this at sea, as the parameters are not well known nor do they remain fixed. The lack of seabed geoacoustic parameter information is a huge problem for field research, while models do use exact values for simulations; these values have to be established carefully, as was done in Appendices C and H.

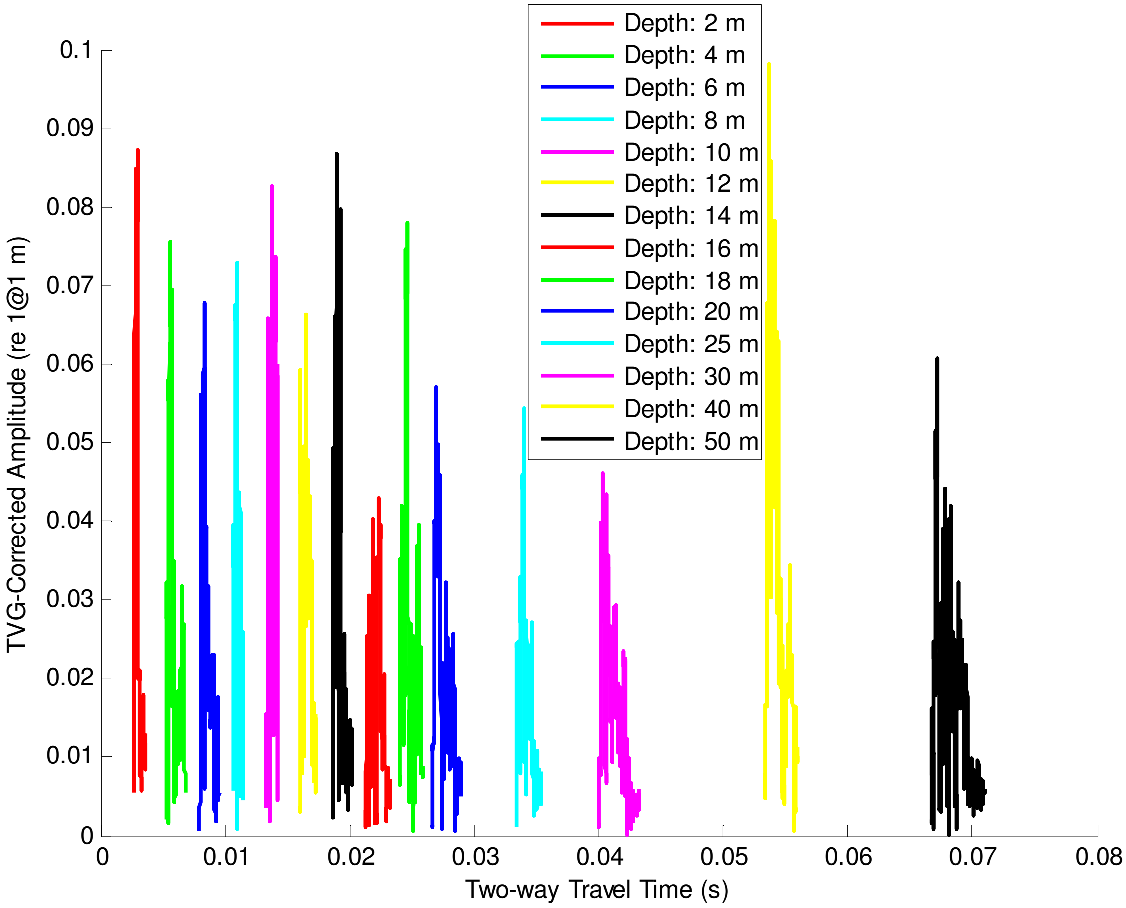

Figure 5-2. BORIS simulated echo amplitude time series (from bottom pick to tail pick) as an example of a controlled experiment for varying water depths. Depths range from 2 to 50 m for the Sand seabed (see Table C-3) with the Airmar38 echosounder (see Table C-1). The exact same seabed realization and positioning is used in every echo. The variability observed accurately depicts the random nature of echoes due to incoherent scattering as the amplitudes are corrected for spreading losses. One can also observe the lengthening effect of depth on the echoes.

To demonstrate the capabilities of BORIS for experimentation, a simulated survey is shown in the remainder of this section. A survey is simulated by moving the transducer along the simulated seabed according to normal survey parameters. For example, surveys are typically run at three to five knots with a ping or repetition rate of 3 Hz, so the transducer may move 0.7 m between simulations. Because of tiling (see section 4.5), the seabed available to BORIS is infinite, as long as the spacing between realizations is not a multiple of the tile size (in that case duplicate echoes have been observed). In comparisons with real data, the variability of a simulated survey is qualitatively the same (further testing with the variability statistics developed in chapter 7 is an avenue for future work).

One of the simplest ways of viewing the results of a set of simulations is to view them as an echogram. As discussed in chapter 1, an echogram is an image of successive pings. The echogram for an example simulated survey is shown in Figure 5-3. A set of simulated surveys like this one was created for the seabed classification repeatability research of Gleason, [Gleason, 2009]. This survey is used to demonstrate the effects of slope and depth on several existing and new echo features in sub-section 5.4.4.

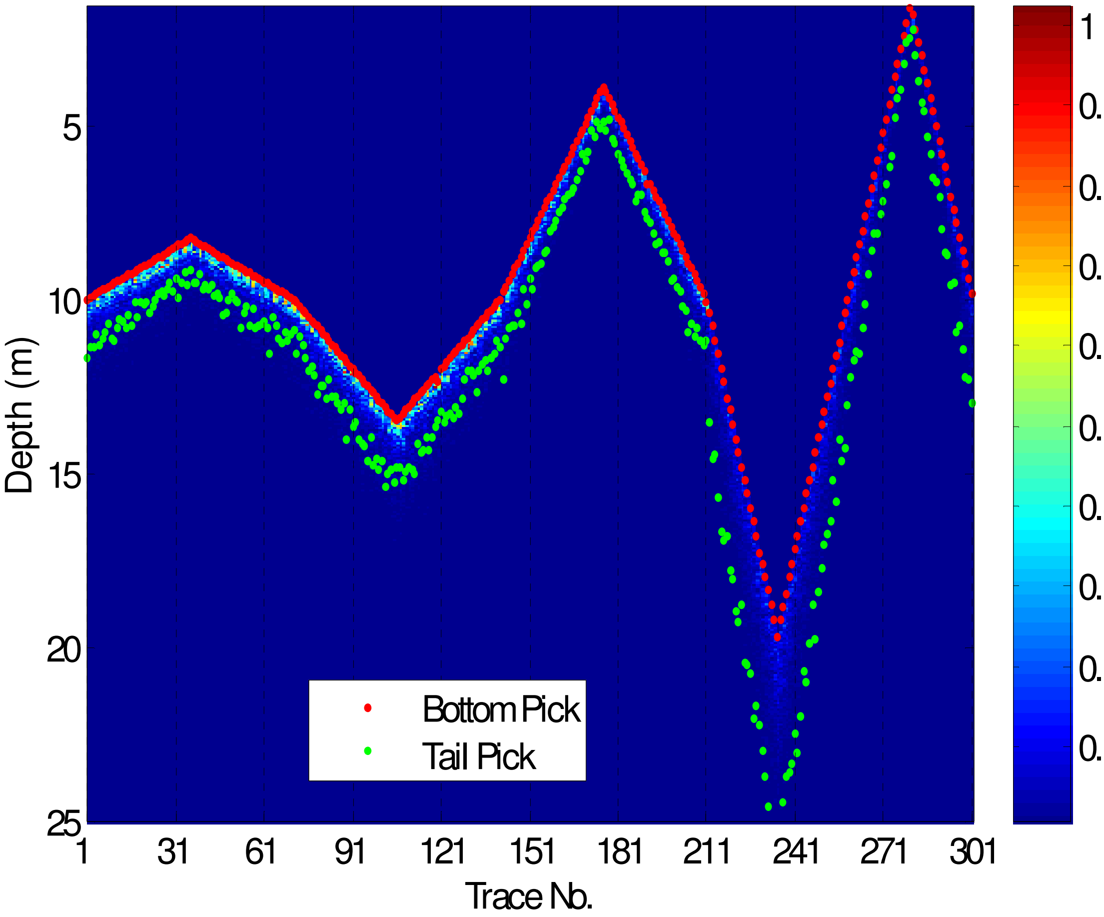

Figure 5-3. An echogram of a simulated survey of the Suzuki50 SBES on the SandAG seabed. The colourbar indicates amplitude (corrected for spreading loss). This survey occurs over pairs of slopes of increasing angles: 3, 6, 10 and 22°. The locations where the slope increases are trace numbers 71, 141 and 211 (the 22 degree section gets an extra 20 echoes). Vessel speed: 3 m/s, ping rate: 2 Hz. The 'vessel' heading, pitch and roll are Gaussian random with a standard deviation of 1, 2 and 1° respectively (to simulate real survey conditions).

5.2 The Analytic Echo Model

The analytic echo model (AEM) for echo time series is an extension of Heald's model [Heald, 2000] described in section 4.4. Similar to other temporal models of seabed echoes, Heald's model integrates the backscatter from a moving annulus of ensonified seabed as depicted in Figure 4-2. Solving the time-dependent integral of backscatter over the seabed (Equation 4-8) is made tractable by applying a Gaussian beampattern

(Equation 4-11), a backscattering coefficient from [Brekhovskikh and Lysanov, 1988] (Equation 4-12), and expressions for the leading and trailing edges of the ensonified annulus. In order to modify any of these terms to improve and extend the model requires that the new terms contribute to a solution of Equation 4-8, which is not often possible. One of the major deficiencies of Heald's theory is that the Brekhovskikh and Lysanov backscattering coefficient, and hence any model derived from it, is frequency independent. Attempts to substitute frequency-dependent backscattering coefficients (from [APL-UW, 1994] and [Medwin and Clay, 1998]) resulted in an integral is not solvable analytically. Other attempts to generalize the problem to three dimensions to handle seabed slope also did not result in solutions. However, this author did find a new solution for modelling a complete echo. It is based on new expressions for the leading and trailing edges of the ensonified area so that the solution represents the entire first echo. The new expressions are:

$$\begin{aligned} \theta_{lead} &= \operatorname{arcsec}\left(\frac{ct}{2D}+1\right) \\ \theta_{trail} &= \operatorname{arcsec}\left(\frac{c(t-\tau)}{2D}+1\right) \text{ for } t \geq \tau, \text{otherwise} \\ \theta_{trail} &= 0 \end{aligned} \tag{5-1}$$

where D is the true water depth (as will be discussed further in the next section), t is defined as zero at the time of onset, so the time since transmission is $t + {}^{2D}\!/_{c}$. When compared to the original expressions for θ_{lead} and θ_{trail} (Equation 4-14), the two sets of expressions are numerically equivalent except for $t \leq \tau$, when the trailing edge is stationary at the origin of the ensonified area. The substitution of Equation 5-1 into the solution of integral of Equation 4-8 leads to the analytic echo model:

$$\begin{aligned} I_{AEM}(t) &= \frac{M}{2\beta^2}\left(1-\exp\left(-\beta^2 \operatorname{arcsec}^2\left(\frac{ct}{2D}+1\right)\right)\right) \text{ for } 0 \leq t \leq \tau, \\ I_{AEM}(t) &= \frac{M}{2\beta^2}\left(\exp\left(-\beta^2 \operatorname{arcsec}^2\left(\frac{c(t\text{-}\tau)}{2D}+1\right)\right)-\exp\left(-\beta^2 \operatorname{arcsec}^2\left(\frac{ct}{2D}+1\right)\right)\right) \text{ for } t \geq \tau \end{aligned} \tag{5-2}$$

$$\text{where: } M = \frac{I_0 \Re_\perp^2 \lambda_0^{\,2}}{16 h_{RMS}^{\,2} D^2}, \beta^2 = \frac{\lambda_0^{\,2}}{4 h_{RMS}^{\,2}} + \frac{2}{\tan^2 \theta_{-4.34dB}}$$

The correlation length, λ_0, and roughness height, h_{RMS}, always appear as a ratio; hence the *roughness ratio*, $h_{Ratio} = {}^{\lambda_0}\!/_{h_{RMS}}$, can replace both parameters. The initial intensity, I_0, and the reflection coefficient, $\Re_{\perp}^2$, are constants that do not affect the shape of the theoretical echoes; they can be replaced with a single constant for scaling purposes (usually taken as unity). In all, there are five geoacoustic/geotechnical parameters in the AEM model: water depth, pulse duration, beamwidth, roughness ratio and the scaling factor.

A major simplifying approximation made in the derivation of the AEM is that of using a Gaussian directivity function in place of a measured or theoretical beampattern. This approximation is evaluated in Figure 5-4. For most of the main-lobe of the beam, the approximation is quite sufficient. This region of the beam represents the vast majority of beam sensitivity. Then, the AEM represents the behaviour of real echoes on terms of angular response; echo shapes and measures of duration produced by the AEM are sufficiently accurate.

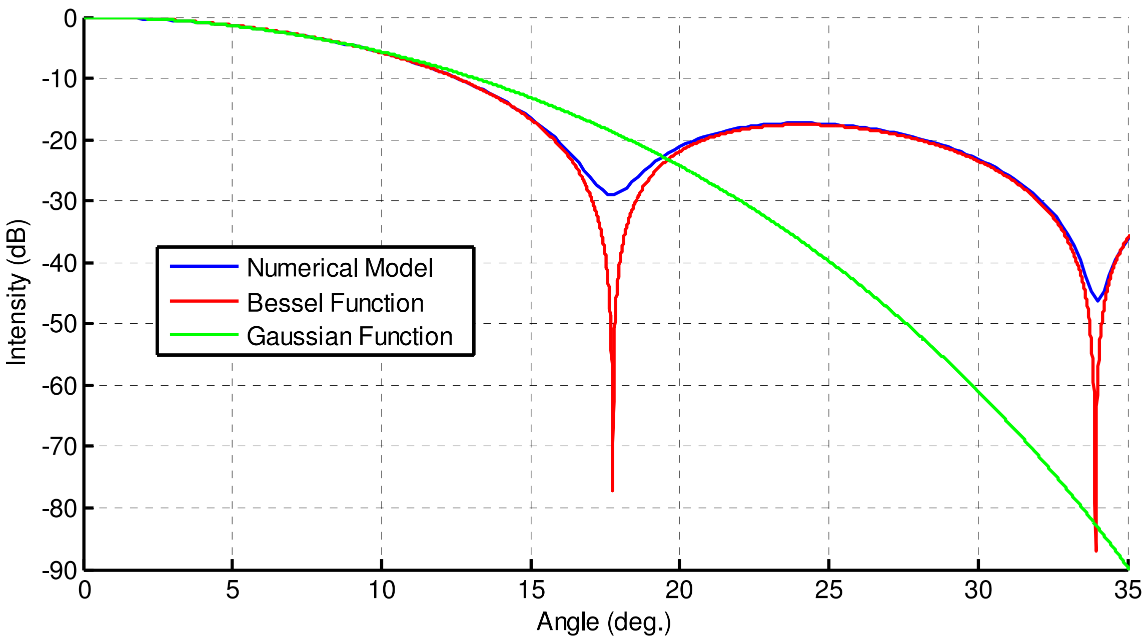

Figure 5-4. A comparison of beampattern models. The numerical model (*blue*) is the Deep50 echosounder beampattern used in BORIS simulations, (see Table C-1) which based on a numerical approximation of a Bessel function. The analytic Bessel function plotted in *red* is from [Lurton, 2002] based on a 12 cm diameter circular transducer. The Gaussian directivity function (*green*) is based on the -4.34 dB beamwidth of the Deep50 echosounder.

It is evident from Equation 5-2 that the AEM lacks direct frequency dependence. Also, echo time series produced by the AEM are time series of the mean echo intensity, not the individual echo amplitude as would be measured by real echosounders and as modelled by BORIS. Thus, the model result is equivalent to the sample-wise square of the stacked echo envelope. Note that the stacked echo envelope is also correctly referred to as the mean echo or the total amplitude, and is not to be confused with the coherent echo or the mean incoherent echo, terms that will be defined in section 5.3.

The AEM only considers surface backscatter. The volume backscatter coefficient, as devised in [Jackson et al., 1986], is proportional to seabed attenuation[27] and independent of incident angle (isotropic). To add elementary volume scattering to the AEM, this author considered adding the term $(1-\Re_{\perp})^2 \alpha_{volume}{}^2 f_c^2 \kappa^2$ to the scattering coefficient $S(\theta_t)$ of Equation 4-12, where α_{volume} is the proportionality between the attenuation and the volume scattering cross-section. Adding this expression to the surface scattering coefficient (Equation 4-12) would still leave Equation 4-8 solvable. However, this simplified volume scattering model would suffer from not considering the local transmission and travel time delay within the seabed. The volume scattering component of the modelled echoes would also have only a rudimentary shape related to the beampattern and ensonification area as it has no dependence on angle; the volume scattering coefficient would only control the amplitude of the volume response and its shape would be the same as the interface echo.

Recently, the Acoustic Remote Sensing Group at the University of Delft published two papers [Simons et al., 2009; Siemes et al., 2010] presenting a simpler version of the Sternlicht model [Sternlicht, 1999]. Both Sternlicht and Simons solve the propagation model numerically, which can be computationally intensive. One of the major differences is that Simons's model utilizes a simple approach to volume scattering similar to that discussed in the previous paragraph. This approach for volume scattering does not consider the local transmission coefficient, in addition to the ignoring the effects of refraction angle, time delay, etc. - effects that were rigorously accounted for by Sternlicht [Sternlicht, 1999]. In [Simons et al., 2009], the forward model is optimized to

---

[27] The volume scatter cross-section usually taken to be $0.04c\kappa$ [Sternlicht, 1999].

SBES data to estimate mean grain size and other properties of the seabed. The optimization is done by a genetic algorithm.

A benefit of using the AEM over numerical models is that the speed of computation facilitates running more advanced, computationally intensive inversion techniques (i.e. Metropolis sampling). These techniques provide much more information than optimization, specifically error estimates. Inversions of SBES data with the AEM will be investigated in chapter 8. For future work, complete accurate models of echo time series such as Sternlicht's model or the BORIS model could be paired with advanced inversion techniques.

5.2.1 Implications and Qualitative Verification of the Analytic Echo Model

One of the goals of developing the AEM is to improve the understanding of the physics of echosounding by working with analytic expressions for echo time series. This information can also be viewed graphically. The following figures show AEM echo time series for variations of the four model parameters that affect the echo shape (Figure 5-5, Figure 5-6, Figure 5-7, Figure 5-8). Each model parameter is varied while holding the others at the default scenario which is:

- 10 m water depth;
- $20°$ beamwidth (re -4.34 dB);
- 0.3 ms pulse duration;
- a roughness ratio of 10 ($h_{ratio} = \lambda_0 / h_{RMS}$);
- and a scaling factor of 0.20 ($I_0 \Re_\perp^2$).

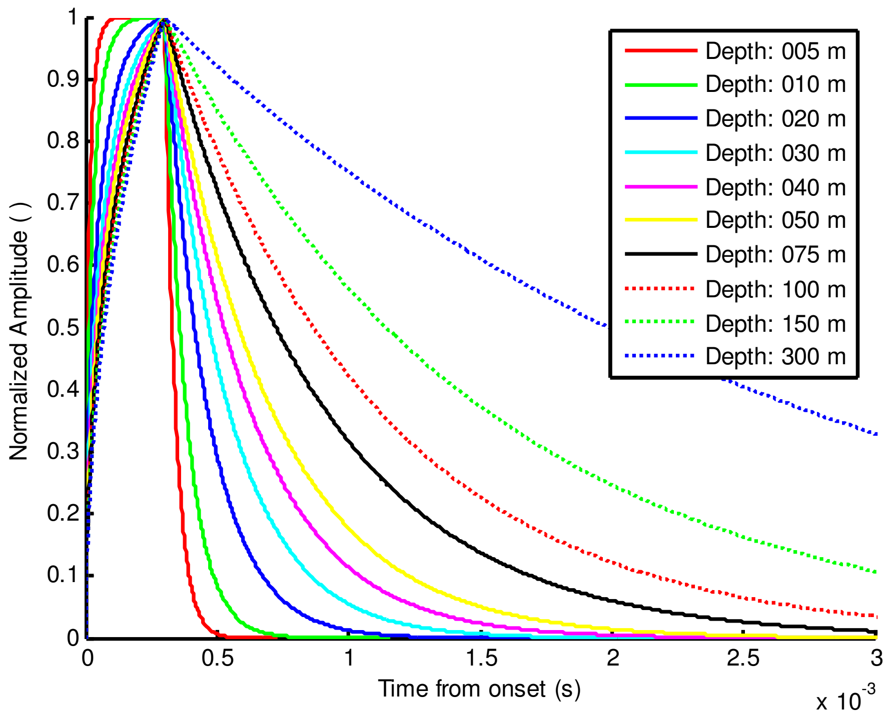

Figure 5-5. AEM simulated echo times series for a number of depths.

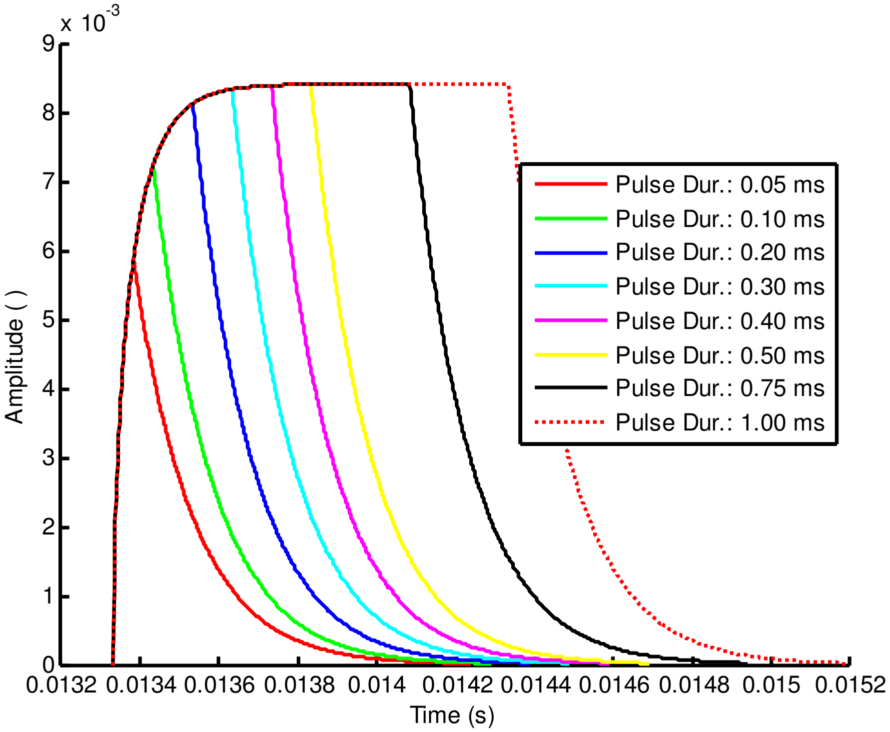

Figure 5-6. AEM simulated time series for a number of pulse durations. Note the amplitudes need not be normalized here, as pulse duration does not affect amplitude.

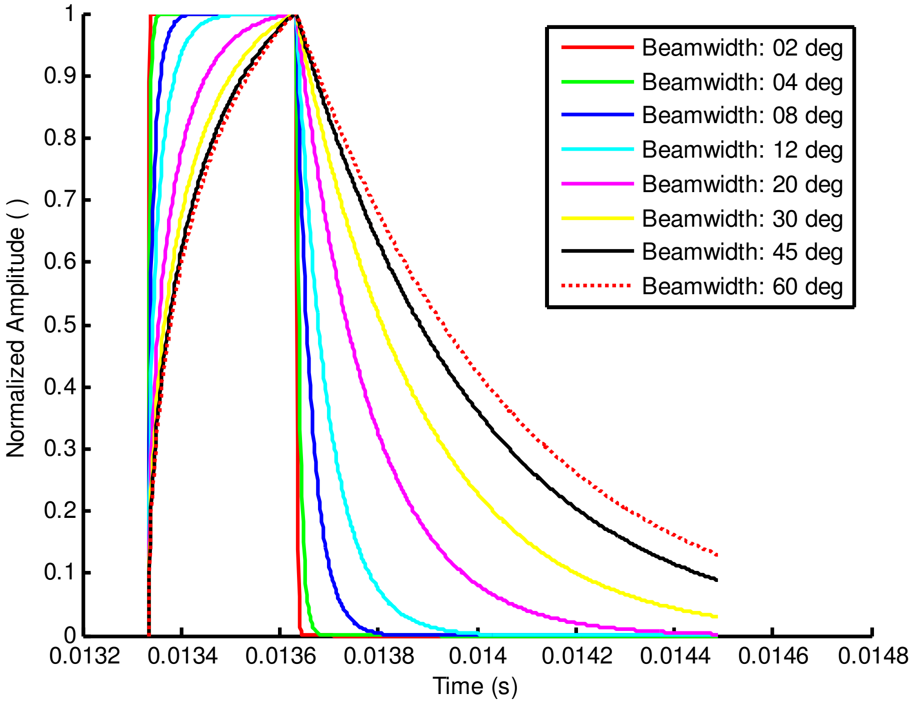

Figure 5-7. AEM simulated echo time series for a number beamwidths.

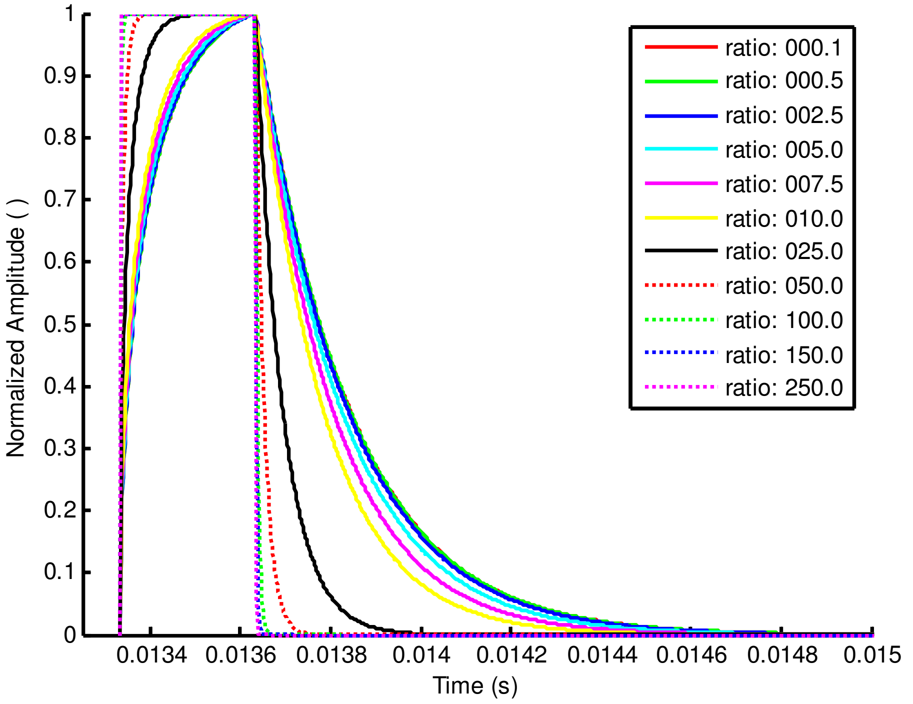

Figure 5-8. AEM simulated echo time series for a number of roughness ratios.

The echoes in the above figures tend to vary in the same manner, progressing from a square pulse (a reflection of the transmit pulse) to a peaked echo. The exception is the case of varying pulse duration (Figure 5-6). Increasing depth does lengthen the echo tail, but does not move the peak amplitude location nor change the shape of the echo before the peak appreciably for depths greater than 75 m (Figure 5-5). This effect is explored further in section 5.4 following the introduction of the echo duration model. Increasing beamwidth and pulse duration both increase the echo duration (Figure 5-6 and Figure 5-7). The roughness ratio does not affect the echo for values less than 5 and greater than 100 (Figure 5-8). The lower limit can be derived mathematically; a value of four was found. This lower limit is consistent with the limits on KA theory (see chapter 4), and observations from a numerical study [Broschat and Thorsos, 1997]. The upper limit of roughness ratio is when the echo becomes purely specular with no roughness backscattering. This limit is closest to smooth compact fine sand, while the lower limit is effectively the high roughness limit where the seabed could consist of cobbles or small boulders. At this high level of roughness or low level of roughness ratio, the echo is of maximum duration, and that maximum is reached because backscattering occurs throughout the acoustic footprint – the footprint is effectively saturated. This is the beam-widening effect that will be discussed further later in this section.

The substitution of roughness height, h_{RMS}, and correlation length, λ_0, by their ratio, $h_{ratio} = \lambda_0 / h_{RMS}$, is further examined in Figure 5-9. In this case, echo duration is used as an indicator of echo behaviour as echoes from the AEM change in shape in a monotonic way. Figure 5-9 shows that echo duration is constant with constant ratios of correlation length to roughness height. This makes sense because, as stated in section 4.4, the scattering coefficient used in the AEM is a function of the RMS micro-slope of the seabed, which is inversely proportional to the roughness ratio, see Equation 4-13. The relationship of the roughness rate to RMS roughness height and seabed type is unknown as measured values of correlation length and RMS micro-slope are extremely rare even though both are important parameters for scattering theory. Correlation length is a difficult parameter to measure because its value is a strong function of the scale on which it is measured. Correlation lengths are measured from BORIS simulations later in section

8.2. For now, roughness ratio has a negative relationship with grain size such that small values are characteristic of gravels, while large values are characteristic of clays.

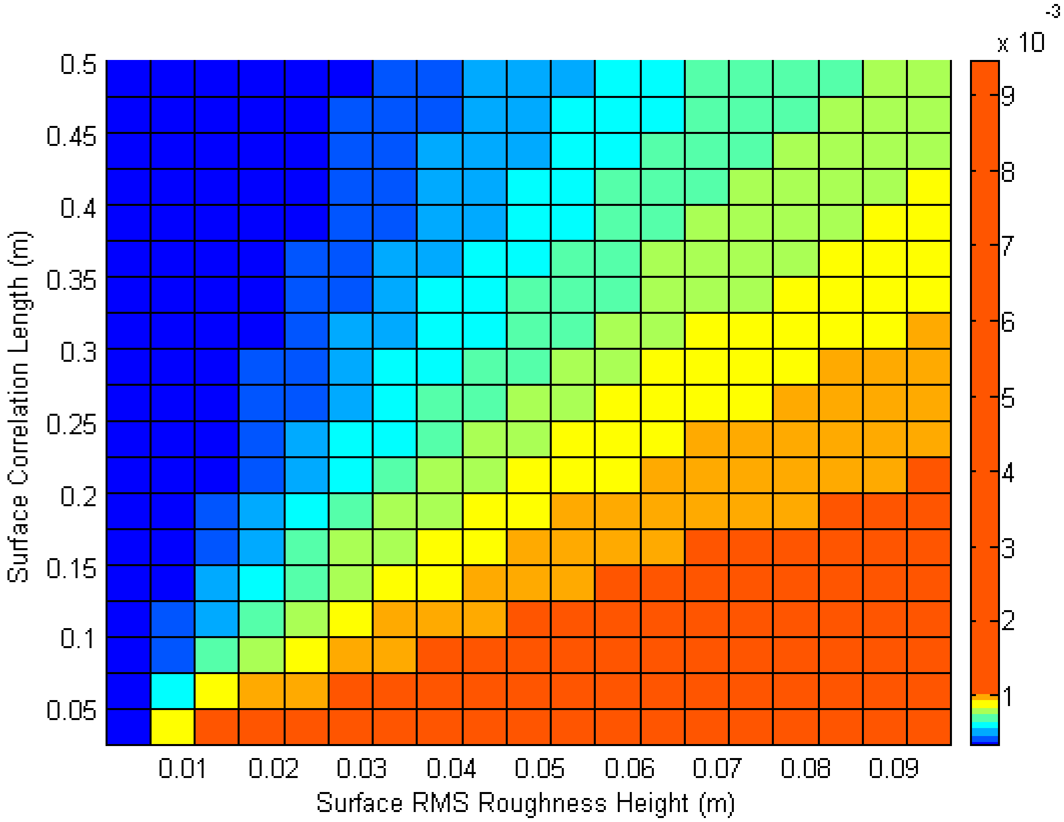

Figure 5-9. Echo duration from AEM echoes for varying values of correlation length and RMS roughness height (other parameters held at the base case values). The colourbar indicates the echo duration (as measured by the 99% cumulative amplitude) in units of seconds (colours are scaled for maximum sensitivity, accounting for extreme values).

Figure 5-10 is another example of the AEM providing useful information on the physics of echosounding pertaining to seabed classification. It shows how SBES with wider beams tend to produce more energetic echoes at larger roughness compared to SBES with narrower beams. When energy is used as a feature for classification, SBES with wider beams would discriminate seabeds over a greater range of roughness.

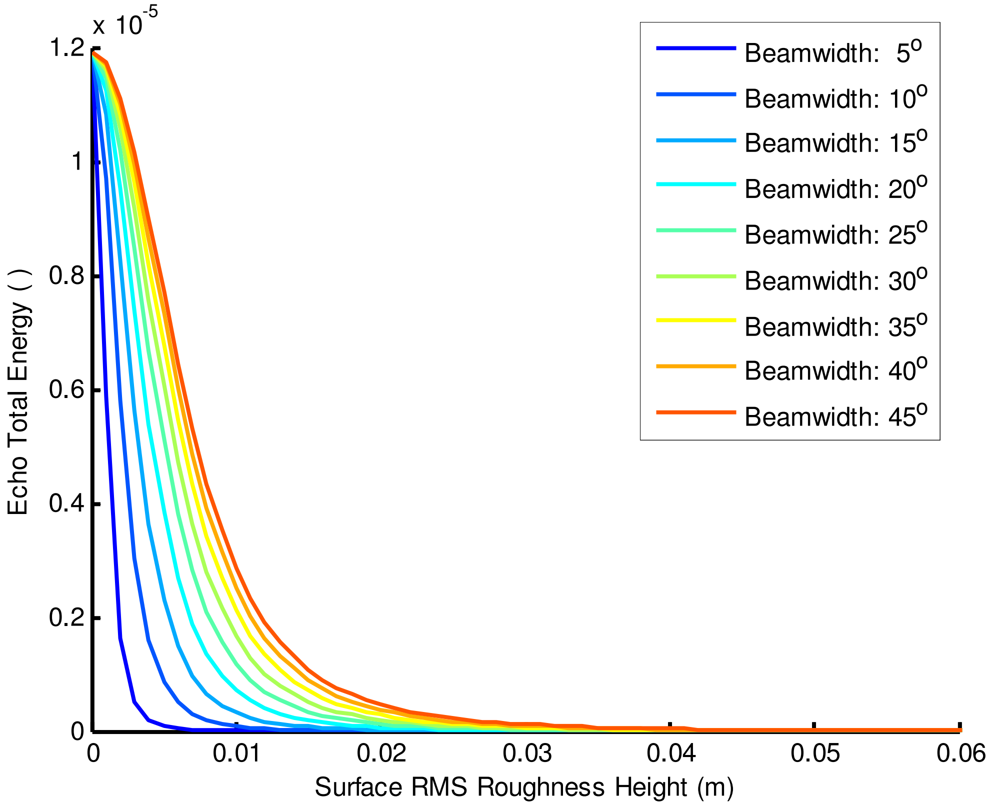

Figure 5-10. Total echo energy measured from AEM simulations for varying roughness height and beamwidth. As was the case for other AEM results in this section, amplitude and energy do not have units, as the initial intensity, I_o, is unity.

5.3 The Echo Duration Model

As presented in [Biffard et al., 2005], the echo duration model (EDM) is based on the geometry of a conical beampattern. Similar geometries have been used to estimate footprint size, define TVG and derive SEL compensation (Equation 2-6, [Preston, 2004]) as discussed in chapter 2.

There are two cases based on whether the specular reflection ray is inside the beam (case I) or is outside the beam (case II). In Figure 5-11, the upslope and downslope rays at the edges of the beam are marked by s+ and s-, respectively. So, when the seabed normal is outside the beam defined by s+ and s-, the specular ray is not possible and the situation is case II. Case I includes the normal incidence situation and zero slope situation. In the extreme case of the slope / angle of incidence being so great that s- never intercepts the seabed, the echo duration goes to infinity; this occurs for $\varphi \geq \pi/2 - \theta/2$. In all cases, echo duration is the sum of contributions from the transmit pulse duration, travel time due to wavefront curvature (taken along the s- ray), penetration depth l_p and macro-roughness height l_H. In Figure 5-11, the travel times of the curvature, penetration and macro-roughness are depicted by C_v, t_p and t_H. Macro-roughness is defined as large objects such as macroalgae, coral-heads and large rocks that cause the bottom pick to occur earlier than a bottom pick triggered by the mean water depth. The curvature term is the time from first contact of the wave front at the seabed to the time of contact along the s- ray. The point of first contact is either along the seabed normal ray (case I) or the s+ ray (case II). Curvature is generally the dominant contributor to echo duration except for the condition that pulse duration dominates for water depths less than ~20 m. Putting all terms together leads to Equation 5-3 for echo duration, with many of the variables as defined in chapter 2:

$$\begin{aligned} t_{ED} &= \frac{2d}{c_w}\left(\sec\left(\phi+\frac{\theta}{2}\right)-1\right)+\frac{2l_p}{c_s}+\frac{2l_H}{c_w}+\tau && \forall \quad 0\leq\phi\leq\frac{\theta}{2} && \text{(case I)} \\ t_{ED} &= \frac{2d}{c_w}\left(\sec\left(\phi+\frac{\theta}{2}\right)-\sec\left(\phi-\frac{\theta}{2}\right)\right)+\frac{2l_p}{c_s}+\frac{2l_H}{c_w}+\tau && \forall \quad \frac{\theta}{2}\leq\phi<\frac{\pi}{2}-\frac{\theta}{2} && \text{(case II)} \end{aligned} \tag{5-3}$$

A few important distinctions must be made here: *d* is the water depth as measured by the echo onset, c_w is the water sound speed and c_s is the sediment sound speed. The true water depth directly below the echosounder is *D*. This quantity can be calculated from the measured depth, *d*:

$$\begin{aligned} D &= d\sec(\phi) && \forall \quad 0\leq\phi\leq\frac{\theta}{2} && \text{(case I)} \\ D &= d\sec(\phi)\cos\left(\phi-\frac{\theta}{2}\right) && \forall \quad \frac{\theta}{2}\leq\phi<\frac{\pi}{2}-\frac{\theta}{2} && \text{(case II)} \end{aligned} \tag{5-4}$$

Equation 5-3 does account for the fact that in case II, the measured water depth will be dependent on the upslope ray and not the seabed normal (as was the situation in case I).

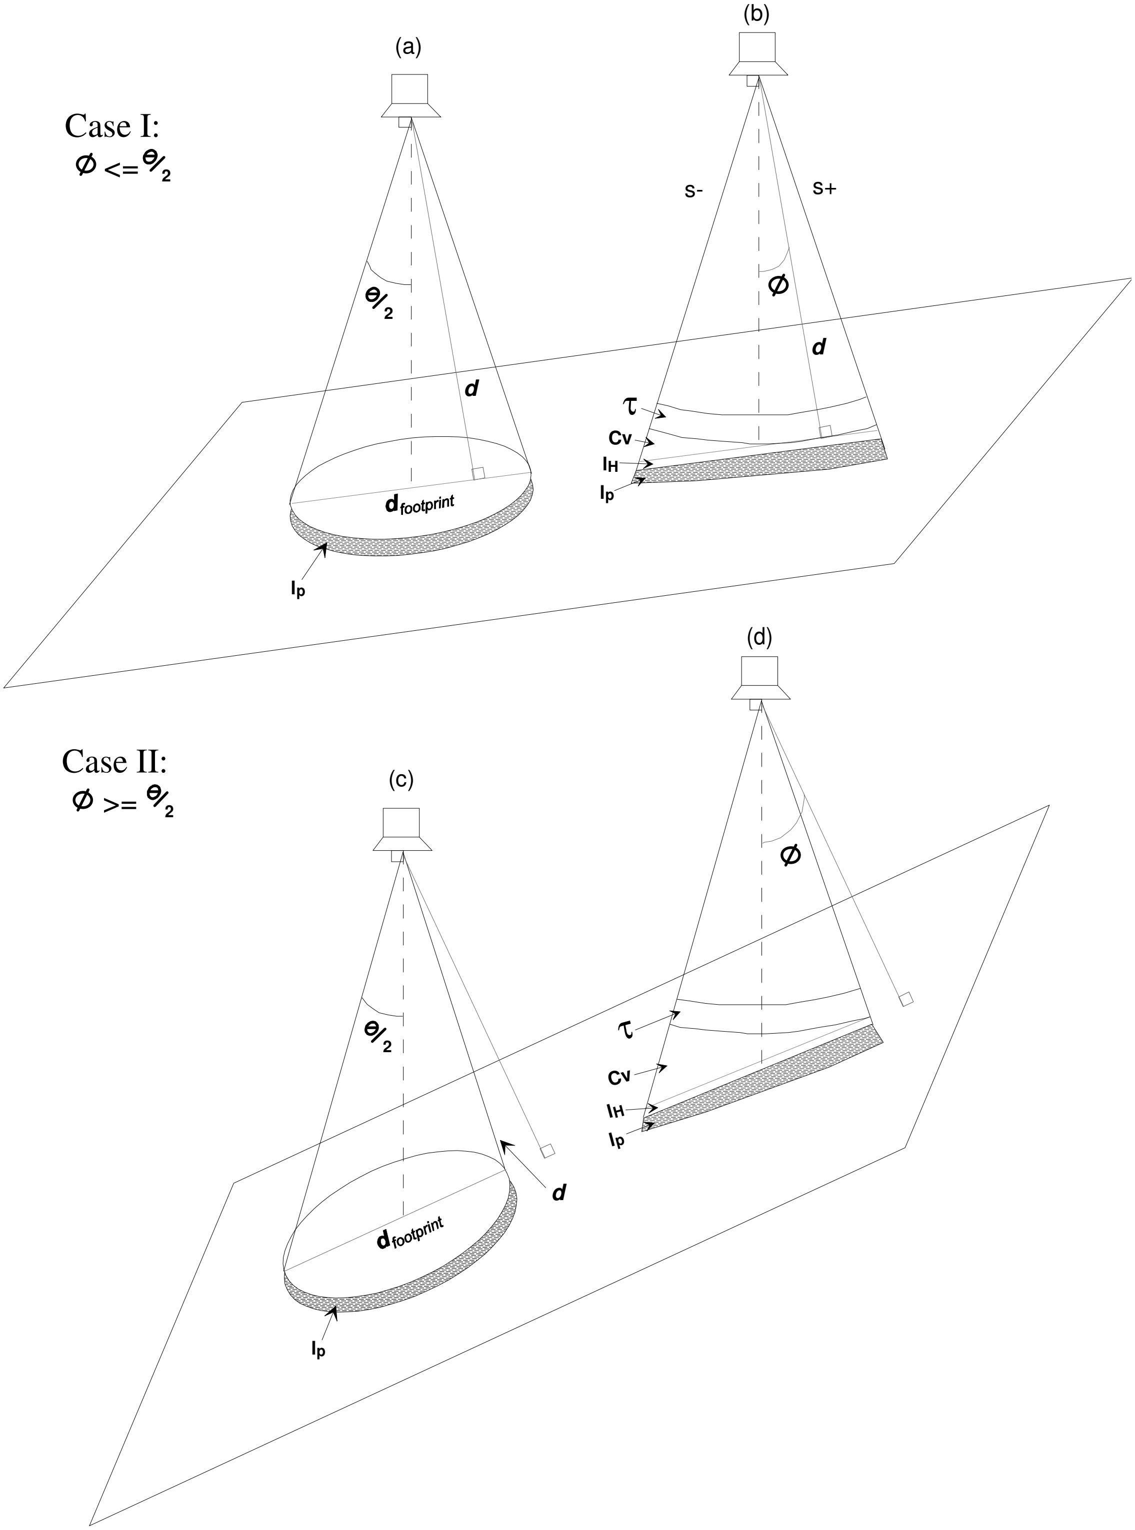

Figure 5-11. Ray-trace diagrams for case I and II: (*a*) and (*c*) illustrate the footprint geometry, while (*b*) and (*d*) show the ray tracing and echo duration terms used to derive Equation 5-3. s- and s+ are the downslope and upslope rays marking the edges of the beam. *Cv* denotes the distance along the s- ray due to wavefront curvature.

One effect evident from Equation 5-3 is that SBES bathymetry in areas of significant seabed slope will be underestimated. The footprint size, $d_{\text{footprint}}$, is also a strong function of seabed slope. As shown in Figure 5-11, the footprint becomes an increasingly eccentric ellipse with increasing slope. The minor semi-axis length remains the same as the zero slope footprint diameter (Equation 2-2), while the major semi-axis length is:

$$
\begin{aligned}
d_{\text{footprint}} &= d\left(\frac{\sin(\theta)}{\cos\left(\frac{\theta}{2}+\phi\right)\cos\left(\frac{\theta}{2}-\phi\right)}\right) \quad \forall \quad 0 \le \phi \le \frac{\theta}{2} \qquad \text{(case I)} \\
d_{\text{footprint}} &= d\left(\frac{\sin(\theta)}{\cos\left(\frac{\theta}{2}+\phi\right)}\right) \quad \forall \quad \frac{\theta}{2} \le \phi < \frac{\pi}{2}-\frac{\theta}{2} \qquad \text{(case II)}
\end{aligned}
\tag{5-5}
$$

Figure 5-12 shows the predictions of echo duration and footprint major semi-axis length from the EDM for several typical echosounder configurations. The EDM predicts a linear relationship between echo duration and depth, a strong effect of slope on echo duration, and moderate lengthening of the footprint major semi-axis due to the slope. The lengthening of the footprint major semi-axis has a relatively minor role in increasing echo duration compared to the displacement of the footprint caused by its tilt, i.e. differences in the travel time to and from the up-tilted and down-tilted sides of the footprint.

The echo duration model makes several approximations. First, seabed slope is measured over distances longer than the acoustic footprint. Smaller-scale variation is defined as seabed roughness, for which a mean seabed level is defined or assumed. The extra duration caused by macro or unusual roughness significantly above the mean seabed level is handled by the term l_H in Equation 5-3. The size of roughness above which l_H becomes non-zero is generally on the order of tens of centimetres, as will be determined empirically from BORIS simulations in the next section. Penetration depth, l_p, is modelled in one of two ways. One may use the extinction relation of Equation 2-7, which is based on an extinction of 10 dB due to the attenuation coefficient. Unlike the

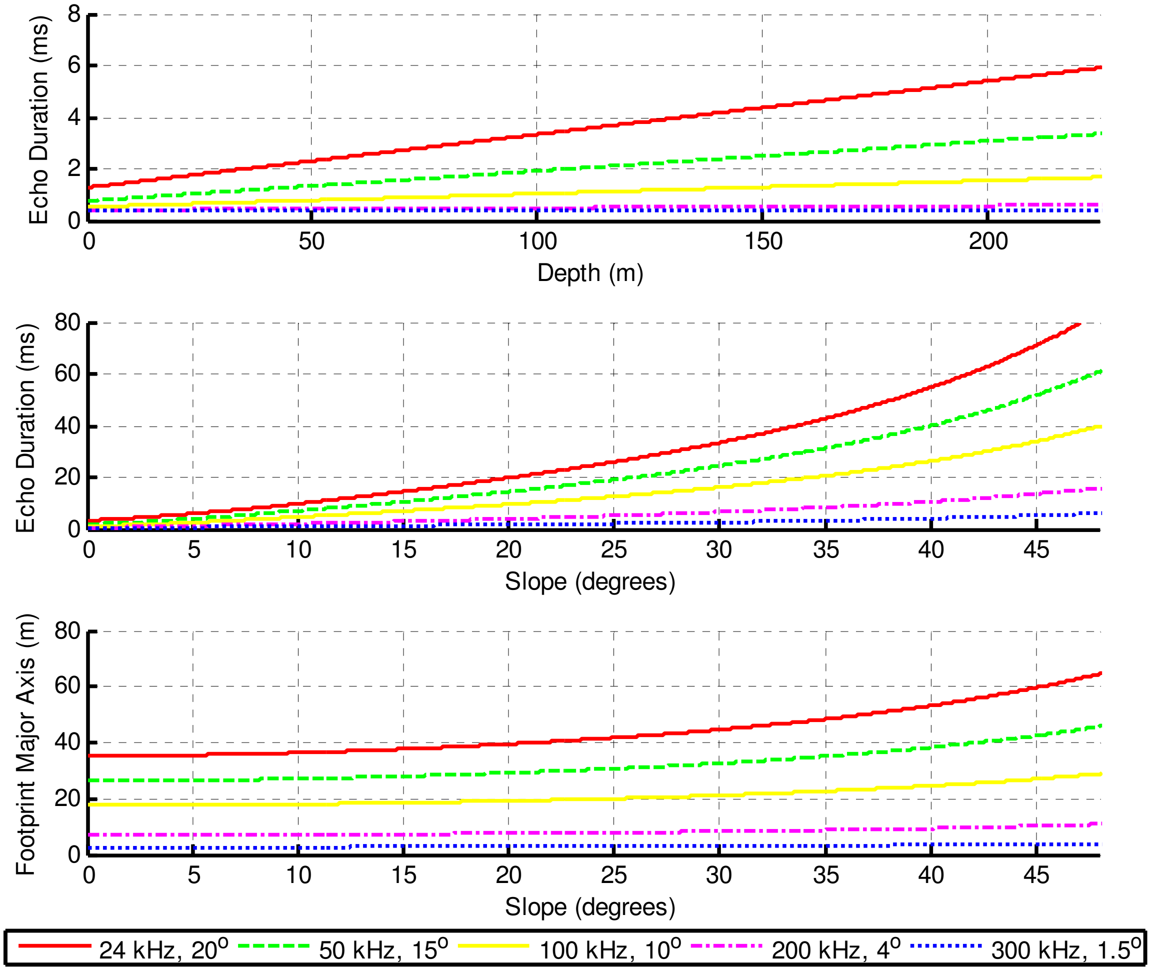

Figure 5-12. Plots of echo duration as a function of depth (*upper*) and slope (*middle*), and footprint major semi-axis length as a function of slope (*lower*) for typical SBES beamwidths and frequencies shown in the legend, based on an attenuation of 0.5 dB/kHz/m and 100 m water depth for the lower two plots.

SEL depth compensation formula (Equation 2-7), Equation 5-3 uses the sediment sound speed to calculate the penetration duration from the penetration depth. A more complete model of penetration depth (skin depth) is discussed in Appendix D, Equation D-1, which includes the reflection coefficient. However, the reflection coefficient is accounted for when measuring penetration time relative to the echo peak amplitude by way of echo duration, so an extinction model is more appropriate. Another issue is that the seabed parameters are generally not known, including the sediment sound speed used in Equation 5-3. Then, the extinction model is used with an *effective* attenuation and an assumed value of the sound speed equal to the water sound speed for consistency with the

SEL depth compensation equation (Equation 2-6). Effective attenuation is a value that gives a reasonable value for the penetration time (the duration of the volume backscatter) – a value that matches the data. The penetration time can be found explicitly from Equation 5-3 and Equation 2-7: $t_p = \frac{20}{c_w \kappa_{eff} f_c}$. This will be discussed further in chapter 6.

The most constraining approximation is that of a purely conical beam with equal sound intensity throughout and zero intensity outside the beam. This is a reasonable approximation for the main lobe of the beampattern, but it certainly does not account for side-lobes. The definition of the beamwidth used for a conical approximation was discussed in chapter 2; conventionally, it is taken as the standard -3 dB beamwidth. However, the -3 dB beamwidth may not be appropriate. Instead, the value that should be applied is the *effective* beamwidth, which is defined as a value that gives a reasonable duration for the wavefront curvature term – a value that matches the data. Results to be presented in chapter 6 will show that the effective beamwidth and the beamwidth that should be used for all calculations (including footprint diameter) is greater than the $1/e$ beamwidth or even the -6 dB beamwidth. As a consequence of the conical beam approximation, the model assumes penetration is equal across the footprint out to its edges, where it will be less but not significantly so unless the incident rays approach the critical angle. As discussed in chapter 4, many studies show that volume backscatter is dominant at the edges of the beam, away from the specular direction. This justifies the idea that volume backscattering does extend the duration of echoes as modelled in Equation 5-3. Also, it is certainly possible that surficial backscatter can determine the end of an echo making the inclusion of the volume term in Equation 5-3 erroneous. This may occur if the volume inhomogeneities are very weak and the attenuation is relatively high: this situation occurs for higher centre frequencies (>200 kHz), high homogeneity, high reflection coefficients and high attenuation. The best naturally-occurring examples are homogenous medium and fine sand, although it is virtually impossible not to create inhomogeneities through natural processes such as layering.

Despite all of these approximations and assertions, the echo duration model is very useful as it provides an estimate of a measurable quantity: echo duration. Echo duration is the basis for depth and slope compensation, improved discrimination and for new characterization methods. It will be shown in the next section and in chapters 6 and 7 that the model is valid when using the appropriate *effective* parameters of beamwidth and attenuation.

5.3.1 Initial Verification of the Echo Duration Model

There are several ways to verify the echo duration model. Field data can be used to compare measured echo durations to echo durations predicted by Equation 5-3. Simulated echo time series data from BORIS can be used in a similar fashion or can be used to test the approximations of the echo duration model directly. Pouliquen successfully used BORIS to test depth compensation by reference depth [Pouliquen, 2004]. The analytic echo time series model can also be used, although it is based on a Gaussian approximation to a full beampattern. The scope of the results presented in this section is expanded through a dedicated experiment described in chapter 6.

An initial verification of the EDM is shown in Figure 5-13, which presents a comparison of measured versus modelled echo duration for field data from the survey of Eagle Bay described in section 3.1. Measured echo duration is the average from stacks of five echoes. Tail picking was done with the robust threshold technique of appendix B[28]. The overall Pearson correlation coefficient between the measured and calculated echo duration was significant at 0.84. Low correlation generally occurred in specific regions of Eagle Bay that have more gravel in the sediment than the overall seabed type which was found (by video and grab-sample) to be muddy sand. The same analysis was also run for nearby Clio Bay and the results were good in a small area where there was no log debris. The overall correlation for Clio Bay was significant at 0.40.

---

[28] The more advanced picking techniques had not yet been devised.

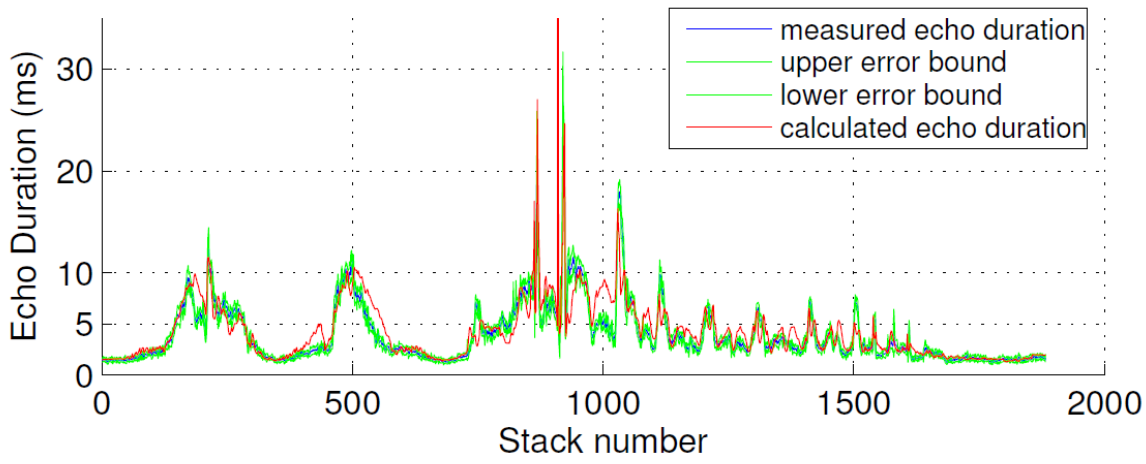

Figure 5-13. Measured and modelled (or calculated) echo duration for the 24 kHz survey of Eagle Bay, British Columbia. The error bounds are the standard deviations of each stack. Only stacks are plotted because navigation data are only available for each stack. EDM parameters used are discussed in the text below.

For both bays, the modelled echo duration was based on the attenuation coefficient of muddy sand only[29]; this could be a source of error. Another source of error is that the effective beamwidth was determined by observing the best correlation in a range of beamwidths between the -3 dB value (19°) and the -6 dB value (28°) – a more objective means of determining the effective beamwidth will be devised later in this sub-section and then improved throughout chapters 6 and 7. The effective beamwidth used was 24°, which is near the $1/e$ beamwidth. When seabed slope data were excluded (i.e. set to zero), the result was significantly degraded: the correlation coefficient was 0.62 (for Eagle Bay). The seabed slope was determined from the initial 200 kHz survey bathymetry and is considered to be of low quality[30]. However, even this low-quality seabed slope information is sufficient to improve the result greatly. Another source of error is tail picking error (bottom picking error is negligible in comparison). Echoes from sloping or low-attenuation seabeds have long, low-amplitude 'tails' that are difficult to distinguish from the noise floor. This causes the measured echo durations to be less the

---

[29] see Table C-2. The skin depth model of penetration depth was used in this calculation.

[30] If a survey does not have sufficient line density and crossing lines to resolve bathymetric features of scale equal to or less than the acoustic footprint, than the bathymetry is considered low quality. The correction of survey bathymetry for the effect of slope is discussed in the next sub-section.

modelled values for slopes greater than 12°. Threshold tail picking is particularly prone to underestimating the duration of these types of echoes. Additional non-systematic variability of the measured echo duration is due mainly to incoherent scattering. Overall, the modelled echo durations match the measured values very well as indicated by the high correlation coefficient and by inspection of Figure 5-13.

The second verification of the EDM comes from the initial testing of the QTC DEEP system aboard ROPOS, presented in [Bloomer et al., 2007]. Ascent and descent data were collected 4 km north of Patricia Bay in the vicinity of Patey Rock, B.C. This is an interesting data set since the seabed did not change while the distance between the transducer and the seabed varied as the ROV moved vertically. Technically the distance measured by the echosounder is the ROV's altitude, but it is identical to the water depth in terms of how echo duration behaves as a function of water depth. One can also see in Figure 5-11 that there is no distinction between altitude and depth for the echo duration model or any model only involving the direct-path echo. Equation 5-3 shows that echo duration is a linear function of water depth with the slope of the line being a function of the beamwidth and seabed slope, and the y-intercept being a function of the pulse duration, roughness height and penetration depth. Then, plotting echo duration data during the ascent and descent as a function of altitude should show a linear relationship, with the linear regression coefficients directly related to an *effective* beamwidth and attenuation (assuming a flat smooth seabed). This is shown in Figure 5-14. The equations used to calculate the effective beamwidth and attenuation are:

$$\begin{aligned} \theta_{eff} &= 2\left(\sec^{-1}\left(\frac{mc_w}{2}+1\right)-\phi\right) \qquad (\text{for } \phi \le \frac{\theta}{2}) \\ \kappa_{eff} &= \frac{20}{(b-\tau)c_w f_c} \end{aligned} \tag{5-6}$$

where *m* and *b* are the first and second regression coefficients, referred to as the regression slope and intercept. Also, note that f_c is in units of kHz to be compatible with the units of attenuation, which are dB/kHz/m, as noted in reference to Equation 2-7.

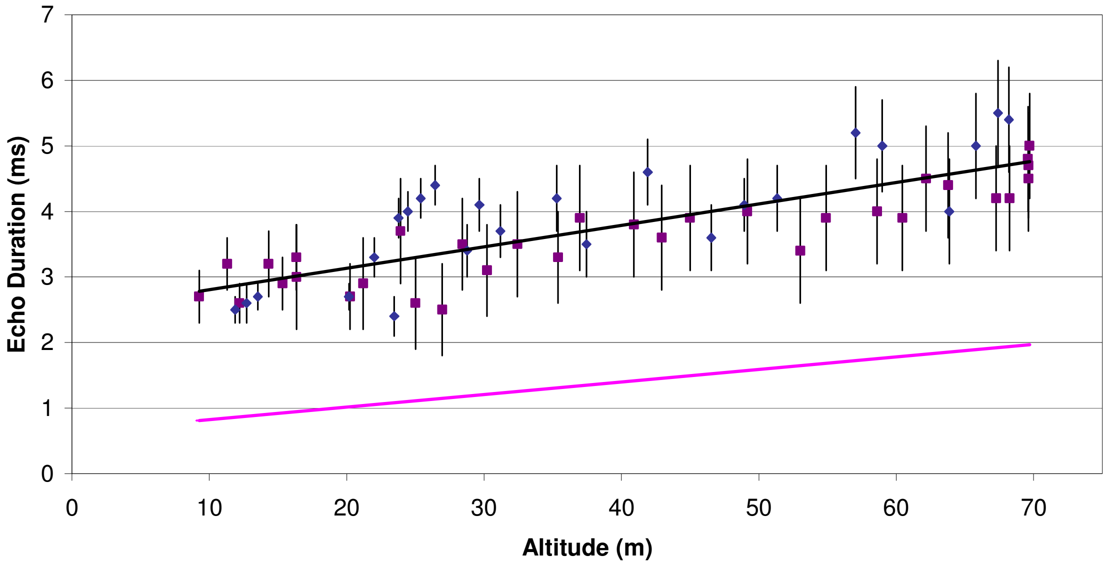

Figure 5-14. Measured echo duration from the vicinity of Patey Rock, B.C. The blue diamonds and purple squares represent the data during ROPOS's ascent and descent, respectively. The data points are the mean of each stack with error bars two times the standard deviation of each stack. The magenta line is the EDM calculated using the manufacturer's 19.3° beamwidth, pulse duration of 0.1 ms and 0.5 dB/kHz/m attenuation, representing one SEL as described in section 2.2.1. The black line is the ordinary least squares linear regression to the data. Based on Equation 5-6, the resulting effective beamwidth is 25.1±0.7° and the effective attenuation is 0.112±0.022 dB/kHz/m.

Figure 5-14 shows that the relationship between echo duration and depth is indeed linear with a high degree of variation explained by the linear fit to the data (with a correlation coefficient of 0.81). The standard -3 dB beamwidth that one would use for SEL depth compensation is significantly less than the value that would produce the linear relationship observed in the data. The effective beamwidth, 25.1° is consistent with the -6 dB beamwidth of 25.6°.

There some notable sources of error in this analysis. First, this was a data set of opportunity, and so the factors of positioning, seabed slope and seabed type are not well controlled. ROPOS video does indicate a homogenous fine to medium silt, which would imply that the attenuation values range from 0.11 to 0.24 dB/kHz/m. Another source of error is that the transducer was observed to 'ring' for 50 to 150 ms after the transmit pulse was broadcast. This is likely due to an impedance mismatch as the system was calibrated to operate in 800 m of water. The effect is minimal compared to the y-intercept in Equation 5-6 and Figure 5-14.

In order for modelled echo durations to match the data precisely, the criteria defining the end of the echo (i.e. the tail pick parameters) need to correspond to model parameters that also define the end of the echo. This is what has occurred here. A reasonable value of both the tail pick threshold (10 % for the threshold algorithm, see appendix A) and the penetration depth extinction parameter (10 dB) leads to an effective attenuation that corresponds well with the estimated seabed attenuation. The effective beamwidth value is also reasonable (it would not be reasonable it were smaller than the -3 dB beamwidth or larger than the first sidelobe). There is no expectation of a good correspondence between these parameters, as there is no direct link between the tail pick threshold and the effective beamwidth and attenuation. Such a link will be established in section 6.2, using flat homogenous seabed testbed sites that are well controlled. The experiment in section 6.2 will address many of the deficiencies of the Patey Rock experiment.

These two preliminary examples, Patey Rock and Eagle Bay, demonstrate that the EDM replicates the trends in the data due to depth and seabed slope. A more comprehensive analysis of the model follows in sub-section 5.4.3 using data from the BORIS model. Further analysis is the focus of section 6.2 using well-controlled field data as mentioned above. The implications of the EDM are discussed in the next sub-section.

5.3.2 Implications of the EDM for Seabed Classification on Sloping Seabeds

The geometry of the EDM shows that the specular reflection is lost at slopes greater than one-half the beamwidth. This significantly changes the nature of the echoes, reducing the coherent component and increasing variability. So as a rule, classification results are not valid for such slopes [Biffard et al. 2005]. This agrees with the observations from surveys in chapter 3 and from several other surveys [von Szalay and McConnaughey, 2002, Cooke et al. 2003].

The other effect of slope is to increase echo duration, as shown by Equation 5-3. Echoes are stretched in time, analogous to the effect of depth on echo time series. As shown in Figure 5-12, the effect is approximately ten times that of depth. From the EDM, the maximum change in echo duration due to seabed type is about 4 times for most systems. An approximately equal change in echo duration due to slope occurs at slopes equal to or slightly less than the beamwidth. Then, for slopes greater than the beamwidth, the effect of slope dominates over the influence of seabed characteristics on the echo.

5.3.3 Implications of the EDM for Bathymetry on Sloping Seabeds

The effect of slope on the accuracy of bathymetry is also significant. Using Equation 5-4 one can show that for slopes of 20° the vertical depth is 6.4%, 1.2% and 0.46% greater than the measured depth for beamwidths of 20°, 10° and 1.5° respectively. This is illustrated in Figure 5-15. The measured depth is always shallower than the vertical depth because the seabed normal ray always moves upslope to the edge of the beam as the slope increases as shown in Figure 5-11. Fortunately, this calculation of the vertical depth can be applied to any data for which slopes can be calculated accurately. However, the slope must be calculated from the vertical bathymetry. To overcome this issue, an iterative scheme was devised and tested on the Eagle Bay and Patricia Bay data sets (the later data set is presented in chapter 7). The iterative process is as follows:

- calculate the seabed slope from the bathymetry;
- correct the original survey bathymetry using Equation 5-4;
- then use the new bathymetry as input to the next iteration to again find the slope and correct the original bathymetry.

At the end of each loop, the corrected bathymetry is compared to the previous loop's corrected bathymetry and if the change is below a threshold (usually 1% of the mean difference) the iteration stops, keeping the latest corrected bathymetry. The calculation of seabed slope requires gridding the bathymetry using kriging[31] and then calculation of the surface gradient. It generally takes three to five iterations for the corrected bathymetry to converge. The absolute difference between the vertical bathymetry and the measured depths averaged 2.3% for Eagle Bay, which is significant considering hydrographic surveys aim for accuracies better than 1%.

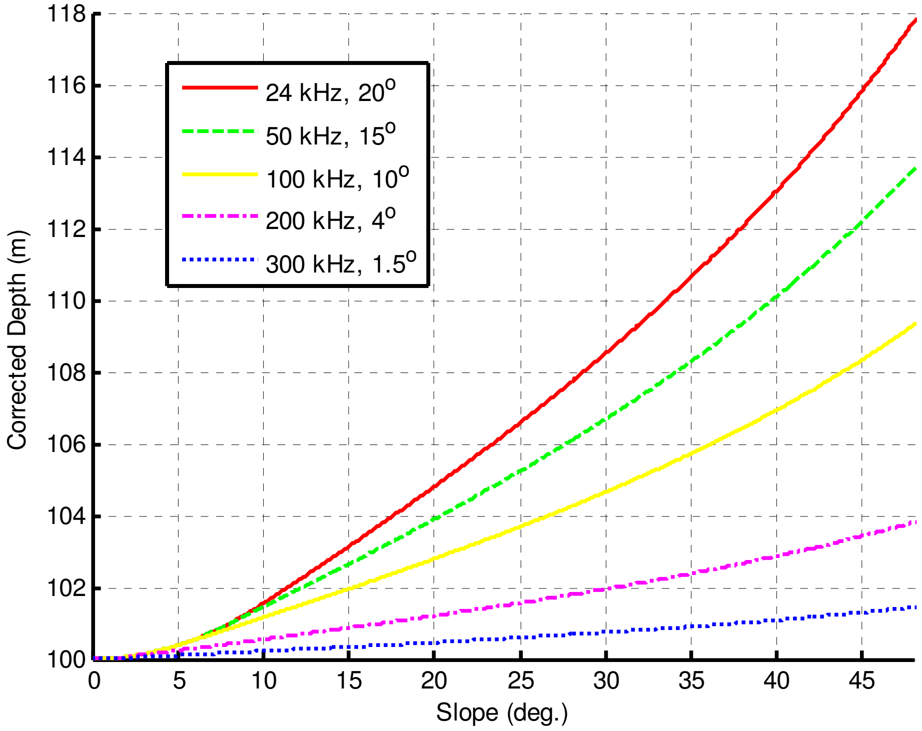

Figure 5-15. The correction of measured bathymetry for seabed slope based on Equation 5-4 for several typical SBES frequencies and beamwidths. The measured depth is constant at 100 m, pulse duration 0.3 ms and 0.5 dB/kHz/m attenuation.

---

[31] Golden Software's SURFER surface mapping software is used for gridding data by kriging; the SURFER user manual refers to: Cressie, N. A. C., *Statistics for Spatial Data*, John Wiley and Sons, Inc., New York, 1991, pp. 900. SURFER is also used to plot the classification maps of chapters 2, 3, 6 and 7.

5.4 Modelling Studies and Demonstration of Concepts

This section presents studies of several aspects of backscattering using the models of chapter 5 together. In addition to exploring the physics of echosounding, another aim is to demonstrate the effects of the theories of chapters 4 and 5 on seabed classification and characterization.

Aside from geotechnical parameters, (SBES centre frequency, beamwidth, water depth, seabed slope, etc.), the geoacoustic parameter that has the greatest influence on the characteristics of echoes is seabed surface roughness as shown in Figure 4-5. The AEM shows that impedance is effectively a scaling factor, while BORIS simulations indicate volume inhomogeneities are only dominant in some cases: inhomogeneous mud, the very tail of most echoes. Therefore, understanding the influence of roughness is vital to many aspects of this dissertation (design of features for seabed classification, characterization, etc.). The second part of this section explores the effects of the instantaneous ensonification of the seabed that is vital to forming the shape of seabed echoes. The third part of this section discusses the estimation of effective beamwidth and attenuation from the AEM and BORIS models, while the last part of this section presents a BORIS study of potential new features for seabed classification.

5.4.1 The Effects of Roughness

The BORIS model facilitates many experiments not possible in the field. For instance, one can vary the roughness height or the water depth separately without the other parameters. The simulations presented here are controlled experiments for the effects of roughness.

For the first controlled experiment, successive simulations with increasing roughness heights are generated while everything else remains fixed. The result is shown as an echogram in Figure 5-16. In the figure, one can observe several effects of roughness: the amplitude quickly decreases then increases again; the speckled appearance of the echogram increases; and the echo duration increases dramatically with both the bottom pick and tail pick moving apart. A similar echogram was constructed for varying the surface height power-law spectral exponent. Over the range of values (two to four), no systematic variation in the bulk characteristics of the echoes were observed so this result is not shown.

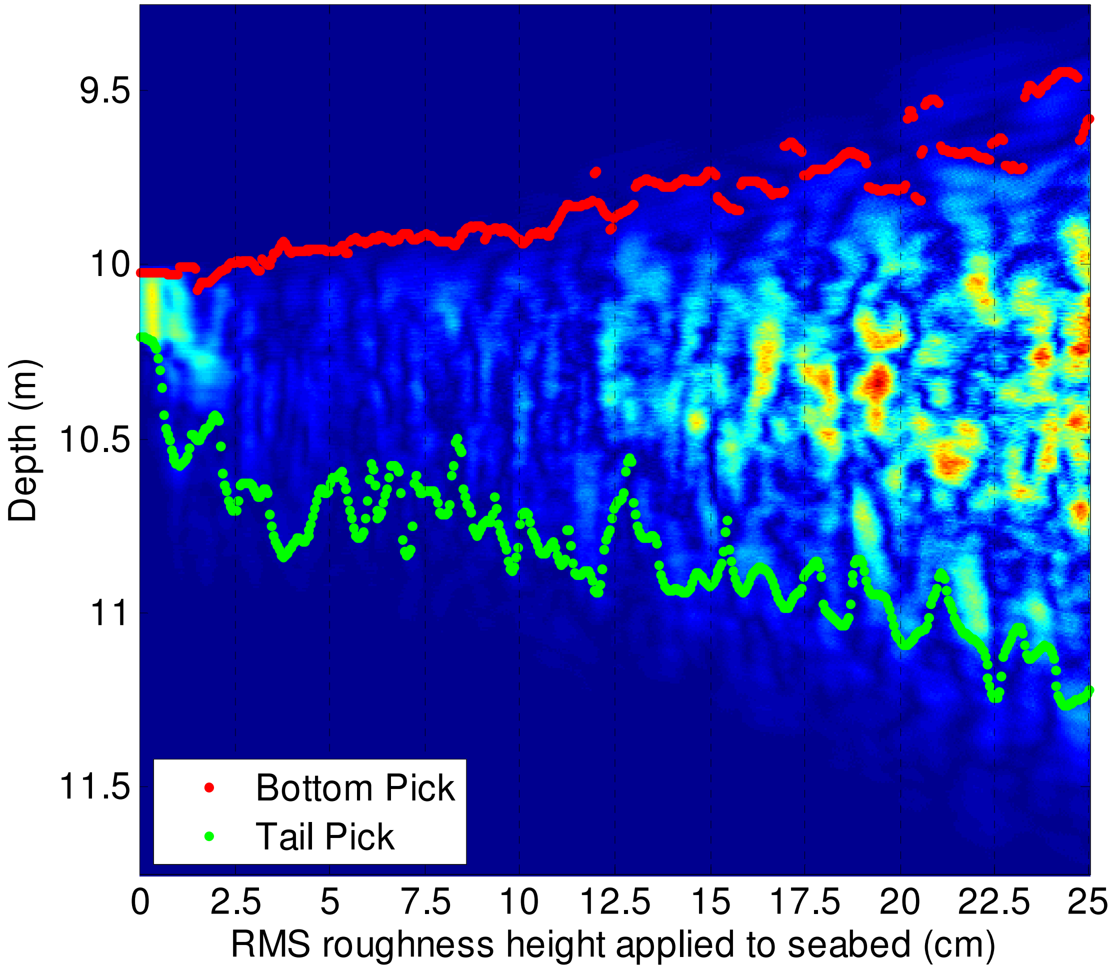

Figure 5-16. An echogram of increasing roughness from BORIS simulations of the circ50_20 virtual echosounder on the HardFlat seabed – see appendix C for the definition of these simulation scenarios. Roughness was increased from zero to 25 cm in 0.05 cm steps, effectively up-scaling the same stochastic seabed for each successive echo. There is no volume scattering. Colours indicate amplitudes corrected for spherical spreading; the progression of colours is the same as the colourbar in Figure 5-3. Red bottom picks are at 10% of peak amplitude, green tail picks are at 99% of the cumulative energy.

The effects shown in Figure 5-16 are consistent with scattering theory as discussed in section 4.7. The best way to discuss this is in relation to the Rayleigh criterion, Equation D-22, which is applied to separate the scale of roughness into three domains: roughness well below the Rayleigh criterion, roughness of order of the criterion and roughness well above the criterion. At normal incidence, the Rayleigh criterion can be simplified to $h_{RMS}^{RC} = \lambda/8$, where h_{RMS}^{RC} is the RMS roughness above which incoherent scattering dominates and the coherent echo is reduced by 10.7 dB. Then the three domains of relative roughness are:

$$\begin{aligned} h_{RMS} &<< \lambda/8 \quad \text{(low)} \\ h_{RMS} &\sim \lambda/8 \quad \text{(moderate)} \\ h_{RMS} &>> \lambda/8 \quad \text{(high)} \end{aligned} \tag{5-7}$$

These domains also represent the relative strengths of the coherent and incoherent components of the echo as seen in Figure 5-17. The relative strengths of these components are often expressed as a ratio of coherent to incoherent intensity; this ratio is γ and will be used extensively for characterization in chapter 8.

The decomposition of the echoes into their components as shown in Figure 5-17 was achieved by a process this author calls *coherent stacking*. It is the literal application of the theoretical field definitions discussed in appendix E, and obtains similar echo components as shown in Figure 4-5. Normal stacking as discussed in chapter 2 finds the total average amplitude by extracting the envelope or analytic signal using the Hilbert transform and then averaging several echoes (five is the standard number, 500 is the number used here). The *coherent stack*, or the average coherent amplitude as seen in Figure 5-17, is found by stacking many full-waveform time series followed by extraction of the envelope or analytic signal with the Hilbert transform. By averaging the full signals first, the 'random-walk' phases of the incoherent component average out and only the coherent component (defined by a constant phase) remains. The average incoherent amplitude can only be defined by subtracting the coherent amplitude from the total amplitude, as described in detail in appendix E. Coherent stacking is possible with BORIS data because the water depth is exactly constant; real data would have to be aligned precisely to minimise the travel-time phase variance. For BORIS data, it was observed that the minimum number of pings required to stabilise the coherent echo is approximately 30.

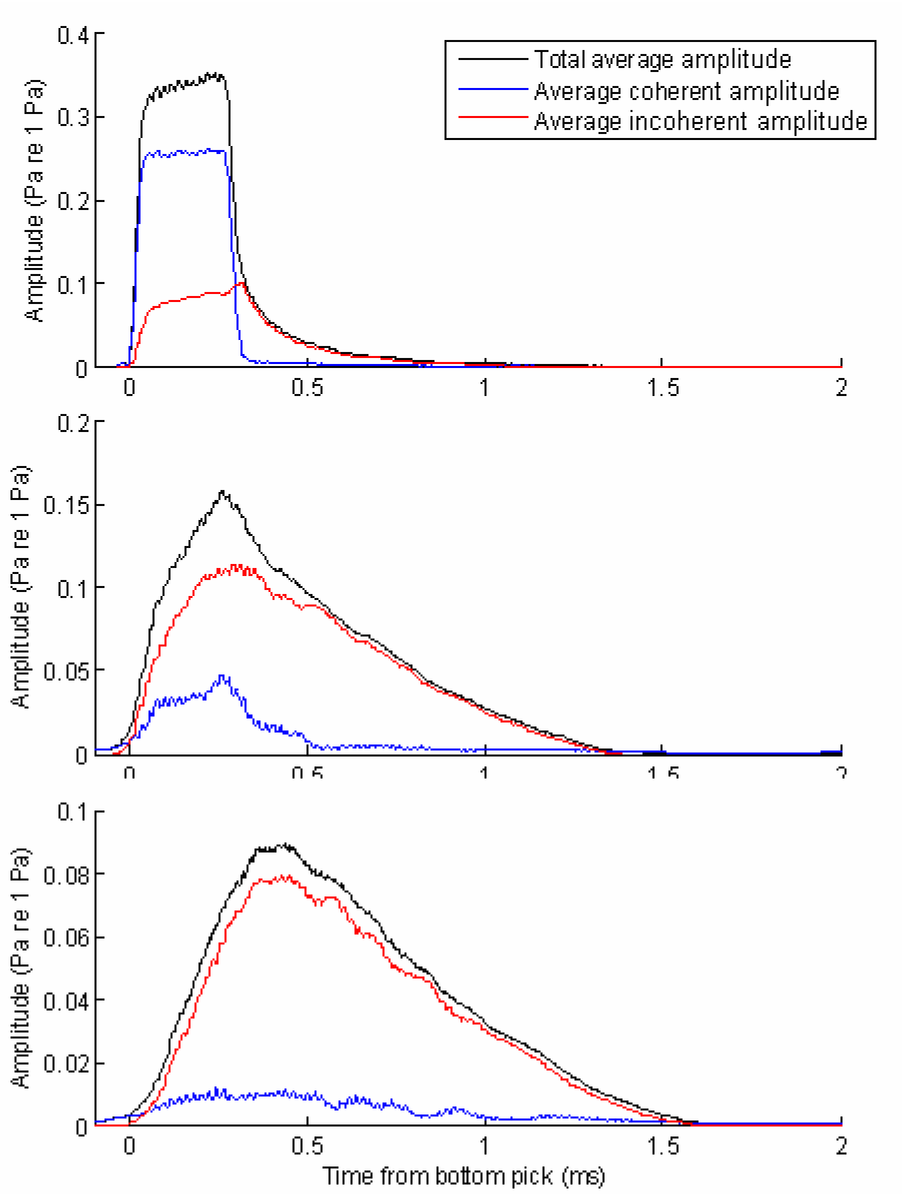

Figure 5-17. Average coherent, incoherent and total amplitudes for the three domains of relative roughness: low (0.25 cm RMS, top), moderate to high (2 cm RMS, middle), high (5 cm RMS bottom). Each mean is a stack of 500 echoes from BORIS simulations. Simulation parameters are as before: Circ50_20 virtual echosounder on the HardFlat seabed, the echosounder moves 1 m between echoes, depth 10 m, no volume scattering.

The total amplitude is the mean or stacked echo envelope. The term 'total' is used as the echo is composed of coherent and incoherent parts. The coherent amplitude can be thought of having two parts: the specular part is a reflection of the transmit pulse (0.3 ms in this case, see Figure 5-17 *upper*) and the non-specular part from off-nadir facets which are orientated normal to the wavefront. The non-specular part increases in duration with increasing roughness due to increasing numbers of facets at higher angles as predicted by KA theory. The non-specular or scattered coherent field can also include coherent volume scatter or reflections from sub-bottom reflectors; however, sub-bottom reflections are rare in the domain of SBES defined here because of the short wavelengths and shallow penetration of these systems. Coherent volume backscatter dominates the echoes observed by lower frequency systems such as sub-bottom profilers and seismics, while incoherent volume backscatter dominates SBES echoes. Table 4-1 summarizes the four basic components of echoes.

The relative contributions of the incoherent, specular and non-specular coherent amplitudes to the total echo can explain some of the observations made in Figure 5-16. For low roughness (<0.375 cm), the echo is specular – it is a reflection of the transmit pulse and exhibits no variation from echo to echo. For medium roughness (0.4 cm to 1 cm), the specular component decreases while the incoherent component increases in the area around the seabed normal so that the echo still looks like a reflection of the transmit pulse. The BORIS results show a peak in the mean amplitude around the Rayleigh criterion roughness (0.35-0.4 cm, Figure 5-16) likely due to the added contribution of the incoherent and non-specular coherent fields. As roughness increases past this level, the non-specular coherent component becomes more important (Figure 5-17, *middle* and *lower*), while the scattering region expands decreasing the mean amplitude. As the roughness heights become several times larger than the acoustic wavelength, the number and size of the facets continues to increase. This is seen in Figure 5-16 with amplitudes increasing above the level that was observed at low roughness (for RMS roughness greater than approximately 12 cm).

Since the RMS roughness height is essentially a scaling factor in the creation of a stochastic BORIS seabed surface, all the simulations of Figure 5-16 are based on scaled versions of the same seabed heights. This causes the ping-to-ping variability to exhibit an artificial coherence. The advantage of this is that one can see the effects of several large facets causing high amplitude non-specular coherent reflections. In high-roughness simulated surveys with natural ping-to-ping variability, short-duration high amplitudes (i.e. spikes in amplitude much shorter than the pulse duration) are observed instead, with much less coherence. In general, coherence increases as scattering decreases. The level of coherence or ping-to-ping variability observed in simulated surveys corresponds to levels observed in field data; quantitative measures of variability are developed in chapters 7 and 8.

For a number of roughness heights, the depth was varied over the same seabed to to observe the effects of roughness on the linear trend predicted by the EDM. The resulting echo duration data were plotted against depth in Figure 5-18. In the exceptional case of zero roughness, the backscatter is purely specular and the effective beamwidth is zero, so that the echo duration is equal to the pulse duration at all depths. As roughness increases, the effective beamwidth increases as more facets are angled to backscatter energy, increasing the backscattering area. This author calls this effect *beam-widening*. From the slope of the linear least-squares fits shown in Figure 5-18, the effective beamwidth varies from 23.8° (0.5 cm RMS roughness) to 36.5° (2 cm RMS roughness). 36.5° is definitely larger than the -3 dB beamwidth of 20°, and is wide enough to include the first sidelobe. There is a roughness height above which there is no further beam-widening, in this case, somewhere between 2 and 5 cm (dependent on frequency, in general). The beam-widening effect likely reaches this maximum because the area from which there is significant backscattering tends to saturate because of the combined effect of the beampattern and the decreasing likelihood of facets being inclined enough to backscatter acoustic rays at the edges of the beam. The effect of the macro-roughness heights term in Equation 5-3 also becomes noticeable around 2 cm RMS roughness. At 20 cm RMS roughness, the linear fit finds that the macro-roughness, l_H, corresponds to a peak roughness of about 2.5 times the RMS roughness (50 cm), which adds 0.7 ms to the intercept, as observed. One can also see this effect in Figure 5-16, with the bottom pick to occurring earlier by an amount that corresponds to the peak roughness. The level of roughness at which the macro-roughness term becomes important is not likely related to frequency, but is instead the resolution of the data, which is likely less with real data that has the complication of volume scattering.

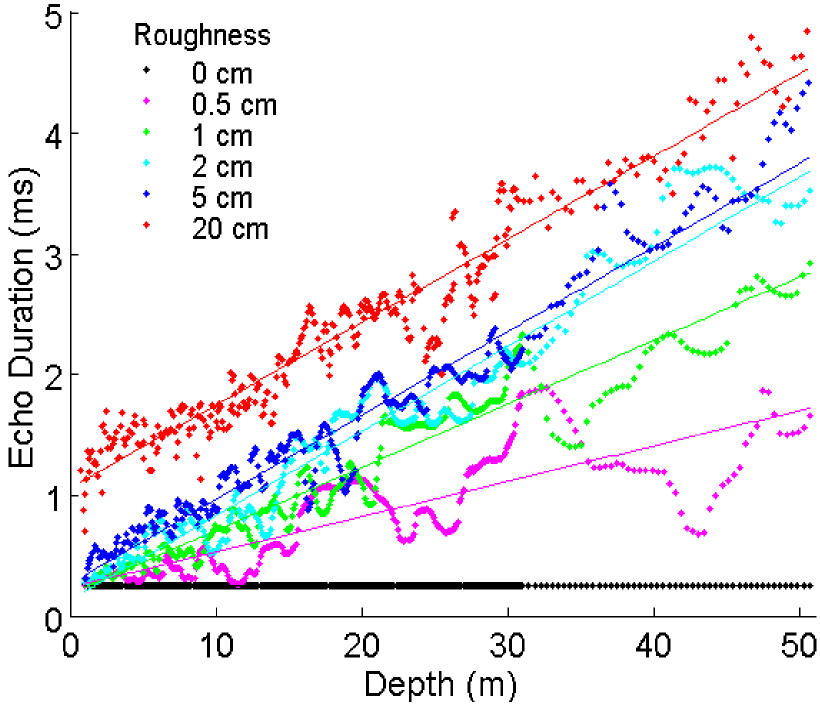

Figure 5-18. Echo duration vs. depth from BORIS simulations of the circ50_20 virtual echosounder on the HardFlat seabed. Straight lines are least-squares fits to each series of roughness. Note that the variation in echo duration exhibits a degree of auto-correlation because the same seabed is used so the backscattering effects are the similar for simulations at similar depths; this effect is slow varying and especially strong for the low-roughness series. This effect can also be seen in Figure 5-16 and is explained in detail in that sub-section (5.4.3).

The AEM also demonstrates the effect of beam-widening, but does not model the effect of macro-roughness. For beam-widening, the saturation point is even clearer because the beampattern is approximated as Gaussian which lacks the sidelobes of a real beampattern. Figure 5-19 shows the saturation point of beam-widening for a range of beamwidths; for the 20° of the simulations above, the saturation point is slightly less than 0.02 m. This is consistent with the findings above because a Gaussian beam has less energy at its edges than a real beam has. A drawback of the AEM is that it is frequency independent and is not able to show the frequency dependence of beam-widening. From the concepts of KA theory, one would expect that the large facets observed in a roughness saturated beam would have to scale with the wavelength. Additional BORIS simulations (at 38 and 24 kHz) confirm this – for moderate sized beams (20 to 24°) the saturation point generally occurs when the RMS roughness height is approximately equal to the wavelength. However, this is a strong function of the beamwidth such that wider beams require greater roughness, as shown in Figure 5-19. Conversely, smaller beams require less roughness (BORIS simulations of 200 kHz echosounders, 4-7° beams confirm this). It appears the limiting condition is that the orientation of the facets has to be normal to the incoming rays, possibly such that the local inclination of the facets (Equation 4-13) has to be equal to the incident angle which would be the half-beamwidth at the edge of the beam. This condition is more important than the size of the facets in relation to the acoustic wavelength. Since practical echosounders' beamwidths decrease with frequency and smaller beams require less relative roughness to saturate the beam-widening effect, the size of the facets relative to the acoustic wavelength is usually significant to facilitate beam-widening. Then in practice, the contribution of frequency to the effect of beam-widening is seen as negligible compared to the beamwidth with the caveat that practically impossibly[32] wide-beamed high-frequency echosounders (e.g. 200 kHz, 20°) may not follow this rule. Since the effect of facet size relative to the wavelength is largely irrelevant and the AEM handles the effect of the facet inclination, the AEM is suitable to model the beam-widening effect. In general, the AEM result shown in Figure 5-19 is a taken as a guide to predict beam-widening for various echosounders.

---

[32] The size of the transducer head required to produce a 200 kHz, 20° wide beam would be so small that it would not produce enough power, without cavitating, for practical applications.

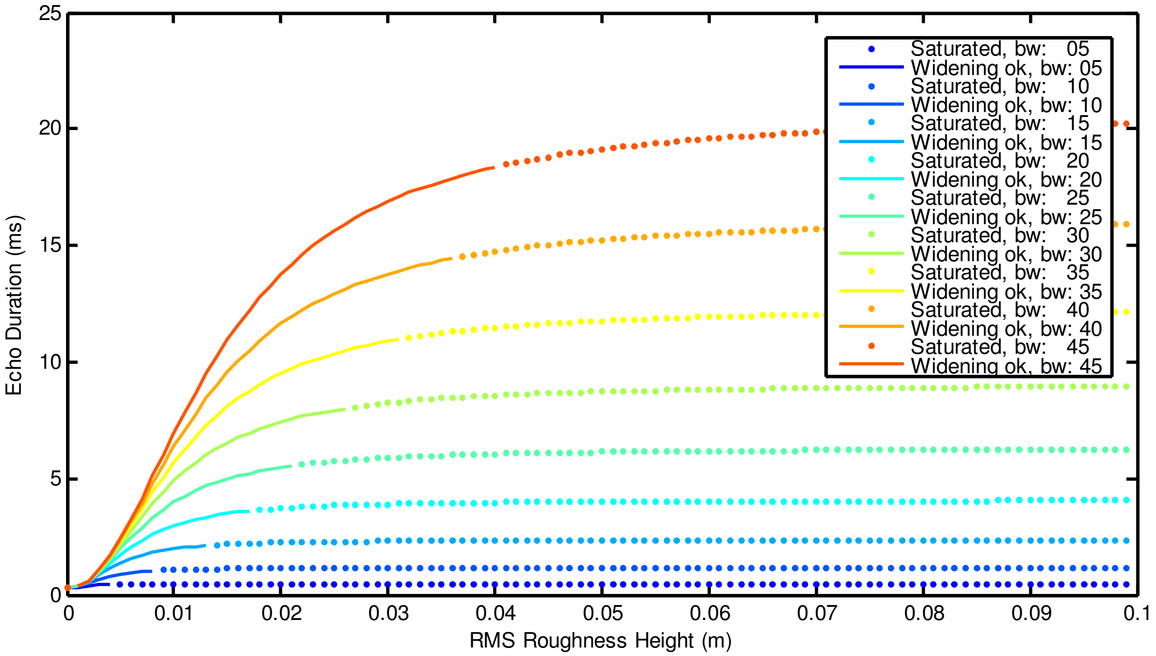

Figure 5-19. AEM simulations of echo duration as a function of roughness (with the default parameters as listed in section 5.2). The effect of beam-widening increases (lines) the echo duration until it saturates (dots), the criterion for which is a change of 0.2 ms or less per 0.001 m RMS roughness.

5.4.2 The Ensonification Regime Effect

The movement of the leading and trailing edges of the transmit pulse as projected onto the seabed has a major influence on the shape of the resulting echo. One can model this using the geometry of the EDM, including the conical beam approximation. As the transmitted pulse spreads along the seabed, the ensonified area is initially a disc (initial ensonification in Figure 5-20). The disc will spread to completely fill the footprint (full ensonification) *or* the trailing edge of the transmit pulse will arrive on the seabed prior to the leading edge reaching the limit of the footprint (annular ensonification). During full ensonification, the ensonified area remains exactly the same and so does the echo amplitude, see Figure 5-21. During annular ensonification, the ensonified area remains the same; however the incidence angle increases, decreasing the echo amplitude. As seen in Figure 5-21, the peak of the echo amplitude is always the point when the trailing edge of the transmit pulse reaches the seabed. This rule applies for average or stacked echoes at normal incidence on normal seabeds. In the final ensonification stage, the ensonified area decreases and the mean incidence angle increases, so the amplitude decays to zero even with a conical beampattern. In all, there are two alternative regimes of ensonification, full or annular, with the changeover point determined by the water depth, beamwidth and pulse. The changeover point is most conveniently expressed as the ensonification critical depth:

$$d_{ec} = \frac{c_w \tau}{2\left(\sec\left(\theta/2\right) - 1\right)} \tag{5-8}$$

The shape of the echo resulting from either of the two regimes is a complex function of many parameters, but as shown in Figure 5-21, the ensonified area and choice of regime have a major influence on echo shape. The square-echo of the full ensonification regime is found in shallow water, while the peaked-echo of the annular ensonification is found in deeper water. For typical SBES, the ensonification critical depth is 10 to 40 m. Wider beams and shorter transmit pulses decrease the depth at which the regime changes, and therefore are preferred for seabed classification as the square-echoes of full ensonification have less seabed information in their shapes than do the peaked-echoes of annular ensonification. This behaviour is observed in real echoes and modelled echoes using realistic beampatterns, further confirming that the conical beam approximation is appropriate if used with an effective beamwidth.

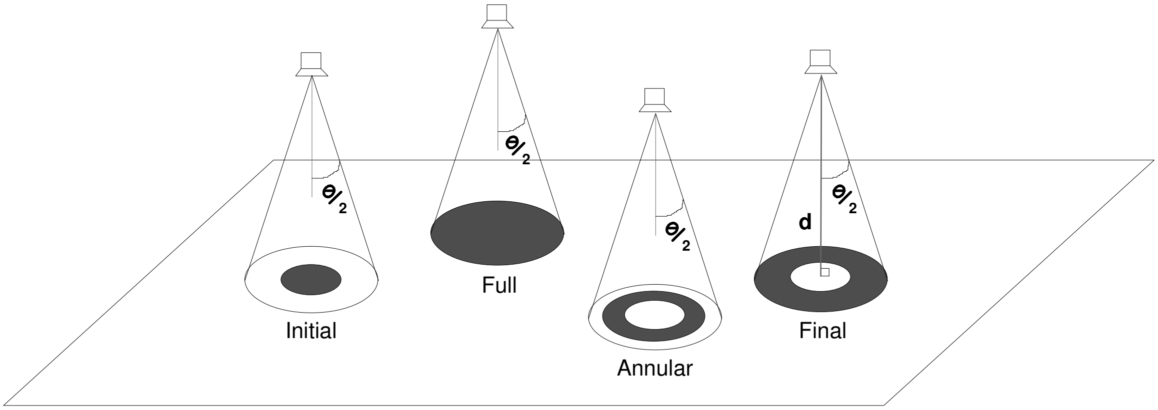

Figure 5-20. The four stages of ensonification visualized with a conical beam. The initial and final stages always occur, while the full and annular stages are exclusive of each other.

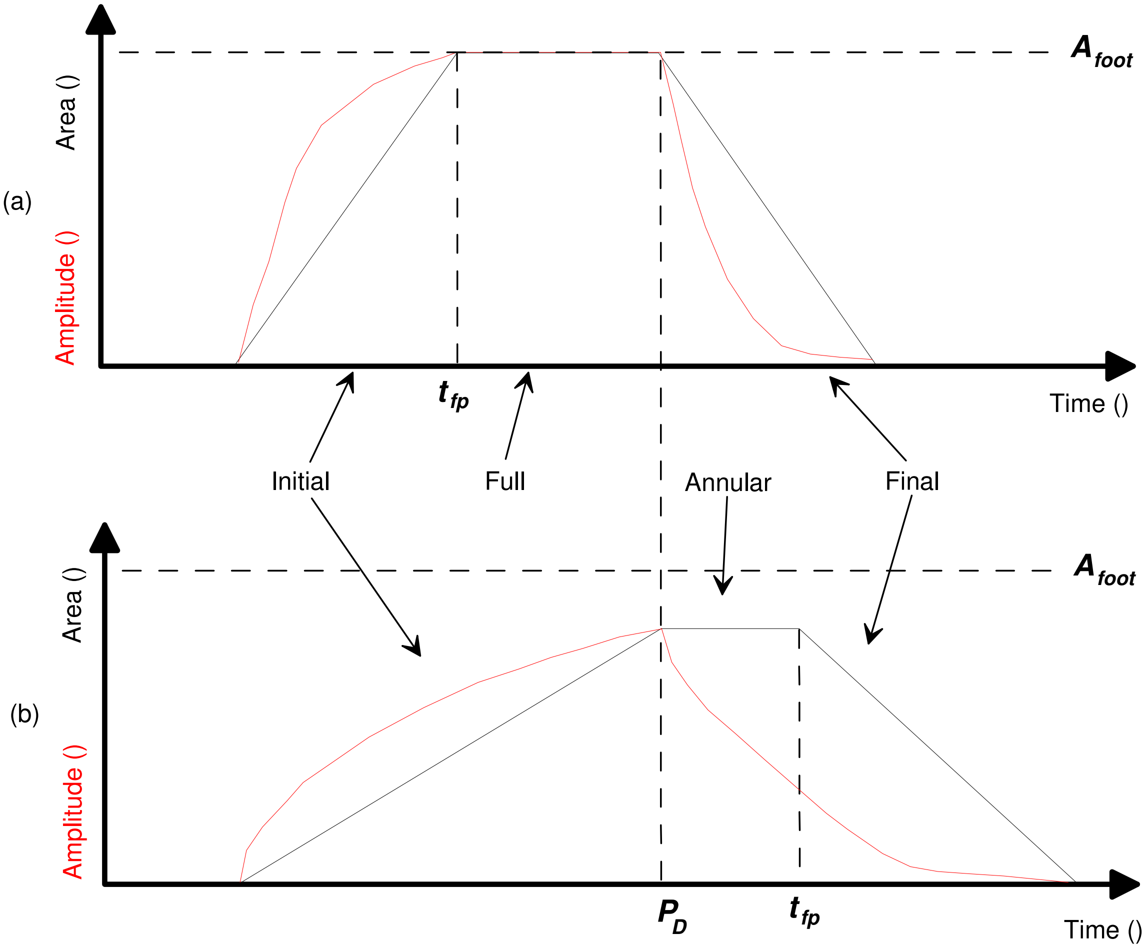

Figure 5-21. Qualitative representation of ensonification area and a typical echo shape as a function of time for full ensonification (*a*) and annular ensonification (*b*). The ensonification area is represented by the black lines, while the red curves represent typical echo amplitudes. t_{fp} is the time at which the leading edge of the transmit pulse reaches the end of the footprint and A_{foot} is the full area of the footprint. Not drawn to scale on any axis.

For sloping seabeds, the effect is more complicated, but made easier if the footprint is broken into upslope and downslope parts around the first arrival ray (the seabed normal that intersects the transducer). For the upslope footprint, the ensonification regime will tend to reach full ensonification at deeper depths, while the downslope footprint will tend to have annular ensonification for shallower depths. One could replace $\theta/2$ in Equation 5-8 with $\theta/2 \pm \phi$ to estimate the critical depth along the slope gradient for case I slopes. However, the shape of the echo will not be dominated by either regime.

Simulation results with the AEM and BORIS models show the effect of ensonification regime on the shape of echoes. The BORIS results, Figure 5-22 (*b*), are aimed at simulating real surveying conditions for the same seabed at various depths. There is a subtle difference in shape before the peak amplitude between the 5 and 10 m echoes and the 30 and 40 m echoes – this is due to the ensonification regime. It is a secondary effect when compared to the effect of depth on the echo shape as seen in the tails of the echoes - the tails of the echoes are stretched out with increasing depth. The ensonification regime does not affect the shape of the tail of any echo as the final stage is the same for all depths (see the 'final' regime in Figure 5-20). The shapes of the echoes from the AEM, Figure 5-22 (*a*), are extremely similar, building confidence in the reliability of the two models and the concept of the ensonification regime effect.

The ensonification regime effect can be observed in single realizations of field echoes as demonstrated by Figure 5-23. Normally, variability caused by the incoherent fields masks the effect on single echoes; the effect is pronounced in stacked echoes, as seen in BORIS results (Figure 5-22 (*b*)). However, stable positioning and low roughness sand reduce the variability seen in Figure 5-23. In Figure 5-23, the upper echo is the square shape characteristic of a fully ensonified footprint, which occurs at very shallow depths[33]. The middle echo rises to maximum amplitude and 'saturates' at full ensonification for a short time, similar to the 10 m echo of Figure 5-22 (*a*). The lower echo is in the annular ensonification regime and exhibits the expected shape: it rises to a point approximately one pulse duration after onset. Although the effect is blurred in real data by variability, the effect is consistent and can be seen in ensemble averages, i.e. stacks; it does affect classification results.

---

[33] Also in the upper time series of Figure 5-23, one can see the second echo and the electrically-attenuated transmit pulse.

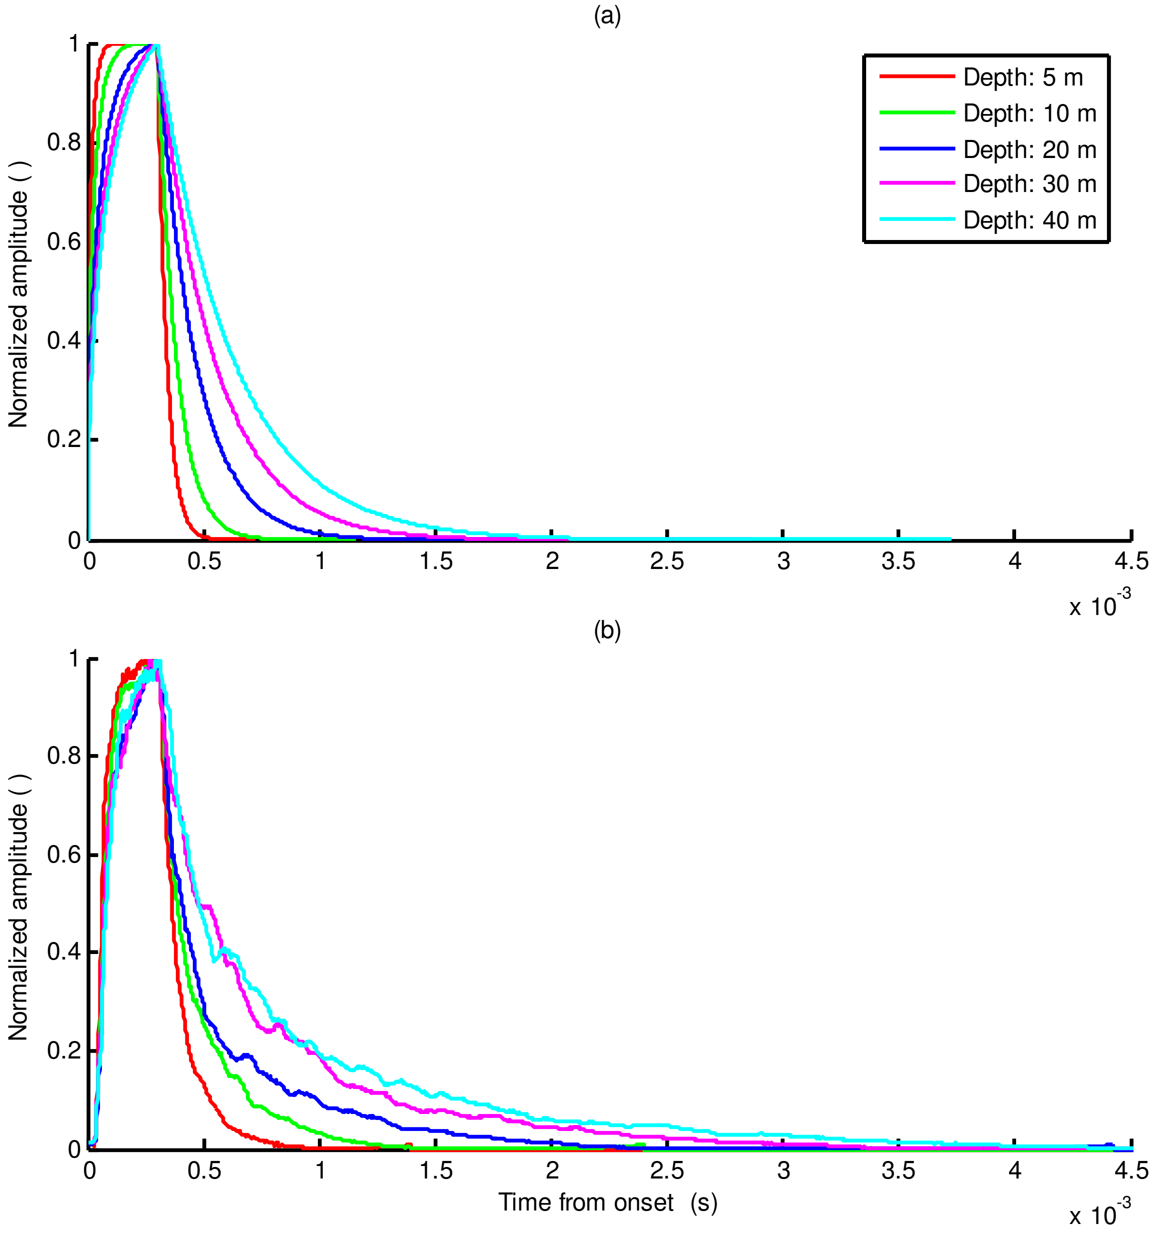

Figure 5-22. AEM (*a*) and BORIS (*b*) simulated echo time series at several depths around the ensonification critical depth of 14.1 m. For the BORIS simulations, the time series are stacks of 50 simulations for the Circ50_20 virtual echosounder on the Sand seabed (see Table C-3, Table C-1), using a simulated survey speed of four knots and three pings per second. The parameters for the AEM simulations (*a*) match those in (*b*) and are very similar to the default scenario used in section 5.2. The time from onset is slightly different in each case: for BORIS it is taken as five samples prior to the 99% CE bottom pick and for AEM it is taken from the travel time of the specified water depth. The difference between the BORIS and AEM simulations seen in the tails of the simulations is due to volume scattering that is not included in the AEM (volume scattering is important in tails of echoes).

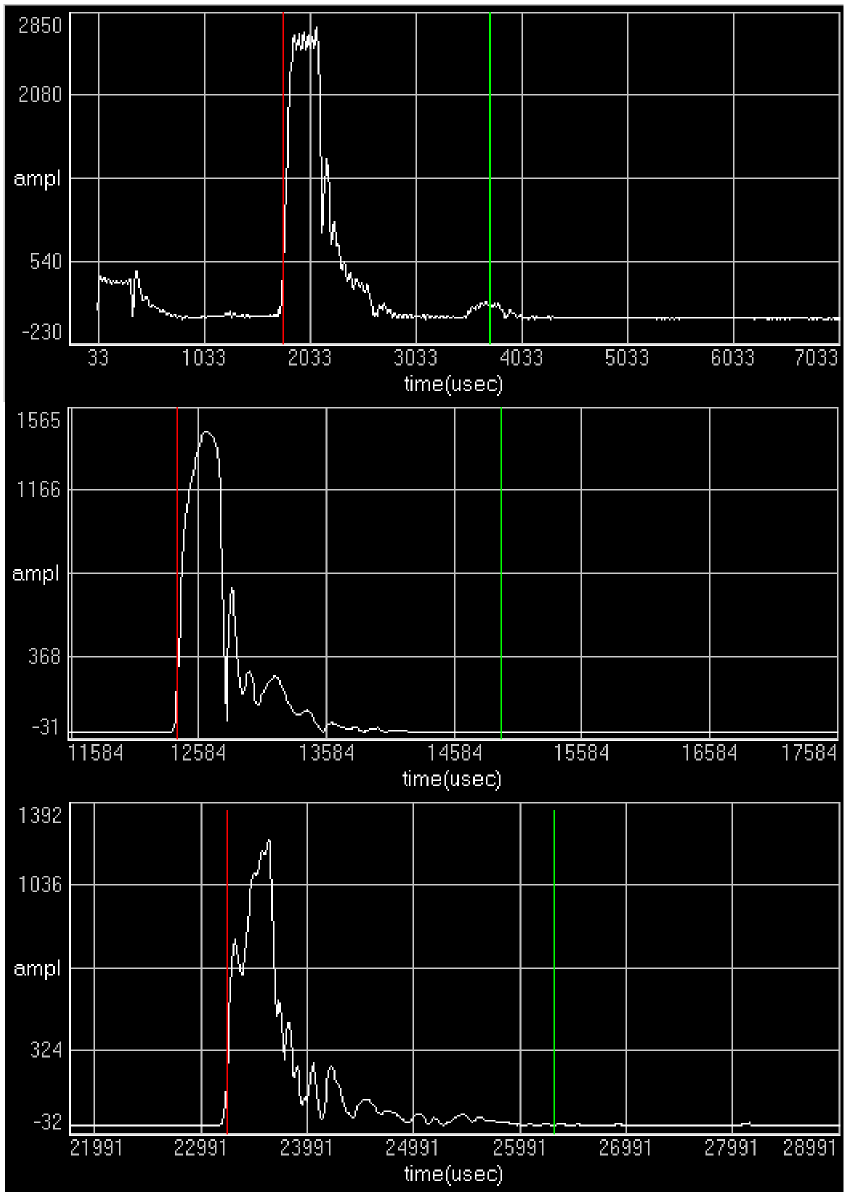

Figure 5-23. Echo envelope time series at three depths (1.35 m (*upper*), 9.40 m (*middle*), 17.41 m (*lower*)) as in QTC IMPACT from the Odom24 echosounder at the sand testbed site during the depth experiment (described in detail in chapter 6). The ensonification critical depth is 15.56 m. Time axis is in microseconds, relative amplitude is measured in units proportional to voltage or pressure and not all digits may be significant.

In chapter 2, the original proportional depth compensation called TNORM was introduced and shown to be inadequate. A proposal to improve proportional depth compensation is to add a proportional pulse length adjustment. This was presented and evaluated by simulation in [Pouliquen 2004]. The idea is to adjust the pulse duration proportional to water depth: for example, if at 50 m the pulse duration is 0.3 ms, at 100 m the pulse duration would be 0.6 ms and at 5 m it would be 0.03 ms. Although many SBES have the capability adjust their pulse lengths automatically based on depth, they usually do so in large steps (0.1 ms) and do so to improve the accuracy of depth measurements in deep water. Other reasons continuous proportional adjustment is not practical include that depth often changes too rapidly for accurate adjustments and that erroneous depth picks would cause irreversible errors. However, an unseen benefit is that the pulse duration adjustment effectively keeps echoes from any depth in the same ensonification regime. This is due to the fact that the ensonification regimes scale proportionally to pulse duration: if the ratio of depth to pulse duration is the same for all echoes, the ensonification regime is the same. One can see this in Equation 5-8. In Figure 5-24, the ensonification regime and relative echo shape is constant. Taking this a step further, Figure 5-25 shows the same series of echoes with and without pulse duration adjustment. In both figures, proportional depth compensation would yield a set of identical echoes.

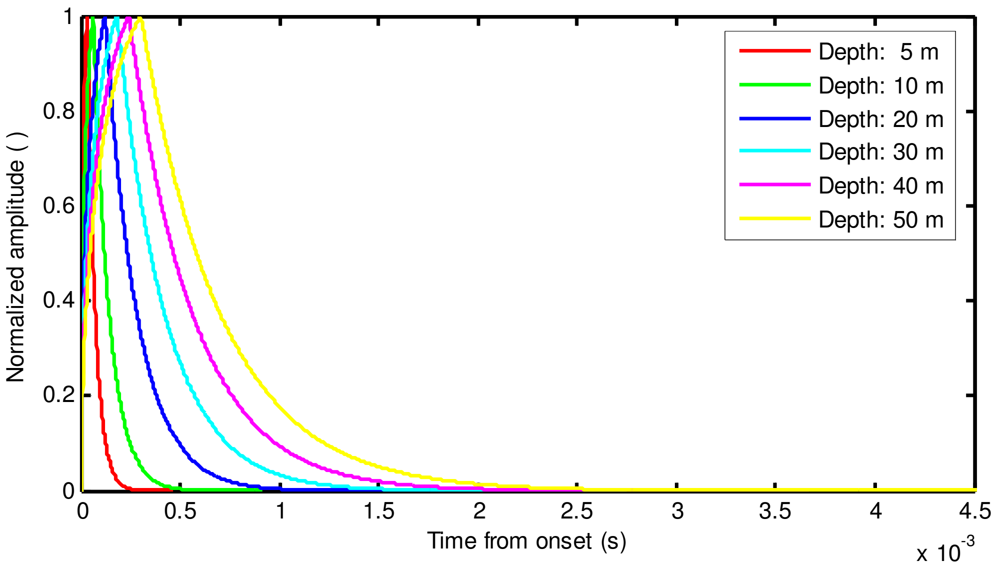

Figure 5-24. AEM simulated echoes with the pulse duration adjustment; reference pulse duration is 0.3 ms at 50 m. The parameters are the same as those in Figure 5-22.

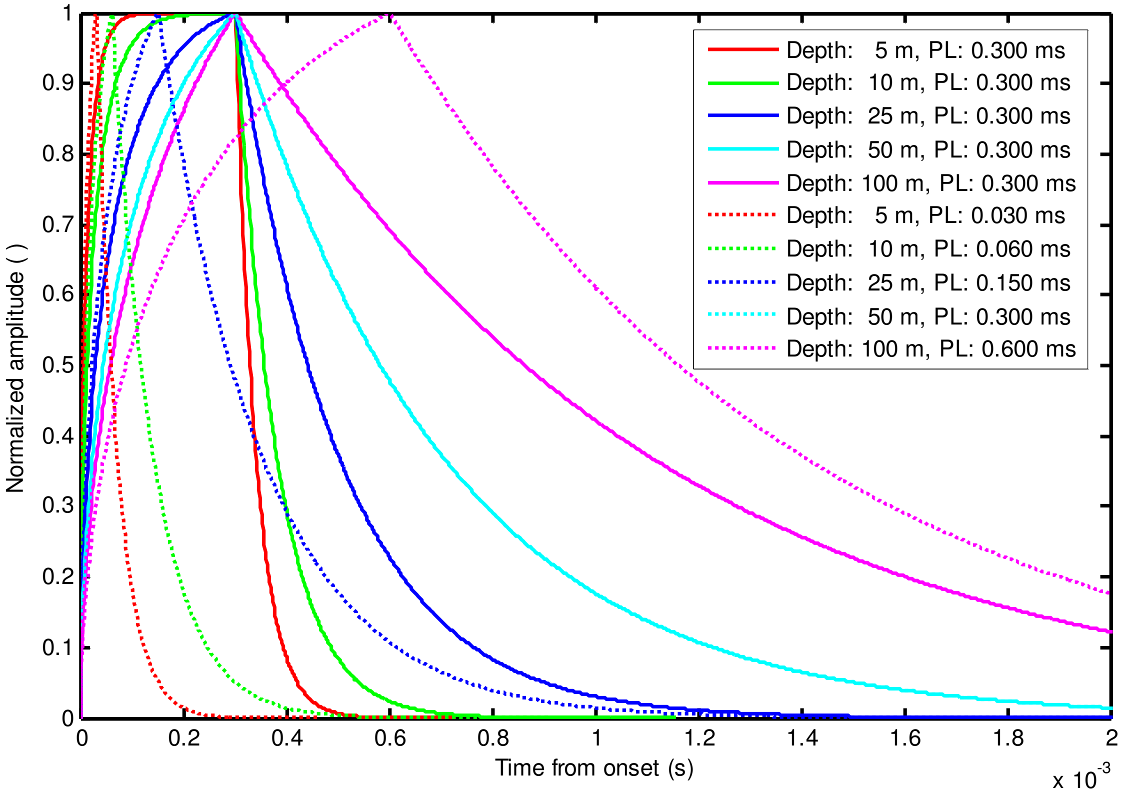

Figure 5-25. AEM simulated echoes with (dashed lines) and without (solid lines) pulse length adjustment for the parameters listed in Figure 5-22. The reference pulse duration is 0.3 ms at 50 m (the two 50 m echoes are identical).

Overall, the effect of ensonification is relatively minor compared to the other non-seabed influences on echoes – those of depth and slope. If depth is compensated for and slope is near zero (or compensated for), it is likely that the ensonification regime effect can cause misclassification. This effect modifies the shape of echoes most noticeably around the ensonification critical depth, so misclassification would occur at those depths. It is also responsible for square-looking echoes from shallow depths that offer limited discriminatory information based on shape. For seabed classification surveys, the best practical solution is to keep the survey data in one regime or the other. In surveys that cross the ensonification critical depth, the classification results may have a shallow-water square-echo class that contains several seabed types. This misclassification has been observed in several surveys, an example appears in chapter 7. If such a classification error is observed, a practical remedy would be to split the data about the ensonification critical depth and analyze both parts separately. A better solution is presented in chapter 7, it is the combination of a new depth and slope compensation method and new features selected to limit this effect.

5.4.3 Effective Beamwidth and Attenuation from Modelled Data – AEM and BORIS Results

As seen earlier in this chapter, the analytic echo model can be used to show the effects of geoacoustic parameters on the echoes. In this case, the AEM is used to evaluate the relationship between echo duration and depth, and to determine if that relationship is linear as predicted by the EDM. Figure 5-26 shows the results.

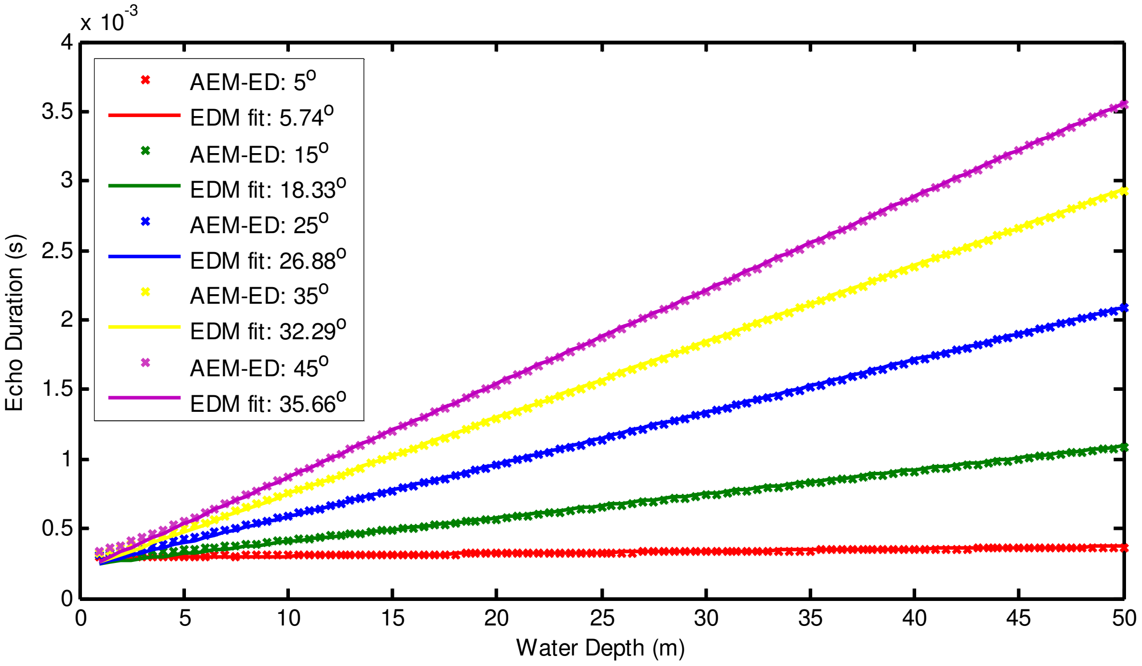

Figure 5-26. Echo duration as function of depth for AEM simulated echoes for a series of different beamwidths as listed in the legend. The legend also indicates the effective beamwidth from a linear fit to the data, calculated from Equation 5-6. AEM parameters: seabed roughness of 0.01 m, correlation length of 0.1 m (a roughness ratio of 10) and a pulse duration of 0.3 ms.

In the above figure, echo duration was measured from AEM simulated echo times series using the 99% cumulative amplitude algorithm (see appendix B). The beamwidths used in the AEM is the $1/e$ beamwidth. As can be seen, the linear fits are excellent. The small deviation for water depth less than 8 m that is due to picking sensitivity. The effective beamwidths as calculated by Equation 5-6 are slightly larger than the $1/e$ beamwidths up to $30°$. It appears the results of simulations with input beamwidths greater than $30°$ begin to suffer from the limit of applicability of the underlying theory; see Equation D-19 or refer to [Jackson et al., 1986] which does not apply KA theory past $30°$ incidence. Since the model does not include volume scattering, the values of effective attenuations are unrealistic and are not shown here.

Overall, the echo durations from the AEM are in excellent agreement with the EDM for conditions within the limits of the underlying theory, described in section 5.2. The BORIS model is more advanced in that it includes volume scattering, a full beampattern and simulates the effects of ping-to-ping variability. It also has a wider range of applicability due to SSA corrections. A similar analysis with BORIS is presented next.

The BORIS model was applied to generate echo duration versus depth plots for various seabeds and echosounders. For a single seabed/echosounder pair, a linear least fit of echo duration to depth is done resulting in point estimates of effective beamwidth and attenuation calculated according to Equation 5-6. The same seabed is used throughout with the transducer position only changing vertically, not horizontally, as would be done in a field experiment in ideal conditions. This means there is some coherence to be expected between echoes, since adjacent simulations will be produced from very similar seabeds with similar geoacoustic properties and scattering facets. This tends to result in local non-linear trends in the data. One can see this in the 'patchy' appearance of Figure 5-16 and the high levels of auto-correlation of series of echo duration in Figure 5-18.

The pattern of echo-to-echo coherence is due to the presence or absence of scattering facets or variation that significantly changes the statistics of the seabed. The growth of the acoustic footprint as a function of depth also effects this process. The seabed statistics are less stable with a smaller footprint, while echo duration is affected more by the absence or presence of a single scattering facet at the edge of a larger footprint. The effect of coherence and varying seabed statistics together could cause local departures from a linear trend in the echo duration versus depth plots, which explains the appearance of Figure 5-18. In the field experiments to be presented in chapter 6, when conditions were very calm (no heave or drift), these effects were observed. These effects are real and not an artefact of the BORIS model. A third possible cause of non-linear trends in plots of echo duration versus depth is simply the variable nature of high-frequency echoes (due to the relative strength of incoherent backscatter).

These three effects that may cause departures from linearity can be diminished, and the simulations and experiments should be designed with this in mind. In particular, the effect of the variability of echoes is quickly diminished by increasing the number of simulations. The coherence effect will be diminished either by simulating different seabed realizations or by moving the virtual transducer relative to the same realization (i.e. a simulated survey).

These effects that may cause departures from linearity were realized once one set of simulations for each echosounder and seabed pair had been completed – this took approximately one year of computation time. To diminish the effects and improve the results, a second set of data for each echosounder and seabed pair was simulated using a different random number seed so that the stochastically simulated seabeds were different from the first run. The two runs were then combined for analysis. The exception is the simulations with Simrad12, which takes much less computation time than the other scenarios. Instead of doing multiple runs for this echosounder, which may introduce unrealistic variability, a small simulated survey was run for each depth interval to provide a larger ensemble of data. As noted in appendix F, the computation time is strongly dependent on centre frequency and angular extent of the simulation. The Simrad12 has the lowest values of both, so for the worst case scenario (water depth of 50 m, SiltyClay seabed) an echo is generated in less than 10 minutes, while in the same conditions, generating an echo for the Simrad200 takes 30 hours.

In order to conserve computation time, simulating echoes in deeper water was avoided (computation time increases with the third power of water depth). For every echosounder seabed pair, simulations were run at the following water depth intervals: every metre between 2 and 20 m, every 5 m between 20 and 40 m and one simulation at

50 m. These simulations, which include volume scattering, are much more time consuming than the roughness experiments shown in Figure 5-16 and Figure 5-18.

Figure 5-27 demonstrates the resulting echo duration versus depth plots for the two types of simulations as discussed above. The Airmar38 plot (Figure 5-27 (*a*)) shows minor coherence within the two separate runs; it is difficult to see coherence in relatively few samples, but a pattern of consecutive points above or below the trend-line is consistent in many results (as observed more clearly in Figure 5-18). The Simrad12 plot (Figure 5-27 (*b*)) shows a large spread in values due the strength of the incoherent volume scattering component. The horizontal variability of the volume inhomogeneities used for the simulations may be too high. More vertical variability may be required, i.e. more simulated layering.

Table 5-1 and Table 5-2 summarize the effective beamwidths and attenuations measured from BORIS simulations of eight echosounders on five seabed types. For each echosounder / seabed pair, an ordinary least squares fit was applied to echo duration and depth data to obtain regression slope and intercept values, and effective beamwidth and attenuation results that follow from Equation 5-6.

To avoid duplication, a detailed examination of the analysis of data residuals of echo duration versus depth for required conditions on normality, zero-mean, homoscedasticity and correlation, is left to section 6.2. Of the forty cases considered, six did not pass the required conditions for the data residuals. These are indicated with asterisks on the effective beamwidth and attenuation point estimates in Table 5-1 and Table 5-2. Of the cases that did pass, their respective 95% confidence intervals / widths are considered valid; otherwise, these values are not shown.

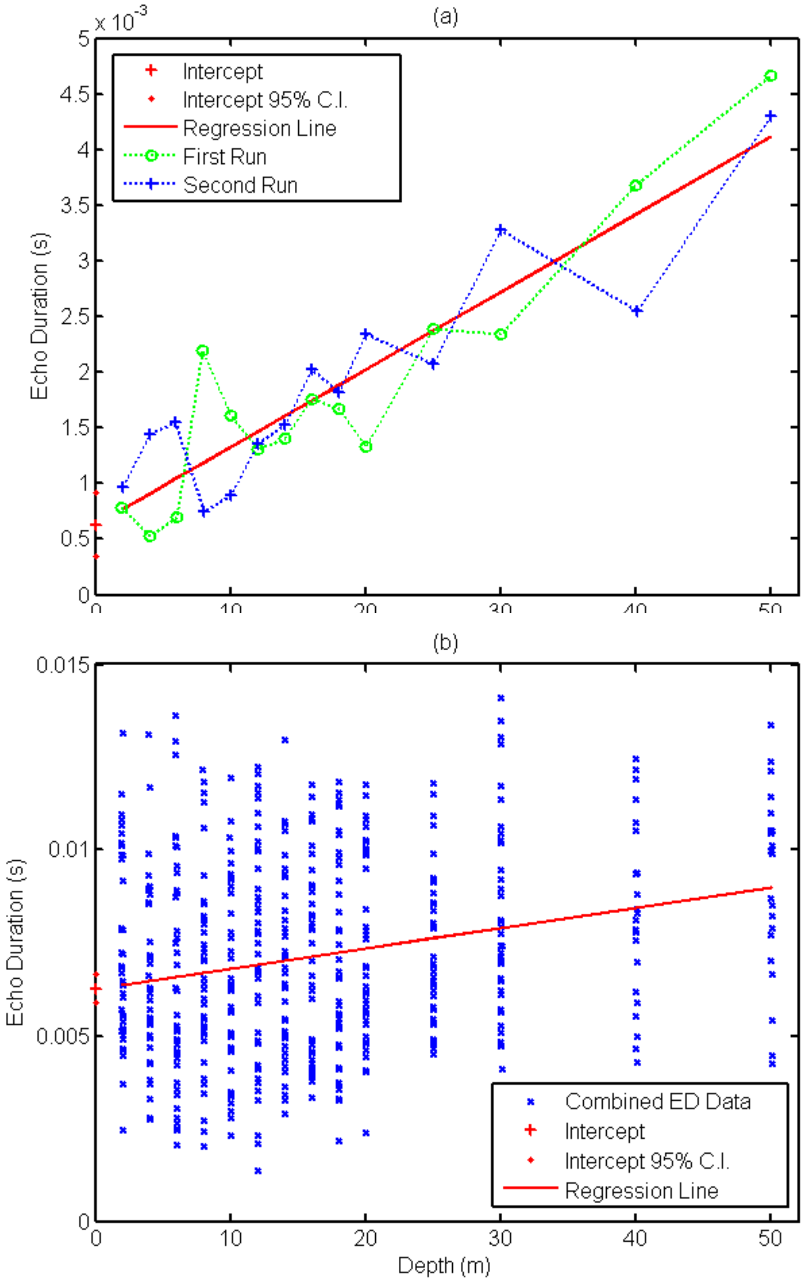

Figure 5-27. Examples of linear regressions of echo duration to water depth from BORIS simulations of: (*a*) the Airmar38 SBES on the Sand seabed and (*b*) the Simrad12 SBES on the SiltyClay seabed. See Table C-1 for echosounder parameters and Table C-3 for seabed parameters.

The 95% confidence intervals for the effective beamwidth and attenuation note above were calculated in the following way. 10^7 normally-distributed[34] random values (based on the means and standard deviations of m and b from the regression) were input to Equation 5-6 to generate sample distributions for θ and κ, from which 95% confidence intervals were measured as the shortest interval containing 95% of the sample distributions. The output sample distributions would be skewed or even become multi-modal if Equation 5-6 behaved non-linearly or there were a singularity within the values of m and b. The singular points of Equation 5-6 are $m = 0$ and $b = \tau$. For most of the results in Table 5-1 and a few in Table 5-2, the output distributions were approximately normal distributions that had confidence intervals matching the confidence intervals calculated in the standard way using t-statistics (n=28) and the propagation of errors formula: $\delta\theta = \frac{d\theta}{dm}\delta m$, $\delta\kappa = \frac{d\kappa}{db}\delta b$. These are the symmetric confidence *widths* marked with a ± symbol in the tables. Otherwise, the confidence *intervals* are listed.

For Table 5-2, the singular point, $b = \tau$, was encountered in many results – these results have absolute values of means and confidence intervals greater than unity as indicated in grey in the table. These values are discounted for the following reason: the CE tail pick threshold[35] used was too low, causing the echo durations to be underestimated. This results in values of b that are close to, and even less than the pulse duration, τ. If b is close to τ, the larger the absolute value of the effective attenuation, and if b is less than τ, effective attenuation becomes negative. Physically, the duration of an echo cannot be less than the pulse duration, and attenuations greater than 1 dB/kHz/m are not found in natural sediments. Note that the tail pick threshold impacts the data as an offset, so the slope, m, and hence the effective beamwidths are not affected.

Effective beamwidth estimates are affected by a bias of the CE tail picking algorithm to echo shape, which is described in detail in appendix A. The effect is difficult to quantify, is known to make effective beamwidth estimates larger for sand and gravel seabeds, and is resolved in the next section with an improved tail picking algorithm.

---

[34] Estimators on m and b are normally-distributed if the data residuals are normally-distributed. This is a concern for the six cases marked with asterisks.

[35] The CE tail pick threshold used was 99%. See appendix A for complete details.

Because of the tail picking errors, the results of these BORIS simulations are considered to be preliminary. It must be emphasized that even with these errors, it is beneficial to show these results as they demonstrate the effects beam-widening, the measurement of effective beamwidth and attenuation and offer exploratory insights on the design, analysis procedure and expected outcomes of the field experiment to be presented in chapter 6. These BORIS simulations will be re-analyzed with cumulative amplitude tail picks with appropriate thresholds for future work. Results would also be improved with additional data to reduce the size of the confidence intervals to avoid the singularities.

Table 5-1. Measurements of effective beamwidth (EBW) compared to model roughness (lower row) and beamwidth (far right column) for simulated seabeds and SBES using BORIS. Point estimates are included for all cases. 95% confidence interval or widths are determined as noted in the text and are valid for cases where the ordinary least squares linear regression model is appropriate (no asterisk).

SBES:		Seabed: Gravel	Sand	SiltySand	MudShells	SiltyClay	-6 dB Beamwidth (°):
Simrad12	EBW (°)	**56.0**	**39.3**	**32.5**	**48.4**	**33.3**	21.8
	95%-CI/W (°)	*±0.7*	*±1.8*	*±2.1*	*±1.3*	*±5.5*	
	R^2	0.69	0.77	0.73	0.71	0.09	
Odom24	EBW (°)	**54.8**	**38.2**	**39.1**	**39.3**	**31.5**	27.7
	95%-CI/W (°)	*±3.0*	*±6.0*	*±4.4*	*±5.4*	*±5.0*	
	R^2	0.94	0.79	0.88	0.83	0.79	
Simrad38	EBW (°)	**24.1***	**21.4**	**14.9**	**26.8**	**15.9**	11.2
	95%-CI/W (°)	-	*±8.8*	*7.2,21.3*	*±6.8*	*4.9,27.1*	
	R^2	0.48	0.43	0.43	0.61	0.18	
Airmar38	EBW (°)	**42.3**	**36.2**	**37.4**	**33.3**	**29.9**	24.0
	95%-CI/W (°)	*±11.9*	*±4.8*	*±3.2*	*±15.8*	*±4.0*	
	R^2	0.50	0.84	0.92	0.48	0.84	
Circ50_20	EBW (°)	**41.0**	**37.2**	**39.1**	**34.1**	**30.6**	27.1
	95%-CI/W (°)	*±11.3*	*±7.6*	*±5.8*	*±8.5*	*±7.9*	
	R^2	0.63	0.80	0.87	0.86	0.83	
Deep50	EBW (°)	**41.6**	**38.0**	**37.1**	**31.9**	**30.2**	25.6
	95%-CI/W (°)	*±8.9*	*±8.6*	*±5.4*	*±7.2*	*±6.5*	
	R^2	0.74	0.79	0.91	0.84	0.82	
Odom200	EBW (°)	**23.6***	**16.5***	**15.3***	**14.9**	**11.1**	5.1
	95%-CI/W (°)	-	-	-	*5.8,22.6*	*3.4,20.0*	
	R^2	0.54	0.55	0.39	0.28	0.08	
Simrad200	EBW (°)	**18.5***	**13.6***	**16.6**	**14.0**	**7.6**	11.0
	95%-CI/W (°)	-	-	*±3.2*	*±2.9*	*2.2,14.9*	
	R^2	0.69	0.77	0.73	0.71	0.09	
Model Roughness (cm RMS):		2.8	2.0	0.7	2.0	1.0	

Table 5-2. Measurements of effective attenuation (EA) compared to the known values for several BORIS simulated seabeds. Entries in grey have absolute values greater than one, indicating these values are approaching the singular point for *b* in Equation 5-6 and are not valid. 95% confidence intervals or widths are determined as noted above. All entries have units of dB/kHz/m.

SBES:		Seabed: Gravel	Sand	SiltySand	MudShells	SiltyClay
Simrad12	EA	**-5.13**	**-6.09**	**-12.62**	**2.15**	**0.21**
	95%-CI/W	*-7.70,-2.95*	*-9.34,-2.97*	*-103.82,-4.05*	*1.56,2.98*	*±0.02*
Odom24	EA	**2.37**	**3.68**	**-2.94**	**2.62**	**-18.51**
	95%-CI/W	*-5.00,8.65*	*-7.06,8.00*	*-8.76,6.12*	*-6.19,8.27*	*-44.31,42.77*
Simrad38	EA	**1.25***	**0.99**	**1.54**	**3.34**	**0.81**
	95%-CI/W	-	*0.38,3.45*	*0.66,4.20*	*-6.78,7.71*	*-3.63,4.92*
Airmar38	EA	**0.38**	**0.99**	**7.07**	**0.35**	**3.42**
	95%-CI/W	*-2.50,2.94*	*0.31,5.55*	*-7.61,8.53*	*0.11,2.29*	*-6.73,8.65*
Circ50_20	EA	**0.31**	**0.52**	**0.51**	**0.27**	**0.21**
	95%-CI/W	*-2.06,2.42*	*±0.28*	*±0.39*	*0.08,2.29*	*0.06,0.38*
Deep50	EA	**0.33**	**0.45**	**0.49**	**0.43**	**0.07**
	95%-CI/W	*-1.47,2.34*	*±0.24*	*0.23,1.23*	*0.14,2.28*	*-0.54,0.62*
Odom200	EA	**0.24***	**0.33***	**0.28***	**0.22**	**0.11**
	95%-CI/W	-	-	-	*0.09,0.69*	*0.06,0.20*
Simrad200	EA	**0.66***	**0.63***	**0.84**	**0.69**	**0.12**
	95%-CI/W	-	-	*-4.04,4.88*	*0.27,2.33*	*0.07,0.20*
Model Attenuation:		0.46	0.49	0.70	0.14	0.06

In a few cases, the variability of echo duration changes with depth. This is due to the changing size of the acoustic footprint. For surface backscatter dominated echoes, as the footprint gets larger, the range of echo duration that can be observed increases. The effect is the opposite for volume scattering as observed for seabeds with strong volume backscatter (Figure 5-27 (*b*) shows this to a small degree as does a result in chapter 6: Figure 6-12 (*e*)). As the footprint area increases, the variability of volume inhomogeneities tends to decrease as the effects of individual scatterers 'average out' over a larger footprint volume. Fortunately, cases of the variability of echo duration changing with depth are few – they occur for high frequencies on gravel seabeds and weakly for low frequencies on silt/clay/mud seabeds. These are the cases marked in Table 5-1 and Table 5-2. For these cases, the procedure of weighted least squares or another maximum likelihood method should be used to account for the changing variance in the data[36]. However, an objective estimate of the variance as a function of depth is difficult to achieve as the only estimate available is from the data itself. In cases where the

---

[36] Ordinary least squares is applicable for non-homoscedastic data, it is just not optimal.

variability of the data varies with depth, the above effects also tend to cause the distribution of the data residuals to become skewed and deviate from normality[37].

Departures from linearity due to coherence are generally limited in their impact. While changing statistics or random variation and departures from homoscedasticity due to footprint size are important to consider, they only appear in a few cases. The consideration of these issues is examined further in section 6.2.

Now that all the issues of this analysis have been presented, the results are more readily discussed. The linear regressions of echo duration to depth generally represent a large fraction of the variance of the data as shown by the R^2 (coefficient of determination) values in Table 5-1. Low R^2 values (less than 0.3) do not indicate that a linear fit is not appropriate; instead, they indicate the linear trend in the data is a small fraction of the total variance. Low R^2 values tend to occur for narrower beamwidth cases (Simrad38, the Odom200 and the Simrad200) mainly because narrower beamwidths decrease the slope of the regression line relating echo duration to depth. Low R^2 values are also symptoms of the effects that diminish linearity: coherence and changing statistics of the footprint as described earlier. These effects are also expected to be more prevalent in the narrower beams. Simulations for the SiltySand and the MudShells seabeds have high variability and moderate coherence (somewhat greater than that seen in Figure 5-27 (*a*)). However, the SiltyClay seabeds with low R^2 exhibit normal variability due to scattering, similar to that seen in Figure 5-27 (*b*). Overall, the R^2 values and qualitative analysis show that echo duration is a linear function of depth, consistent with the EDM.

The next result to discuss is the relationship of effective beamwidth to the measured beamwidth and to the effect of beam-widening. The results of Table 5-1 indicate a strong general beam-widening effect. The effective beamwidths for each echosounder should correspond to the measured beamwidths, and they do except that the beamwidths are much wider than the -3, -4.34 or even the -6 dB beamwidths. Many beamwidths found for the Gravel seabeds correspond to the width of the first sidelobe. The large beamwidths are primarily due to beam-widening, mostly from roughness scattering with some widening in cases of strong volume scattering (as seen in the

---

[37] Although not shown here, these residuals will be shown as part of the analysis of field data in chapter 6.

SiltyClay seabeds). The bias in the cumulative energy tail pick may enlarge effective beamwidths by artificially shortening echoes at shallow depths. However, large effective beamwidths occurred in the Patey Rock and AEM results, and are likely to be seen in the field experiment in section 6.2.

The correspondence between the effective attenuation values and the expected values can be observed in Table 5-2. The valid results (in black) are generally consistent with the expected values. As discussed earlier, erroneously large negative or positive effective attenuation values (shown in grey in Table 5-2) indicate regression intercepts near the pulse duration so that the duration of the echo within the seabed, the penetration time, is not realistic. The confidence intervals reflect this and show large uncertainly in these erroneous effective attenuations. The erroneous attenuation values occur quite often for lower frequencies. It appears the appropriate tail picking threshold level required to produce more accurate echo durations may need to be higher for lower frequencies than for higher frequencies. The aforementioned tail picking bias may also contribute to this by reducing the regression intercepts for lower frequencies with strong coherent components and shallow depths. As discussed in the Patey Rock example, there is no means of determining what the appropriate threshold level would be in order to tail-pick the time within the echo when the volume component decreases by 10 dB, matching the extinction model of Equation 5-6. Then, a method of relating the picking threshold level to the EDM must be devised. Such a method will be presented in chapter 6, based on the findings of this sub-section and appendix A, so that the final tail picking result is also sensitive, consistent and robust to noise and echo shape.

Overall, the simulation also demonstrates the effect of beam-widening and the measurement of effective beamwidth and attenuation for use in depth compensation parameters. The reanalysis of this BORIS data is left for future work, as the main goals of validating the EDM and exploring the relationship between echo duration and depth were achieved.

5.4.4 Conclusion: Physics-Bases Features and Effective Depth Compensation

In many approaches to acoustic seabed classification, the extraction of information from echo time series is phenomenological and not based on the physics of echosounding. Statistical techniques such as PCA or neural networks are relied upon to select the features that provide discriminatory information. The feature algorithms in the original QTC seabed classification method are one of the best examples [Caughey, 1996]. As discussed in appendix B, features that have no discriminatory information, or worse contain noise, tend to degrade the performance of PCA / cluster analysis. Seabed classification by statistical segmentation may be improved by applying a physical basis to justify existing features, and, more importantly, to develop new physics-based features that offer improved discrimination.

In short, the physical basis is contained in chapters 4 and 5; the combination of scattering theory and the new models shows how echoes respond to roughness (beam-widening, relative strength of echo components), impedance (scales overall amplitude), and volume scattering (increases echo duration); for a summary see Table 4-1. Instead of studying the effects of new features on the results of seabed classification, it is easier to study how the geoacoustic parameters affect features, which also facilitates improving and devising new features. To that end, a study with trial new features is presented in appendix G, which makes use of the example BORIS simulated survey shown in Figure 5-3. The initial results presented in appendix G are further developed in chapter 7 with the operational set of new features.

In all, this chapter contains models and developments essential to the advancement of physics-based features. Although interesting on their own, these findings are most useful in developing an improved seabed classification method (chapter 7) and a new seabed characterization method (chapter 8). As part of an improved seabed classification method, features must operate on echo time series that are free of non-seabed influences or must be designed to be independent of these influences. To that end, the EDM and the concepts of effective beamwidth and attenuation are developed further in the next chapter, resulting in effective compensation for depth and seabed slope.

Chapter 6. Controlled Experiments in the Patricia Bay Testbed

The objective of this chapter is to verify the model studies of chapter 5 and to provide a basis for the new methods presented in chapters 7 and 8. In particular, the aim is to quantify the effects of depth and slope on echoes in order to compensate for them. This will be done by two controlled field experiments in Patricia Bay. The concept of a controlled experiment in this case is to hold all parameters constant while varying just one, as was done in three simulation experiments in chapter 5. The first experiment will vary depth with the aim of refining the method for estimating effective beamwidth and attenuation for depth compensation of echo time series. The second experiment will vary seabed slope to show the effects on features used for segmentation. A method of compensating for seabed slope is developed and tested on the same data. In order to carry out both experiments a suitable *testbed* must be established. This testbed must have a high diversity of seabed types, yet have homogenous areas that meet specific conditions for the experiments. The testbed is presented in section 6.1, followed by the depth experiment in section 6.2 and the slope experiment in section 6.3. The testbed is also applied to evaluate the new classification methods (chapter 7) and the new characterization methods (chapter 8).

6.1 The Patricia Bay Testbed

The objective of this section is to create an accurate seabed map of the testbed area and then use it to establish testbed sites for the controlled experiments. Both experiments require relatively small (200 m by 200 m) homogenous areas with no bedforms, vegetation or seabed slope. The initial plan for the testbed was outlined in [Biffard et al., 2006]; more complete results are presented here.

The testbed area is Patricia Bay, an 8 km^2 bay with water depths up to 120 m. Its location is shown in Figure 3-1. Patricia Bay is adjacent to the Institute of Ocean Science, it is home to a VENUS[38] cabled observatory, and is only 30 km from the University of Victoria. In addition to being a convenient location, there are data and information available from previous work that can be drawn upon, including that Patricia Bay is known to have to a complex seabed that represents a challenge to standard seabed classification methods[39]. The complexity is due a wide range of depths, seabed slopes and seabed types. This makes Patricia Bay a perfect site to test the main hypothesis of this dissertation – that new and improved seabed classification methods developed here can successfully remove non-seabed influences and improve the discrimination of seabed types. The data available within the testbed consist of:

- EM3000 MBES full bay survey and classification (courtesy of the Canadian Hydrographic Service (CHS))
- Airmar38 / Simrad200 full bay survey (acquired in collaboration with J. Galloway and the CHS)
- Extensive ground-truth data: 203 grab-samples with simultaneous video, of which 178 are expert classified and 45 have grain size analysis, 34 STING profiles and an interpreted side-scan and video seabed class map, found in [Bloomer et al., 2007]
- Odom24 / Odom200 experiment site verification survey
- A Simrad12 pass through the experiment sites
- Odom24 / Odom200 depth experiment
- Odom24 / Odom200 slope experiment

All of the data are presented here in chapter 6 except for the Airmar38 / Simrad200 full bay survey; that data set is presented in section 7.3 and used in chapters 7 and 8 to evaluate the new classification and characterization methods.

---

[38] see [Barnes and Tunnicliffe, 2008] or http://venus.uvic.ca/

[39] Personal communication with K. Rhynas, QTC marine technician and other members of the QTC marine division. Mr. Rhynas in particular advised not to use Patricia Bay because previous SBES survey attempts had poor results. These results will be duplicated and then improved upon in chapter 7.

6.1.1 EM3000 Seabed Classification Survey and the Selection of Experimental Sites

In 2003, Patricia Bay was surveyed with a Simrad EM3000™ echosounder. Bathymetry and seabed classification data were processed using standard methods (Caris software for bathymetry and QTC MULTIVIEW for classification). QTC MULTIVIEW™ produces the same type of segmentation result as does QTC IMPACT for SBES; after categorical interpolation with QTC CLAMS, the classes are interpreted and labeled with ground-truth data. Ground-truth data initially consisted of 129 expert interpreted grab samples with simultaneous video for confirmation. This ground-truth data set was used to interpret the results. Table 6-1 is a ground-truth interpretation table used to verify the seabed type labels applied to the five multibeam classes shown in Figure 6-1. The five seabed types range from rock to mud over a wide range of depths and slopes, with some small flat homogenous patches. Of the 98 grab-samples of Table 6-1 only 4 differ from the multibeam classes. Overall, the agreement is excellent between the ground-truth data and the multibeam map. The diversity of seabed types and range of depths and slopes indicate Patricia Bay meets the requirements for the testbed.

It is interesting to note that the multibeam results are independent of seabed slope. A map of the seabed slope is shown as Figure 6-2. QTC MULTIVIEW corrects the amplitudes of each beam for incident angle and range using a 'compensation table', which accounts for the local slope in addition to providing accurate compensation for transmission losses (i.e. TVG) [Preston and Christney, 2005]. Depth compensation of time series (known as snippets) is not done since MULTIVIEW classification is based on amplitudes not on measures related to time. Due to the method of compensation tables MULTIVIEW classification is free of depth and slope artifacts, unlike standard SBES classification.

The only area of concern in the multibeam seabed map is an area of 'striping' consisting of stripes of gravel classes amongst the predominately mud area, as indicated by the black box in Figure 6-1. This area is curious as there are no obvious patches of gravel in the bathymetry, Figure 6-3, which is very smooth in this area. There is an edge to the mud area that consists of gravel and rock that is visible at the western extreme of the survey (Figure 6-3) where the bay drops into the deeper Saanich inlet. Survey coverage in this area is not complete, although coverage is 100% in the rest of the bay. Inspection of the raw (un-interpolated) class map shows that there are gravel outliers on the edges of the multibeam swaths in this area. Such outliers would normally be interpolated out, but in this case the outliers are interpolated into the spaces between the survey lines. The gravel outliers maybe related to the presence of a crust on the mud. The *Mw* grab sample class (Table 6-1) indicates mud with a white crust (as observed in simultaneous video). The crust is a bacterial mat that covers some of this area. Video observations of ROPOS's arm pushing into the crust confirm its rigidity in some locations. At high angles of incidence (the beams at the edge of the swaths), these crusts may appear much thicker, resembling higher impedance seabeds, so that they would be classified as gravel, causing the confusion. Lower frequency MBES or SBES echosounders will be less affected by the crusts. Analysis in this area will be avoided.

Based on the multibeam classification results, five sites were selected for the experiments; see Figure 6-3 for their designations. Three of these sites have little slope (1±0.3°) and, according to Figure 6-1, are homogenous examples of mud, sand, and gravel. The other two sites have multiple seabed types; the first is not sloped and transitions from fine sand to gravel, while the second is sloped (varying from 4 to 15°) with a complex transitions among the gravel, medium and fine sands.

Table 6-1. A summary of ground-truth data for each multibeam class from Patricia Bay. Grab samples had to be located at least 10 m from any additional multibeam classes. Grab sample labels are based on the ternary classification scheme found in [Folk, 1980], also [Jackson and Richardson, 2009]. Additional information: capital letters indicate the class comprises 50% or more of the sample, C is for cobble, w is a white crust, FS is for fine sand.

Multibeam Class:	Grab Sample Class with Fractional Ocurrence:						Total Number of Samples:
Gravel	R/cR	G/sG/sC	gS	gmS	gmS	mS	
	0.19	0.34	0.13	0.16	0.16	0.03	32
Med. Sand	gS	gmS	S	mS	MS		
	0.08	0.04	0.58	0.29	0.04		25
Fine Sand	FS	mS	MS				
	0.38	0.54	0.08				24
Sandy Mud	mS	gmS					
	0.50	0.50					2
Mud	Mw	M	sM				
	0.69	0.19	0.13				16

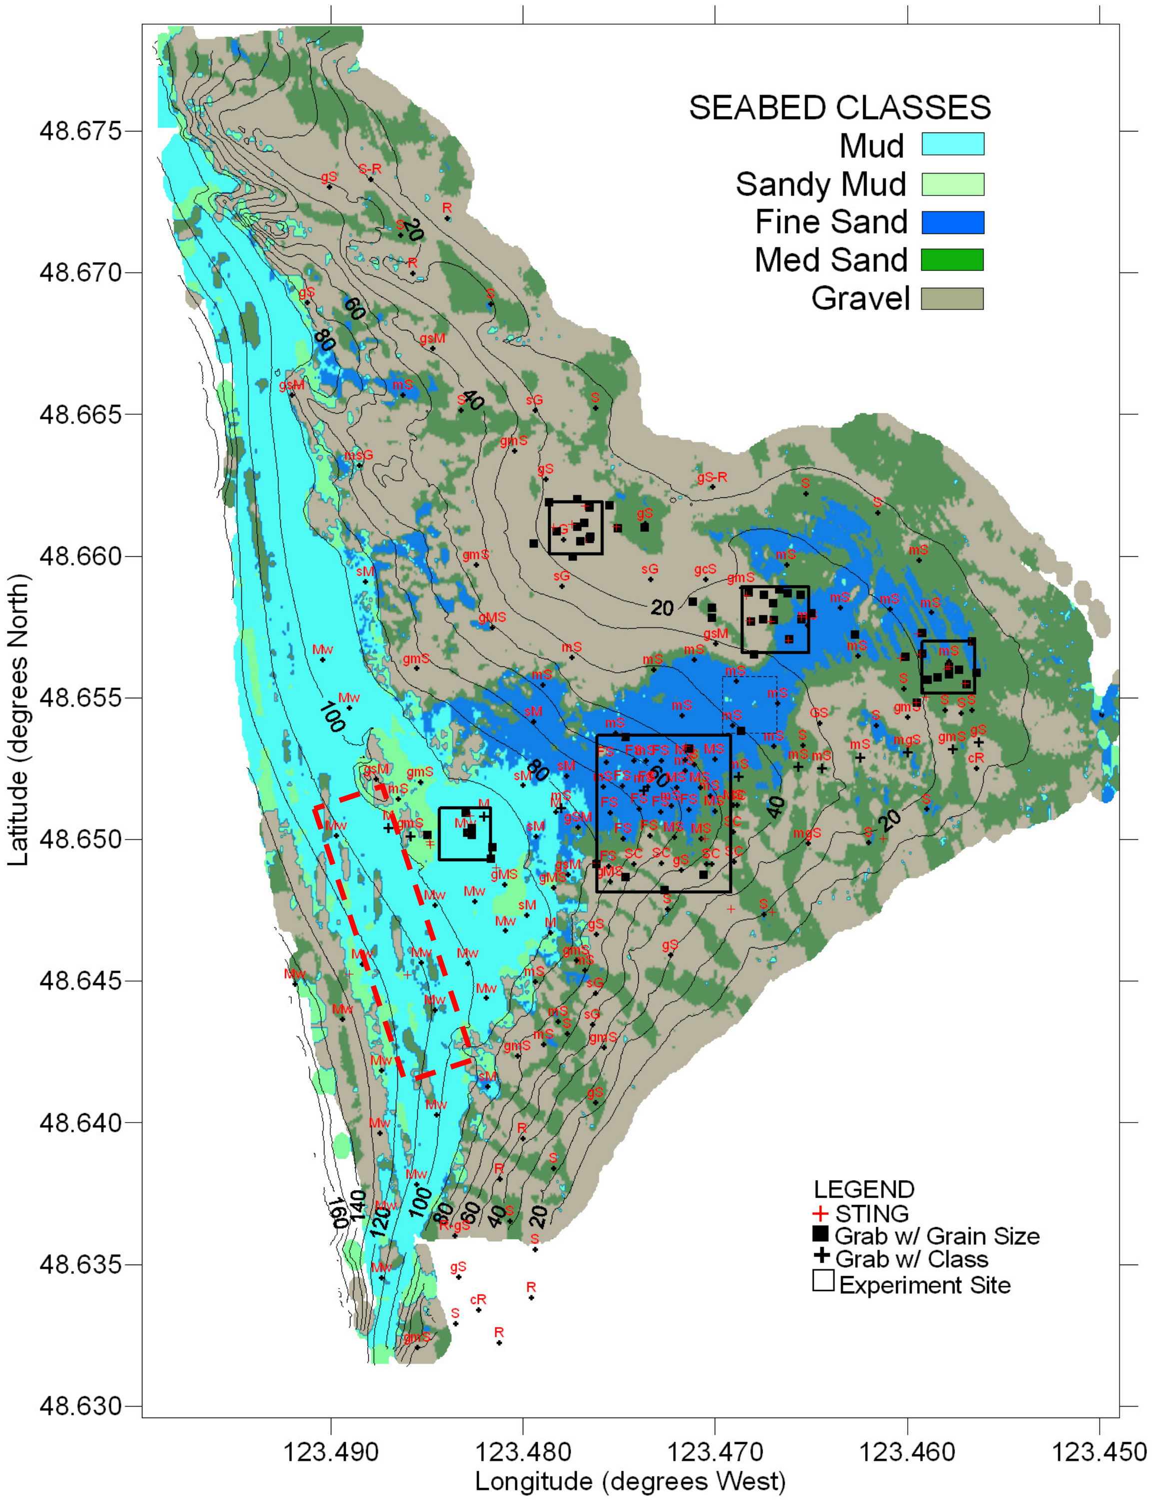

Figure 6-1. The MBES seabed classification map of Patricia Bay, B.C. processed from Simrad EM3000 survey data with QTC MULTIVIEW. The map is overlaid with depth contours, ground-truth data and the boundaries of the experiment sites (the lightly dotted boundary is a training site used later in chapter 7). The red dashed-line box highlights an area of concern.

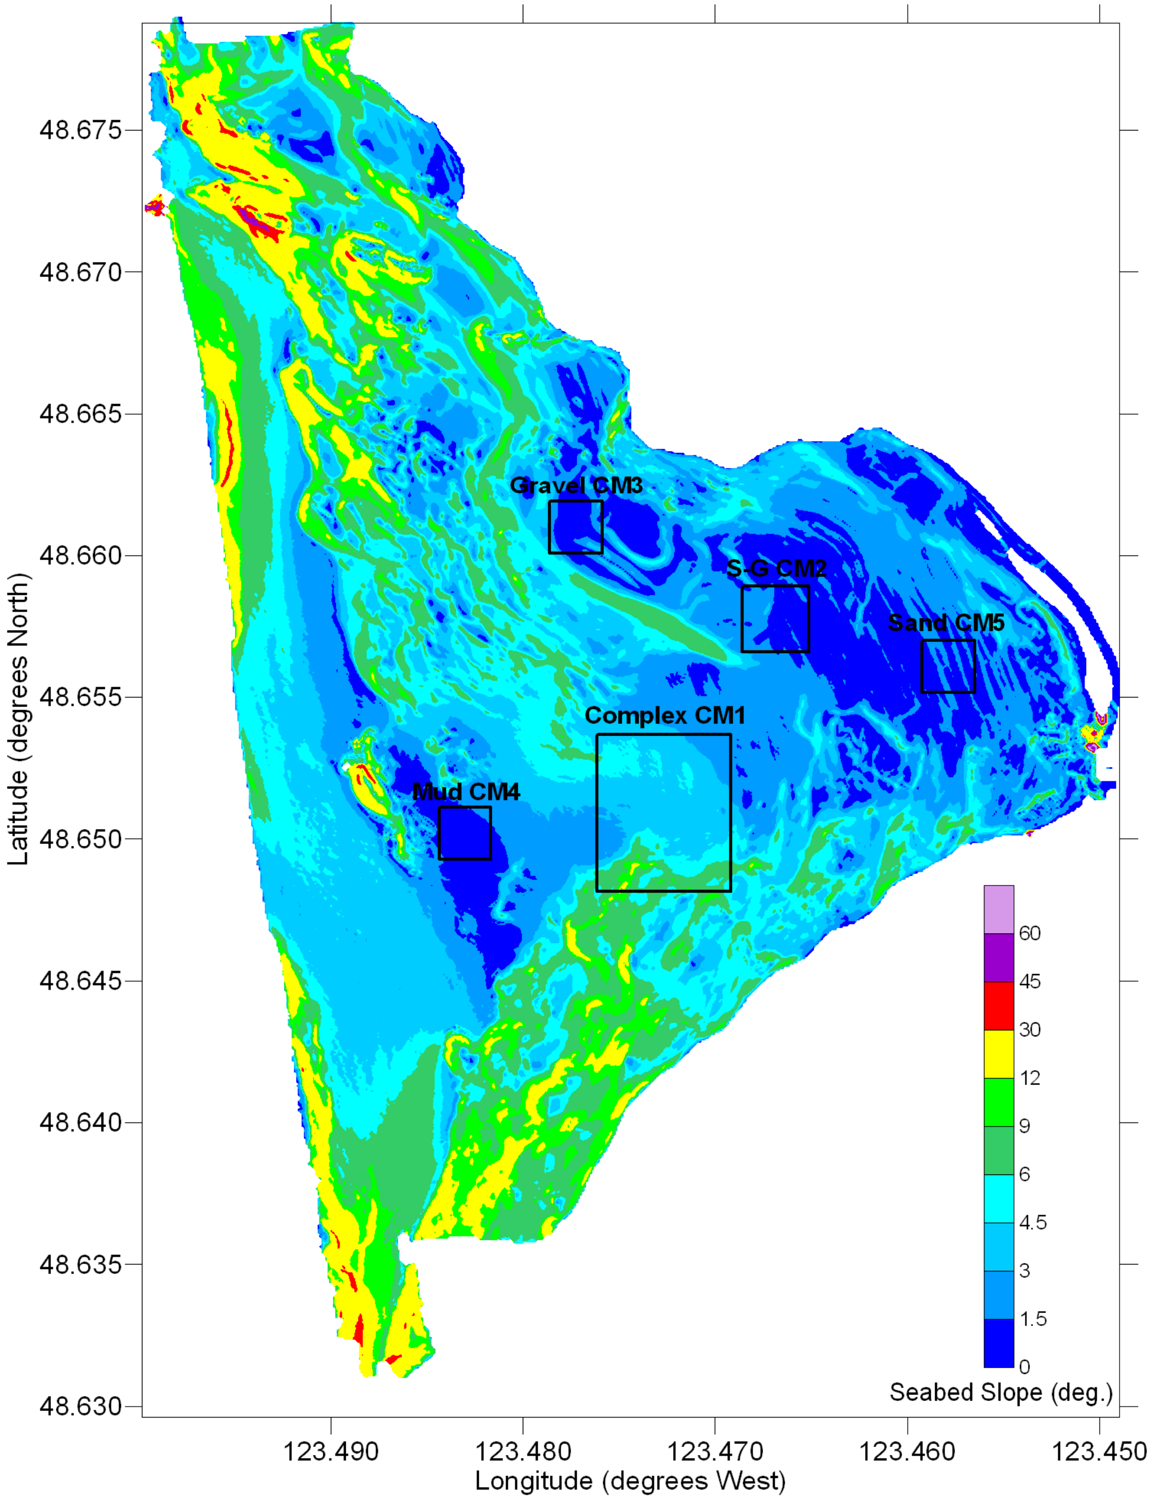

Figure 6-2. Seabed slope, as extracted from EM3000 multibeam bathymetry for Patricia Bay, overlaid with the boundaries and names of the experiment sites.

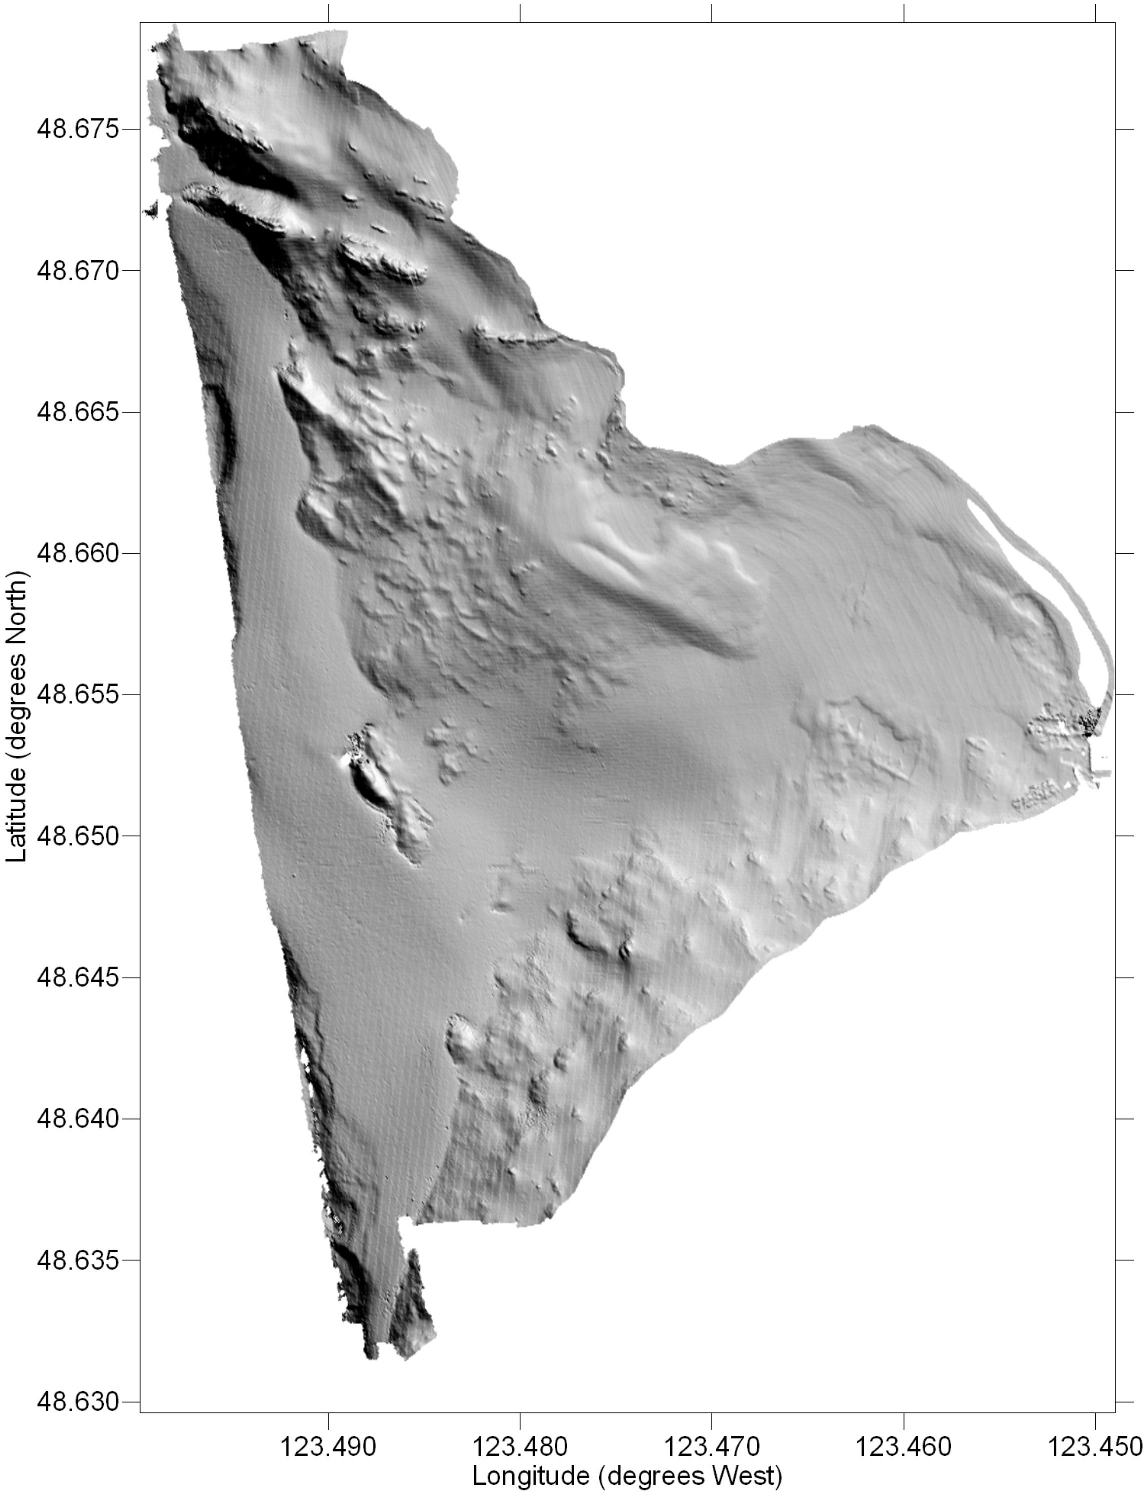

Figure 6-3. Sun-illuminated bathymetry from the EM3000 multibeam survey of Patricia Bay.

6.1.2 Experimental Site Seabed Classification Survey

To confirm the homogeneity or diversity of the experiment sites, all sites were surveyed with the Odom24 and Odom200 SBES for seabed classification via QTC IMPACT. A regular grid of 20 m line-spacing was run. However, the resulting classification using standard methods failed to separate the gravel and sand classes. Instead the OCPR results are shown as Figure 6-4. The OCPR results confirm the homogeneity of the mud, sand and gravel sites, and the transitions in the sand-gravel site and the complex site. However, the failure of standard classification methods (for both frequencies) is concerning. The standard methods appear to suffer from depth and slope artifacts and fail to provide adequate discrimination, as was observed in other surveys shown in chapter 3. In the next section, an extensive ground-truth survey was undertaken to confirm the quality of the sites selected and determine the geoacoustic parameters of the sites.

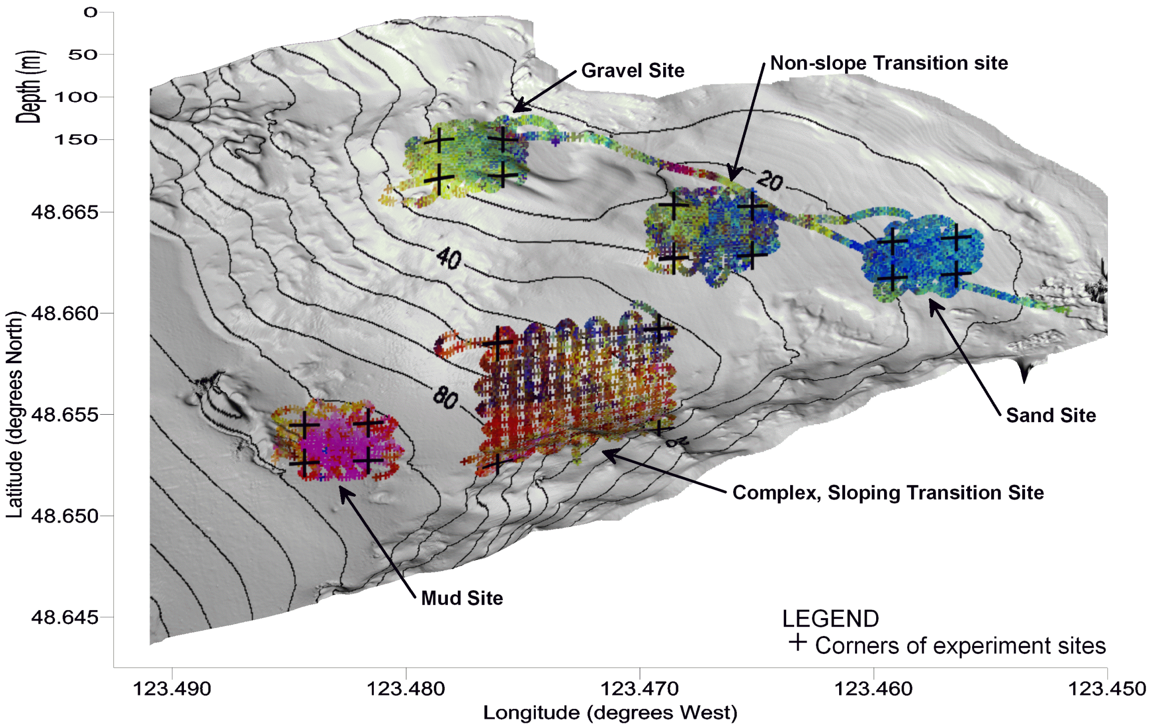

Figure 6-4. Seabed classification result for the Odom24 survey of the experiment sites – displayed as OCPR overlaid on the 3-D surface based on the multibeam bathymetry (2 m resolution).

6.1.3 Experiment Site Ground-truth survey

The ground-truth survey consisted of grab samples, video and STING penetrometer data collected thorough 2005 and 2006 aboard the vessels *Otter Bay* (CHS), *Strickland* (UVIC) and CGR's survey boat. The specific data collection and interpretation methods are described first, followed by a discussion of results for each of the five sites.

Grab sample data collection and analysis was done with the same procedure to that of the CHS Patricia Bay survey of chapter. A Shipek™ grab sampler and grab-mounted video system was deployed onboard the CCGS *Otter Bay*. A slightly smaller, yet safer Ponar™ grab sampler and grab-mounted video system was employed aboard the RV *Strickland*. In addition to the video, visual and tactile expert classification of the grab samples, a dry sieve analysis was completed. Samples were kiln-dried and then broken into loose sediment if cemented during drying. The grain-size distributions of samples were measured by falling through a series of seven sieves of decreasing size in a vibrating stack. The sieves were weighted to determine the fraction of the samples they contained.

The STING free-fall penetrometer was deployed from the CGR boat and the MV Strickland. The advantages of the STING are: it samples the seabed in situ, profiles the upper 0.1 to 2 m of seabed, (much deeper than a grab sampler), and provides a direct measure of seabed bearing strength in Newtons per square metre. The STING is a seabed characterization instrument in that it provides a direct estimate of a seabed property that is used to estimate the seabed type. The accelerometer at the heart of the STING is calibrated by gravity. The bearing strength measurements are calibrated for the foot size and drag (only important during penetration). The nature of the flow of the sediment material around the penetrating foot adds uncertainty, especially when using different diameter feet [Mulhearn, 2002]. An advantage of the STING is that it is deployed by hand from any vessel, while other penetrometers require a winch [Osler et al., 2005]. The version used here has an autonomous recording time of 50 seconds, which is enough time for five impacts. After the first impact, the STING is pulled up ten metres then released to free-fall to each addition impact. Upon retrieval, the data are downloaded, the deceleration curves are picked manually and the STING software calculates the bearing strength as a function of depth. Bearing strength values are then interpreted for a seabed type or sediment class using a lookup table provided with the STING operator's manual [DREP, 2001]. Variables such as porosity cause uncertainly in the relationship between bearing strength and seabed type/grain size; this is accounted for with error bars.

The analysis of ground-truth data as follows, listed by site:

Gravel Site – CM3

From the seven sieve-analyzed grab samples within the gravel site (see Table 6-2), the mean grain size is greater than 3.8 mm (using the moment calculation and mid-range size values). In phi units, the upper limit to the mean grain size is -1.92 Φ. Only the upper limit can be found, as 79% of the sample by weight had a diameter greater than the largest size sieve of 4 mm. Grains of size 4 mm and larger are pebbles through gravels, cobbles, boulders, etc. The upper limit (in phi units) was found by assuming the mean diameter of the fraction greater than the 4 mm sieve size was 4.5 mm. Visual inspection of the samples confirm the classification of a sandy gravel for this site. The samples within the site had little variability. STING measurements showed high relative bearing strength with no layering present or observable. The average penetration depth was 23 cm. Bearing strengths ranged from 400 to 1000 kN/m2 corresponding to medium dense to dense gravel [DREP, 2001]. Video records also confirm these findings with observations of pebbles, cobbles and some sand. Occasional macro-algae were observed - estimated to be less than 5% coverage. Occasional wood debris was observed (one ~1 m branch was seen). Of interest is a sandy area immediately to the East of the site. This area appears on both the multibeam and single-beam classification maps, and as a shallow depression in bathymetric maps. A STING sounding in this area resulted in a deeper 50 cm penetration with bearing strengths characteristic of muddy sand. Only one sounding was taken as the STING became stuck on the first impact and required much effort to retrieve it. Upon retrieval, marine clay was observed on the STING foot and gas bubbles were observed in the water. This sandy clay area adjacent to the gravel site was avoided in the acoustic experiments. Overall, the gravel site is a very homogenous sandy gravel seabed, which is likely glacial till. It falls in between the gravel and sandy gravel seabed types in the APL-UW geoacoustic parameters table (Table C-2).

Table 6-2. Grain size analysis for the Gravel site (CM3), listed by grab location within the site. 'cm' grabs are Shipek grabs, 'fs' grabs are Ponar grabs.

Grain Size (μm)	Class	% sample fraction, by location:							
		cm3-1	cm3-2	cm3-3	cm3-7	fs3-1	fs2-1	fs2-3	Mean
>4000	pebbles	67.0	74.3	54.5	80.3	94.6	93.8	88.3	79.0
2000	granules	6.9	7.7	8.5	3.6	1.6	0.4	0.8	4.2
1000	v. coarse sand	8.1	4.7	11.6	3.2	1.2	0.9	0.9	4.4
500	coarse sand	6.3	3.6	13.8	3.0	1	1.6	1.5	4.4
250	medium sand	5.2	3.7	6.3	4.3	0.6	1.4	2.7	3.5
125	fine sand	4.4	3.6	3.6	3.7	0.6	1.1	3.7	3.0
62.5	v. fine sand	1.6	1.5	1.4	1.3	0.3	0.5	1.6	1.2
<63	silt/clay	0.6	0.7	0.2	0.4	0.1	0.2	0.5	0.4

Table 6-3. Grain size analysis for the Sand site (CM5), listed by grab location within the site.

Grain Size (μm)	Class	% sample fraction, by location:					
		cm5-1	cm5-7	fs2-6	fs3-4	fs3-3	Mean
>4000	pebbles	0.0	0.3	0.1	0	0.1	0.1
2000	granules	0.0	0.2	0	0	0	0.0
1000	v. coarse sand	0.2	1.6	0.6	0.4	0.4	0.6
500	coarse sand	2.3	13.9	3.1	3.2	3.7	5.2
250	medium sand	13.6	35.7	30.7	37.7	31.6	29.9
125	fine sand	64.1	40.2	54.3	48.4	51.8	51.8
62.5	v. fine sand	17.0	6.6	8.8	8.6	9.9	10.2
<63	silt/clay	2.8	1.6	2.2	1.7	2.8	2.2

Sand Site – CM5

Over the five sieve-analyzed grab samples within the sand site (see Table 6-3), the mean grain size was found to be 269 μm or 1.89 *Φ*. Referring to the APL-UW geoacoustic parameters table (Table C-2) this is finer than a medium sand, but is not a muddy gravel, it is between a medium sand and a fine sand. Most of the sample (51.8%) was caught by the fine sand sieve, followed by the medium sand sieve, supporting the classification of the site as a fine sand. Only 2.3% of the collective sample (by weight) was either above or beyond the size range of the sieves, therefore this grain size measurement is considered accurate. Visual inspection of the samples classified the samples as muddy sand (which is nearly the same mean grain size). Muddy sand is sand with a minimum 10% fines component, easily mistaken for very fine sand. STING measurements showed moderate relative bearing strength and some layering - the bearing strength tended to increase after 15 cm of penetration. The average penetration depth was 42 cm. Bearing strengths averaged around 200 kN/m^2 for the first 15 cm of penetration - this corresponds to medium dense sand or stiff clay. After 15 cm of penetration, the bearing strength averaged around 600 kN/m^2 indicating either a dense sand or medium gravel. Video records confirm the classification as fine sand. Sea pens and some macro-algae were observed covering less than 5% of the surface. Some bioturbation was observed - some burrows, but only one or two per square metre. Within the sediment samples, few shells or organisms were observed. The multibeam classification result, Figure 6-1, classifies roughly 85% of the site as medium sand with some fine sand in North-West corner of the site. A grab sample in this area, cm5-1, confirms a subtle shift corresponding to the multibeam classification - the shift consists of moving about 15% of the sample from medium to very fine sand, with the proportion of fine sand remaining the same (~50%). The fine sand areas should be avoided to improve homogeneity within the site.

Mud Site – CM4

Six grab samples were taken in the mud site and with eight grab samples taken in the vicinity. All six grab samples were completely too fine and passed through sieves. The mean grain size is less than 63 μm or greater than 4.0 in phi units. No sand or gravel was present. Visual inspection of the samples finds the sediment as cohesive 'ooze', with extremely fine grains. Such sediments or materials are often classified as a colloid. When dried the sediment cements together and when ground the sediment has the consistency of dust. This fits the finest classes on the grain-size scale and the corresponding classes in the APL-UW geoacoustic parameters table (Table C-2), from 'very fine silt' to 'clay' with a mean grain size of 7.8 μm to 1 μm (7.5 to 9.0 in phi units). The STING penetrated more than the 2 m length of the extended leg with the standard 3.5 cm (diameter) foot. This made it difficult to retrieve the probe. An attempt with a larger 5 cm diameter foot was made but the result was the same, except that the probe was nearly impossible to retrieve. Average bearing strength was 50 kN/m^2, corresponding to soft clays and silts. Video records confirm a flat feature-less sediment with little life or bioturbation. Occasional depressions made by a flatfish (possibly Dover Sole) were observed. A couple of shells were found in the six samples along with very thin layers of slightly darker material. As mentioned earlier, a white crust was observed on top of the mud, but only in waters deeper than the mud site at 100 m. Multibeam classification shows this to be a homogenous mud area and does not distinguish the crusty areas from the non-crusty areas. Our research group has two possible theories on the formation of this unique sediment: a drowned beach, although the lowest sea level in this area was likely only 65 m below present[40]. Or, more likely, it is a mixture of river silt and organic colloids forming laminae in the presence of anoxic bottom water that is formed by a sill (rising to 70 m water depth) at the entrance of Saanich Inlet. Anoxic conditions persist year round with the boundary varying around 90 m [Gucluer and Gross, 1964].

Sand - Gravel Transition Site – CM2

Twelve grab samples were taken within the sand-gravel transition site (CM 2) and analyzed for grain size distribution. Grain size distributions are shown as Figure 6-6, indicating how the sediments grew coarser moving from east to west over the transition from fine sand to sandy gravel. None of the grabs from the fine sand and medium sand classes had any gravel component. Two of the grabs from the sandy gravel class had low levels of gravel (~10%) – this could be due to sampling error, positioning error or is accurate. Five STING measures show a strong progression from moderate bearing strength to high bearing strength. STING profiles also show how the gravel basement dips below the sand through the transition. In the STING profile shown in Figure 6-7, the gravel basement lays below 16 cm of overlaying fine to medium sand. The gravel basement was confirmed with observations from a pass through the site with a Simrad12 SBES onboard the CCGS Tully. The pass traversed the site east to west through the middle, and the echoes were viewed in the standard echogram display of the Simrad EA400 software and Myriax EchoView™ software[41].

---

[40] Personal communication, Dr. Brian Bornhold, Adjunct Professor, University of Victoria

[41] For information on this software, see http://www.km.kongsberg.com/ and http://www.echoview.com/, respectively.

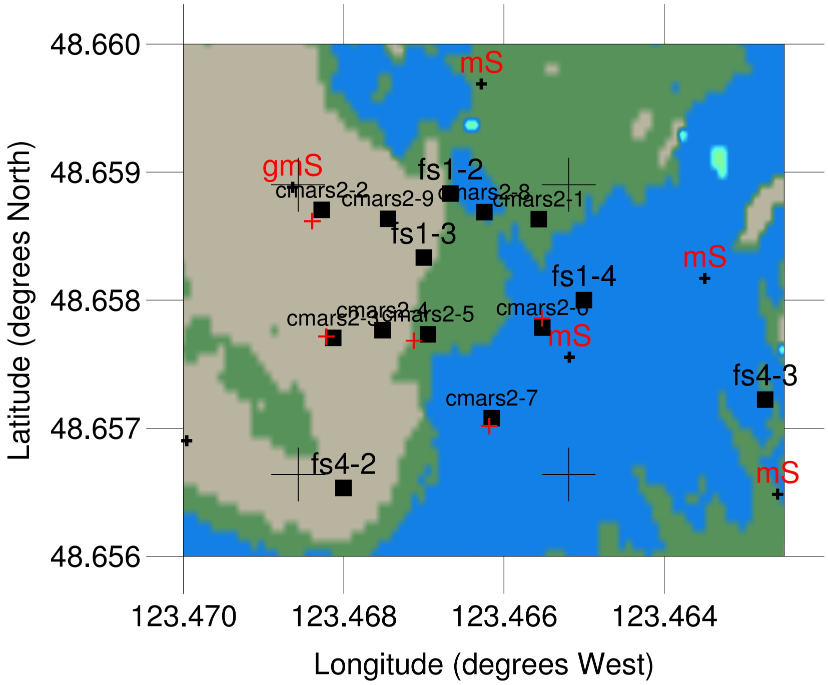

Figure 6-5. An enlargement of the Patricia Bay multibeam seabed classification map for the sand - gravel transition site (CM2). Class colours and ground-truth are the same as Figure 6-1. The four large black crosses are the four corners of the testbed site. Sieved grab samples are the black squares.

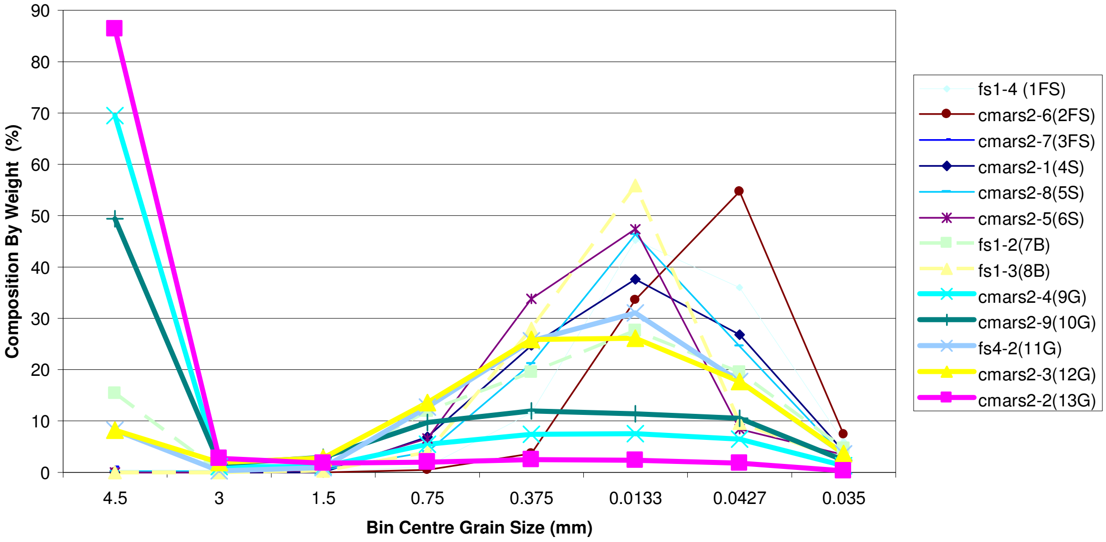

Figure 6-6. Grain size distributions for the Sand - Gravel Transition Site (CM2). The legend denotes, in brackets, the location of each grab ranked by distance from the transition, while the letters FS, S, B, G stand for the multibeam class at the grabs' location (B indicates the grab was on the border). Grabs from the gravel areas are shown in thick lines, those from the sand areas are shown as thin lines and those immediately on the border are dashed thick lines.

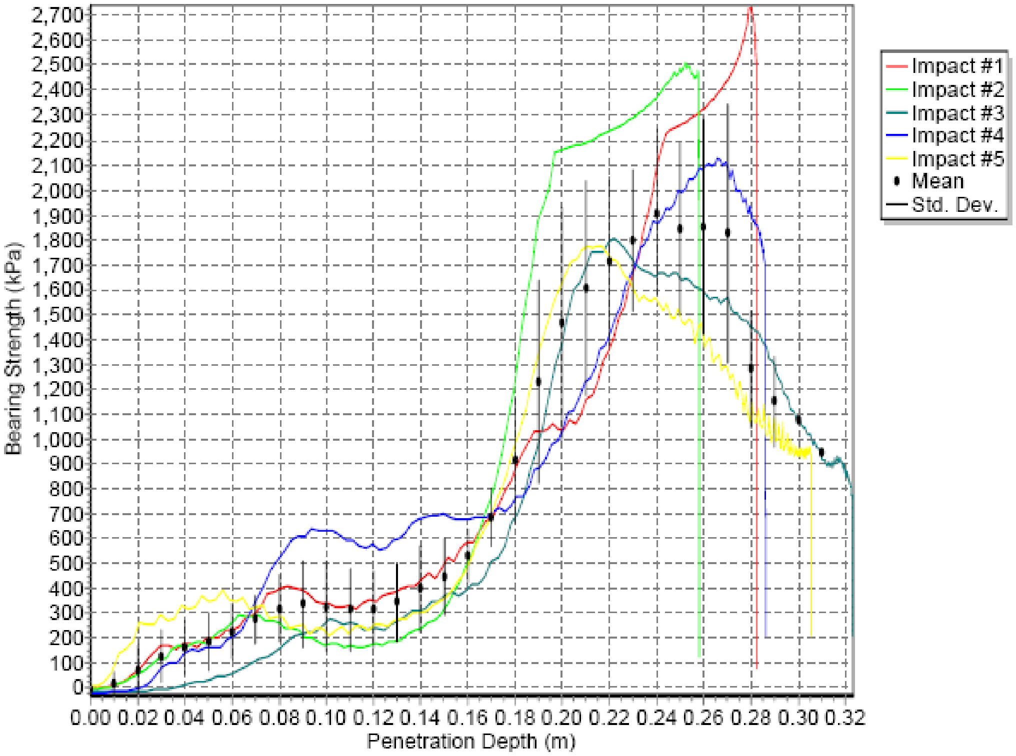

Figure 6-7. The STING sounding result at location CMARS 2-5.

Complex Sloping Site – CM 1

This site was originally a CHS sonar systems group standard catalogue site. As such, additional survey and ground-truth data exists for this area. This site has depth and slope variation and three major seabed classes, all in a small area. Grab samples, STING measurements and video data all confirm the multibeam classification of Figure 6-1. An enlargement of the classification for complex site is shown as Figure 6-8 with more ground-truth detail. The grab samples collected in the recent fieldwork analyzed by dry sieve confirm the previous grabs that were classified by expert interpretation.

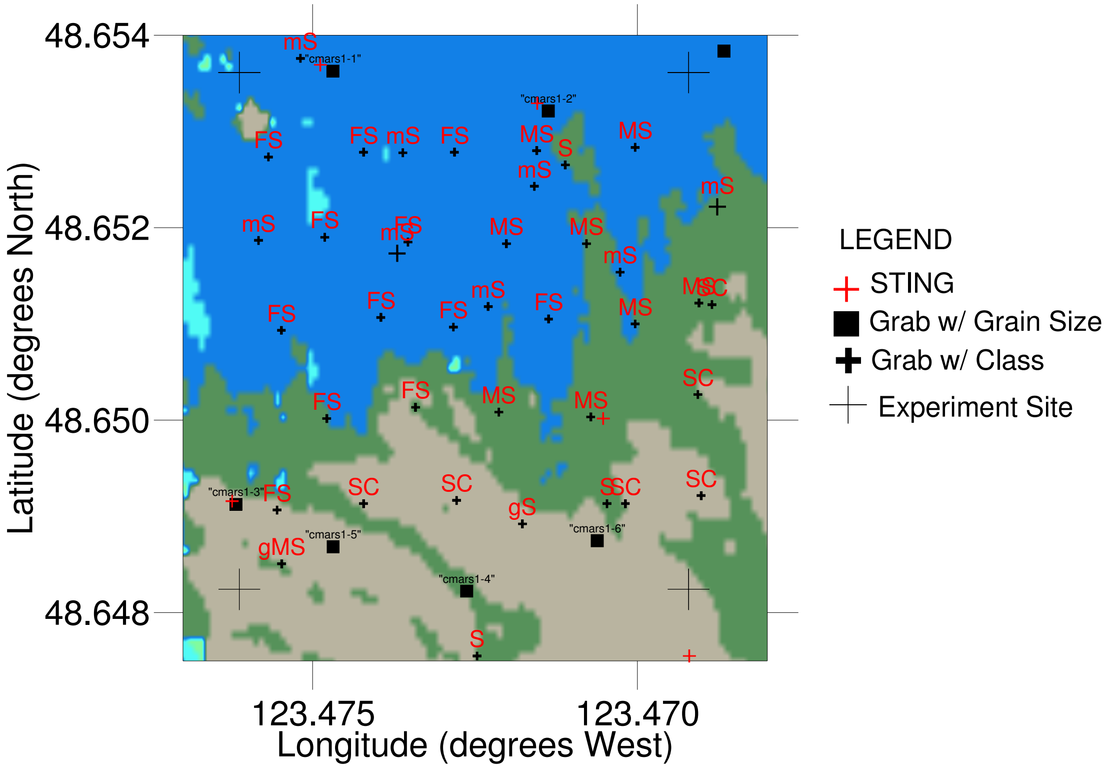

Figure 6-8. An enlargement of the Patricia Bay multibeam seabed classification map for the complex sloping site (CM1). Class colours and ground-truth interpretation are the same as Figure 6-1.

6.2 The Depth Experiment

The most prominent non-seabed influence on seabed echoes is the water depth. The aim of this experiment is to vary only the water depth without varying any other physical or system parameter in order to isolate the effect of depth on echoes as observed by echo duration as was done with model data in chapter 5. As noted earlier, echo duration represents many characteristics of the echo and responds strongly to the effect of depth and seabed slope, in addition to responding to roughness and seabed volume effects. The EDM (Equation 5-3) predicts a linear relationship between echo duration and depth, where the regression slope is function of the beamwidth and regression intercept is the sum of the pulse length, penetration time and macro-roughness, as shown in Equation 5-6.The field data shown here improve upon earlier efforts with model data by making using of improved tail picking and greater sample sizes. The field data are carefully controlled to avoid variation caused by other factors. The aim is to validate further the EDM and the new concept of effective measures of beamwidth and attenuation.

In the current practice of SEL depth compensation, the parameters of beamwidth and attenuation are poorly understood; most often operators use the manufacturer's beamwidth and a default volume attenuation value of 0.5 dB/kHz/m. These values are almost always inappropriate as will be seen here and as was noted in chapters 2 and 3. This is especially true for the choice of beamwidth. The effective beamwidth and attenuation estimated here essentially calibrates the depth compensation, thereby improving it by removing the largest fraction of the effect of depth as is possible. The resulting improvement in classification performance will be assessed in chapter 7. In addition, the improved depth compensation should also demonstrate a reduction in the number of echoes truncated by the end of the analysis window. Standard values for the beamwidth and attenuation used in standard SEL depth compensation tend to truncate many echoes as will be shown in this section. The methods applied to estimate the effective beamwidth and attenuation are designed to offer the best, most reliable performance, while being viable on an operational basis. This ensures that the methods are generally applicable and robust to all survey scenarios.

The EDM is also the basis of one of the three main characterization methods to be presented in chapter 8. To confirm the validity of the approach, the EDM needs to be verified as mentioned above in conjunction with establishing picking parameters that produce accurate measurements of echo duration. The details on picking the onset and tail of echoes are presented in appendix A. The appendix also evaluated the various methods, finding that QTC's threshold detection algorithm was suitable for bottom picking, while the new RANLS cumulative amplitude algorithm generated the best tail pick. Then, one of the main objectives of this experiment is to select the best picking threshold parameters. The methods of doing so and results obtained will be applied in chapter 7 for depth and slope compensation and in chapter 8 for seabed characterization via the EDM, a method dubbed *ABEL*.

6.2.1 Experimental Setup

The premise of the experiment is to simulate varying the water depth by instead varying the transducer's altitude above the seafloor while controlling all other parameters, including location. To this end, a dual frequency transducer was moved vertically through a range of altitudes above the seabed in a series of steps, collecting ~200 echoes at each ~1 m step over the gravel (CM3), sand (CM5) and mud (CM4) sites. The transducer was mounted in an apparatus that included a sensor package, wings to stabilize the apparatus and approximately 20 kg of lead ballast, see Figure 6-9. The sensor package consisted of a two-axis tilt sensor and a compass that were used to level and then monitor the apparatus to ensure it kept level to within ±1°. In additional to aiding in stability, the wings limited the rotation of the apparatus, although rotation was not expected to be an issue as both transducers have symmetric beampatterns. Position was maintained to within 5 m, (exceptions will be noted), by three or four point anchoring. As in chapter 2.6, the dual-frequency Odom echosounder was used; the characteristics of this echosounder appear in Table C-1.

The three experimental sites used were the flat, homogenous sites (CM3, CM4, CM5) that were established with throughout ground-truth in section 6.1. The experiment was conducted over four days on site, which required about two months to set up and to wait for intervals of calm weather. Seas greater than approximately 20 cm were avoided. 8,000 to 40,000 echoes were collected for each site at each frequency – two gigabytes of data in all.

Figure 6-9. Apparatus setup aboard the CGR research boat for the depth dependence experiment. The red arrow points to the transducer. The ballast had not yet been attached. The apparatus was balanced in the water after adding ballast prior to deployment.

6.2.2 Data Preparation

In survey situations, acoustic noise is generated by vessel motion, engines and machinery, and it can be significant. Inspection of the data in QTC IMPACT showed the traces to be virtually free of noise, but contaminated with multi-path echoes. Low noise was achieved because the survey vessel was stationary; its engines were turned off. Power was provided by a small 2500 Watt generator. The generator was a remarkably clean source of power, with little electronic noise. Instead, the biggest source of noise in this case was electronic noise generated within the echosounder itself[42] that primarily affected the Odom200 data over the Mud site. At this site, the primary echoes were of very low amplitude due to the low reflection coefficient and modest volume scattering. The high base gain required to acquire echoes at this site boosted the electronic noise. This limited the useful data from the mud site at Odom200 to a maximum altitude of 25 m. The Odom24 data and data from the other sites were not affected.

The multi-path echoes affected both frequencies, but were a problem only for the sand and gravel sites as the total backscatter at the mud site was of much lower strength. To address this issue, data were excluded from further processing that had a multi-path echo occurring within the tail picking analysis window (which includes the cases when the multi-path echo overlapped with the primary echo). Figure 6-10 shows multi-path echoes in an example trace from the gravel site. In the example, the primary echo overlaps with both the transmit pulse ring-down and the second echo and therefore would have been excluded from processing. The criteria to exclude data were based on altitude. For altitudes greater than 2 m, there is no overlap with the transmit pulse. The primary and second echoes will overlap if the primary echo's duration is greater than two times the time of its onset (i.e. twice the altitude). The water-depth multiple (see Figure 6-10 for the definition) will also overlap with the primary echo as the altitude nears the water surface. The air-sea second-echo (or real second echo) would only affect data in cases of low water depth or impossibly long duration echoes; neither is the case here. The multi-paths can be avoided by restricting the range of altitudes permitted for further analysis. The maximum altitude permitted is $d - {}^{lc}\!/_{2}$ and the minimum is ${}^{lc}\!/_{2}$, where l is the duration of the tail picking window. To avoid excluding a majority of the data, l was chosen as small as practical without sacrificing picking quality (picking quality improves with a larger window, see appendix A). The exception to the above altitude rule is the mud site as there were no observable multiples, presumably due to water depth and the weak reflection coefficient and total backscatter strength of the mud.

---

[42] The source was the thermal paper chart printer, which was discovered in a later survey.

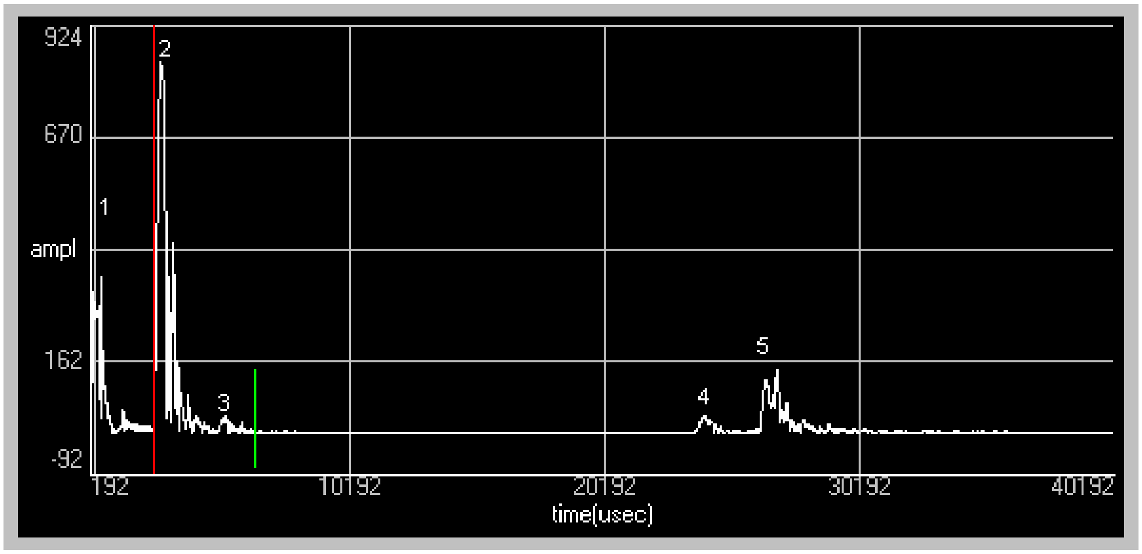

Figure 6-10. An example trace with multi-path echoes: Odom 24 kHz at the Gravel site at 1.90m altitude, water depth 17.84m. Annotations on this figure correspond to those of Figure 6-11. (1) is the transmit pulse. (2) is the first-order or primary echo. (3) is the second-order echo formed by the reflection off the transducer apparatus. (4) is dubbed the water depth multiple as its path length is always twice the water depth. (5) is the second-order echo formed by the reflection off the air-sea interface. The vertical green line is the 99% cumulative amplitude tail pick which is erroneous because of the overlap of (2) and (3). The usual picking window limit to exclude second echoes was overridden for this demonstration.

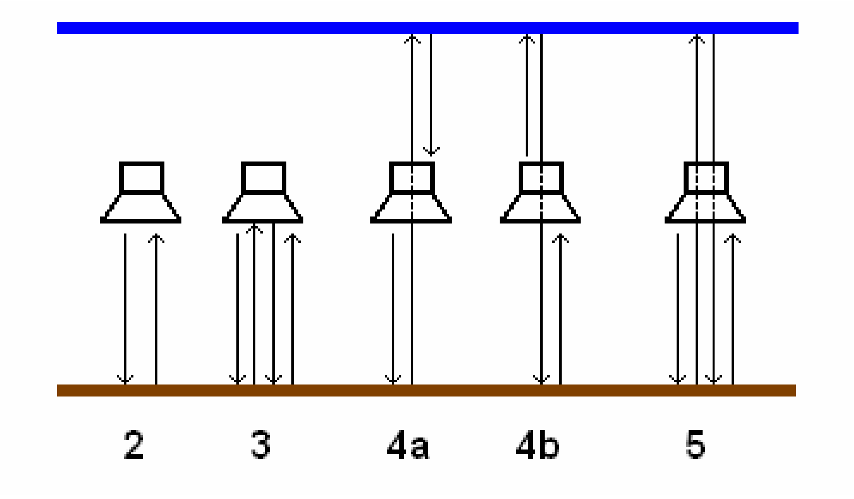

Figure 6-11. Ray diagram for multi-path echoes in the depth experiment, the numbering matches the observed echoes in Figure 6-10. The top line (blue) represents the air-sea interface, while the bottom line (brown) represents the seabed. For each case, rays start from the left and propagate while advancing to the right. Note that (3) requires a reflection off of the transducer apparatus, while (4a&b) require the transducer to receive or transmit with its top-side. Beampattern diagrams for Odom transducers show the top-side sensitively to be ~30 dB less than the peak sensitivity at nadir. For these reasons, the long-path second echo (5) is the strongest echo (aside from the first echo).

6.2.3 Data Analysis Methods – Fitting the EDM to Echo Duration Data

In to order verify the EDM and estimate effective beamwidths and attenuations for real echosounders and seabeds, the measurement of echo duration must be consistent, robust and respond to volume backscatter present at the ends of echoes. Echo durations are measured as the time from bottom pick to tail pick, both of which are described in appendix A. The bottom picks were done with QTC IMPACT using the parameters indicated in appendix A. The tail picks were done using the optimal RANLS CA tail pick method. The usual method of selecting the picking parameters is to maximize the sensitivity of the picks by using the most sensitive thresholds possible without becoming so sensitive that the picks trigger on any artefacts or noise. In this way, the tail picks should respond to the volume backscatter at the end of the echo (Figure 4-5 illustrates the volume component in the tails of echoes as discussed in section 4.7). The issue here is how to link the tail pick threshold precisely to the penetration time (t_p of the EDM, see Figure 5-11). Only an empirical link is possible – no theory has been devised, although future work on including volume scattering in the AEM could lead to such a link. To address this issue, the analysis is run for a series of different RANLS CA tail pick thresholds. In appendix A, it was found that the 99% threshold for the RANLS cumulative amplitude tail picking algorithm generated the best result in terms of the most consistent and reasonable picks. In this experiment, additional, more rigorous, criteria are available. For each threshold, linear regression analysis is run to quantify how well the linear EDM model fits the echo duration. The tail pick thresholds are evaluated on the quality of the fits, the number of erroneous picks, and how well the effective beamwidths and attenuations correspond to known values. In addition, the sediment sound speed estimated from the mean grain size for each site is used to calculate estimated attenuation values measured from echo duration. The tail pick threshold that has the best performance for all three sites, (per frequency) and will be used to measure echo duration in the analysis of chapters 7 and 8.

As mentioned above, the analysis of the water depth (as simulated by altitude) and echo duration is done primarily by linear regression as the EDM predicts a linear relationship between depth (or altitude) and echo duration. However, there is another issue to consider - the effect of seabed slope on echo duration, which is not linear. All of the sites are extremely low slope: 1°±0.3°, maximum 2° (combined over all the sites). As will be shown in section 6.3, seabed slope can be ignored here. However, because an aim of this experiment is to develop and test a generally applicable method for measuring the effective beamwidth and attenuation for any data set, seabed slope will be accounted for.

To estimate effective beamwidth and attenuation, a series of methods have been developed and tested. They will be described in order of increasing complexity. The basic method is the linear regression of echo duration to depth; with the effective parameters calculated from Equation 5-6 (sub-section 5.3.1) assuming the seabed slope is zero. If the seabed slope is known, one can account for it by first calculating an effective beamwidth from Equation 5-6 (*a*) with zero slope and than subtracting two times the mean slope from that zero slope value – this is called the slope subtraction method. The mean slope is the mean of the slope estimate for every data point, which is calculated from the gradient of the seabed over an area the matching the acoustic footprint (see Equation 2.2 for the definition of footprint semi-major axis). The source of the seabed slope data can be bathymetry from the survey itself (corrected for slope as described in section 5.3.2) or bathymetry from another source such as multibeam, LIDAR or even coarse resolution data from a marine chart. If the mean seabed slope is greater than half the beamwidth then the direct measure of effective beamwidth using Equation 5-6 as derived from Equation 5-3 case I does not apply. The form of the EDM that does apply, Equation 5-3 case II, cannot be solved to find the beamwidth. To address this problem there are two approaches. The first is to exclude data that occur on slopes greater than one-half the -3 dB beamwidth; this creates a low-slope subset of the data. The second approach is to remove the extra echo duration caused by the slope, so that Equation 5-6 can be applied assuming zero slope.

This second, more complex, method is done iteratively and numerically to overcome the non-linear response of echo duration to slope – this is the iterative slope removal method. Here is the processing flow described step by step. To aid in understanding, one may think of it as similar to Newton's iterative numerical method. Starting estimates for the beamwidth and attenuation are required; usually the manufacturers' or -3 dB beamwidth with a default attenuation value of 0.5 dB/kHz/m. For each echo, two modelled echo durations are calculated: one with the echo's mean slope and one with zero slope. The difference of the two is the extra duration caused by the slope. This difference is then subtracted from the measured echo durations, then the effective beamwidth and attenuation are calculated assuming zero slope. In the next iteration, these new effective beamwidth and attenuation estimates are used as the starting point to model the extra duration caused by the slope, which leads to new effective beamwidth and attenuation estimates. When the change in the effective parameters (of beamwidth and attenuation) drops to less than 1%, the process is considered converged and the final effective parameters are produced.

This iterative process can fail to converge when the starting effective parameters are significantly inaccurate or if the tail picking method underestimates the echo duration in cases of high slope. To reduce the chance of the latter, the slope values used can be capped at a maximum – usually 30°. If the iterative method fails to converge, the low-slope subset result is taken instead. In the 36 effective beamwidth and attenuation estimates shown in Table 6-4, this iterative slope removal method failed 6 times: 4 times for the mud with high tail picking thresholds and twice for the sands with low tail pick thresholds. In these instances, the failures were due to poor starting points (convergence was achieved by changing the starting point). Convergence was also improved by running the iterative effective parameters method on the low-slope subset. In all, there are four methods: slope subtraction on the full survey or low-slope subset, and the iterative slope removal method on the full survey or low-slope subset. These four methods are applied in chapter 7, and of the four, the *iterative LS* (low-slope) method is generally considered the most reliable. In 34 of the 36 runs, the differences between the four methods of estimating effective parameters were less than the 95% confidence intervals – indicating that the result from the four methods were consistent. For the full Patricia Bay survey presented in section 7.3, the methods also agreed within their 95% confidence limits (although the iterative effective parameters method resulted in tighter confidence intervals for the beamwidth than the other methods: 0.15° versus 6.0° - on average).

6.2.4 Results from Field Data

Figure 6-12 shows plots of measured echo duration as a function of simulated water depth (altitude) for both echosounders over the three experimental sites. Table 6-4 summarizes the results of ordinary least squares linear regression analysis of the depth experiment data. The table lists the effective beamwidth and attenuation, estimated attenuation and expected values for comparison. Estimated attenuation is calculated from the regression intercept using the sediment sound speed determined from the mean grain size for each test site; this differs from the effective attenuation, as that calculation has to assume a sound speed of 1500 m/s. The plots of Figure 6-12 are for the tail pick thresholds highlighted by dotted lines in Table 6-4 – 99% and 98% for the Odom24 and Odom200 respectively. These are the most suitable tail pick thresholds that fit the criteria discussed in section 6.2, and will be established after further verifying EDM and that the regression methods used are sufficient.

The first expectation of the results to be tested is that the linear EDM fits the data. Qualitative inspections of echo duration versus depth plots, some of which are shown in Figure 6-12, show that echo duration clearly has a linear dependence on depth (altitude)[43]. The nature of the relationship, in terms of regression slope and intercept, will be shown to be consistent with the EDM and ground-truth information derived from the mean grain size of each experiment site. As discussed in section 5.4.3, the R^2 values are provided to indicate the strength of the linear relationship of echo duration to depth relative to the overall variance, and do not indicate the goodness of fit. In this case, the R^2 statistic reacts as expected. It increases with the effective beamwidth and decreases as the overall variance of the data increases due to incoherent scattering (gravel and mud seabeds, see section 4.7). In the case of a narrow-beam echosounder, such as the Odom200, the variability in echo duration due to scattering is greater than that caused by the effect of depth. However, even cases with the lowest R^2 values, linearity can still be observed and the fits yield reasonable values.

---

[43] The results of hypothesis testing for linearity and significance are not presented here, as the strict conditions required for their validity are not satisfied in all cases. These conditions are discussed below.

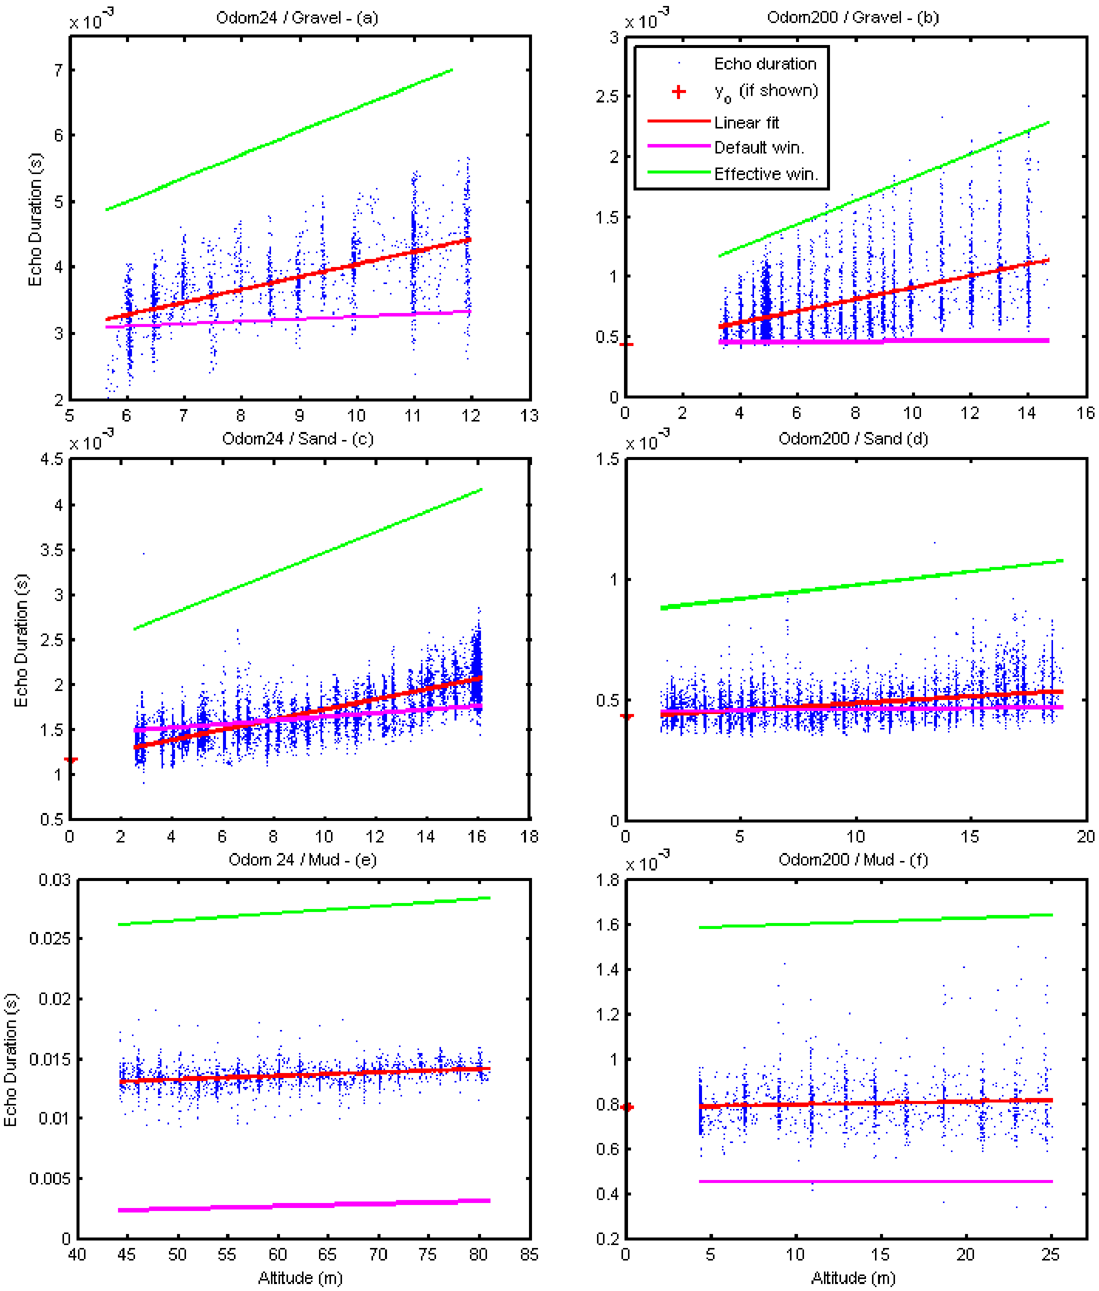

Figure 6-12. Linear regressions of echo duration to 'depth' (as simulated by transducer altitude) for the depth experiment. Echo duration is based on the 99% (Odom24) or 98% (Odom200) CA RANLS tail picks. Odom24 on the left (*a,c,e*), Odom200 on the right (*b,d,f*). Sites as follows: gravel (*a,b*), sand (*c,d*), mud (*e,f*). The green lines represent the duration of the SEL analysis windows calculated using the effective parameters, while the magenta lines are the duration of the standard SEL analysis windows. Note that data points can overlap each other and that the vertical striations are due to the way the data were collected – in depth interval steps of approximately 0.5 to 3 m.

Table 6-4. Table of results for the depth experiment for both echosounders over all sites with a variety of tail picking thresholds applied. Sample sizes ranged from 4000 to 15000 data points (after filtering). 95% confidence intervals / widths are calculated as described in sub-section 5.4.3 – asterisks indicate cases that did not pass the conditions of validity for ordinary least squares linear regression for which confidence interval estimates are not meaningful. Entries in grey are approaching the singular point in Equation 5-6 and are not valid. Expected beamwidth values are based on the -4.34 dB beamwidth (see Table C-1). Expected attenuation values based on mean grain size for each site.

CA Tail Pick	**ODOM24 / Gravel - (a)**					
Threshold (%)	Eff. Beamwidth (°)	95% CW (°)	Eff. Atten. (dB/kHz/m)	95% CW (dB/kHz/m)	Est. Atten. (dB/kHz/m)	R^2
90	28.7	1.12	1.07	0.06	0.80	0.25
95	34.5	1.64	0.61	0.04	0.46	0.18
98	49.6	1.58	0.41	0.02	0.30	0.29
99	56.1	1.35	0.33	0.01	0.25	0.40
99.5	57.2	1.37	0.23	0.01	0.17	0.40
99.9	48.3	1.80	0.14	0.00	0.10	0.23
Expected Values	23.8	-	0.27 - 0.46	-	0.27 - 0.46	-

CA Tail Pick	**ODOM24 / Sand - (c)**					
Threshold (%)	Eff. Beamwidth (°)	95% CW (°)	Eff. Atten. (dB/kHz/m)	95% CW (dB/kHz/m)	Est. Atten. (dB/kHz/m)	R^2
90	20.2	0.28	2.08	0.05	1.81	0.46
95	24.2	0.25	1.29	0.02	1.13	0.57
98	28.4	0.23	0.86	0.01	0.75	0.67
99	30.2	0.24	0.66	0.01	0.58	0.67
99.5	31.0	0.28	0.53	0.01	0.46	0.61
99.9	28.4	0.47	0.33	0.00	0.28	0.33
Expected Values	23.8	-	0.49 - 0.52	-	0.49 - 0.52	-

CA Tail Pick	**ODOM24 / Mud - (e)**					
Threshold (%)	Eff. Beamwidth (°)	95% CW (°)	Eff. Atten. (dB/kHz/m)	95% CW (dB/kHz/m)	Est. Atten. (dB/kHz/m)	R^2
90	23.6	1.12	0.09	0.00	0.09	0.78
95	26.9	1.18	0.08	0.00	0.08	0.77
98	27.1	1.25	0.06	0.00	0.06	0.66
99	27.2	1.55	0.05	0.00	0.05	0.50
99.5	28.2	1.61	0.05	0.00	0.05	0.39
99.9	27.5	1.69	0.03	0.00	0.03	0.15
Expected Values	23.8	-	0.05 - 0.07	-	0.05 - 0.07	-

Table 6-4. Continued.

CA Tail Pick Threshold (%)	**ODOM 200** / **Gravel - (b)** Eff. Beamwidth (°)	95% CW (°)	Eff. Atten. (dB/kHz/m)	95% CI / W (dB/kHz/m)	Est. Atten. (dB/kHz/m)	R^2
90	9.3*	-	1.36*	-	1.02*	0.12
95	15.9*	-	0.90*	-	0.67*	0.19
98	28.5*	-	0.40*	-	0.30*	0.28
99	35.6*	-	0.28*	-	0.21*	0.49
99.5	47.0*	-	0.21*	-	0.15*	0.44
99.9	40.4*	-	0.06*	-	0.04*	0.36
Expected Values	4.9	-	0.27 - 0.46	-	0.27 - 0.46	-

CA Tail Pick Threshold (%)	**ODOM200** / **Sand - (d)** Eff. Beamwidth (°)	95% CW (°)	Eff. Atten. (dB/kHz/m)	95% CI / W (dB/kHz/m)	Est. Atten. (dB/kHz/m)	R^2
90	3.5*	-	2.82*	-	2.45*	0.08
95	4.9*	-	1.13*	-	0.98*	0.11
98	7.9*	-	0.60*	-	0.52*	0.18
99	10.3*	-	0.43*	-	0.37*	0.21
99.5	11.9*	-	0.29*	-	0.25*	0.16
99.9	7.4*	-	0.11*	-	0.09*	0.01
Expected Values	4.9	-	0.49 - 0.52	-	0.49 - 0.52	-

CA Tail Pick Threshold (%)	**ODOM200** / **Mud - (e)** Eff. Beamwidth (°)	95% CI/W (°)	Eff. Atten. (dB/kHz/m)	95% CW (dB/kHz/m)	Est. Atten. (dB/kHz/m)	R^2
90	4.8	1.3	0.27	0.01	0.27	0.01
95	5.3	1.3	0.19	0.01	0.20	0.03
98	5.2	1.7	0.14	0.00	0.15	0.04
99	5.2	2.0,8.0	0.12	0.00	0.12	0.03
99.5	5.6	1.9,9.2	0.09	0.00	0.09	0.02
99.9	13.2	6.4,18.7	0.06	0.01	0.06	0.02
Expected Values	4.9	-	0.05 - 0.07	-	0.05 - 0.07	-

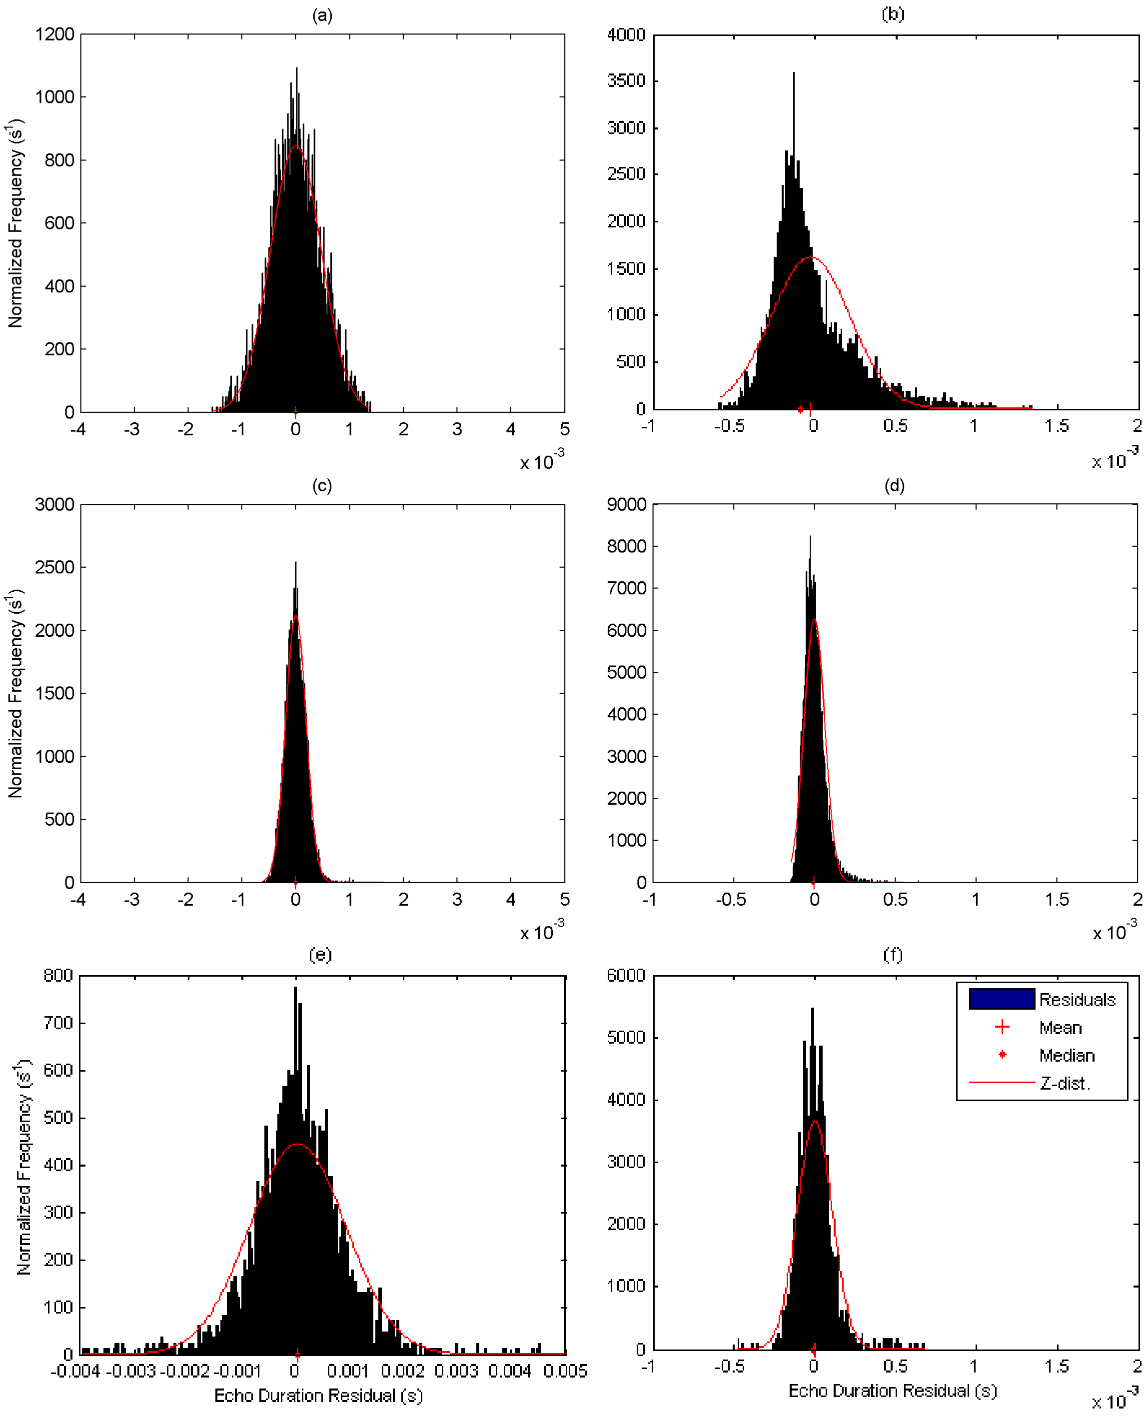

Figure 6-13. Histograms of echo duration residuals – the data minus the echo duration model, with each panel corresponding to those in Figure 6-12, i.e. Odom24 on the left Odom200 on the right. Echo duration is based on the 99% (Odom24) or 98% (Odom200) CA RANLS tail picks. A normalized frequency scale is used to facilitate better comparisons between data sets by accounting for the number data points and bin widths. The red curves are normal ('Z') distributions calculated from the sample size, mean and standard deviation. Bins are plotted in blue, but the density is too high to see the colour.

There are four conditions required for linear regression analysis [Walpole and Myers, 1985], they are:

1) the data (echo duration) residuals are normally distributed,
2) the mean of the residuals is zero,
3) the standard deviation is constant (homoscedasticity),
4) and the data are not significantly correlated.

These conditions will be addressed in the following paragraphs.

The issue of data normality was investigated by inspection of histograms of the echo duration residuals that are shown in Figure 6-13. Departures from normality in these data are likely due to the following: increasing variance with increasing footprint size, which leads to positively skewed results (Odom200 gravel and sand); the possibility of sub-bottom reflectors lengthening or shortening echoes (the mud cases); residual noise spikes (as seen in Odom200 mud). Inspection shows that all of the Odom24 data and the Odom200 mud data are approximately normal with zero mean (conditions 1 and 2)

In sub-section 5.4.3, effects that could lead to non-linear trends were discussed at length; these effects tend to break the condition of homoscedasticity (condition 3). The effect of coherence and varying statistics within the footprint are seen here in the variations about the regression line (Figure 6-12). They are most evident in the data collected from the gravel sites, which show increasing variance with depth, and are strongest for the Odom200 gravel data. This is as expected, since the higher frequency sonar has a much narrower beamwidth, and so is more sensitive to the impact of isolated scattering centres within the beam footprint.

The echo duration residuals for the Odom24 sand data are plotted versus altitude in Figure 6-14. Inspection of these data shows that the standard deviation does not change significantly over the altitude variation in this experiment. This result is typical of the other cases except the Odom200 gravel and sand data. The condition of homoscedasticity is only challenged in the case of high frequency, narrow beamwidth and high roughness.

The last condition for the validity of linear regression analysis is that the data are not significantly correlated (condition 4). Correlated data would mean that the sample size used in the tests is effectively smaller than the stated number of samples. This can be roughly accounted for by dividing the sample size by number of samples over which the data are correlated (known as the auto-correlation length). Figure 6-14 shows the echo duration residuals and the auto-correlation of these residuals for the Odom24 sand data set. This figure shows that the normalized auto-correlation drops very quickly, indicating there is no significant correlation. Any correction for correlated data would not change the confidence interval results. This conclusion is the same for all six cases and regardless of the ordering of the data (sorted by altitude or in order in which the data were collected). The case shown in Figure 6-14 (Odom24 / sand) is ideal in that ping-to-ping coherence or breaking of homoscedasticity was not expected on the basis of the physical reasoning discussed earlier. Coherence is observed in other data sets, as expected, with auto-correlation lengths of 5 to 20 samples, which is again not significant (a value of 100 would be, for a data set of 10,000 samples). Figure 6-14 (*upper*) is effective at showing that this particular data set is homoscedastic (condition 3) and results in residuals that are zero mean throughout (condition 2).

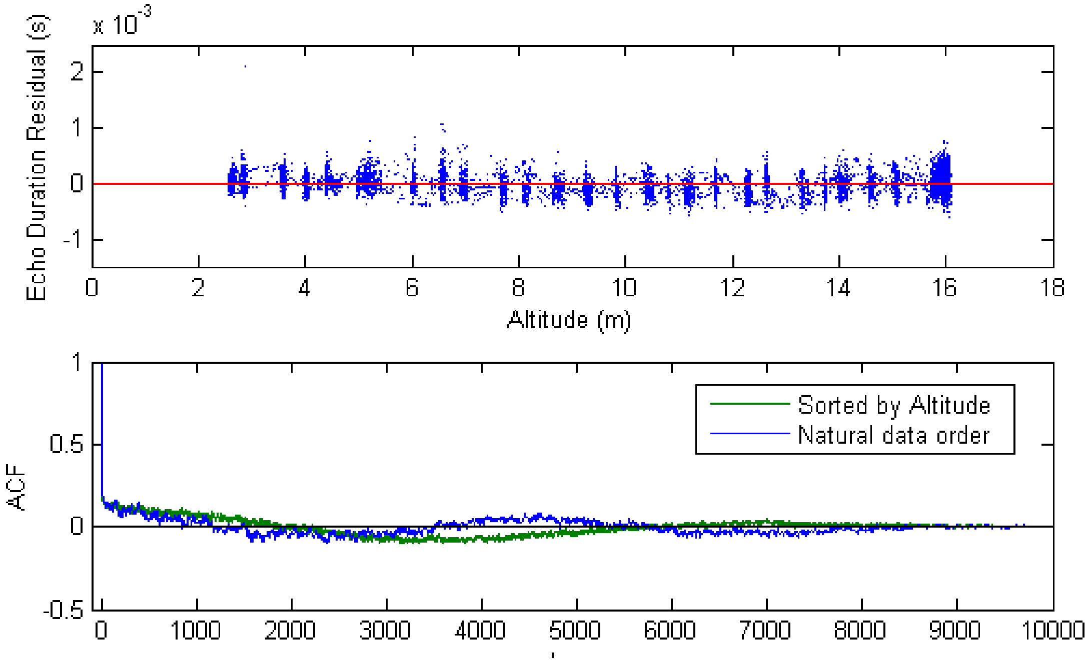

Figure 6-14. Echo duration residuals plotted against altitude (*upper*) for the Odom24 / sand / 99% case. The red line is a linear regression of the residuals that had zero slope and intercept. (*lower*) Auto-correlation analysis for the Odom24 / sand / 99% case.

From the results of the preceding paragraphs, the linear regression model was found to be valid for all the Odom24 cases and the Odom200 / mud case shown in Figure 6-12. Subsequent results can now be discussed. As shown in Table 6-4, a series of tail picking thresholds were applied in order to establish the most appropriate thresholds to use in this and subsequent data sets. The criteria for selecting the most appropriate thresholds were discussed in the previous sub-section (6.2.3). The selected thresholds (99% and 98% for the Odom24 and Odom200 respectively) have the best performance in terms of sensitivity to noise, reasonable placement within the echo signals (defined in appendix A) and reliability of data provided to the linear regression analysis. These thresholds generally result in higher R^2 values (with consideration of the regression slope) and EDM-data residuals that are more symmetric, more closely resemble normal distributions and have fewer outliers than the other thresholds. These thresholds also result in effective beamwidth values that respond reasonably to the beam-widening effect (as will be discussed further in the next sub-section). The effective beamwidth values are also consistently larger than the -4.34 dB beamwidth, as expected.

A goal of this experiment is to establish an empirical link between the effective attenuation (calculated from a 10 dB extinction assuming the sediment sound speed is 1500 m/s) and the tail pick threshold. In usual survey scenarios, the sediment speed will not be known, so the selection of the thresholds should consider both the effective attenuation and the estimated attenuation (based on an estimated sediment sound speed)[44]. The best thresholds will produce attenuations consistent with the ground-truth data. For both echosounders, the selected thresholds overestimate the attenuation of the sand seabed and underestimate the attenuation of the gravel seabed (see Table 6-4). However, both of these differences are not completely unexpected as some sands have higher attenuations than that listed in Table 6-4 (i.e. 'Very Fine Sand' in Table C-2, has an attenuation of 0.634 dB/kHz/m), and some gravels have lower attenuations (i.e. 'Cobbles, Gravel, Pebble' seabed has an attenuation of 0.273 dB/kHz/m). The attenuation of the mud is estimated accurately for the Odom24, while the Odom200 shows a higher

---

[44] The difference between the effective attenuation and the estimated attenuation as seen in Table 6-4 is only significant (greater than 0.1 dB/kHz/m) for the gravels, which have higher a sound speed (2000 m/s) than the sands (1720 m/s) or the muds (1470 m/s).

attenuation that is likely due to the thin crust/bacterial mat observed on the seabed surface. The Odom24 likely penetrates this thin crust, while the shorter wavelengths of the Odom200 experience more attenuation for this short distance. The empirical relationship between mean grain size and attenuation [APL-UW, 1994] has uncertainty that is unaccounted for here. Measurement uncertainty for the grain size measurements is also unaccounted. For these reasons, it is reasonable to conclude that the estimated attenuations from the 99% and 98% thresholds (Odom24, Odom 200 respectively) are approximately consistent with attenuation values characteristic of sand, gravel and mud. For both frequencies, the tail picking thresholds immediately lower or higher than the two selected thresholds result in considerably more disparity between estimated and expected attenuation values. Overall, the two selected thresholds generated the best results.

In addition to variability and physical processes that tend to break homoscedasticity and linearity, there is also some systematic error to be considered that would affect the attenuation estimates only. There are two sources or error: the water sound speed, and the pulse duration. The true water sound is approximately 1484 m/s[45], and this difference is negligible. Pulse duration can suffer from ring-down, and the effect of the shape of the transmit pulse is unknown. However, the transmit pulses of the Odom echosounder are known to be square-shaped with little ring-down. The pulse duration was measured from data originally presented in [Park, 2003]. The error on the pulse duration is -0.025 to +0.050 ms, which translates to an uncertainty of approximately 5% for the Odom200 results and 0.5% for the Odom24 results.

After considering all of the sources of error and upon evaluating the results, it is reasonable to conclude that the EDM is a valid model of echo duration for varying depth in most cases. This validates the SEL depth compensation approach, but more importantly, it facilitates improved depth compensation by using effective beamwidth and attenuation.

---

[45] Calculated based on temperature (10℃) and salinity (30 ppt), for reference see [Lurton, 2002].

6.2.5 Discussion and Comparison to Previous Results

The effective and estimated attenuations produced by the selected tail picking thresholds shown in Table 6-4 exhibit similarity between frequencies and vary relative to the expected values. The values were only approximately consistent to the expected bounds as was discussed in the preceeding sub-section. The range of expected attenuations shown in Table 6-4 reflects the uncertainly in the ground-truth and the uncertainty in the relationship between attenuation and grain size. The effects of echo-to-echo coherence and varying footprint size were observed in the BORIS results of section 5.4.3 were also observed here.

The beam-widening effects on the effective beamwidth observed here were similar to those seen in chapter 5, particularly the BORIS results. The consistently large effective beamwidth results seen here and in chapter 5 are not entirely unexpected given the phenomenon of beam-widening, and they are also not unrealistic given the properties of the beampatterns. For instance, beampattern cross-sections of the Odom24 and Odom200 echosounders are shown in Figure 6-15. Based on these beampattern data, the amplitude levels corresponding to the effective beamwidths for the sand site were found to be -7 dB (Odom24) and -12 dB (Odom200). This indicates that both of these effective beamwidths are well within in the main-lobes of the respective beams as these main-lobes are 47° and 9° wide, respectively. However, in the high roughness case, the gravel site, the beam-widening effect is strong enough so that the sidelobes influence the echo duration for both frequencies – these effective beamwidths are located within the first side-lobes. This tends to explain why the effective beamwidths for sufficiently rough seabeds are dramatically larger. The effect of sidelobes does not appear to be a problem for the model as the echo duration is still linear with depth. It is interesting to speculate that beam-widening via volume backscattering is likely contributing to the wide effective beamwidths of the mud site and possibly the sand site, two sites that have little surface roughness. This is made possible by transmission into the volume on high-angle rays that do not contribute any surface backscatter. Since volume scattering at these frequencies is isotropic, a volume return may be present where there is no surface return as seen in some BORIS observations (BORIS separates volume and surface echoes). In comparing the beam-widening effect between the echosounders on the basis of the relative increase in the *slope* of the regression line, it was found that the Odom200 echosounder experiences a stronger beam-widening effect (a factor of 5.0 increase for the Odom24 versus a factor of 34.3 increase for the Odom200 for the gravel site over the sand site). This is due to the characteristics of the beampatterns: partially the strength and spacing of the sidelobes relative to the main beam, but mainly just the width of the beam overall. It appears that the beam-widening effect is much stronger for narrower beams, which explains why narrow-beam high-frequency SBES can generate meaningful classification on the basis of roughness. However, wider-beam SBES tend to respond to a wider range of roughness (beam-widening saturates at higher-levels of roughness), in addition to generating more responsive shape features (due to the length of the tail of the echo, see Figure 2-7). Narrow-beamed echosounders may produce more classification errors if the effective beamwidth is not used.

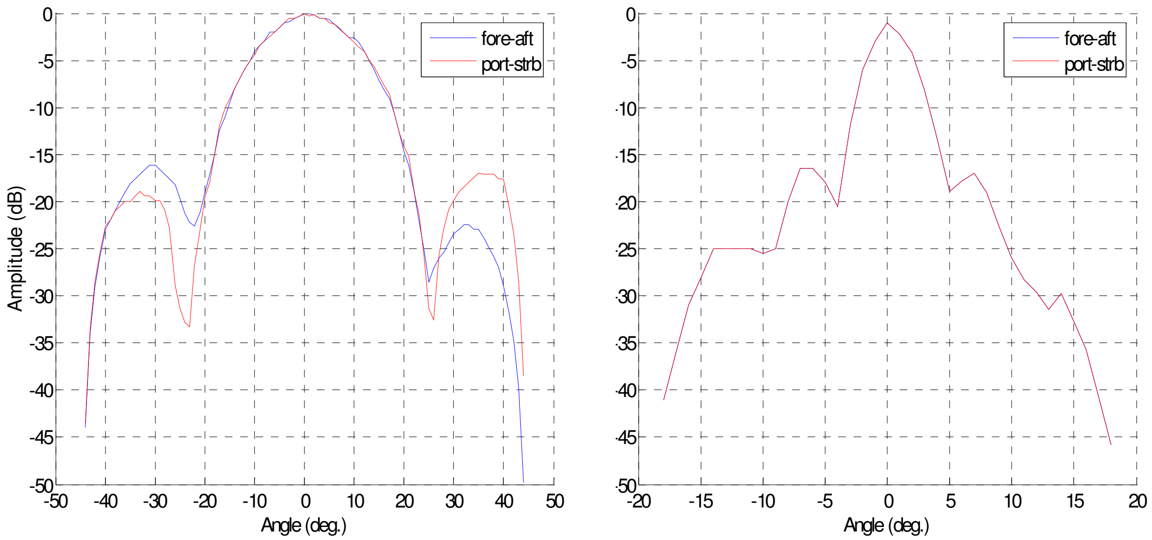

Figure 6-15. Beampattern profiles for the Odom dual-frequency echosounder. The Odom24 is (*left*) and Odom200 is (*right*). Note that the two profiles for Odom200 are equal (the blue curve is covered by the red), see Figure F-2 for a 3-D representation.

A possible issue with this analysis is that the EDM assumes that volume scattering and the penetration time adds to the echo duration at times after the spreading time has past. In other words, the volume scattering has to contribute strongly to the tail of the echo for its effect to be measurable. As mentioned above, there are mechanisms for volume scattering to contribute at the edges of the beam and the tail of the echo. Results from chapters 4, 5 and from [Sternlicht, 1999; Jackson et al., 1986] indicate that volume scattering does contribute to the tail of the echo, especially for wide-beam, low-frequency SBES. There are scenarios such as narrow-beam high-frequency, where the contribution of volume scattering is slight. However, the accurate measurement of seabed attenuation from the narrow-beam high-frequency Odom200 SBES demonstrates that attenuations can be measured even in the least likely scenario. Then, it appears likely that the contribution of volume scattering need not dominate the tail of the echo. This is likely a consequence of the *cumulative* amplitude tail pick used. Volume scattering almost always contributes to the tail of the echo, and because it can persist for a long time past the tail pick even at a very low level, it will shift the position of a high-sensitivity cumulative amplitude tail pick. The key is to have long analysis windows, low noise and/or noise compensation. For the range of frequencies used in this work, observations from BORIS simulations indicate that the volume contribution can be orders of magnitude less than the surface contribution in peak amplitude and still be measurable with a sensitive tail pick. When the volume scattering is insignificant, the measured effective attenuation will be erroneously large. Indeed this is what is seen with tail picking thresholds at 95% or less.

Overall, the beam-widening effect is quite pronounced; manufacturer's beamwidths should not be used as the basis for depth compensation. The effective attenuation values obtained are significantly different from the fixed value of 0.5 dB/kHz/m used in SEL compensation. The method of estimating effective beamwidth and attenuations presented here objectively optimizes the EDM to the data so that depth compensation based on these values will remove the greatest portion of variance caused by depth as possible. This improved method reliably limits the number of truncated echoes as indicated by the green lines in Figure 6-12 as compared to the large number of echoes that are longer in duration than the standard SEL analysis window (magenta lines in Figure 6-12). Measurement of the effective beamwidth and attenuation will remain an empirical process due to the lack of a reliable way to account for the many factors, including shape of the beam, scattering, noise level, picking, etc. So instead of running a depth experiment for every seabed type and echosounder possible, a much more pragmatic approach is to measure the effective parameters from the survey data as will be shown in chapter 7. This approach lends itself to being to be robust to any combination of echosounder and seabed type and does not require calibration now that suitable picking thresholds have been established.

6.3 The Slope Experiment

As discussed in chapters 3 through 5, seabed slopes strongly affect echoes, by lengthening duration and by modifying the relative strengths of the echo components. Seabed slope and transducer tilt are equivalent as shown by the EDM; they affect echoes in the same way. This has two consequences. One is that acoustic surveys done in rough water, with vessel roll or pitch often significant compared to half the sonar beamwidth, can give poor results. Gleason et al. showed that classification consistency declines markedly for only modest values of vessel roll and pitch [Gleason et al, 2009]. A much older paper by Stanton revealed errors in echo integration measurement caused by vessel roll were significant and increased for smaller beams [Stanton, 1982]. The other is that experiments on seabed slope can be done by varying transducer tilt, which is far more practical than trying to find the same sediment on sites with a range of slopes. This is precisely how this experiment was done: by repeatedly transecting back and forth along the centre-line of each of the three homogenous seabed sites while increasing the transducer tilt on every other transect. The same sonar and sensor package used in the depth experiment was used here. Features extracted from echo time series drive classification results, so the most direct way to observe the effect of slope on classification is to observe the features. The effectiveness of slope compensation on classification can also be seen in the feature values.

6.3.1 Compensation of Slope – The SELw Method

Compensation for the lengthening effect of slope on echoes is certainly possible, especially for low slopes where the relative strengths of the echo components are not strongly affected. Such compensation could make use of the resampling approach applied in TNORM and SEL depth compensation. Indeed, simply adding a slope term in Equation 2-6 based on the form of the EDM (Equation 5-3) would achieve this. For this experiment, the slope is taken as the tilt, however for real surveys slope data could be obtained from survey bathymetry (after compensation for the bathymetry error caused by slope as discussed in section 5.3.2) or from an outside source such as multibeam bathymetry. As seen in Figure 5-12, slopes can lengthen echoes by up to 20 times their original duration, which means that these echoes would have to be downsampled by that amount. For envelope data from many echosounders, such downsampling is likely to reduce the Nyquist frequency to values about the same as the bandwidth. The only logical alternative is to fit the analysis window to the data, instead of the usual practice of resampling the data to fit the analysis window. This new method is called variable length window compensation or SELw. The implications of this are as follows:

- the feature algorithms have to account for varying analysis window size (the original QTC features always assume 256 samples in an analysis window);
- resampling is the most computationally time consuming part of SBES classification; SELw completely skips it, dramatically reducing computation time;
- the bandwidth of the original signal is retained;
- there is no restriction on the size of the window, so any time series may be analyzed.

SELw analysis windows consist of a number samples before and after the bottom pick. The number of samples in each case is determined as a fraction of the expected echo duration, as calculated from the EDM using the effective beamwidth and attenuation, pulse duration and zero macro-roughness. The exact formulation of SELw is shown as Equation 6-1:

$$\begin{Bmatrix} s_{bbp} \\ s_{abp} \end{Bmatrix} = \begin{Bmatrix} n \\ m \end{Bmatrix} f_{samp} \left[\frac{2d}{c_w} \left(\sec\left(\phi_{local} + \frac{\theta_{eff}}{2} \right) - 1 \right) + \frac{20}{c_w \kappa_{eff} f_c} + \tau \right] \quad \forall\ 0 \le \phi \le \frac{\theta}{2} \text{(I)}$$

$$\begin{Bmatrix} s_{bbp} \\ s_{abp} \end{Bmatrix} = \begin{Bmatrix} n \\ m \end{Bmatrix} f_{samp} \left[\frac{2d}{c_w} \left(\sec\left(\phi_{local} + \frac{\theta_{eff}}{2} \right) - \sec\left(\phi_{local} - \frac{\theta_{eff}}{2} \right) \right) + \frac{20}{c_w \kappa_{eff} f_c} + \tau \right] \ \forall\ \frac{\theta}{2} \le \phi < \frac{\pi}{2} - \frac{\theta}{2} \text{(II)}$$

$$\text{(6-1)}$$

where s_{bbp} and s_{abp} are the number of samples before and after the bottom pick and n and m are fractional amounts of expected echo duration before and after the bottom pick. The effective beamwidth and attenuation are normally used along with the local seabed slope as measured by the slope of the mean seabed plane fitting an area the size of the footprint. As is the case for SEL compensation, the SELw analysis window is usually at least twice the echo duration to allow for long echoes due to high roughness and/or penetration (i.e. $n + m = 2$). A good choice for n and m are 0.125 and 1.875, respectively. These values place one cumulative integral feature before the pick, given that the cumulative integral feature algorithm places 16 features evenly throughout the analysis window. Of the original feature set, only FFT and wavelet feature algorithms are not easily adaptable to a variable length window as will be examined further in section 7.4.

6.3.2 Data Analysis

Three of the new features that are related to features developed in section 7.5 are used here to study the effect of slope; these features are chosen as they are representative of the shape of echoes, and are taken to represent the general behavior of shape features. The three features are: cumulative histogram (CH) at 40% and 80% and the timespread. The CH features are similar to the cumulative integral features and the cumulative threshold features of appendix G. The algorithm for the cumulative histogram feature family is defined as the position within the analysis window (valued from zero to unity) where the normalized cumulative amplitude reaches a certain threshold, which is either 40 or 80% in this case. The timespread feature is as used in [van Walree et al., 2005];

except that amplitude rather than the amplitude squared (energy) is used. This feature, known as the 168th feature[46], is:

$$F_{168} = \sqrt{\frac{4\int p(t)(t-t_0)^2\,dt}{\int p(t)dt}} \tag{6-2}$$

where $p(t)$ is the echo amplitude time series, the integrals are over the analysis window, and t_0, the 167th feature, is:

$$F_{167} = t_0 = \frac{\int p(t)t\,dt}{\int p(t)dt} \tag{6-3}$$

The motive for using amplitude rather than energy was to emphasize the smaller amplitudes in the tail of the echo.

These three features are calculated from variable length windows and the usual stacks of five. It is important to note that the values of the window length parameters, n and m, are not constant between the sand and gravel sites (the sand site used the usual values, while the gravel site used double those values, see Equation 6-1 and associated discussion). This means that in the four figures about to be presented, it is the trends in the feature values are important instead of the values themselves.

Figure 6-16 and Figure 6-18 show these features calculated from stacks of echoes that have been compensated for depth but not for tilt. The features are displayed in the sequence in which the stacks were acquired. Each column, demarcated by vertical black lines, is from two transects through the centre of either the homogeneous gravel or sand site as indicated. Each feature has a monotonic trend as tilt is increased, showing the expected effect of slope on these features. Repetitive variations and trends seen within each column (where slope is constant) are due to varying seabed properties, seabed slope and some depth effect since the -3 dB beamwidth was used with SELw compensation. Seabed slope was 0.8° for the gravel site and 1.1° for the sand site. Seabed slope interacts with the tilt of the transducer if the two are aligned – the transducer was always tilted up toward the bow of the survey vessel. Since the gravel transect was oriented parallel to the

---

[46] The feature numbers will given here will be used later in chapter 7.

depth contour, the effect of slope (and depth) was constant and not noticeable, while the sand transect crossed back and forth across depth contours maximizing the effect of slope. Specifically, the first transect was downhill so the tilt and the slope add, while the second transect has the tilt and slope counteracting. The turnaround between the two transects coincides with the discontinuity seen in the middle of many of the columns in Figure 6-18 and Figure 6-19. If this explanation is accurate, the jump in the feature values (the discontinuity) should be equivalent to the change seen at approximately two degrees tilt and it is.

Figure 6-17 and Figure 6-19 are the tilt compensated features. Since tilt (simulating seabed slope) is the only variable in the experiment, compensating for it should render all features constant as slope varies. Slope is compensated by using the transducer tilt as the seabed slope for SELw compensation (Equation 6-1). In most cases compensating for tilt in this manner was very effective. With gravel echoes, the effects of tilt on CH at 80% and on timespread have been fully erased – the feature is unchanged at all tilts. This is true even for tilts well in excess of half the -3 dB beamwidth, 10°. With CH at 40% tilt compensation is beneficial, even though some variation remains. With the sand echoes, the conclusions are similar. Data were acquired but plots of features are not shown for the Odom200 and the mud site (the Odom200 data for the mud site was irreparably noisy and excluded). Table 6-5 shows a summary of the results as the limits of the tilt at which the feature values become significantly different from the zero tilt case, with and without compensation for tilt. The criterion for two samples being significantly different was 20% or less misassignment. Percent misassignment is basically the percentage of data points that overlap between two distributions. Figure 6-20 shows an example and a full description of this non-parametric statistic. For tilts or slopes greater than these tilt limits, classes may be segmented on the basis of slope, causing misclassification. The tilt limit without compensation is significantly less than one-half the beamwidth, while the compensated tilt limit is much larger. For scenarios that lack a strong coherent echo, such as the Odom200 sand and gravel, Odom24 gravel, slope compensation works for tilts up to the effective beamwidth (11 and 33 °, see chapter 7). The Odom24 sand has a strong coherent component which is lost at the half-beamwidth. Slope compensation cannot compensate for such an effect, it can only account for the stretching of the echo in time, which means measures of amplitude need to be compensated in another way as will be discussed in chapter 7. Slope compensation for the Odom24 mud site overcompensates for slope/tilt and it appears volume backscatter is less affected by slope than is the surface backscatter. This is not a problem because mud seabeds are not found at any significant slope (these sediments have a low angle of repose). Operationally, slope compensation can probably be applied for slopes at least as large as the effective beamwidth; as the compensation, although not perfect, is still preferable to uncompensated echoes. This will be evaluated further in chapter 7.

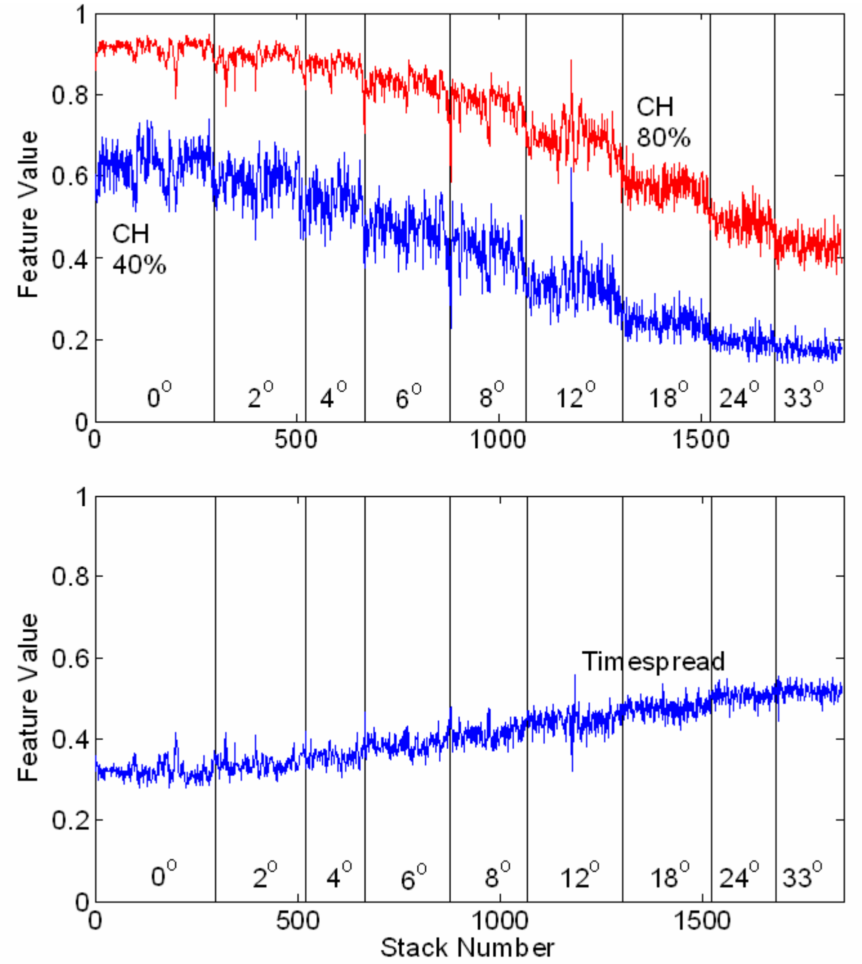

Figure 6-16. Three features calculated from Odom24 kHz echoes collected over the gravel site without slope compensation. The upper plot shows the cumulative amplitude threshold (CT) at 40% and 80%, while the lower plot shows the timespread. The stack number is just an index; the transducer tilt is indicated above it.

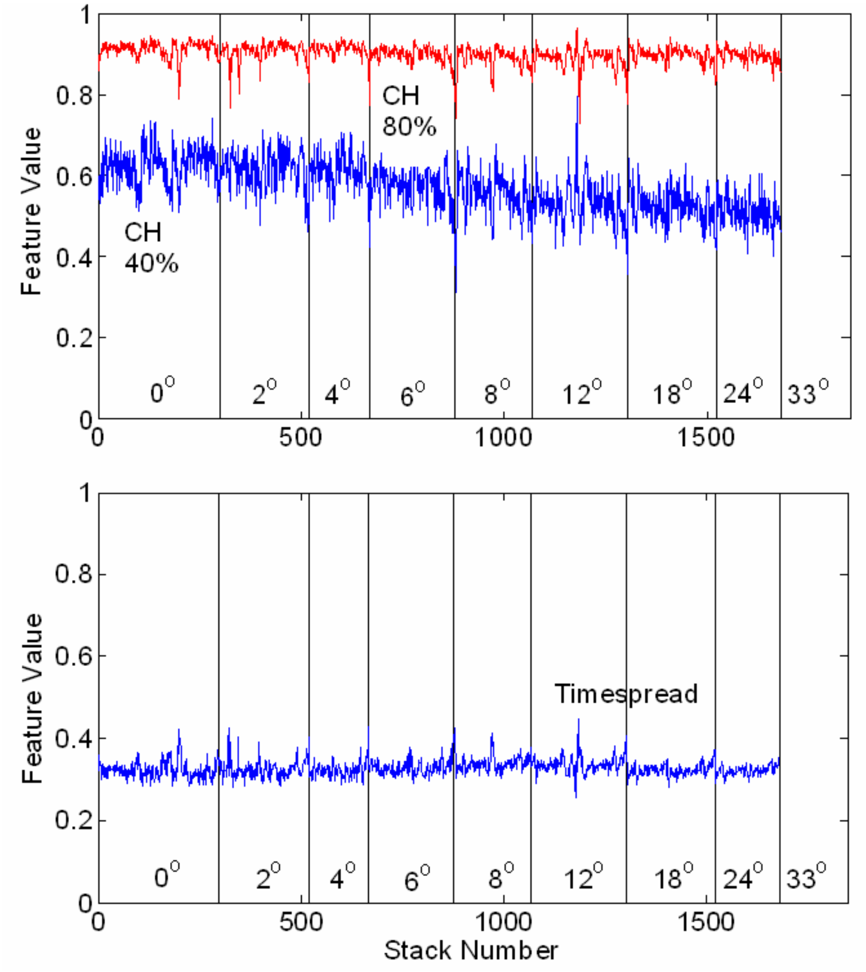

Figure 6-17. The same three features calculated from the same echoes as in Figure 6-16 (Odom24, gravel site), except with tilts compensated by using the transducer tilt values as the seabed slope in Equation 6-1.

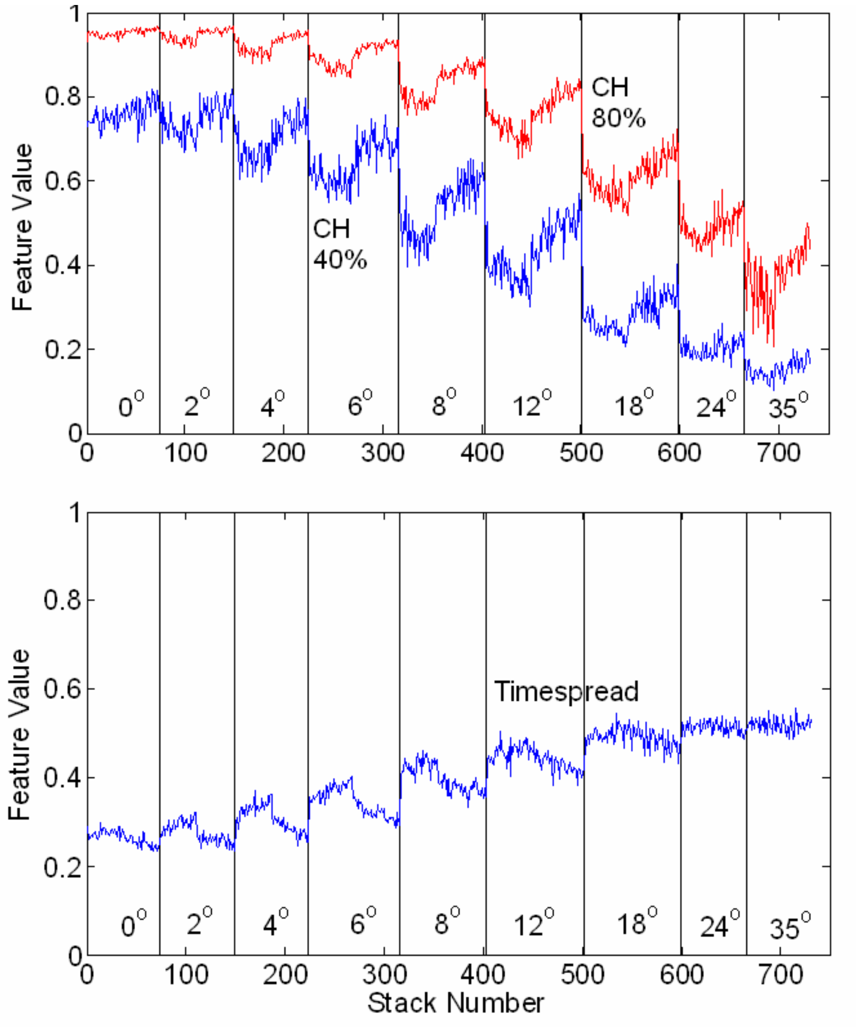

Figure 6-18. The same three features as before calculated from echoes collected over the sand site (Odom24). This site has slope of 1° on average. At each tilt, two passes in opposite directions were made over the area, illustrating the combined effects of tilt and slope.

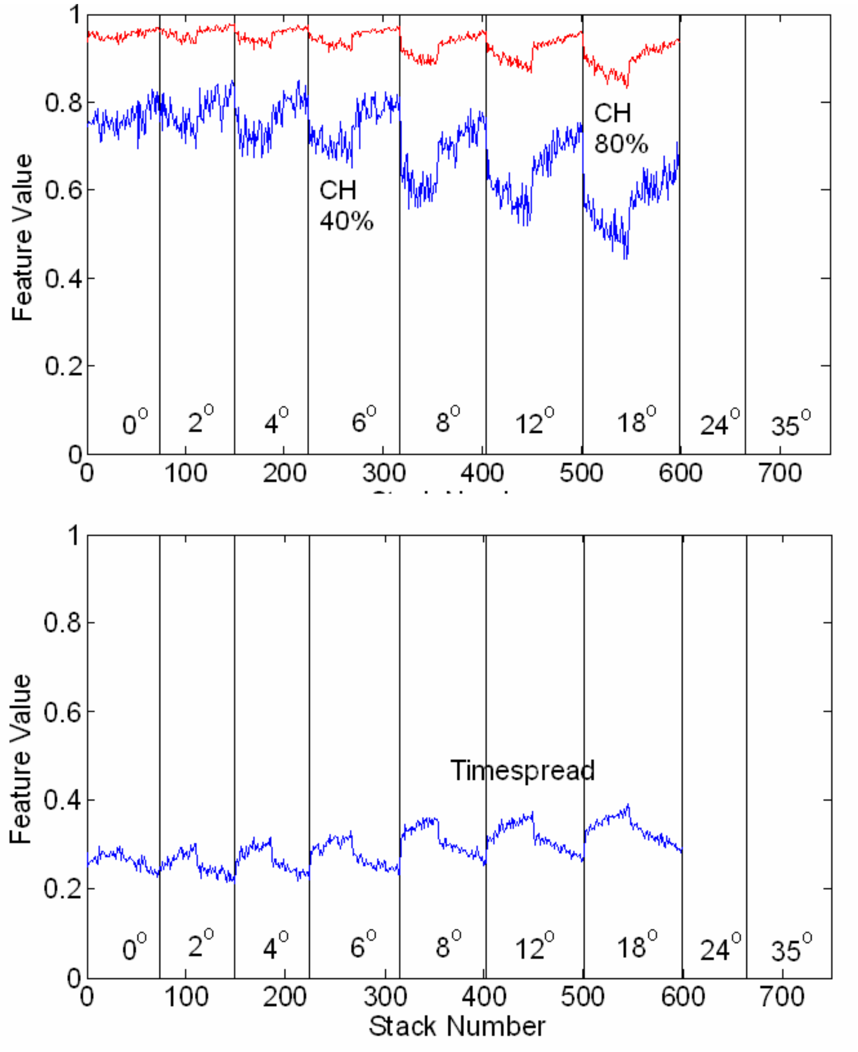

Figure 6-19. The same three features calculated from the same echoes as in Figure 6-18 (i.e. Odom24 over the sand site), except with tilts compensated for by using the transducer tilt values as seabed slope in Equation 6-1. No compensation was attempted for the actual incidence angle (seabed slopes plus tilt).

Table 6-5. Summary of tilt limits for reliable classification, with and without including a value for the seabed slope (tilt) in the SELw calculation for depth and slope compensation.

Site:	System:	Fraction before bottom pick, *n*	Fraction after bottom pick, *m*	Tilt limit	
				No tilt compensation	With tilt compensation
Sand	Odom24	1/8	15/8	4°	8°
Gravel	Odom24	1/8	15/8	6°	24°
Mud	Odom24	1/8	15/4	6°	8°
Sand	Odom200	1/8	15/8	2°	8°
Gravel	Odom200	1/8	15/8	2°	12°

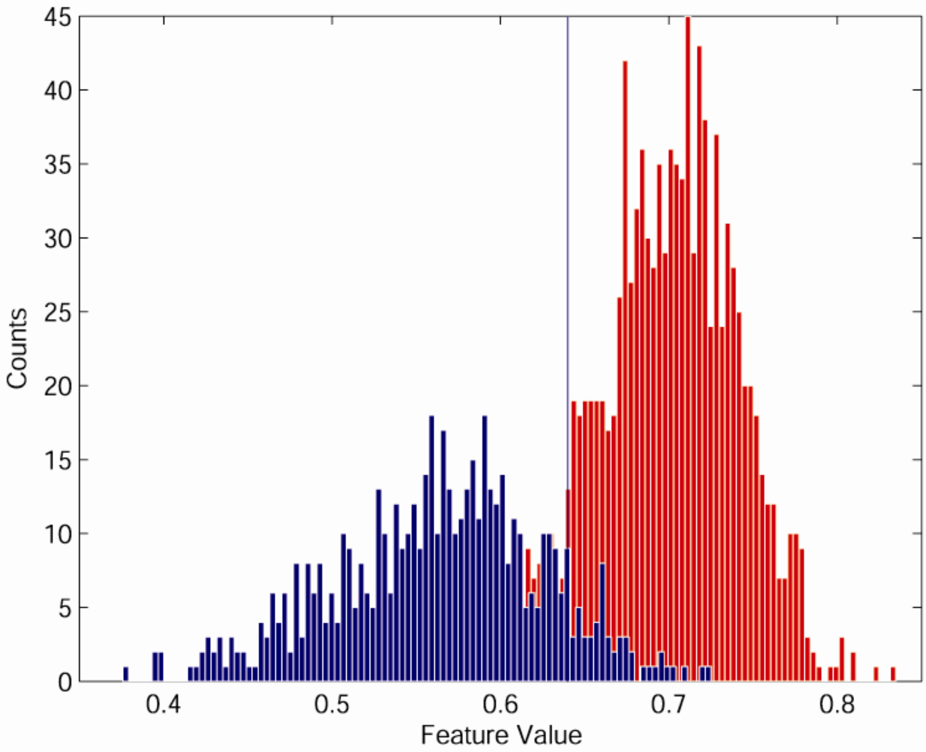

Figure 6-20. An example showing the calculation of percent misassignment. For any given feature, histograms of two samples are plotted in red and blue. The vertical line represents the point where the percentage of points of the blue class to the right of the line is equal to the percentage of points of the red class to the left of the line. That percentage is the probability of misassignment for classification of data to either sample that could be achieved for known sample distribution type, means and covariances. As such, this statistic is non-parametric.

6.3.3 Discussion and Conclusions

This experiment has shown that echo features can be compensated for seabed slope and transducer tilt, within limits. As a rule, slope can be compensated for up to one-half the effective beamwidth and above this, compensation is preferable to not compensating at all. Slopes above the one-half beamwidth cause a change in the scattering regime that cannot be compensated for in the method, i.e. the specular reflection no longer backscatters to the echosounder as discussed in section 5.3. Such data can be discarded. Seabed slope information can be difficult to get and to use. Survey bathymetry has to be of sufficient resolution and then has to be processed carefully. Accurate SELw compensation for slope would require the slope at each echo location, estimated by the average slope value over the acoustic footprint. However, even crude slope data would be a great improvement over none at all. In an actual survey, data from a pitch/roll sensor on the vessel could be used in addition to slope data to compensate for vessel attitude to improve survey reproducibility, which as Gleason observed, was reduced in modest seas causing modest tilts [Gleason, 2009].

6.4 A Discussion of Classification Errors Caused by Seabed Slope and Standard Depth Compensation Parameters

An interesting consequence of these experiments is that one can explain certain classification errors caused by depth compensation with the standard parameters. Figure 6-21 shows the ratio of echo duration modelled with default parameters to echo duration modelled with effective parameters. The effective parameters are from the depth experiment results, so there is a different set of parameters for each of the experiment sites. Since standard classification is primarily based on shape features, which are all correlated to echo duration, these ratios are useful for illustrating the causes of classification artifacts related to improper depth compensation and a lack of slope compensation. The ratio is nearly unity for echoes from sand seabeds, but is much greater for mud and gravel echoes. The effect increases for sloping seabeds. Class confusion will result if any of the lines cross. Figure 6-21 illustrates two cases of class confusion: the (*upper*) plot shows that slightly sloping sand can classify as gravel, while the (*lower*) plot shows that slightly sloping gravel can classify as mud. Using effective parameters with depth and slope compensation will prevent such misclassifications.

In general, a potential problem with processing data in analysis windows is the danger of truncating echoes. The analysis window is usually only twice the expected echo duration, i.e. a maximum ratio of 2.0 in Figure 6-21. Without slope compensation, truncation occurs for modest slopes. As demonstrated in the depth experiment, Figure 6-12, depth compensation using the effective beamwidth and attenuation avoids truncation errors (compare the green lines to the magenta lines). Even with depth and slope compensation with effective beamwidth and attenuation, it is important to have the capacity to automatically tag truncated echoes as part of data quality assurance. This capacity is built into the data analysis methods utilized in chapter 7 as part of the new seabed classification suite about to be presented.

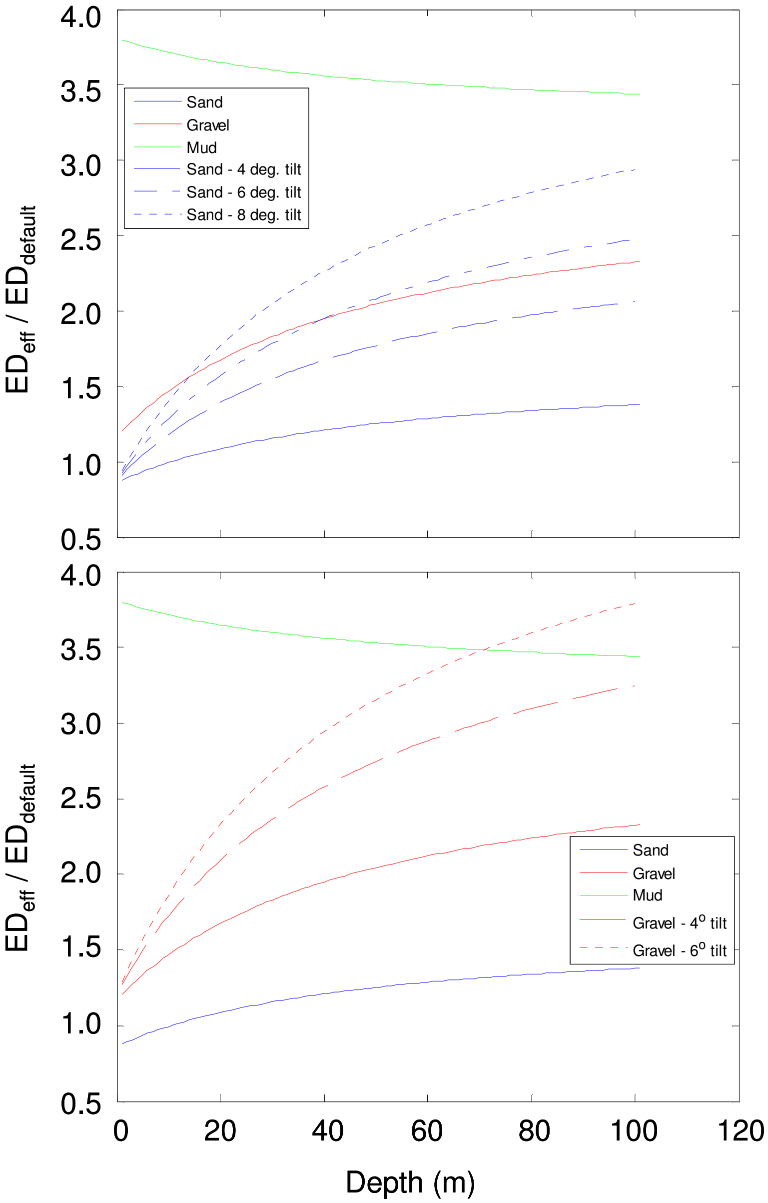

Figure 6-21. Ratio of echo duration calculated with effective parameters to echo duration calculated from the standard or default parameters (beamwidth: 24°, attenuation: 0.5 dB/m/kHz) for the three homogenous testbed sites (for the Odom24 only). The dashed / dotted lines show the ratios with added tilt for the sand site (*upper*) and the gravel (*lower*) site.

Chapter 7. New and Improved Seabed Classification Methods

The overall objective of this chapter is to present and evaluate new methods of seabed classification based on the knowledge gained through the research presented thus far. The chapter focuses on two major issues that are shown to generate significant improvements in classification performance: the removal of non-seabed influences and poor discrimination. Two new methods are presented here to address these issues: the SELw effective depth and slope compensation algorithm and the introduction of physics-based features. The expectation for both is that classification will be improved as will be evaluated both qualitatively and quantitatively. An additional benefit of physics-based features is that these features facilitate the development of three feature classification and relative characterization (section 7.8) and feature interpretation, which will be presented as part of post-class characterization in chapter 8.

This chapter begins with an overview of the new classification and characterization methodology (section 7.1) presented in chapters 7 and 8. Following that is a description of all of the tools used to evaluate the various classification methods (section 7.2). The evaluation is based on a comparison to the results of standard QTC IMPACT processing on a data set from the Patricia Bay testbed; these baseline results are shown in section 7.3. The SELw effective depth and slope compensation algorithm is introduced and evaluated as part of section 7.4, followed by the introduction and evaluation of new physics-based features that are used to generate the final classification method and results (section 7.5). Section 7.6 presents supervised classification results using the new features. Section 7.7 is the conclusion for the new classification method.

7.1 An Overview of the New Seabed Classification Methodology

Prior to proceeding with the remainder of this chapter, it is useful to show how the new methods will work within the new classification methodology that is the subject of both chapters 7 and 8. To that end, Figure 7-1 shows a flowchart summarizing the new methodology, which will serve as a reference for this chapter and for chapters 8 and 9.

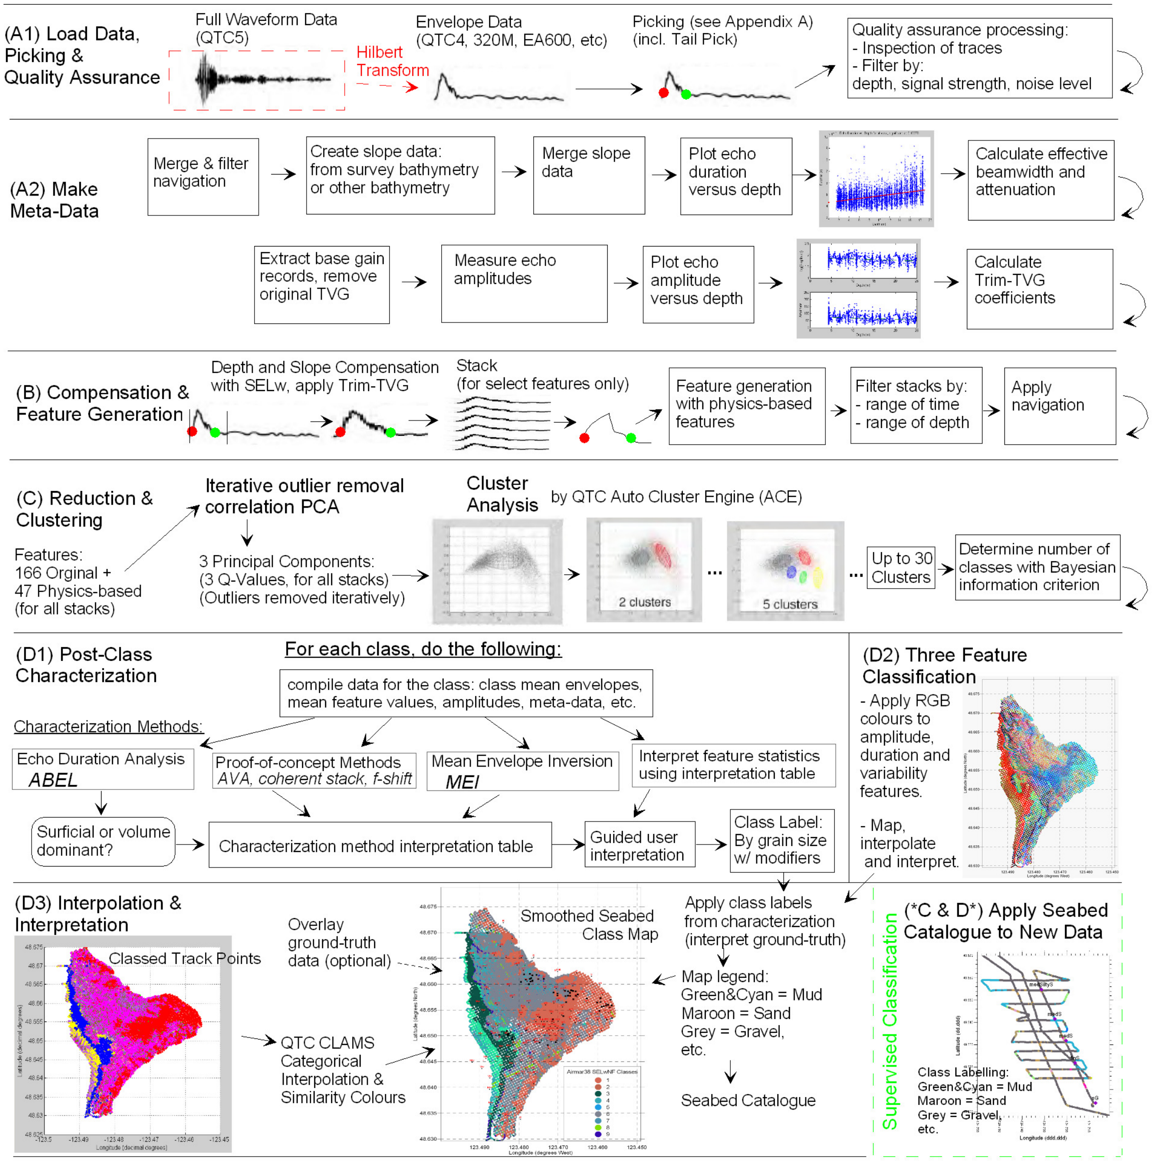

Figure 7-1. The new seabed classification methodology, including SELw depth and slope compensation, physics-based features, three-feature classification and post-class characterization.

The methodology is based on the original QTC VIEW seabed classification methodology (Figure 2-4), and has the same four basic steps: picking and quality assurance (A), compensation (removal of non-seabed influences) and feature generation (B), reduction and clustering (C), interpolation and interpretation (D). In the new methodology, new methods either have replaced old methods or have been added as shown by Figure 7-1. The SELw effective depth and slope compensation algorithm (section 7.4) involves improvements to steps (A1), (A2) and (B). Section 7.5 involves aspects of stacking, feature generation and PCA in steps (B) and (C). The subject of section 7.6, supervised classification is shown as the alternate to steps (C) and (D) – (*C & D*). Three feature classification (D2) serves as a relative characterization method, providing information to label classes with seabed type information similar to the role of post-class characterization (D1), which is presented in chapter 8.

The methods of chapters 7 and 8 are intended to be applied together as shown in Figure 7-1, providing a complete seabed classification methodology that is robust to any condition or survey scenario and provides reliable seabed type information without the need for extensive ground-truth. The methodology is also designed to provide a breadth of tools that can be used together to satisfy many survey objectives. The robustness of unsupervised classification is partially due to fact that the number of classes produced is not fixed, with the end result reflecting the natural boundaries in the data. However, if a survey objective was to classify on the basis of the major seabed types, i.e. gravel, sand, mud, additional classes may be merged into the main classes using the tools presented in the next section. This requires the labelling of classes with characterization information provided by the methods of section 7.8 and chapter 8. Although the flow chart of Figure 7-1 presents an overview the methodology, its use and flexibility will only become fully developed once the methods of both chapters 7 and 8 are presented.

7.2 Analysis Methods – Tools for Evaluating the Performance of New and Improved Classification Methods

There are a number of tools required to evaluate the performance of new and improved classification methods. These tools are either qualitative or quantitative and can apply directly to the methods or apply to evaluate the results of the classification. The analysis methods are summarized in Figure 7-2; a flowchart showing the data sources for the methods and how some methods are related. Following the figure, the methods will be discussed in the order shown in Figure 7-2, starting with "Classification maps and qualitative comparisons" and proceeding counter-clockwise.

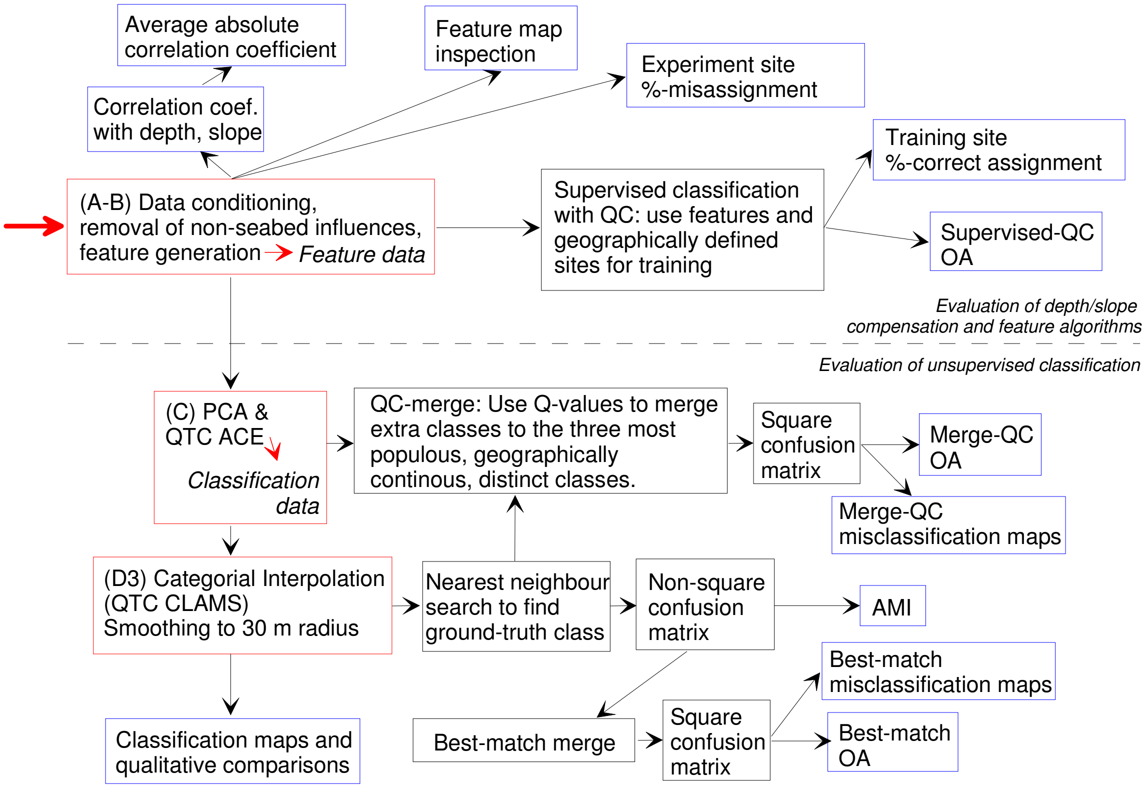

Figure 7-2. A summary of the methods used to evaluate the new and improved classification methods. The red boxes are the starting points, which refer to steps in Figure 7-1. Black boxes are intermediate steps, while the blue boxes are the statistics or qualitative results used for the evaluation.

In general, it is useful to start with the qualitative comparisons of classification maps to help understand the results of the confusion matrix and the resulting OA and

AMI statistics. In this chapter, the comparisons will be made to a ground-truth seabed class map, and are made in reference to certain geographic areas that are either problem areas or areas that are very well known, such as the experimental sites discussed in section 6.1.3. It should be noted that evaluating classification performance is a difficult task, complicated by the nature of the process: classification is phenomenological, while ground-truth data are often from different frequency echosounders, video or grab-samples that are difficult to compare with SBES classification results. This type of analysis is rarely done for seabed classification and has not been applied to evaluate improvements in classification methodology. The closest examples of objective analysis of classification performance appear in the following: [Preston et al., 1999; Gleason, 2009; Gleason et al, 2009; Wilmut et al., 2009].

The qualitative comparisons need to be verified quantitatively. All quantitative measures are derived from confusion matrices. Each row of a confusion matrix represents the instances in a predicted class, while each column represents the instances in an actual or ground-truth class. In this case, each cell is the percentage of the total data that has both the survey class (row) and the ground-truth class (column) of the corresponding row and column headings. The assignment of ground-truth classes to the survey data points is done spatially. Each point in the testbed survey is assigned the nearest ground-truth seabed class grid point. If a survey point is more than 5 m away from a ground-truth class grid point, it is excluded from further analysis.

In general, confusion matrices are applied to compare separate sets of results with the same number of classes, so the confusion matrix is ideally square. A common measure of classification performance, the **o**verall **a**greement (OA), is only applicable to square matrices. For acoustic seabed classification, square matrices can be achieved by choosing a consistent number of classes in ACE, which would result in poor classification results, or by merging classes upon identifying which classes represent the same seabed. Since the number of ground-truth classes in this application is either five or three[47], the number of testbed survey classes must be made to match.

---

[47] Two of the ground-truth classes can be merged into others, as will be discussed later in this section.

There are two methods of merging classes to be presented; both are shown in Figure 7-2: best-match merge and merge-QC. To discuss these methods, example confusion matrices are required. The survey classification data used are from standard QTC IMPACT processing for envelope data (TNORM) from a survey of Patricia Bay to be presented in section (7.3). Because the multibeam map (Figure 6-1) compared so well to the ground-truth (grab-sample, video and STING data, see section 6.1), it is taken as the ground-truth seabed class map of Patricia Bay. As such, these classes are referred to as ground-truth classes. The multibeam classification map and the survey maps have been interpolated with QTC CLAMS [QTC, 2004c] to a grid on a radius of 40 m so that each assignment of a ground-truth class to a survey point is approximately the modal class within the acoustic footprint. The best-match merge procedure is described next.

From a seabed classification point of view, every survey class must be assigned a best-match ground-truth label, i.e. each survey class is assigned a ground-truth label based on which ground-truth class occurs most frequently within that survey class. This is usually done with ground-truth interpretation tables; examples of this procedure were discussed in reference to Table 6-1, and can be found in [Biffard et al., 2009]. These tables are essentially confusion matrices, so the same procedure is applied here. The criterion for merging survey classes would be by their best-matches as highlighted in green in Table 7-1 and Table 7-3. The survey classes (rows) that have best-matches with the same ground-truth classes are merged by row addition. Using this procedure for merging the classes in Table 7-1 would result in only two classes, matching the gravel and the fine sand ground-truth classes, leaving three rows of the matrix as zeros (Table 7-2). The alternative to leaving empty rows is to choose classes not to merge, though doing so is rather arbitrary. Either way, this would leave three ground-truth classes without best-matches to any survey classes. The situation is better for the Simrad200 survey with only two empty rows (Table 7-4), but still not satisfactory.

In general, if the survey classes are not consistent with the boundaries of the ground-truth seabed map, merging classes is difficult and may lead to unreliable results. It appears that merging is only viable when the survey class map matches the ground-truth class map sufficiently well such that there are no empty rows after merging the survey classes by best-matches; these results can be evaluated with the OA statistic defined in the following paragraphs. However, it is clear that a statistic such as the **a**verage **m**utual **i**nformation (AMI) that can evaluate non-square matrices would be useful. The AMI statistic will follow the introduction of OA, while a second method of merging classes; merge-QC will be presented later.

Table 7-1. Confusion matrix for the Airmar38 TNORM processed survey using the multibeam seabed class map as the ground-truth classification. The green numbers indicate the best-matches between the two sets of classifications. All confusion matrix results are presented as percentages of the total number of data points in the survey.

True class ►	Gravel	Med. Sand	Fine Sand	Mud	S. Mud
TNORM class ▼	1	2	3	4	5
1	5.1	6.8	**7.6**	2.5	1.4
2	**16.0**	9.0	3.5	0.4	0.1
3	**3.0**	0.6	0.0	2.1	0.3
4	**6.1**	1.9	0.4	2.5	0.7
5	**11.3**	1.2	0.1	9.7	1.0
6	**1.5**	0.1	0.0	0.1	0.4
7	0.2	0.6	**2.3**	0.7	0.2
8	**0.2**	0.1	0.0	0.1	0.1

Table 7-2. Square confusion matrix for the Airmar38 TNORM (TN) testbed survey, as found by merging classes within the full confusion matrix (Table 7-1) by best-matches.

True class ►	Gravel	Med. Sand	Fine Sand	Mud	S. Mud
Merged TN class ▼	1	2	3	4	5
1	38.1	12.9	4.0	14.9	2.7
2	0	0	0	0	0
3	5.3	7.3	9.9	3.2	1.6
4	0	0	0	0	0
5	0	0	0	0	0

Table 7-3. Full confusion matrix for the Simrad200 TNORM testbed survey compared to the ground-truth seabed map. Best-matches are highlighted in green.

True class ►	Gravel	Med. Sand	Fine Sand	Mud	S. Mud
TNORM class ▼	1	2	3	4	5
1	0.2	0.3	0.1	**3.0**	0.1
2	**2.8**	0.6	0.2	0.2	0.3
3	**4.3**	2.4	1.8	0.4	0.3
4	0.1	0.2	0.2	**1.3**	0.1
5	**4.3**	0.5	0.2	0.2	0.4
6	**3.1**	0.5	0.0	0.4	0.4
7	**0.6**	0.1	0.1	0.1	0.1
8	**0.3**	0.0	0.0	0.1	0.1
9	**1.6**	0.1	0.1	0.1	0.4
10	0.3	0.3	0.3	**4.9**	0.6
11	0.2	0.4	0.6	**3.2**	0.8
12	0.1	0.1	0.1	**0.4**	0.1
13	**2.6**	0.6	0.2	0.2	0.2
14	3.8	5.6	**8.1**	2.9	0.7
15	**5.1**	1.3	0.3	0.1	0.3
16	**10.5**	6.3	1.6	0.7	0.6
17	**1.3**	1.1	0.6	0.1	0.1

Table 7-4. Square confusion matrix for the Simrad200 TNORM (TN) testbed survey, as found by merging classes within the full confusion matrix (Table 7-3) by best-matches.

True class ►	Gravel	Med. Sand	Fine Sand	Mud	S. Mud
Merged TN class ▼	1	2	3	4	5
1	36.4	13.6	5.1	2.7	3.1
2	0	0	0	0	0
3	3.8	5.6	8.1	2.9	0.7
4	0.9	1.2	1.2	12.8	1.8
5	0	0	0	0	0

For a square matrix organized so the largest values or best-matches fall on the main diagonal, the **o**verall **a**greement (OA) is the matrix's trace divided by the sum of the matrix [Gleason et al., 2009; Legendre and Legendre, 1998]. When comparing survey classes to the ground-truth classes, the OA statistic expresses the fraction of the data set that classified correctly. When the best-match merge procedure results in empty rows, the OA values tend to be artificially high due to the concentration of co-occurrences onto the main diagonal of the matrix. This effect is compounded if there are empty rows despite having many more survey classes than ground-truth classes, as was the case for Table 7-4.

Overall agreement is also biased upwards by the chances of random correct classification. A square confusion matrix of random classifications results in an OA of 100% divided by the number of classes (row or column). OA can be corrected for this effect, and the result is the Kappa statistic. Kappa is potentially useful here, but OA values are easier to interpret so are used instead. For more information on Kappa, see [Gleason et al., 2009; Finn, 2003].

The AMI statistic can be calculated from non-square or rectangular confusion matrices directly. It is a measure of the information common to two nominal data sets that was derived from information theory, specifically the concept of entropy [Finn, 1993] - see that paper for the formulas used to calculate AMI from confusion matrices. Finn evaluates AMI with several case studies, with idealized and real examples of square and rectangular matrices [Finn, 1993]. AMI is used here to express the percentage similarity between the survey class map and the multibeam ground-truth class map. It ranges from 0%, indicating the class maps share no information, to 100%, when the class maps are identical. Unlike OA, AMI accounts for the entries off of the main diagonal, so that it does not require the confusion matrix to be organized with the largest entries (best-matches) on the main diagonal. For instance, consider an identity matrix with all the rows swapped so that no entries lie on the main diagonal; OA would give zeroes, while the AMI would be 100%. Now consider a confusion matrix with a trace of 100% and then subtract a fraction from the diagonal elements evenly, and redistribute that fraction equally throughout the entire matrix – when that fraction is 100% the matrix is filled with equal values (i.e. 11% for a 3 by 3 matrix). This is the idea behind Figure 7-3, which shows how AMI and OA are related when there are evenly-distributed off-trace elements throughout the matrix. Kappa is shown here as it exactly measures the fraction subtracted from the trace. This figure shows OA is higher than Kappa and much higher than AMI in that generally applicable scenario. For instance, for a three-class OA of 70%, only 24% AMI may be observed. As the off-trace elements within the confusion matrix become non-evenly distributed, AMI will increase from 24%. AMI accounts for complex relationships between classes, while OA is only concerned with the matrix trace and is influenced by the number of classes in the matrix.

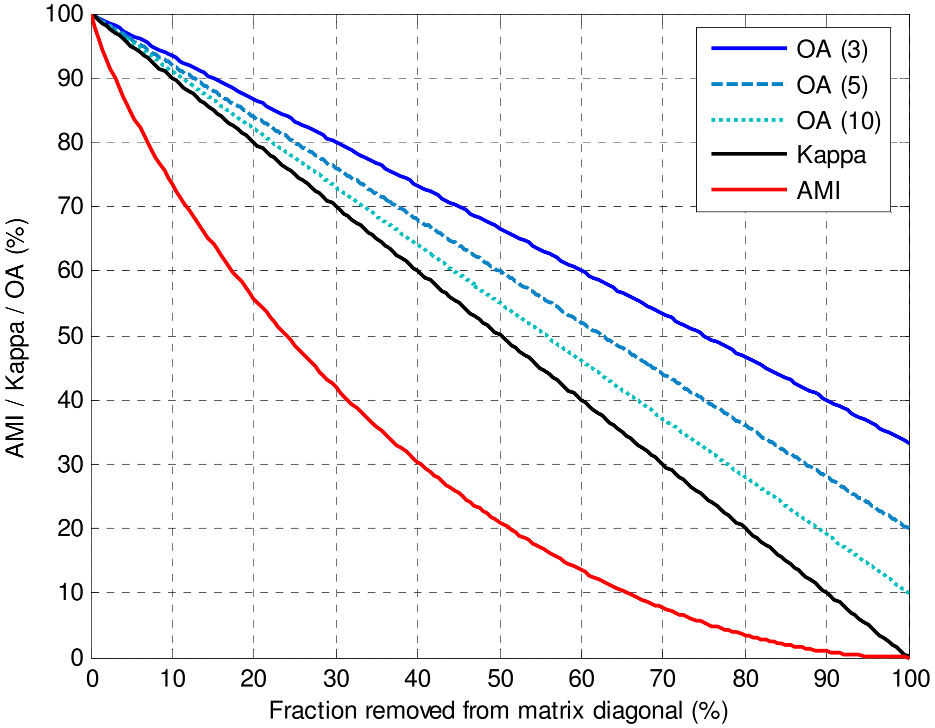

Figure 7-3. The relationships between the three statistical measures of classification performance. Three scenarios of square confusion matrices are shown: 3, 5 and 10 classes. The Kappa statistic is shown for reference, showing the fraction removed from the matrix diagonal and redistributed evenly.

Confusion matrices can also be used to create misclassification maps. These maps show the spatial distribution of correct and incorrect classification. The proportion of correctly classified data of the total is equal to the OA. Unfortunately, the approach of the AMI statistic cannot be used to create misclassification maps, or significance tests. The best-match procedure can be used to make these maps, as well as the other class merging method, merge-QC, which will be discussed next.

The process of quadratic classification (QC) applied here is described in [Wilmut et al., 2009], and for more background information on the method refer to [Rencher, 2002]. There are two types of QC used here, merge-QC and supervised-QC (S-QC) and it is best to describe the standard application of QC first that is common to both applications. In short, QC is a general supervised classification method. It assigns new data to any of the known classes based on a metric. In this case, the Mahalanobis distance is used, because no Bayesian prior information exists, which is different from [Wilmut et al., 2009] where a Bayesian metric is used. Data of known classes are used initially as training sets to establish the mean and covariance of each class. These training sets can be defined spatially, or can be defined from existing class definitions (see sub-section 2.2.5 for a description of supervised classification methodologies). For each new data point after training, QC computes the metric to each of trained classes, and assigns the datum to the closest class. The Mahalanobis metric can incorporate many features, so an *n*-dimensional feature space is generated.

Here is where the merge-QC and supervised-QC (S-QC) differ. S-QC skips PCA and cluster analysis (step C in Figure 7-1). It is very similar to standard QTC supervised classification (see Figure 2-4) except that it is easier to define training sites spatially with S-QC. One major difference is that in S-QC only the features that best separate the training sets are used as defined by percent correct assignment[48], as poor features reduce the performance of S-QC. This evaluation of feature discriminatory power will be applied to evaluate the new features in section 7.5. Of the best features, uncorrelated features are selected to maximize the effectiveness of the classification. In addition to statistics on

---

[48] Percent correct assignment is distinct from OA as it is not derived from a confusion matrix. It is the proportion of data points of the training sites that are correctly assigned to those sites using QC. This statistic also appers in Figure 7-2.

feature selection, the other products from S-QC include class maps, the OA between the S-QC classification and the ground-truth seabed class map and misclassification maps (see Figure 7-2). One of the benefits of using S-QC is that it can objectively create square confusion matrices if the training sets are selected to represent ground-truth classes (i.e. the homogenous experiment sites in sub-section 6.1.1).

As discussed in detail earlier in this sub-section, it is difficult to merge acoustic classes to create a square confusion matrix. QC offers an alternative to the best-match merge procedure: merge-QC (see Figure 7-2). Instead of directly classifying on the features and training sites, QC is applied to the Q-values from PCA and trains on a number of classes. The number of classes used is equal to the number of ground-truth classes in order to generate a square confusion matrix. This process effectively merges the data from the superfluous classes into the training or main classes according to the position of the data in Q-space. The main classes are selected, without knowledge of the ground-truth, from the largest classes (by number of data points) that also represent homogenous areas of seabed and span Q-space[49]. For an example of this procedure, see the Q-space plot in Figure 2-9 that shows three classes that were selected to train on, to which all the other data were re-classified. To improve the results in Patricia Bay, three classes (gravel, sand, mud) were used instead of five for QC analysis[50]. Merging the two sand ground-truth classes and the two mud ground-truth classes into broader sand and mud classes is acceptable as these ground-truth classes are quite similar (see ground-truth seabed map, Figure 6-1 and the grab-sample analysis of sub-section 6.1.3). With the main classes selected, survey data is assigned to the main classes based on the Mahalanobis distances calculated from the Q-values, instead of the features. In all, merge-QC is an objective way of creating square confusion matrices to evaluate the results of unsupervised classification by PCA / cluster analysis.

S-QC and its resulting statistics (OA and percent correct assignment) are a few of the statistics used to evaluate the performance of the new depth and slope compensation

---

[49] Note that other means of selecting the main classes are available, but are not used here: on the basis of ground-truth data or, most notably, classes characteristic of specific seabed types, i.e. gravel, sand, mud, can be identified without prior knowledge using the three-feature classification method (section 7.8) or the post-class characterization method (chapter 8).

[50] OA statistics were calculated from the best-match merge using the five original ground-truth classes.

algorithms and physics-based features. The others, which are shown in Figure 7-2, are described below. Percent misassignment, percent correct assignment and feature inspections are only used for the evaluation of feature performance in section 7.5. It is more convenient to discuss these tools in detail in that context. Percent misassignment was introduced in sub-section 6.3.2, where an example was shown - Figure 6-20. The difference between percent misassignment and percent correct assignment is that percent misassignment is a non-parametric measure of discrimination between two seabed types, while percent correct assignment is calculated for multiple training sites. Although feature map inspection is qualitative, objective criteria are applied to evaluate feature performance. These maps are very useful as a diagnostic tool for features.

The final tool shown in Figure 7-2 to be described is the correlation coefficient. The effects of the different depth and slope compensation methods on individual features are inferred by the relative values of the correlation coefficient of each feature to depth and slope. To evaluate the compensation of depth and slope as a whole, an absolute average of correlation coefficients is taken over a number of features used in every classification method. The features used are the eight cumulative integral and quantile features listed as 'Used' in Table 7-17 and Table 7-18. These features influence the classification outcome as they are weighted strongly by PCA in all cases (see appendix B). As depth and slope compensation is improved from the baseline result, the average absolute correlation coefficients are expected to decrease. However, some correlation between the features and depth and slope is expected because seabed types tend to be correlated with depth and slope. This is why depth and slope biased classification can perform relatively well. In the case of Patricia Bay, the mud seabed is entirely below 90 m on a nearly flat seabed. Rocky seabeds usually have high slope. The correlations should not to go to zero with the most effective depth and slope compensation.

7.3 The Baseline to Compare the New Methods to: QTC IMPACT (TNORM) Processing for Envelope Data

In order to evaluate the improvements in classification performance made by the new methods, a baseline must be established for comparison. That baseline is the seabed classification result produced by standard QTC IMPACT processing of a dual-frequency SBES data set that covers most of Patricia Bay. The remainder of this section presents this SBES survey of Patricia Bay that will serve as the testbed data set, followed by the TNORM classification result which is evaluated relative to the ground-truth multibeam seabed map, using the qualitative and quantitative methods of the previous section.

7.3.1 The Patricia Bay Testbed Full SBES Survey

The dual-frequency Airmar38/Simrad200 SBES survey of Patricia Bay testbed is used for all methods of chapter 7 and 8. This data set was chosen because it offers full coverage of the bay, covering the same area as the multibeam survey. The survey includes all of the varied terrain and seabed types that make Patricia Bay one of the most challenging classification scenarios one could face. If the new methods are effective here, they should be similarly effective anywhere else. The echosounder systems used are the Airmar38 and Simrad200 as shown in Table C-1, except with slightly longer pulse durations of 0.400 and 0.316 ms respectively. The data set was collected with a Knudsen 320M echosounder head unit, which uses analogue circuitry to band-pass the signal and create envelopes – FWF data were not collected. The amplitude envelopes were sampled at 24 kHz. The survey consists of eight survey days spaced (due to weather) over four weeks in November and December of 2004. With this amount of survey time, a high survey line density was achieved, with survey lines spaced at 50 m in an orthogonal grid with additional grids offset by 45° to provide added diagonal crossings over ~60% of the survey, as shown in Figure 7-4. Also, since the footprint size of the Airmar38 is 58% of the water depth, the 50 m line spacing gives complete coverage for water depths greater than 86 m. However, complete coverage is not achieved with the Simrad200. The survey was carried out by J. Galloway and his group aboard the CCGS Otter Bay with some assistance from the author. The survey appears here courtesy of the CHS.

7.3.2 Standard QTC IMPACT Processing with TNORM

The only depth compensation method normally available in QTC IMPACT for envelope data is TNORM or reference depth compensation, as described in detail in chapter 2 (Figure 2-4 in particular). In this case, the reference depth was 40 m. Based on the restrictions noted in chapter 2, limits on the data were establish between 16 and 125 m water depth, and for uniformity this limit is held throughout chapter 7. The standard QTC bottom picking routine was used, tuned to maximize sensitivity while minimizing poor picks. After depth compensation and feature generation, clustering with ACE finds 8 classes for the Airmar38 survey and 17 classes for the Simrad200 as selected by the lowest Bayesian information criterion. The resulting seabed classification maps are shown in Figure 7-4.

In these figures, one can see the multibeam survey extends to waters both deeper and shallower than the single-beam survey. This is due to the aforementioned limits imposed by TNORM depth compensation[51]. However, the multibeam class data files were available so the data were interpolated on the same search radius as the single beam results. All classification results are shown in similarity colours to best show the relationships among classes. Note that attempting to colour match the classes between different results may lead to inappropriate conclusions. It is better to study the borders of the classes, with additional information on which classes are similar as provided by similarity colours.

---

[51] Unfortunately, the depth filtering was not available on the multibeam survey because the MULTIVIEW database was lost. Recovering the data would be time consuming.

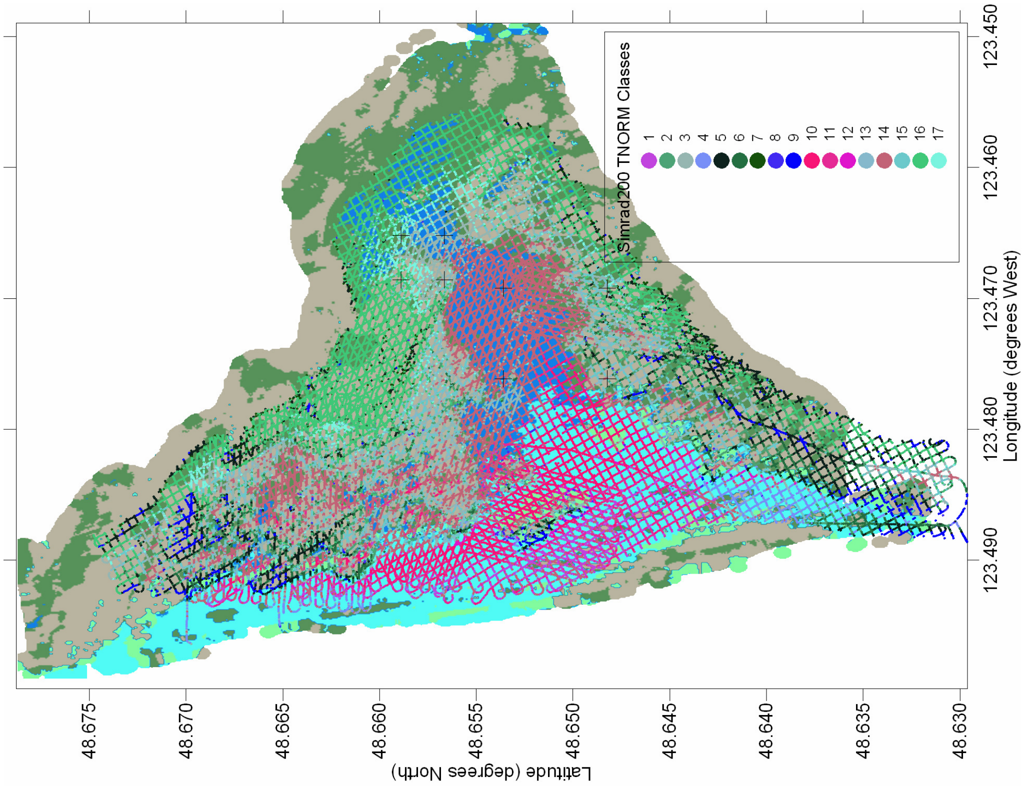

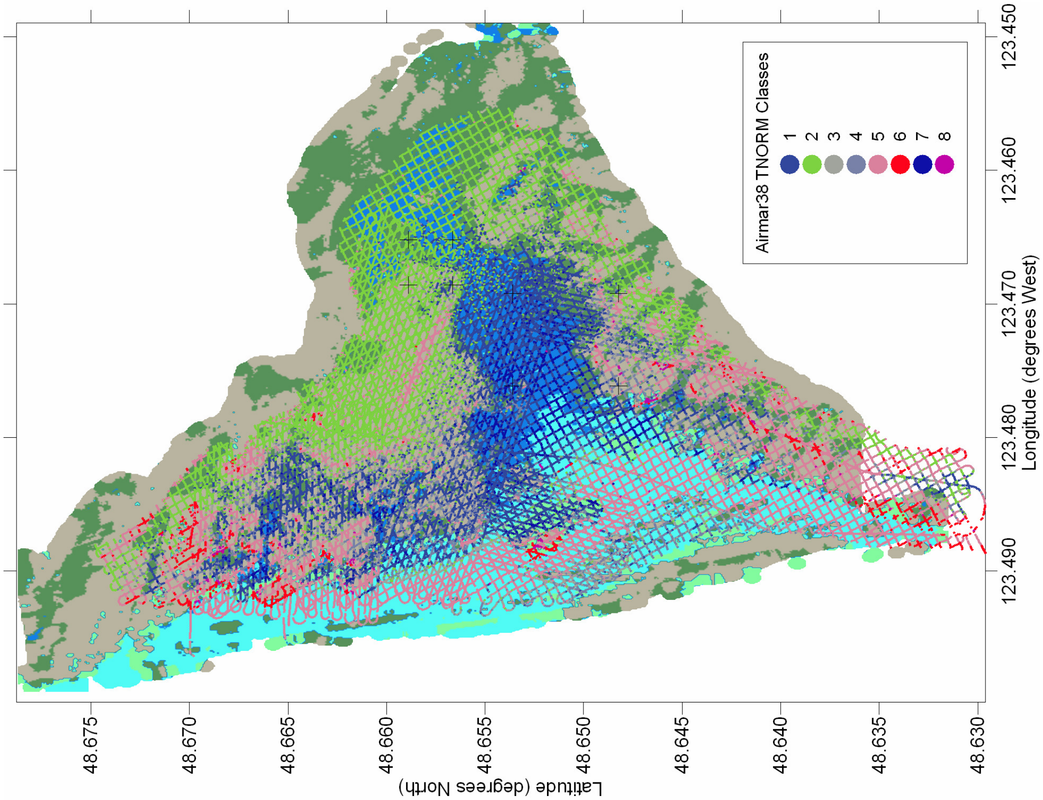

Figure 7-4. Airmar38 (*upper*) and Simrad200 (*lower*) seabed classification from the standard processing method overlaid on the EM3000 multibeam seabed map. Only the corners of CM1 and CM2 are marked (thin crosses) to avoid obstructing the map.

7.3.3 Evaluation of the TNORM Classification Results

The techniques described in section 7.2 to evaluate the quality of classification results are applied here, starting with qualitative comparisons of the classification maps.

At both frequencies, the overlay of the TNORM classification on the multibeam ground-truth map (Figure 7-4) shows the TNORM result modestly represents the class boundaries. The testbed transition sites are marked in these figures, where one can observe that the sand-gravel transition is not evident at either frequency because the site is encompassed in a shallow water class. The boundary between the medium sand and gravel is evident for the complex transition site, although the consistency of the classification is not satisfactory. (See Figure 6-1 for a map of all the testbed experiment sites).

More comparisons are available in the side by side comparisons with the ground-truth seabed class maps: Figure 7-5 and Figure 7-6. The single-beam classes represent the boundary between the ground-truth fine sand and gravel classes well, except in the areas shown in the green and yellow boxes. (These are the blue classes (1 and 7) for Airmar38 and the burgundy class (14) for Simrad200; see the confusion matrices for more details: Table 7-1, Table 7-3). However, the distinction between the ground-truth medium sand and fine sand is also not made. The classification errors in the yellow boxes are due to classification by water depth – a shallow water class encompasses the medium sand, fine sand and gravel seabeds for water depths less than approximately 25 m for both frequencies. The green boxes indicate areas that are predominately gravel but are classified with the predominately fine sand classes at both frequencies. This is likely due to a depth artefact in that the fine sand classes have the same depth as this area. The seabed slope in this area may also contribute as it is moderate (6° to 12°, more or less in isolated areas, refer to the slope map in Figure 6-2). The red boxes show examples of high slope (>12°) causing further class confusion. For the Airmar38, high slope areas and mud areas are confused for each other, which is understandable as BORIS simulations and real echoes show that echoes from mud sediments look like those from high slopes: both are long, low amplitude and variable. For the Simrad200, this area is its own high slope classes – the dark blues and greens (classes 5 to 9), which are also present at shallow depths at the limit of 16 m. This extreme shallow water and high slope class identifies signals of truncated echoes, caused by either high slope or erroneously high upsampling of the echo time series by TNORM in shallow water. In TNORM compensation, shallow water echoes are erroneously up-sampled to infinite length at zero depth (see Equation 2.5) – this leads to severe truncation of the echo in the analysis windows. As indicated by the purple boxes in Figure 7-5 and Figure 7-6, the boundary of the mud seabed is not followed particularly well in either result, which is surprising as the mud is so different (in terms of geoacoustic properties) from the sands and gravels. The boundary of the mud seabed is very evident for most of the bay on the multibeam bathymetry (Figure 6-3). For the Airmar38, the misclassified area (purple box in Figure 7-5) could be a symptom of the differing penetration depths of the frequencies being compared here. For the Simrad200, the classes that represent mud in the shallow mud area (classes 10, 11, 12) extend into the fine sand seabed. The deeper mud area is split among several classes with these classes appearing elsewhere in non-mud areas. An example is highlighted by the more southern purple box in Figure 7-6. It is clear that depth and slope artefacts are causing a majority of the misclassifications observed. They may be responsible for additional problems because of shifts of class boundaries to accommodate these artefacts, as the effects of depth and slope can be much stronger than the variation in echoes due to the seabed.

The OA and AMI quantitative statistical measures of classification performance are shown in Table 7-5 to compare with the above discussion and the misclassification maps that appear below. Merge-QC OA was not obtained for the Simrad200 data set, because training classes were difficult to define for the TNORM result and for the SELwNF result to be presented, merge-QC tended to concentrate all the data into one class.

Table 7-5. Classification performance for the TNORM results. These are the baseline results, with values in percent.

TNORM	Airmar38	Simrad200
Best-match OA (%)	48	57
No. classes:	8	17
Merge-QC OA (%)	53	-
AMI (%)	15	13

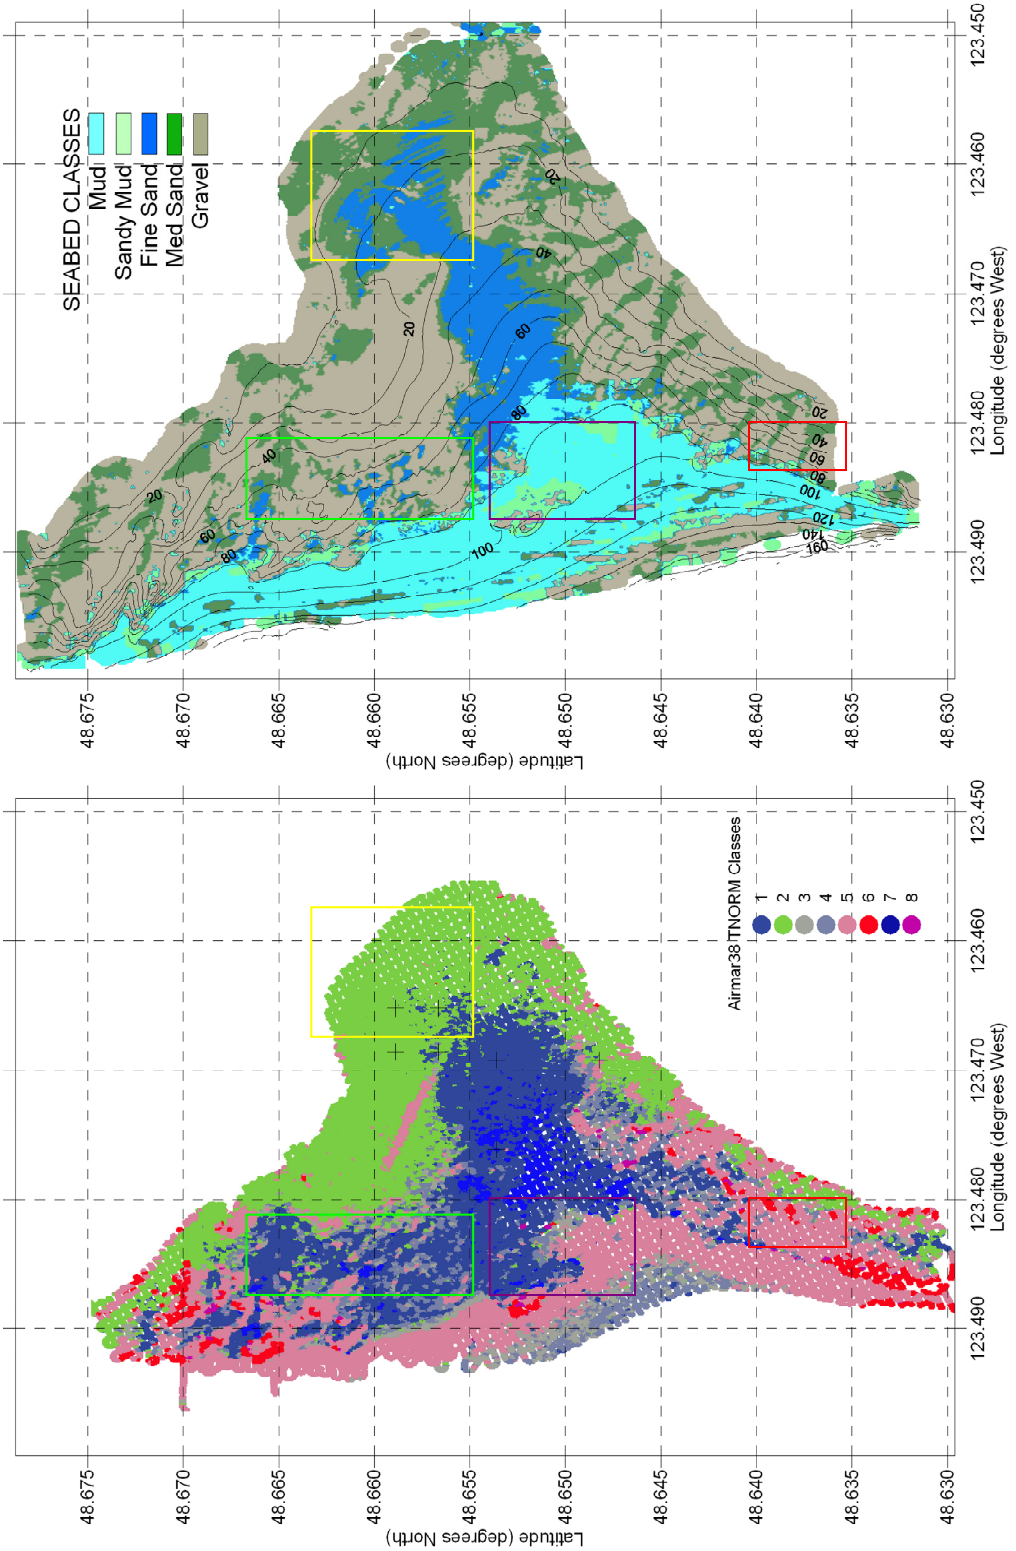

Figure 7-5. A comparison of the TNORM classification map for the Airmar38 testbed survey (*lower*) to the multibeam derived ground-truth class map (*upper*). The boxes highlight areas of misclassification. Both maps were generated by categorical interpolation with QTC CLAMS on a search radius of 30 m and displayed in similarity colours.

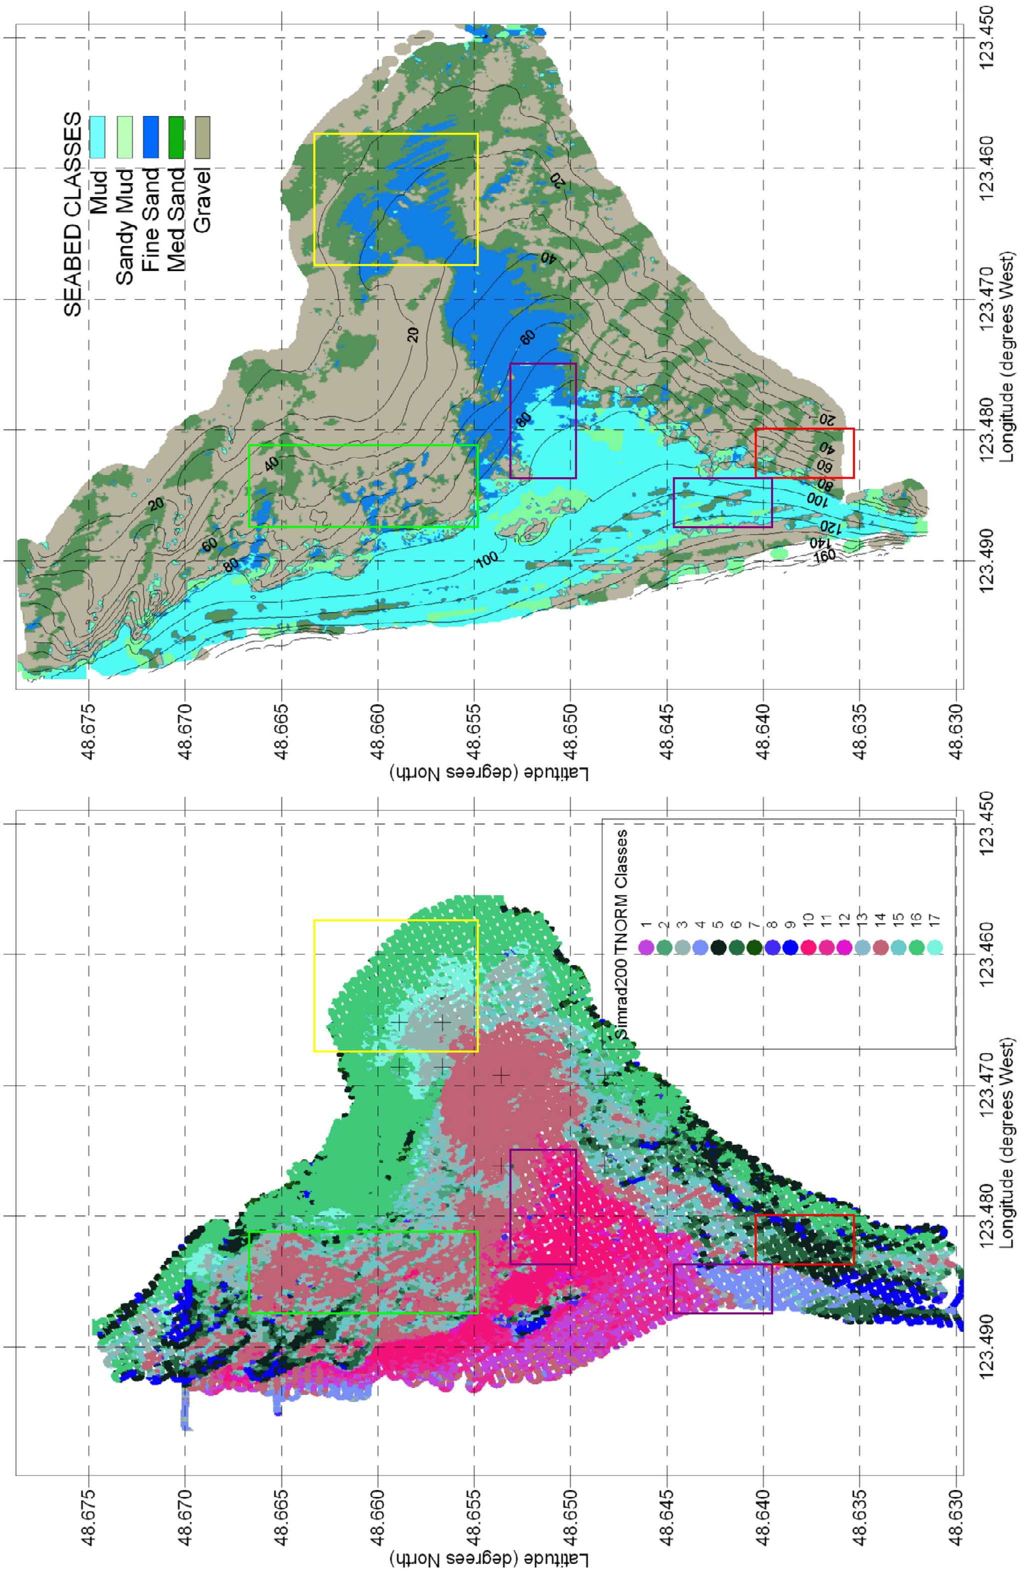

Figure 7-6. A comparison of the TNORM classification map for the Simrad200 testbed survey (*lower*) to the multibeam ground-truth class map (*upper*). The boxes highlight areas of misclassification. Both maps were generated with a search radius of 30 m.

Figure 7-7 shows overlays of the correctly and incorrectly classified survey data points associated with the best-match merge procedure. The correctly classified points are those where the points agree with their class labels and vice versa. These maps are useful to confirm the observations made of the class maps. Indeed, the locations of the boxes highlighting areas of misclassification (see Figure 7-5, Figure 7-6) are also areas of high rates of misclassification shown here. The Simrad200 misclassification looks substantially better than the Airmar38, mostly because the Simrad200 has a class that is assigned as mud, while the Airmar38 does not. The Airmar38 does better in most other areas. This illustrates that misclassification maps and the associated OA values can be misleading when comparing different numbers of acoustic classes to ground-truth classes.

It is clear that the standard processing of echo envelopes via TNORM in QTC IMPACT results in many misclassifications. With this analysis and the knowledge base built up through chapters 2 through 6, causes of poor classification are clear: depth and slope biases and poor discrimination.

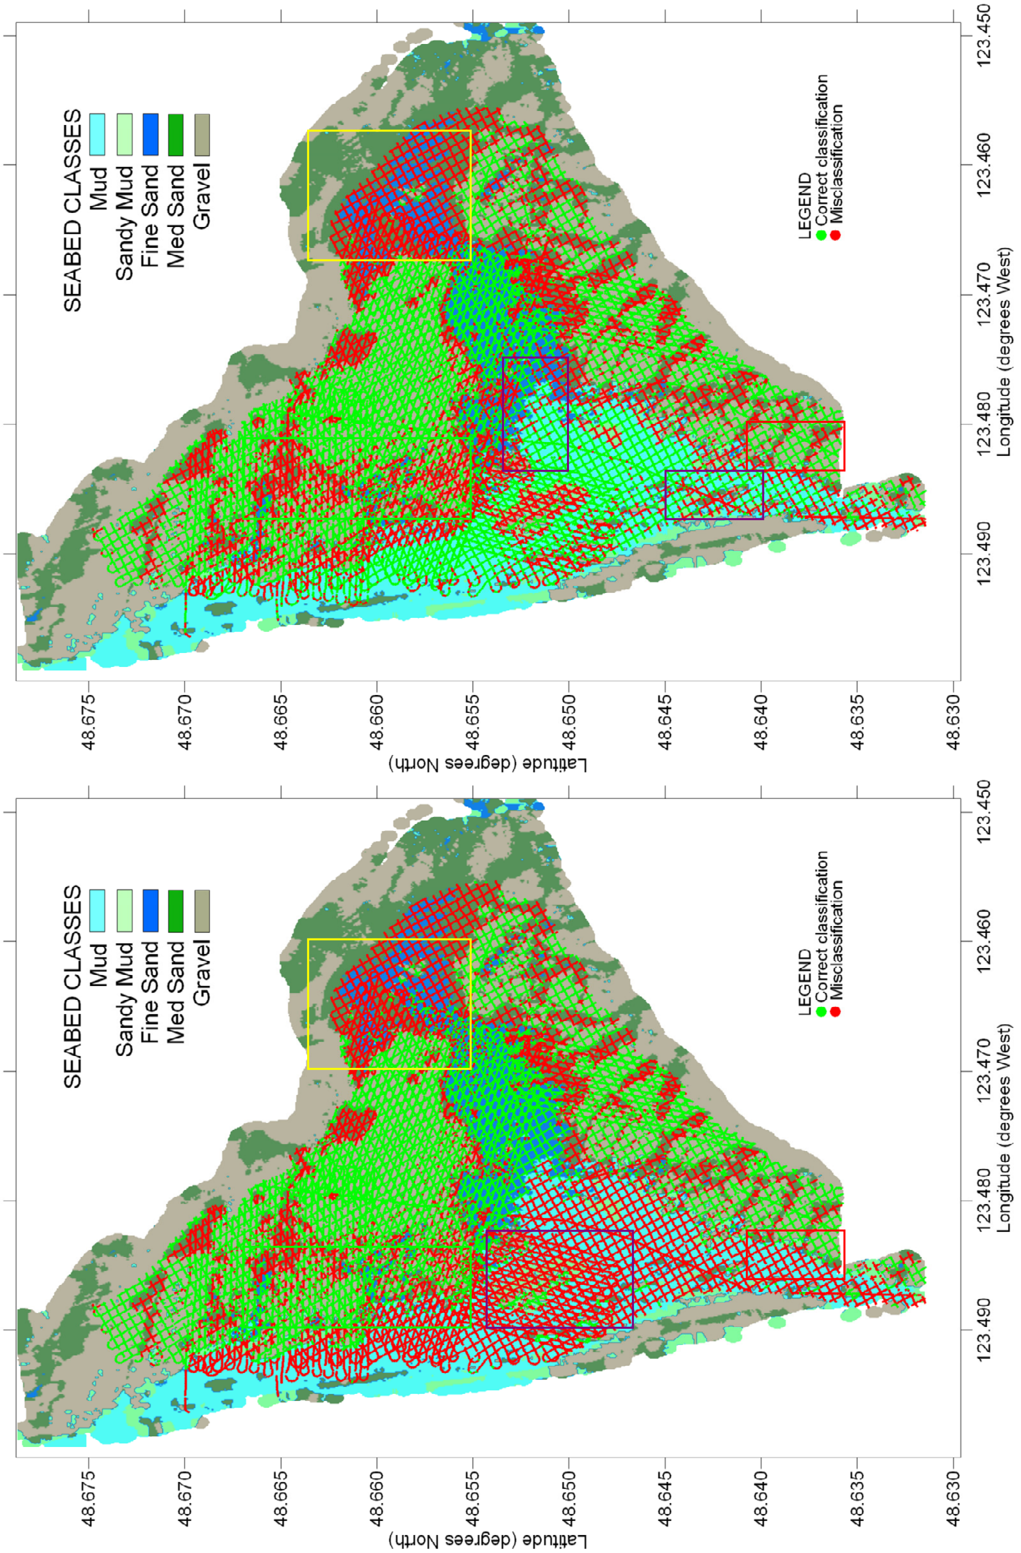

Figure 7-7. Maps of correct and incorrect classification for TNORM processed Airmar38 (lower) and Simrad200 (upper) testbed surveys. Correct classifications are determined from the best-match procedure - the best-match OA represents the percent correct on these maps. The boxes highlight the same areas as those in Figure 7-5 and Figure 7-6.

7.4 Removal of Non-Seabed Influences: Effective Methods of Depth and Slope Compensation

The MATLAB suite developed for this work was designed to accommodate different versions of the processing steps in order to facilitate research and development. Several different depth and slope compensation methods are available. However, this section will concentrate on the final method SELw, presenting only a synopsis of the other, less advanced methods.

7.4.1 A Synopsis of The Intermediate Depth and Slope Compensation Methods

The first advancement over the TNORM method is QTC's SEL depth compensation as introduced in sub-section 2.2.1. SEL only applies to QTC5 FWF data and not to envelope data. As part of the development of more advanced methods, SEL was adapted to process envelope data. The author applied this improved method for ROV-mounted SBES seabed classification surveys as presented in [Bloomer et al., 2007]. In that case, the problems with seabed slope and poor discriminatory ability of the standard feature set were still evident. The result of applying SEL for envelopes to the Patricia Bay testbed surveys is strong depth and slope artefacts similar to or even worse than TNORM classification. This is caused by the values used for beamwidth and attenuation. The -3 dB beamwidths and assumed attenuations of 0.5 db/kHz/m were used, as is common practice for SEL depth compensation. As shown in sub-sections 5.4.3 and 6.2.4, these values are erroneous.

The next advancement is to apply effective beamwidth and attenuation to SEL compensation. To do this, a linear regression analysis is run on the echo duration data for the full survey, using Equation 5-6, to calculate effective beamwidth and attenuation and a similar procedure to that described in section 6.2. There are two differences between the analysis in section 6.2 and the analysis here: slope data are not used and the effective beamwidth and attenuation represent an average over the different seabed types within the survey. This procedure to find the effective beamwidth and attenuation is the one that would be used for surveys that do not have seabed slope information[52]. In the classification methodology flow chart, Figure 7-1, step (A2), the make meta-data step, is now required. The tail picks were generated with the RANLS method based on the thresholds and parameters established in section 6.2. With accurate and consistent tail picks, analysis of the echo duration data resulted in effective beamwidth and attenuation estimates as shown in Table 7-6. The effective beamwidths are similar to those found in sub-sections 5.4.3 and 6.2.4, and are quite wide compared to the -3 dB beamwidths of 16° and 7° for the Airmar38 and Simrad200 respectively. This is due to the effects of high roughness and low attenuation seabeds, and seabed slope. Seabed slope is responsible for inflating these values by approximately 10° (the average value of seabed slope throughout Patricia Bay), and it is accounted for in the next advancement.

Table 7-6. Regression results for effective beamwidth and attenuation without any slope information – Patricia Bay testbed survey.

	Airmar38	Simrad200
Effective beamwidth (deg.)	47.8	22.4
Eff. beamwidth 95% C.W.	0.1	0.1
Effective attenuation (dB/kHz/m)	0.23	0.32
Eff. attenuation 95% C.W.	0.01	0.02
r^2	0.52	0.30
Mean residual (%)	19.9	15.8

As discussed in section 6.3.1, applying seabed slope to SEL compensation for envelopes with effective beamwidth and attenuation (EBA) is not practical on an operational basis because the resampling required in SEL tends to cause aliasing (violating the Nyquist sampling theorem). This occurs for data from water depths greater than a critical depth that is:

$$d_{C\phi} = \frac{c_w}{2\left(\sec\left(\frac{\theta}{2}+\phi\right)-1\right)}\left(\left(\frac{n_{SEL}}{s}-1\right)\tau-\frac{20}{c_w \kappa f_c}\right) \tag{7-1}$$

---

[52] The only situation where slope information would not be obtainable is if the survey was not run on a grid – i.e. a single line.

The various slopes of Patricia Bay cause the critical depth for the Airmar38 to range from 20 to 340 m. To address this, the new method of SELw (section 6.3.1) was applied, which skips the resampling step.

7.4.2 The Implementation of SELw Effective Depth and Slope Compensation

As described in section 6.3.1, the SELw depth and slope compensation method adjusts the size of the analysis window to accommodate the echo duration as predicted by the EDM, using effective beamwidth and attenuation combined with slope data.

To apply SELw to survey data, the effective beamwidth and attenuation are found from the whole survey; accounting for slope with the two techniques described in subsection 6.2.3: slope subtraction and the iterative slope removal technique. Both are applied on either the full survey or a low-slope subset of the survey, and all four results for both frequencies are shown in Table 7-7. The *final* values for the effective beamwidth and attenuation in Table 7-7 are from a grid search to minimize the mean percent residual for the whole survey. The size of the grid was 0.1 ° for the effective beamwidth and 0.01 dB/kHz/m for the effective attenuation.

Table 7-7. Measures of effective beamwidth and attenuation with slope information, using the four different methods as introduced in section 6.2.3 – the slope *subtract*ion and *iterative* slope removal methods and their low-slope (*LS*) variants.

SBES ▶	Airmar38				
Method ▶	Subtract	Subtract LS	Iterative	Iterative LS	Final
Effective beamwidth (deg.)	37.6	37.4	35.9	36.8	33.0
Eff. beamwidth 95% C.W.	7.3	6.3	0.1	0.1	-
Effective attenuation (dB/kHz/m)	0.23	0.22	0.22	0.23	0.30
Eff. attenuation 95% C.W.	0.00(5)	0.00(1)	0.00(4)	0.00(4)	-
r^2	0.52	0.54	-	-	-
Mean residual (%)	17.4	17.5	12.2	13.7	5.6

SBES▶	Simrad200				
Method ▶	Subtract	Subtract LS	Iterative	Iterative LS	Final
Effective beamwidth (deg.)	12.3	14.0	did	12.8	11.0
Eff. beamwidth 95% C.W.	7.4	2.4	not	0.1	-
Effective attenuation (dB/kHz/m)	0.32	0.37	converge	0.26	0.30
Eff. attenuation 95% C.W.	0.02	0.04	-	0.01	-
r^2	0.30	0.38	-	-	-
Mean residual (%)	10.4	21.0	-	17.2	2.5

The echo duration residual maps ((*c,d*) in Figure 7-8) generated with the final effective beamwidth and attenuation values display less variation throughout the fine sand seabed (centre of map) as compared to the *iterative LS* maps ((*a,b*) in Figure 7-8. These maps also have improved the contrast between the sands and gravels, and the gravels and muds. The echo duration residual is also improved in areas of high seabed slope. The echo duration residuals are ideally positive in cases of high roughness (gravel) and low attenuation (mud), assuming a median seabed of medium sand. The final effective beamwidth and attenuation values also marginally improve the classification results over those obtained by using the *iterative LS* estimates (not shown for brevity).

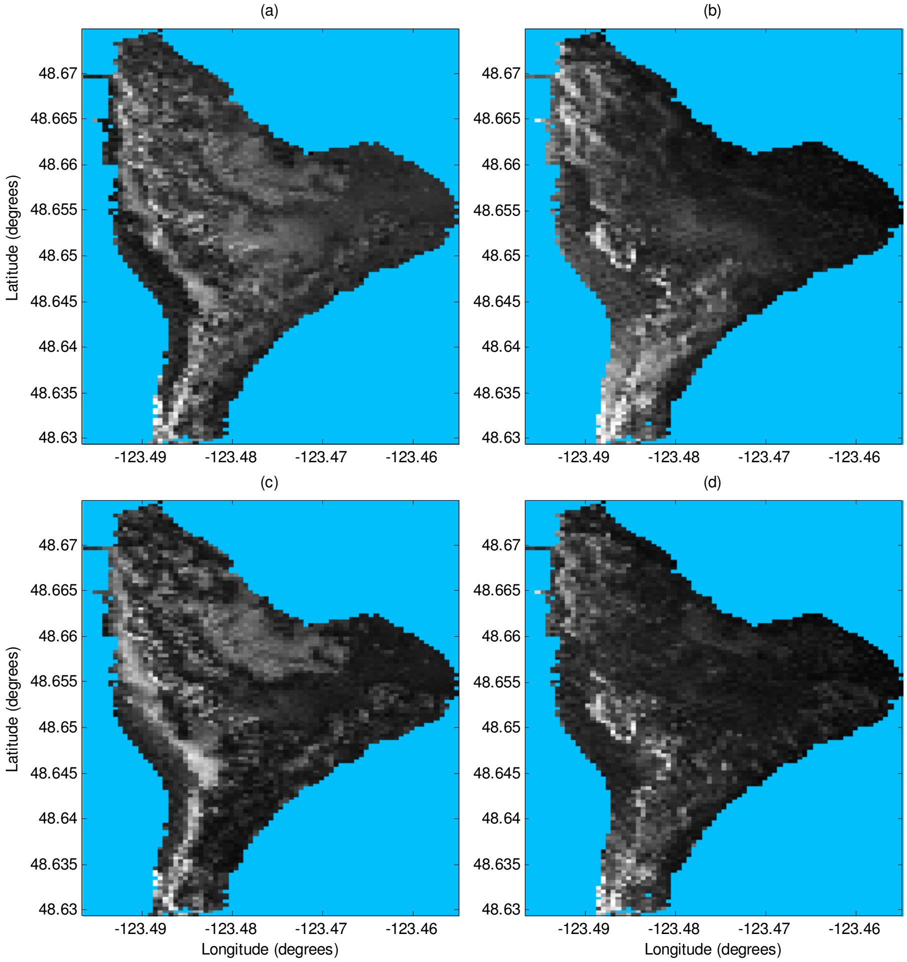

Figure 7-8. Echo duration residual maps generated with slope information. Based the iterative LS parameters (*a,b*) or the final parameters (*c,d*) for either the Airmar38 (*a,c*) or the Simrad200 (*b,d*) testbed survey. The grey scale is linear from minimum to maximum values.

7.4.3 Qualitative Evaluation of SELw Depth and Slope Compensation

Now that the effective beamwidth and attenuation values are established, the classification results for SELw can be observed.

There are no visible depth or slope artefacts for the Airmar38 SELw classification map, Figure 7-9. The classification differs from the ground-truth along the mud-fine sand boundary within the purple box. This effect will be observed consistently in later results and is likely due to the penetration of a thin mud layer over gravel by the 38 kHz SBES, while 300 kHz MBES is more sensitive to the thin layer. The Airmar38 SELw classification is not very proficient at drawing clear boundaries between the gravel and sand classes that are seen in the ground-truth map. New features need to be added to improve the discrimination.

Airmar38 SELw result is in contrast to that for the Simrad200 SELw result (Figure 7-10). Reasons for this will be discussed below.

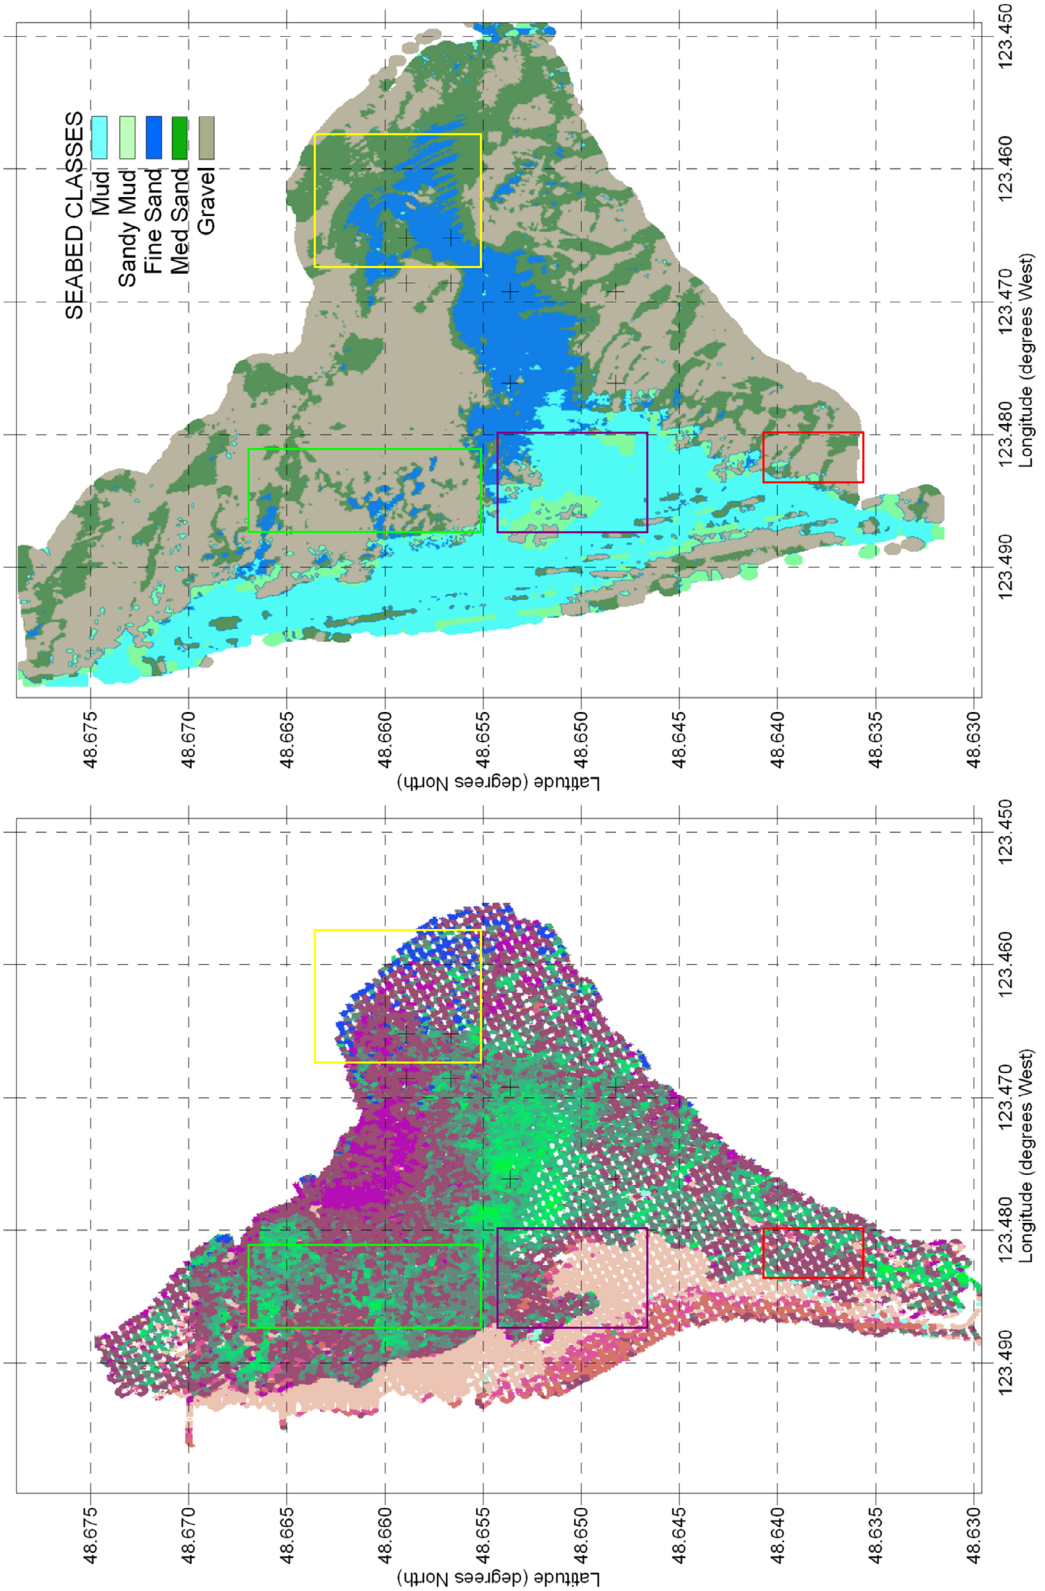

Figure 7-9. A comparison of the SELw classification map for the Airmar38 testbed survey (*lower*) to the multibeam derived ground-truth class map (*upper*).

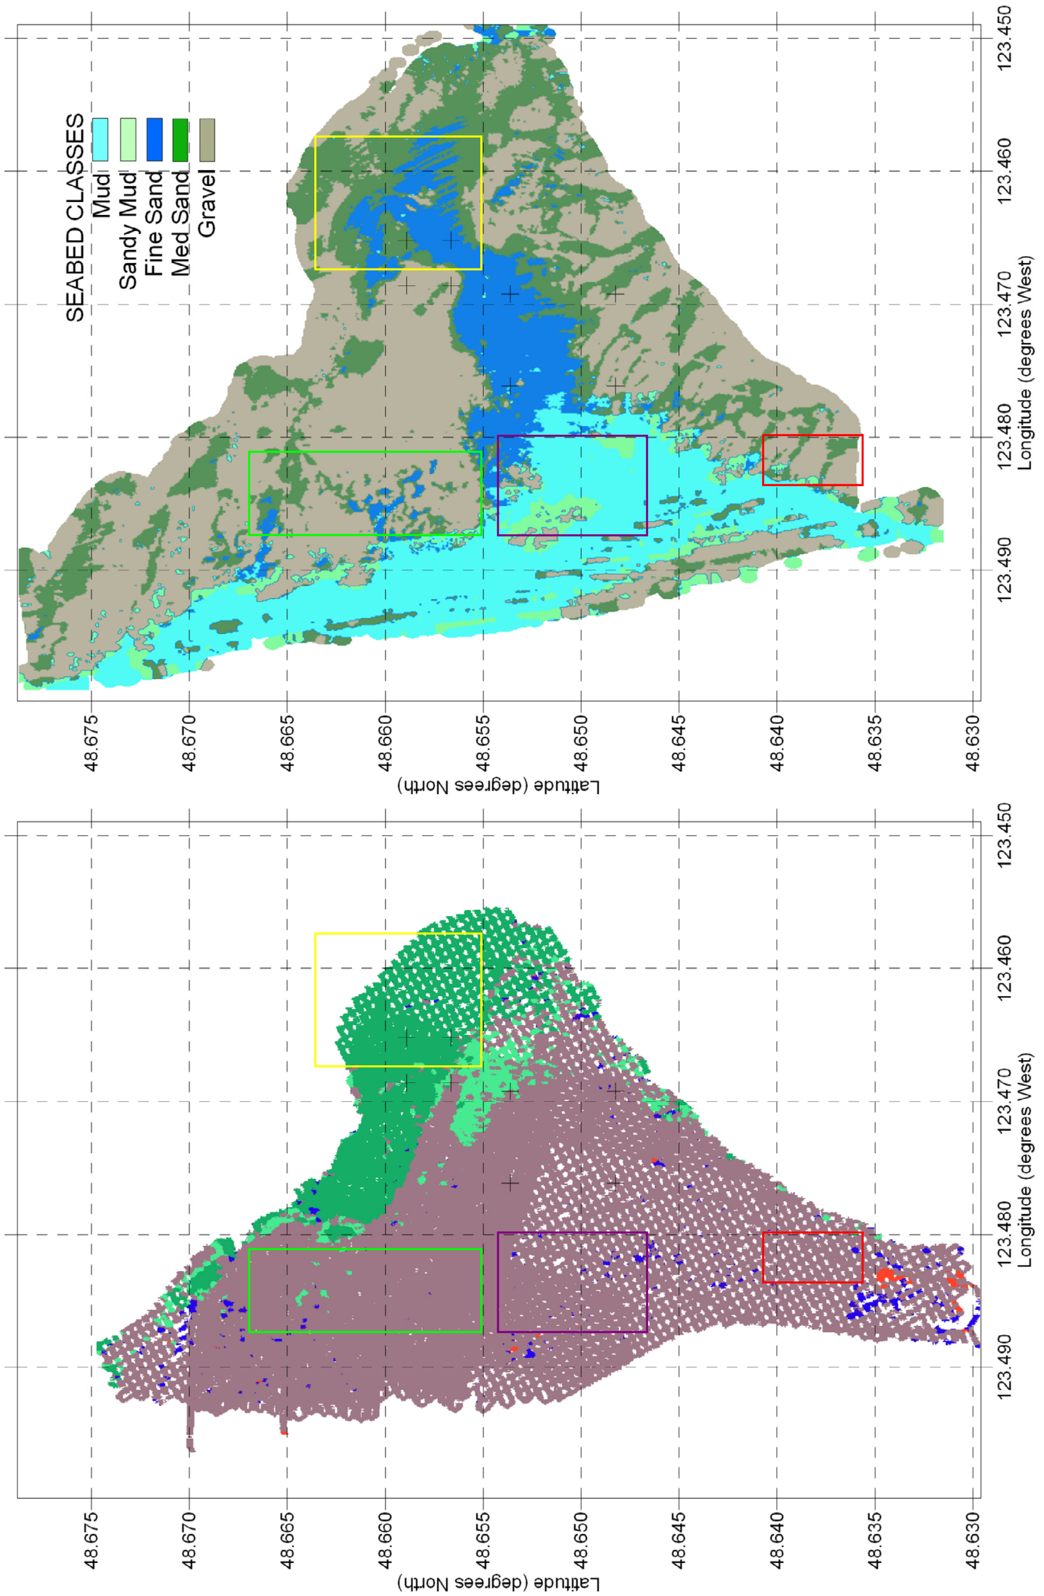

Figure 7-10. A comparison of the SELw classification map for the Simrad200 testbed survey (*lower*) to the multibeam derived ground-truth class map (*upper*).

7.4.4 Quantitative Evaluation of SELw Depth and Slope Compensation

Overall, the quantitative measures of classification and feature performance confirm what was seen in the qualitative comparisons of the SELw class maps to the ground-truth seabed class map. In comparing the classification performance of the TNORM and the SELw results, the AMI and OA statistics show significant improvement for the Airmar38 survey as presented in Table 7-8. However, the Simrad200 results do worse because of the effect of ensonification regime which will remain after compensation for depth and slope. The comparison is made worse from the fact that the TNORM result actually benefits from being depth and slope biased since the seabed types in Patricia Bay are partially correlated with depth and slope. The best-match merge OA value for the Simrad200 TNORM result also benefits from a large number of classes. In general, seabed classification results with the Simrad200 are generally poor because of its narrow beamwidth and lack of volume penetration; this echosounder generates far less discriminatory information than does the wide-beam low-frequency Airmar38.

Table 7-8. Quantitative measures of classification performance for the improvements in the removal of non-seabed influences. merge-QC OA is more reliable than the best-merge OA as best-merge OA is biased upwards for higher numbers of classes, due to the merge process. Relatively low AMI values are expected as was discussed in reference to Figure 7-3.

TNORM	Airmar38	Simrad200
Best-match OA (%)	48	57
No. classes:	8	17
Merge-QC OA (%)	53	-
AMI (%)	15	13

SELw	Airmar38	Simrad200
Best-match OA (%)	63	44
No. classes:	13	11
Merge-QC OA (%)	65	-
AMI (%)	29	6

Improvements in the removal of non-seabed influences are not necessarily always reflected in the classification results as was seen in the case of the Simrad200 testbed survey. In addition to seabed types being correlated with depth and slope, there is some complexity introduced by the process of classification. PCA can weight poor features if such features are highly covariant, and cluster analysis can make apparently arbitrary decisions on class boundaries if the distribution of points in Q-space is a continuum. For these reasons, the average absolute correlation coefficients to depth and slope are useful means of evaluating improved methods for the removal of non-seabed influences. The results of Table 7-9 show the improvement is substantial, especially when compensating for seabed slope for the wide-beam Airmar38 as one would expect from the results of the slope experiment, section 6.3.

Table 7-9. Average absolute correlation coefficients to depth and seabed slope. The calculation of these values was described in section 7.2.

TNORM	Airmar38	Simrad200
Depth	0.49	0.48
Slope	0.38	0.45

SELw	Airmar38	Simrad200
Depth	0.32	0.28
Slope	0.15	0.26

7.5 Improved Discrimination: Physics-Based Features

In the first sub-section that follows this introduction, the new feature algorithms will be described in detail, making reference to the theoretical basis, connections with geoacoustic properties of the seabed and some initial testing with the BORIS simulations. Despite a having a good basis, it is unknown how well each feature will perform; the new features are relatively untested. So after presenting the new features in sub-section 7.5.1, the features will be evaluated for performance in terms of non-seabed influences and discriminatory ability in sub-section 7.5.2. Features that are found to provide biased, correlated, poor or erroneous information should be eliminated, as they will reduce classification performance. The final classification result of the combined and reduced feature set is presented and evaluated in sub-section 7.5.3.

7.5.1 The New Features

Table 7-10 lists all of the new features as grouped by *family*, defined by their parent algorithm. Each family will be described in order and in detail following the table.

Table 7-10. A summary of the new feature families.

Feature Family Name	Feature Number	Data Source:	Indicator of:
Moments of echo amplitude	167-169	Stacked echo	Specific shape
Cumulative amplitude threshold	170-174	Stacked echo	General shape
Cumulative energy threshold	175-179	Stacked echo	General shape
Relative amplitudes	180-185	Stacked echo	Impedance, roughness
Measures of time	186-189	Stacked echo	Roughness, attenuation
Amplitude variability	190-192	Echoes of 5 stacks	Roughness
Envelope variability	193-195	Echoes of 5 stacks	Scattering
Fractals	196-197	Echoes of 1 stack	Scattering
Residual echo duration	198	Echoes of 1 stack	Roughness, attenuation
AEM cross-correlations	199-205	Stacked echo	Specific shape related to roughness
Spectral moments	206-213	Concatenated echoes	Spectal effects of scattering

The first new family of feature algorithms is the set of moments of echo amplitudes (features 167 to 168). These are usual mathematical definitions of moments applied to the echo envelope time series. van Walree applied moments of the echo intensity time series, which are just the envelope values squared [van Walree et al., 2005]. The first two moments (centre of gravity and time-spread) are as calculated by Equation 6-2 and Equation 6-3 (see sub-section 6.3.2). The echo skewness (feature 169, F_{169}) is:

$$F_{169} = \frac{8}{F_{168}^{\ 3}} \frac{\int p(t)(t-t_0)^3\,dt}{\int p(t)dt} \tag{7-2}$$

The echo moments are all-purpose indicators of general shape. These three features, the centre of gravity and time-spread in particular, express the two major degrees of freedom of echo shape – how 'pointy' and how 'long' an echo is. The scattering theory, AEM and BORIS results show that high impedance, low scattering echoes are 'pointy', while scattering lengthens echoes (via the beam-widening effect).

The next two feature families, cumulative amplitude thresholds (features 170-174) and cumulative energy thresholds (features 175-179), are related. Both families are related to the cumulative picking methods of appendix A and to the within-echo cumulative amplitude test features of appendix G. In this case, the features are the normalized locations within the analysis window (from zero to one) where the normalized cumulative amplitude or energy reaches a prescribed threshold. The thresholds are shown in Table 7-11. The threshold values were chosen to capture variation in echo shape in regions of typical echoes where the echo responds to macro-roughness heights (features 170 & 175), rise time (171 & 176), and decay of the echo, which is strongly determined by backscattering (172-174 & 177-179).

Table 7-11. Thresholds for the cumulative amplitude and energy threshold features.

Feature Number	Threshold (%)	Amplitude or Intensity?
170	1	Amplitude
171	20	Amplitude
172	60	Amplitude
173	75	Amplitude
174	90	Amplitude
175	0.5	Intensity
176	10	Intensity
177	95	Intensity
178	99	Intensity
179	99.9	Intensity

The relative amplitude features (180-185) are more complicated than those of the two previous feature families. Many SBES are not calibrated so that the absolute backscatter strength is not known; instead, a compensated relative measure of backscatter amplitude is used in all cases. Usually amplitude is measured in raw units from the analogue to digital converter, which is in units proportional to received voltage and hence proportional to received pressure. Without calibration, the relative echo amplitudes in raw units should be sufficient to generate amplitude features after accounting for the varying gains applied by the echosounder. Spherical spreading and water column losses are compensated for by TVG (see section 2.1). However, if the TVG is not exactly correct, these relative amplitudes will be functions of depth as shown in Figure 7-11 and Figure 7-12. The spreading term is determined by type of target and the way the target is ensonified by the echosounder. The transmit pulse either fully ensonifies the acoustic footprint or only ensonifies an annulus that moves outward towards the edge of the footprint. This is another consequence of the ensonification regime as described in Sub-section 5.4.2; for an illustration, see Figure 5-20. In all, there are three forms of TVG:

$$\begin{aligned} TVG_{full} &= 20\log_{10} r + 2\alpha r \\ TVG_{annular} &= 30\log_{10} r + 2\alpha r \\ TVG_{point} &= 40\log_{10} r + 2\alpha r \end{aligned} \qquad (7\text{-}3)$$

In Equation 7-3, r is the range as related to the two-way travel time, and α is the water column absorption coefficient. The full derivation of Equation 7-3 may be found in [Lurton, 2002]. A derivation from KA theory showing that the *30logR* rule is appropriate for seabed echosounding in the annular regime can be found in [Pouliquen, 2004].

The variance in water column absorption shown in Figure 7-11 and Figure 7-12 is likely enough to cause class confusion if classification were based solely on amplitude features. The solution is a best fit compromise, similar to the idea behind effective beamwidth and attenuation. After removing the original *TVG*, the spreading term coefficient, absorption and an offset value are found by fitting the form of Equation 7-3 to smoothed peak amplitude data. This *trim-TVG* is then applied to the time series. The benefits of this include first-order accommodation for varying spreading and absorption as functions of depths and minimization of the effect of depth on relative amplitudes. Trim-TVG can be calculated for data compiled from a multi-day survey or, as was the case here, the trim-TVG was recalculated for each day in the survey in case oceanographic conditions change day to day. As was the case for the effective beamwidth attenuation estimates, trim-TVG can be affected by the distribution of seabed types in the seabed. However, these effects are relatively small compared to the propagation loses. Preston conceived the concept of trim-TVG by adapting his method for MBES [Preston and Christney, 2005] to SBES. trim-TVG is formally presented here for the first time.

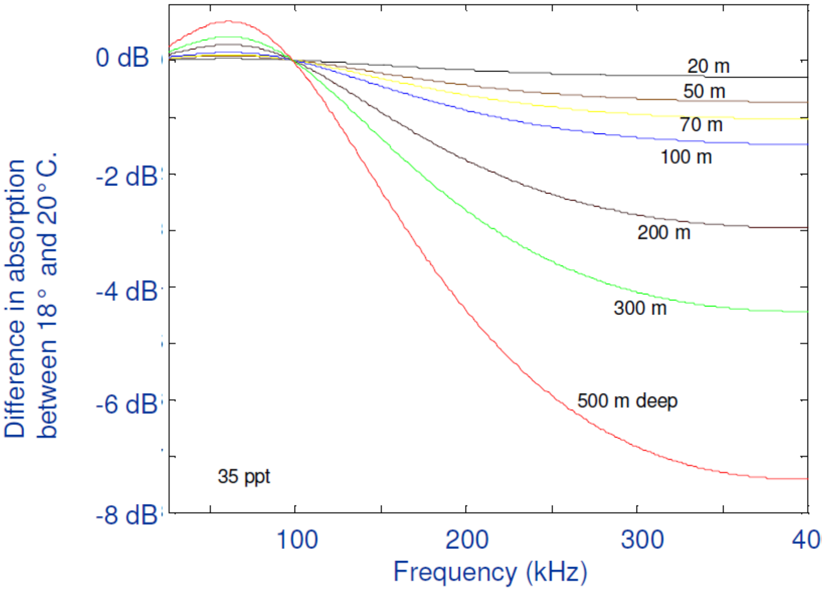

Figure 7-11. TVG inaccuracy at various depths and echosounder frequencies for a two degree difference in temperature, based on Equation 7-3 and the Francois-Garrison model as presented in [Lurton, 2002]. The null at 97 kHz arises because absorption is a complicated function of temperature and frequency.

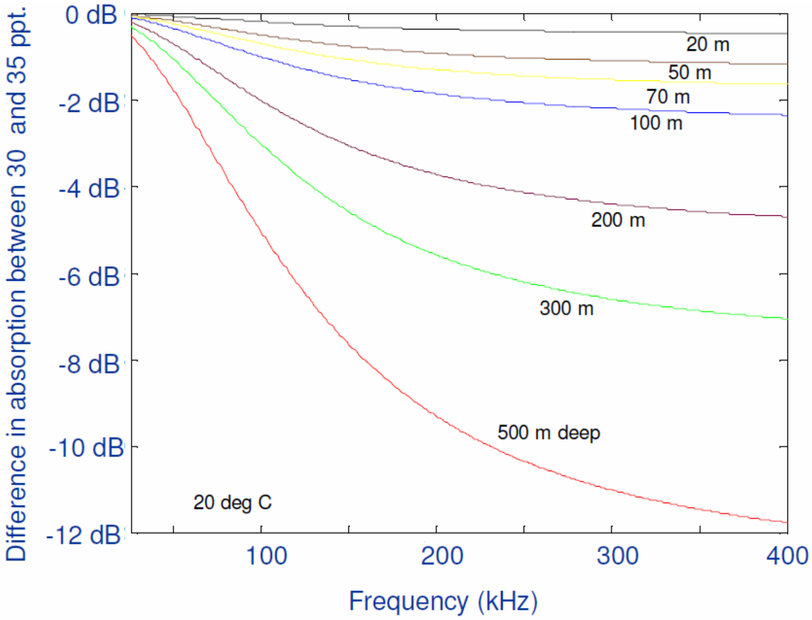

Figure 7-12. TVG inaccuracy at various depths and frequencies for a 5 ppt difference in salinity, based on Equation 7-3 and the Francois-Garrison model as presented in [Lurton, 2002].

The relative amplitude features are listed in Table 7-12. Computational details are not necessary as the names indicate the general nature of the features. Feature 182 is similar to RoxAnn's E1; it increases with backscattering (both surface and volume).

Feature 183 could also be called the relative tail energy; it increases with surface and volume backscattering. For feature 184, the 'half maximum' means the region defined by mean locations of points near 50% peak amplitude on either side of the echo peak. Feature 185 is an energy ratio as experimented with BORIS in appendix G; it will decrease with impedance and increase with surface and volume backscattering. Features 180, 181 and 184 are sensitive to impedance; however roughness backscattering will decrease amplitudes, while additional volume inhomogeneities can increase amplitudes.

Table 7-12. The relative amplitude feature family.

Feature Number	Feature Name
180	Smoothed peak amplitude
181	Mean amplitude from bottom pick to tail
182	Mean amplitude from end of pulse to tail
183	Mean intensity peak to tail divided by peak amplitude
184	Mean amplitude of the half maximum
185	Ratio of energy after pulse to energy during pulse

Features 186 through 189 are measures of time as summarized in Table 7-13. The half maximum duration (186) should approximate the pulse duration for a seabed of less than Rayleigh critical roughness. It will increase slowly for rising roughness backscattering where the echo is still largely peaked like AEM echoes. For seabeds with low impedance, such as the mud of Patricia Bay, this feature will be extremely long. All other measures of time will increase with surface and volume backscattering. Fall time is the fraction of the analysis window occupied by an echo decaying from its peak to a 99% cumulative amplitude threshold found within the analysis window. Fall fraction is the proportion of the fall time to the total time of an echo within the analysis window. Significant echo duration is the proportion of samples in the window with amplitudes greater than 10% of the peak amplitude.

Table 7-13. The measures of time feature family.

Feature Number	Feature Name
186	Half maximum duration
187	Fall fraction
188	Fall time
189	Significant echo duration

The next three feature families are based on measures of variability. Echo statistics, particularly those of variability, are used primarily in hydroacoustics to estimate fish or zooplankton type, density and even orientation [Medwin and Clay, 1998; Stanton and Clay, 1986]. There are two approaches that apply measures of variability to characterize the seabed [Stanton, 1984; Dunsiger et al., 1981]. Stanton's method is to fit a theoretical distribution to histograms of peak amplitude to determine the ratio of coherent to incoherent intensity[53], which is symbolized as γ. Using KA theory, Stanton linked the Rice distribution to γ and then derived an expression for RMS roughness height as a function of γ [Stanton, 1984]. In general, as roughness increases the distribution of peak amplitudes changes from approximately Normal to a much wider, positively skewed distribution. Dunsiger arrives at γ via a coherence function which is the ratio of the power spectral densities of a stacked echo and the average of the echoes in the stack [Dunsiger et al., 1981]. The relative strength of the incoherent component of an echo determines the variability of echoes as discussed in chapters 4 and 5. The variability manifests within single echoes and as the variability between echoes, often called *ping-to-ping variability*.

The amplitude variability feature family (190-197) consists of measures of the peak amplitude distribution consisting of the span, standard deviation and skewness. These are features 190 through 191 in Table 7-14. The peak amplitudes of the echoes that make up the standard stack of five echoes are insufficient to sample a peak amplitude distribution. Instead, 25 trim-TVG corrected echoes from three stacks are used – this is the 'echoes from 5 stacks' data source listed in Table 7-10. The four extra stacks consist of two stacks before and after the current stack in the order in which they were collected. The number of stacks can be modified – increasing this number would increase the performance of the features but degrade their spatial resolution. Geographic searching to find the nearest stacks was also considered but is not yet implemented.

---

[53] Although γ is given a ratio of intensity, the histogram is of amplitude. This is as repoted in [Stanton and Clay, 1986].

Table 7-14. The features of the three feature families that measure aspects of variability.

Feature Number	Feature Name
190	Amplitude stack span
191	Amplitude stack standard deviation
192	Amplitude stack skewness
193	Stack mean coherence
194	Stack mean correlation
195	Stack relative deviation
196	Fractal - path length
197	Fractal - box count

The envelope variability features are features 193 through 195 as listed in Table 7-14. These features are applied to the same ensemble of nine stacks as the amplitude variability features, but do not require trim-TVG. Because the number of samples in the variable length windows will be different, the echoes are upsampled to the window with the greatest number of samples, creating a 'time series matrix'. The stack mean coherence is a simplified version of Dunsiger's spectral coherence [Dunsiger et al., 1981]. The stack mean correlation is the mean of the correlation matrix calculated on the time series. This feature also measures the strength of the coherent component relative to the incoherent component of the echo. It will respond to ping-to-ping variability and common shapes or scatterers between echoes, for example a sub-bottom reflector in 12 kHz echoes should increase the stack correlation. The stack relative deviation is calculated as the mean of the sample-wise standard deviation divided by the mean of the sample-wise mean. This feature is the most direct measure of ping-to-ping variability.

Two fractal features are implemented to measure the variability within each echo. They are features 196 and 197 as listed in Table 7-14. Fractals are calculated from echoes that comprise the stack and then averaged to get one value per stack. The first fractal feature is the path length. It is not a classic fractal dimension, but is related to the idea that a fractal can have infinite length between any two points [Mandelbrot, 1982]. In this case, the maximum length is determined by the resolution of the measurement (sample frequency and A/D bits), so variability will increase the path length up to a maximum (which is unrealistic to reach with real data). After scaling the analysis window to range from zero to one in amplitude and time, the path length dimension is calculated by:

$$F_{196} = \log\left(\sum_{1}^{n-1} \sqrt{dx^2 + dy_i^2}\right) \tag{7-4}$$

where n is the number of samples, dx is $\frac{1}{n-1}$ and dy_i is the change in amplitude from the i^{th} sample to the next sample. The second fractal feature is the box count Hausdorff dimension, as calculated by and discussed in [van Walree et al., 2005].

The next feature, residual echo duration (198), is a family all to itself. Although it is a measure of time, it is unique in that it is calculated directly from the meta-data (step A2 of Figure 7-2). It has already been shown and used in sections 6.2 and 7.4. It is calculated based on measured echo durations and the final effective beamwidths and attenuation of Table 7-7. The feature values are in milliseconds and are mean values from echoes within a stack. As noted throughout chapters 5 and 6, echoes lengthen due to beam-widening, macro-roughness or penetration.

Residual echo duration is based on the EDM, and the next feature family is based on the other model developed in chapter 5, the AEM. The AEM cross-correlation features (199-205), as listed in Table 7-15, are various measures from cross-correlations of AEM simulated echoes with stacked echoes. Stacked echoes are required as the AEM calculates mean or stacked echo envelopes. The first four AEM features are the maximum cross-correlation values between AEM simulated echoes and stacked echoes. The four simulated echoes are based on the parameters of the echosounder, the water depth and the roughness parameters as listed in Table C-2[54] and correlation lengths determined later in chapter 8. These features indicate the similarity between the survey echo and the respective AEM echoes. The maximum cross-correlation is used so there are no alignment issues, especially in case of discrete scatterers that occur before the main seabed echo. The cross-correlation lag caused by discrete scatterers, such as macro-algae or unusual heights above the seabed, is also useful as feature 203. This feature scales with the heights of these objects above the seabed if the bottom pick picks on them. The next two features attempt to replace the wavelet features. For these features, the AEM medium

---

[54] The use of the clayey sand seabed here is to represent a seabed with near-zero roughness, and not one with significant volume scattering.

sand echoes are used like match-filter wavelets: they are scaled up or down by linear interpolation and then cross-correlated with the stacked echoes. The best interpolated cross-correlation of any of the AEM echoes is listed as feature 204, while the interpolation used to get that maximum cross-correlation is listed as feature 205. These two features are conceptually equivalent to using the peak wavelet scale and shift values.

Table 7-15. The AEM cross-correlation feature family

Feature Number	Feature Name
199	Max cross-corr clayey sand
200	Max cross-corr medium sand
201	Max cross-corr sandy gravel
202	Max cross-corr rock/cobble
203	Cross-corr lag medium sand
204	Best interpolated cross-corr
205	Best interpolation rate

The final feature family is the spectral moments (206-213). Many of these features were employed in BORIS experimentation in sub-section 5.4.4 and appendix G. These features are designed to characterize the behaviour of the spectral content of echoes as discussed throughout chapters 4 and 5. Selective and coherent surface backscattering favours lower frequencies, so the mean frequencies, bandwidth, etc., will reflect this by shifting lower with increasing roughness scattering. Volume scattering from large inhomogeneities in addition to dispersion in the volume will also act to reduce the relative strength of higher frequencies in the echo. The spectral moment features are listed in Table 7-16.

Table 7-16. The spectral moments feature family

Feature Number	Feature Name
206	Centre of Gravity Frequency
207	Standard deviation PSD
208	Skewness PSD
209	Kurtosis PSD
210	Peak frequency
211	Median frequency
212	Cumlulative sum bandwidth
213	Entropy PSD
(214)	Mean instantaneous frequency
(215)	Standard deviation instananeous freq.
(216)	Standard deviation phase angle

The features are their native values (units of kilohertz where applicable), which are not normalized, nor should they be because correlation PCA will be used with the new feature set. To improve resolution, echo time series comprising the stack are knit together end to end in a processed called echo concatenation as discussed in [Preston and Biffard, 2007]. Basically a Hann window is used to smooth the discontinuities at the ends of the time series. Like the amplitude and envelope variability features, the stacks are expanded by adding neighbouring stacks to make a total of 25 echoes.

The first four features in the table are the moments of the power spectral density. The following four features are other bulk measures. The cumulative sum bandwidth was described in appendix G. The entropy PSD is calculated according to information theory (i.e. Shannon entropy); it is another indicator of bandwidth. The last three features are only available with FWF data. Instantaneous frequency and phase are derived from the Hilbert transform. The mean instantaneous frequency is a measure of centre frequency and was shown in Figure G-7.

7.5.2 Evaluation and Reduction of the Combined Feature Set

The properties of a feature set have a strong influence on PCA and cluster analysis. As discussed in appendix B, covariance PCA is preferable when dealing with data of the same scale[55]. However, the new features and the original features all have different scales. Even though the original features were restricted in range from zero to one, their means and standard deviations were not equal, (as will be shown in Table 7-17 and Table 7-18). This should cause poor results using covariance PCA in the QTC View method. However, the original features that have good discriminatory ability also have high covariance which causes covariance PCA to weight them favourably, resulting in generally useable results, as discussed in sub-section 2.2.3 and appendix B.

The new features are not designed to guarantee such a fortunate outcome. To address this issue, correlation PCA is used for processing data from the combined feature set. To optimize correlation PCA, outliers and large groups of correlated features should not be used, as they will be favoured in the result. Especially degrading to the

---

[55] For example, covariance PCA of temperature time series data is often used to establish the spatial patterns of many climate oscillatory patterns. The discovery of the pacific decadal oscillation is a good example [Mantua et al., 1997], while the most famous example is that of El Nino Southern Oscillation.

performance would be groups of features correlated with an artefact such as depth, slope or the ensonification regime. Features can also have correlated noise or erroneous information and outliers. Outliers are removed by a new method called iterative outlier removal applied to correlation PCA as described in appendix B. The purpose of this sub-section is to eliminate features that detract from the performance of correlation PCA.

A related risk to diminish the performance of PCA is redundant or correlated information. In section 4.7, it was found that echoes vary in a limited number of ways, so having an exorbitant number of features will not improve discrimination. As redundant features are added, the risks of adding artefacts or errors increases. The combined feature set contains 213 features, (146 excluding wavelets and constant features). Considering that many SELw analysis windows for the Patricia Bay testbed survey contain less than 100 samples, this large number of features will undoubtedly contain redundant information. The feature set needs to be reduced to include features free of non-seabed influences, errors, noise and features that contain the highest quality of information, i.e. discriminatory ability.

To satisfy the aims discussed above, a number of features from both the original and new features will be discarded on the basis of biases to depth, seabed slope, sensitivity to the ensonification regime effect, correlation among feature families (i.e. redundancy) or poor discriminatory ability. This is an important step as the new features are essentially prototypes with limited testing behind them, while the original features have never been tested in this way. Two reduced feature sets are established here, accounting for the differences between high-frequency SBES (100-200 kHz) and low-frequency SBES (12-100 kHz). An important aspect of establishing reduced feature sets applicable for general surveying is that the evaluation of the features spans all major seabed types over a wide range of depths and slopes. Since the Airmar38 and Simrad200 Patricia Bay testbed surveys have all of the above in abundance, these surveys are again an ideal testbed.

Four main tools are applied here to assess feature performance in the testbed surveys: S-QC training site percent correct assignment, percent misassignment, inspection of feature maps and correlation coefficients of the features to depth and slope.

See Figure 7-2 for a flowchart showing how these tools are related. After combining the results of the above tools, the last step is to assess the correlation matrix of the features to reduce redundancy and improve PCA performance. The process of reducing the feature set is further elaborated in the following paragraphs.

The two means of quantitatively evaluating feature discriminatory ability are the percent misassignment and the S-QC percent correct assignment as described in section 7.2. In this case, percent misassignment was calculated using survey data drawn from three homogenous testbed sites – gravel, sand and mud (see section 6.1, Figure 6-1). The percent misassignment between the sand site and the other two sites, gravel and mud, was calculated for all available old and new features for both the Airmar38 and Simrad200 SELw processed data sets. The criterion for a feature to have useable discriminatory ability is to have less than 20% misassignment in both comparisons. The 20% criterion is somewhat arbitrary, although such a level of misassignment (or lower) will facilitate depicting an accurate boundary between two seabed types. The S-QC percent correct assignment statistic is the percentage of data points in the training sites that were classified correctly, using one feature. Four training sites were used: the aforementioned homogenous sites plus a fine sand site; also shown in Figure 6-1. See section 7.6 for more details on the S-QC analysis applied here. It is important to note that the testbed sites are comprised of seabeds that span the range of echo characteristics. It is hypothesized that any feature that adequately responds to these seabeds, will adequately respond to any naturally occurring seabed – this is supported by the physics of chapters 4 and 5.

Qualitative inspections of feature maps are useful and informative in facilitating a diagnosis to understand the results of the quantitative measures. There are four main reasons for poor performance: depth, slope or ensonification biases or poor discrimination. Poor discrimination is the most common problem, for which feature maps look 'fuzzy' with no patterns that could represent boundaries of seabed types. The final decision on whether to discard features resides on the clarity of feature maps; features are not discarded on their ability to discern the known seabed types in Patricia Bay. This avoids tuning the feature sets to Patricia Bay. Biased features resemble the depth or slope maps of Patricia Bay (see Figure 6-1 for depth contours, see Figure 6-2 for the slope map). Several example feature maps of quality new features are shown as Figure 7-13.

The depth and slope biases are apparent in the correlation coefficients, which are listed for each feature along with results of the qualitative inspections and quantitative measures in Table 7-17 and Table 7-18 (for the Airmar38 and Simrad200 surveys, respectively). All the information for each feature[56] is combined to make a decision on whether or not the feature is of sufficient quality to be applied in a reduce feature set, as shown under the column heading 'Used?' in Table 7-17 and Table 7-18. Although this decision is based on passing all the quantitative criteria (less than 20% misassignment, greater than 80% QC correct assignment, less than 0.6 correlation to depth / slope), there is some allowance for arbitrary cut-off criteria in cases of features that contribute unique information. As note above, the qualitative feature map inspection is used to make the final decision by either confirming the decision made by quantitative criteria or by making exceptions. Overall, making exceptions is rare, but this does occur for some physics-based features that offer unique seabed information.

Some features perform well but are highly correlated with others, so these features do not add unique discriminatory information. Such groups of features can be replaced by one or two features without losing much information. This occurs for features from the same feature family, especially the original shape features. Correlation PCA will also reduce groups of features like these into one principal component. However, large blocks of correlated features will bias correlation PCA to weighting these blocks favourably and ignoring other features with useful discriminatory information. To reduce this effect, correlation matrices (Figure 7-14 and Figure 7-15) of the feature values are used to identify groups of three or more correlated features that would otherwise pass the selection criteria. Features are then rejected from such a correlated group until the average correlation coefficient of the group to itself decreases below 0.9, or there are only two features remaining. The features that belong to correlated groups are listed prior to the final decision shown in the 'Used?' column in Table 7-17 and Table 7-18.

---

[56] For brevity, the FFT features are summarized on one line in Table 7-17 and Table 7-18. Wavelet features were excluded for computational reasons. Features that have constant values throughout the survey were also excluded. These include features 16, 31 of the original feature set and features 205 (best interpolation ratio) and 210 (peak frequency) of the new feature set. Feature 205 is constant due to an error in the code that was later corrected, while feature 210 is constant due to the nature of the spectra of echo envelopes – this feature will be much more useful for FWF data.

Table 7-17. Feature statistics from the Airmar38 testbed survey, feature evaluation and the reduced feature set for low frequency data. 'y' – yes, 'm' – marginal, no otherwise.

Feature Name			Feature Descriptive Statistics				Quantitative Tests					Feature Map Inspection					Used?
	Feature Number	Sequential Number	Min.	Max.	Mean	Std. Dev.	Corrcoef w. Depth	Corrcoef w. Slope	QC %-correct	Sand to Gravel %-mis.	Sand to Mud %-mis.	Depth Biased?	Slope Biased?	Enson. Biased?	Poor Discrimination?	Correlated?	
Cumulative integral direct (first 8) 1	1	1	0.001	0.049	0.007	0.003	-0.483	-0.322	58.0	14.4	5.3				y		
Cumulative integral direct (first 8) 2	2	2	0.008	0.633	0.273	0.083	-0.170	0.102	68.6	13.1	3.5					y	
Cumulative integral direct (first 8) 3	3	3	0.019	0.833	0.530	0.117	-0.507	-0.153	74.1	8.0	0.0					y	y
Cumulative integral direct (first 8) 4	4	4	0.040	0.907	0.695	0.117	-0.638	-0.212	81.4	4.3	0.0					y	
Cumulative integral direct (first 8) 5	5	5	0.048	0.949	0.797	0.103	-0.700	-0.214	82.6	2.9	0.0					y	y
Cumulative integral direct (first 8) 6	6	6	0.070	0.967	0.863	0.085	-0.741	-0.210	81.6	2.4	0.0					y	
Cumulative integral direct (first 8) 7	7	7	0.079	0.975	0.905	0.068	-0.762	-0.210	82.2	2.4	0.0	m			m		
Cumulative integral direct (first 8) 8	8	8	0.086	0.980	0.933	0.052	-0.765	-0.213	78.0	4.8	0.0	m			m		
Cumulative integral direct (next 7) 1	9	9	0.099	0.986	0.951	0.038	-0.733	-0.208	66.2	8.3	0.0				m		
Cumulative integral direct (next 7) 2	10	10	0.147	0.988	0.963	0.027	-0.666	-0.202	53.5	10.1	0.0					y	y
Cumulative integral direct (next 7) 3	11	11	0.276	0.990	0.972	0.019	-0.567	-0.197	45.9	9.3	0.0					y	
Cumulative integral direct (next 7) 4	12	12	0.496	0.994	0.979	0.012	-0.454	-0.193	39.1	8.0	0.9					y	y
Cumulative integral direct (next 7) 5	13	13	0.714	0.996	0.985	0.008	-0.351	-0.184	35.8	8.3	7.5					y	
Cumulative integral direct (next 7) 6	14	14	0.824	0.997	0.990	0.004	-0.279	-0.178	35.4	9.1	16.8					y	y
Cumulative integral direct (next 7) 7	15	15	0.934	0.999	0.995	0.002	-0.233	-0.175	33.6	9.9	29.2					y	
Cumulative integral ratio (first 3) 1	17	16	0.258	0.896	0.544	0.075	0.034	0.177	43.2	33.6	40.3				m		
Cumulative integral ratio (first 3) 2	18	17	0.120	0.751	0.413	0.081	-0.065	0.119	48.0	22.9	20.4				m		
Cumulative integral ratio (first 3) 3	19	18	0.058	0.705	0.360	0.082	-0.144	0.070	54.8	20.0	8.8					y	y
Cumulative integral ratio (last 4) 1	20	19	0.028	0.680	0.334	0.084	-0.196	0.041	59.6	17.9	4.4					y	
Cumulative integral ratio (last 4) 2	21	20	0.018	0.675	0.319	0.085	-0.231	0.025	62.3	16.8	3.5					y	y
Cumulative integral ratio (last 4) 3	22	21	0.013	0.670	0.311	0.087	-0.255	0.015	63.7	16.0	2.7					y	
Cumulative integral ratio (last 4) 4	23	22	0.011	0.666	0.306	0.087	-0.270	0.010	65.0	15.7	2.2					y	y
Quantile (first 7) 1	24	23	0.000	0.289	0.006	0.010	0.363	0.260	60.0	7.4	0.8					y	
Quantile (first 7) 2	25	24	0.000	0.308	0.010	0.019	0.459	0.220	72.2	3.3	0.0					y	y
Quantile (first 7) 3	26	25	0.000	0.400	0.019	0.035	0.566	0.168	80.7	1.7	0.0					y	
Quantile (first 7) 4	27	26	0.001	0.501	0.039	0.057	0.628	0.129	83.5	2.5	0.0					y	y
Quantile (first 7) 5	28	27	0.003	0.597	0.086	0.084	0.584	0.146	73.2	8.3	0.0					y	
Quantile (first 7) 6	29	28	0.014	0.706	0.190	0.110	0.478	0.153	74.5	9.1	0.0					y	y
Quantile (first 7) 7	30	29	0.069	0.867	0.428	0.122	0.167	-0.003	48.6	22.3	11.6				m		
Histogram (first 7) 1	32	30	0.060	0.980	0.691	0.111	-0.501	-0.140	73.5	6.7	0.0						y
Histogram (first 7) 2	33	31	0.007	0.560	0.092	0.037	0.502	0.057	48.0	35.5	10.7				y		
Histogram (first 7) 3	34	32	0.000	0.486	0.059	0.035	0.508	0.141	51.2	35.5	5.8				y		
Histogram (first 7) 4	35	33	0.000	0.277	0.047	0.032	0.483	0.146	48.2	36.4	7.4				y		
Histogram (first 7) 5	36	34	0.000	0.211	0.037	0.026	0.435	0.138	50.9	32.2	6.6				y		
Histogram (first 7) 6	37	35	0.000	0.182	0.029	0.020	0.244	0.046	43.9	33.9	13.2				y		
Histogram (first 7) 7	38	36	0.000	0.149	0.023	0.015	-0.116	-0.153	33.5	47.9	44.6				y		
Histogram (last)	39	37	0.001	0.128	0.022	0.013	-0.575	-0.379	39.9	47.1	19.5	y					
FFT 1-65 (only 1-59* applicable)	102-166	38-96*	-	-	-	-	-	-	-	-	-	y			y		

Table 7-17. Continued.

Feature Name			Feature Descriptive Statistics									Feature Map Inspection					Used?
	Feature Number	Sequential Number	Min.	Max.	Mean	Std. Dev.	Corrcoef w. Depth	Corrcoef w. Slope	QC %-correct	Sand to Gravel %-mis.	Sand to Mud %-mis.	Depth Biased?	Slope Biased?	Enson. Biased?	Poor Discrimination?	Correlated?	
Envelope centre of gravity	167	97	0.128	0.733	0.220	0.045	0.512	0.145	78.7	0.8	0.0						y
Envelope time-spread	168	98	0.133	0.583	0.264	0.054	0.673	0.251	86.3	1.7	0.0						y
Envelope skewness	169	99	0.158	1.443	0.908	0.125	-0.694	-0.099	72.5	31.2	0.0						y
Cumulative amplitude thres. 1%	170	100	0.009	0.170	0.065	0.005	0.112	-0.033	45.6	37.2	1.7			m	y		
Cumulative amplitude thres. 20%	171	101	0.077	0.673	0.115	0.021	0.042	-0.195	64.2	12.4	4.1						y
Cumulative amplitude thres. 60%	172	102	0.116	0.773	0.211	0.053	0.493	0.143	73.7	7.4	0.0					m	y
Cumulative amplitude thres. 75%	173	103	0.142	0.831	0.271	0.069	0.594	0.188	76.8	5.0	0.0					m	y
Cumulative amplitude thres. 90%	174	104	0.194	0.912	0.380	0.090	0.670	0.190	83.3	1.7	0.0					m	y
Cumulative intensity thres. 0.5%	175	105	0.020	0.368	0.070	0.005	-0.555	-0.449	19.6	27.3	44.6			m	y		
Cumulative intensity thres. 10%	176	106	0.064	0.680	0.095	0.016	-0.311	-0.378	40.7	20.7	27.3				m		
Cumulative intensity thres. 95%	177	107	0.151	0.946	0.302	0.088	0.560	0.176	73.0	8.3	0.0					m	y
Cumulative intensity thres. 99%	178	108	0.196	0.990	0.398	0.107	0.648	0.177	77.8	4.1	0.0					m	y
Cumulative intensity thres. 99.9%	179	109	0.255	0.999	0.526	0.126	0.684	0.162	79.2	2.5	0.0					m	y
Smoothed peak amplitude	180	110	0.292	3.076	1.183	0.211	-0.457	-0.212	52.2	45.6	7.1						y
Mean amplitude from bottom pick to tail	181	111	0.054	0.801	0.213	0.054	-0.187	-0.404	47.5	15.7	13.3						y
Mean amp from end of pulse to tail	182	112	0.050	0.807	0.199	0.051	-0.063	-0.343	47.1	15.5	19.5						y
Mean intensity peak to tail / peak amp.	183	113	0.002	0.228	0.038	0.019	0.085	-0.285	45.2	18.4	47.8				y		
Mean amplitude of the half maximum	184	114	0.180	2.739	0.631	0.188	-0.601	-0.286	59.4	38.4	2.2						y
Ratio of energy after pulse / during pulse	185	115	0.929	10.000	5.988	3.117	0.659	0.363	61.2	25.6	0.0		y		y		
Half maximum duration	186	116	0.010	0.806	0.149	0.083	0.402	0.120	57.1	22.3	3.3						y
Fall fraction	187	117	0.254	0.986	0.901	0.059	0.268	0.354	34.4	21.5	35.5	m		m	y		
Fall time	188	118	0.222	0.905	0.575	0.117	0.386	0.298	65.2	5.0	4.1	m		m	y		
Significant echo duration	189	119	0.127	0.943	0.345	0.119	0.566	0.171	73.6	8.3	0.0						y
Amplitude stack span	190	120	0.173	4.442	1.066	0.421	-0.550	-0.312	54.3	34.4	5.3				y		
Amplitude stack standard deviation	191	121	0.043	1.099	0.303	0.114	-0.578	-0.327	55.8	29.6	4.0						y
Amplitude stack skewness	192	122	-1.518	3.005	0.511	0.540	0.004	-0.022	27.6	42.1	42.5				y		
Stack mean coherence	193	123	0.129	0.763	0.488	0.092	-0.540	-0.151	62.3	21.6	0.4						y
Stack mean correlation	194	124	0.123	0.934	0.749	0.073	-0.737	-0.346	68.7	21.5	0.0						y
Stack relative deviation	195	125	0.342	1.257	0.529	0.053	0.116	0.193	28.0	45.1	45.5				y		
Fractal - path length	196	126	0.984	3.675	1.611	0.459	0.866	0.436	66.0	44.6	0.0	y					
Fractal - box count	197	127	0.132	0.779	0.431	0.083	0.450	0.185	53.0	33.9	6.6						y
Residual echo duration	198	128	-17.100	38.650	1.360	3.420	0.122	-0.122	52.1	2.5	1.7						y
Max cross-corr clayey sand	199	129	0.038	0.890	0.555	0.140	-0.761	-0.452	69.4	20.8	0.0					m	y
Max cross-corr medium sand	200	130	0.355	0.995	0.896	0.091	-0.452	-0.522	64.9	12.0	0.4					m	y
Max cross-corr sandy gravel	201	131	0.442	0.996	0.948	0.052	-0.499	-0.312	56.5	31.4	11.1					m	y
Max cross-corr rock/cobble	202	132	0.541	0.994	0.937	0.035	-0.053	0.385	53.9	7.4	15.7						y
Cross-corr lag medium sand	203	133	0.249	0.502	0.260	0.013	-0.225	-0.106	46.2	31.4	31.4				y		
Best interpolated cross-corr	204	134	0.601	0.990	0.880	0.063	0.535	0.533	73.0	7.4	0.0						y
Centre of gravity frequency	206	135	0.197	0.758	0.375	0.075	-0.762	-0.430	66.3	23.7	0.0	y			y		
Standard deviation PSD	207	136	0.423	1.148	0.703	0.069	0.017	-0.069	25.4	47.7	47.8				y		
Skewness PSD	208	137	2.371	5.691	3.872	0.383	0.341	0.258	44.5	33.1	14.0				y		
Kurtosis PSD	209	138	6.774	46.621	20.696	4.190	0.152	0.167	36.2	34.7	25.6				y		
Median frequency	211	139	0.000	0.372	0.098	0.064	-0.854	-0.426	84.9	10.1	0.0	y	m		y		
Cumlulative sum bandwidth	212	140	0.012	0.892	0.295	0.126	-0.740	-0.394	58.8	24.3	0.0	m					y
Entropy PSD	213	141	-4.441	-1.617	-2.597	0.262	-0.222	-0.325	66.5	14.0	1.7	m	m		m		

Table 7-18. Feature statistics from the Simrad200 testbed survey, feature evaluation and the reduced feature set for high frequency data. 'y'– yes, 'm' – marginal, no otherwise.

Feature Name			Feature Descriptive Statistics				Quantitative Tests					Feature Map Inspection					Used?
	Feature Number	Sequential Number	Min.	Max.	Mean	Std. Dev.	Corrcoef w. Depth	Corrcoef w. Slope	QC %-correct	Sand to Gravel %-mis.	Sand to Mud %-mis.	Depth Biased?	Slope Biased?	Enson. Biased?	Poor Discrimination?	Correlated?	
Cumulative integral direct (first 8) 1	1	1	0.001	0.060	0.007	0.005	0.456	0.583	50.9	44.1	6.0			y	m		
Cumulative integral direct (first 8) 2	2	2	0.005	0.658	0.081	0.046	0.696	0.606	60.1	40.5	0.9			y	m		
Cumulative integral direct (first 8) 3	3	3	0.056	0.899	0.277	0.094	0.608	0.460	55.7	32.7	13.8			y	m		
Cumulative integral direct (first 8) 4	4	4	0.082	0.950	0.506	0.110	0.478	0.387	51.0	36.9	12.9			y	m		
Cumulative integral direct (first 8) 5	5	5	0.105	0.962	0.688	0.100	0.333	0.352	54.1	26.4	36.2			y	y		
Cumulative integral direct (first 8) 6	6	6	0.129	0.977	0.826	0.075	0.111	0.292	51.0	22.2	40.8			y	y		
Cumulative integral direct (first 8) 7	7	7	0.165	0.983	0.905	0.054	-0.105	0.233	48.2	19.5	21.1			y	y		
Cumulative integral direct (first 8) 8	8	8	0.219	0.988	0.948	0.038	-0.295	0.169	48.2	13.5	9.9					y	y
Cumulative integral direct (next 7) 1	9	9	0.283	0.991	0.966	0.027	-0.356	0.143	49.9	9.3	6.7					y	
Cumulative integral direct (next 7) 2	10	10	0.495	0.994	0.977	0.020	-0.405	0.070	54.7	8.1	6.3					y	
Cumulative integral direct (next 7) 3	11	11	0.543	0.995	0.983	0.014	-0.440	0.013	54.9	9.3	7.6					y	y
Cumulative integral direct (next 7) 4	12	12	0.585	0.997	0.987	0.010	-0.504	-0.042	57.1	5.7	3.6	m			m		
Cumulative integral direct (next 7) 5	13	13	0.634	0.998	0.991	0.007	-0.560	-0.065	54.7	6.9	4.0	m			m		
Cumulative integral direct (next 7) 6	14	14	0.713	0.999	0.994	0.004	-0.615	-0.082	56.1	12.3	2.7	m			m		
Cumulative integral direct (next 7) 7	15	15	0.862	0.999	0.997	0.002	-0.666	-0.128	56.2	9.3	0.9	m			m		
Cumulative integral ratio (first 3) 1	17	16	0.101	0.811	0.414	0.080	-0.019	0.018	39.9	41.7	13.9			y	m		
Cumulative integral ratio (first 3) 2	18	17	0.038	0.773	0.222	0.056	0.454	0.362	41.1	33.0	27.4			y	m		
Cumulative integral ratio (first 3) 3	19	18	0.028	0.764	0.163	0.049	0.598	0.440	32.5	39.3	49.1			y	m		
Cumulative integral ratio (last 4) 1	20	19	0.020	0.756	0.137	0.048	0.634	0.461	46.9	36.3	34.5			y	m		
Cumulative integral ratio (last 4) 2	21	20	0.018	0.749	0.126	0.047	0.629	0.462	46.5	34.8	31.9			m	y		
Cumulative integral ratio (last 4) 3	22	21	0.017	0.742	0.121	0.046	0.615	0.462	45.3	33.3	34.5			m	y		
Cumulative integral ratio (last 4) 4	23	22	0.016	0.736	0.118	0.046	0.604	0.462	44.3	30.9	37.9			m	y		
Quantile (first 7) 1	24	23	0.000	0.320	0.003	0.007	0.153	-0.094	35.2	24.1	23.3					y	y
Quantile (first 7) 2	25	24	0.000	0.470	0.007	0.014	-0.113	-0.264	34.1	18.1	29.3					y	
Quantile (first 7) 3	26	25	0.001	0.619	0.015	0.024	-0.071	-0.315	47.4	31.9	33.6					y	
Quantile (first 7) 4	27	26	0.001	0.728	0.044	0.045	0.235	-0.057	57.7	19.8	12.1					y	y
Quantile (first 7) 5	28	27	0.002	0.783	0.143	0.081	0.137	-0.009	47.8	26.7	25.9				y		
Quantile (first 7) 6	29	28	0.004	0.908	0.406	0.128	-0.308	-0.254	44.7	46.6	20.6			y			
Quantile (first 7) 7	30	29	0.023	0.994	0.721	0.133	-0.590	-0.521	53.8	45.6	6.3			y			
Histogram (first 7) 1	32	30	0.063	0.921	0.611	0.064	-0.084	0.073	42.5	34.8	20.6				y		
Histogram (first 7) 2	33	31	0.000	0.395	0.074	0.030	0.224	-0.022	31.7	39.7	39.7				y		
Histogram (first 7) 3	34	32	0.000	0.370	0.050	0.031	0.376	0.268	53.8	37.1	14.7				y		
Histogram (first 7) 4	35	33	0.000	0.405	0.048	0.029	0.315	0.292	44.1	35.3	12.9				y		
Histogram (first 7) 5	36	34	0.000	0.314	0.046	0.028	0.286	0.290	35.9	35.3	23.3				y		
Histogram (first 7) 6	37	35	0.000	0.270	0.044	0.028	0.244	0.261	40.0	42.2	17.2				y		
Histogram (first 7) 7	38	36	0.000	0.275	0.048	0.028	-0.074	-0.045	24.8	45.7	45.7				y		
Histogram (last)	39	37	0.003	0.270	0.078	0.051	-0.632	-0.593	60.1	44.7	4.5	y	m				
FFT 1-65 (only 1-17* applicable)	102-166	38-54*	-	-	-	-	-	-	-	-	-	y	m		y		

Table 7-18. Continued.

Feature Name			Feature Descriptive Statistics				Quantitative Tests					Feature Map Inspection					Used?
	Feature Number	Sequential Number	Min.	Max.	Mean	Std. Dev.	Corrcoef w. Depth	Corrcoef w. Slope	QC %-correct	Sand to Gravel %-mis.	Sand to Mud %-mis.	Depth Biased?	Slope Biased?	Enson. Biased?	Poor Discrimination?	Correlated?	
Envelope centre of gravity	167	55	0.131	0.608	0.273	0.034	-0.461	-0.471	59.9	20.7	29.1				y		
Envelope time-spread	168	56	0.119	0.693	0.227	0.034	0.423	0.048	58.1	17.2	8.6						y
Envelope skewness	169	57	0.260	1.830	0.726	0.081	0.377	0.132	50.5	24.1	8.6				y		
Cumulative amplitude thres. 1%	170	58	0.009	0.119	0.062	0.010	0.243	-0.050	62.5	25.0	6.0			y			
Cumulative amplitude thres. 20%	171	59	0.081	0.521	0.166	0.022	-0.617	-0.471	53.8	18.1	22.4			y			
Cumulative amplitude thres. 60%	172	60	0.110	0.833	0.280	0.039	-0.461	-0.383	58.9	24.1	19.3				y		
Cumulative amplitude thres. 75%	173	61	0.130	0.911	0.332	0.045	-0.351	-0.363	55.8	30.2	33.2				y		
Cumulative amplitude thres. 90%	174	62	0.158	0.956	0.415	0.053	-0.164	-0.342	48.3	31.9	38.8				y		
Cumulative intensity thres. 0.5%	175	63	0.022	0.179	0.084	0.010	-0.320	-0.318	40.5	22.4	8.6			y			
Cumulative intensity thres. 10%	176	64	0.074	0.584	0.150	0.023	-0.681	-0.486	49.6	33.6	24.2			y			
Cumulative intensity thres. 95%	177	65	0.130	0.989	0.383	0.055	-0.315	-0.328	58.2	26.7	29.6				y		
Cumulative intensity thres. 99%	178	66	0.155	1.000	0.451	0.061	-0.145	-0.321	47.7	19.8	41.4					m	y
Cumulative intensity thres. 99.9%	179	67	0.180	1.000	0.533	0.069	-0.014	-0.340	54.4	18.1	24.1					m	y
Smoothed peak amplitude	180	68	0.361	2.909	1.044	0.273	-0.279	-0.143	51.3	34.5	20.2					m	y
Mean amplitude from bottom pick to tail	181	69	0.048	1.497	0.441	0.133	-0.509	-0.273	57.0	37.1	6.7		m			m	y
Mean amp from end of pulse to tail	182	70	0.036	1.025	0.383	0.102	-0.217	0.127	40.0	39.7	27.8					m	y
Mean intensity peak to tail / peak amp.	183	71	0.001	0.893	0.157	0.086	-0.558	-0.301	53.3	48.0	9.4		y				
Mean amplitude of the half maximum	184	72	0.128	2.557	0.813	0.244	-0.426	-0.303	53.3	36.2	14.8		m				y
Ratio of energy after pulse / during pulse	185	73	0.154	10.000	3.249	3.329	0.646	0.725	63.1	34.5	0.9		y				
Half maximum duration	186	74	0.018	0.912	0.228	0.057	-0.283	-0.208	51.1	31.9	22.9				y		
Fall fraction	187	75	0.081	0.970	0.699	0.100	0.436	0.081	34.8	48.3	33.6				y		
Fall time	188	76	0.078	0.835	0.375	0.083	0.431	0.018	39.3	44.8	21.6				y		
Significant echo duration	189	77	0.088	0.953	0.410	0.067	0.161	-0.025	51.9	25.0	19.0						y
Amplitude stack span	190	78	0.263	4.601	1.508	0.495	-0.342	-0.320	40.7	40.5	26.9				m		y
Amplitude stack standard deviation	191	79	0.065	1.358	0.434	0.138	-0.381	-0.348	43.3	43.8	21.1						y
Amplitude stack skewness	192	80	-1.845	2.944	0.514	0.517	0.030	-0.057	30.8	32.4	49.3				y		
Stack mean coherence	193	81	0.012	0.931	0.456	0.098	-0.545	-0.490	56.8	33.6	8.1				y		
Stack mean correlation	194	82	0.082	0.986	0.777	0.091	-0.588	-0.583	57.4	27.9	4.5				y		
Stack relative deviation	195	83	0.245	1.553	0.566	0.096	0.137	0.222	33.9	39.6	44.8				y		
Fractal - path length	196	84	0.884	3.145	1.143	0.232	0.662	0.586	36.2	37.2	36.2	y					
Fractal - box count	197	85	0.000	1.090	0.496	0.127	-0.517	-0.428	49.8	19.0	27.4				y		
Residual echo duration	198	86	-15.820	14.670	0.019	1.650	-0.206	-0.493	37.0	19.8	25.0						y
Max cross-corr clayey sand	199	87	0.229	0.983	0.799	0.133	-0.647	-0.729	57.0	31.2	2.7		y				
Max cross-corr medium sand	200	88	0.342	0.997	0.924	0.090	-0.378	-0.826	46.1	27.6	11.2		y				
Max cross-corr sandy gravel	201	89	0.359	0.997	0.933	0.083	-0.384	-0.796	44.4	41.4	7.6		y				
Max cross-corr rock/cobble	202	90	0.373	0.997	0.937	0.079	-0.390	-0.765	50.0	46.6	8.1		y				
Cross-corr lag medium sand	203	91	0.248	0.504	0.280	0.017	-0.325	0.337	38.7	39.7	28.7		m		m		
Best interpolated cross-corr	204	92	0.573	0.990	0.833	0.067	0.291	0.362	53.8	27.6	31.8		y				
Centre of gravity frequency	206	93	0.244	1.135	0.521	0.119	-0.711	-0.651	54.9	39.6	2.2	m			y		
Standard deviation PSD	207	94	0.498	1.503	0.799	0.102	-0.051	-0.080	27.8	49.8	45.3				y		
Skewness PSD	208	95	2.329	6.908	4.339	0.578	0.461	0.454	45.0	41.4	15.5	m			y		
Kurtosis PSD	209	96	8.291	71.093	28.170	7.165	0.300	0.316	41.9	42.2	19.0				y		
Median frequency	211	97	0.244	0.706	0.120	0.120	-0.771	-0.672	73.9	26.1	0.0	m	m		y		
Cumlulative sum bandwidth	212	98	0.033	1.283	0.433	0.154	-0.700	-0.668	58.1	34.8	1.3	m			y		
Entropy PSD	213	99	-3.801	-1.638	-2.586	0.278	0.730	0.656	63.8	36.2	0.0	y			y		

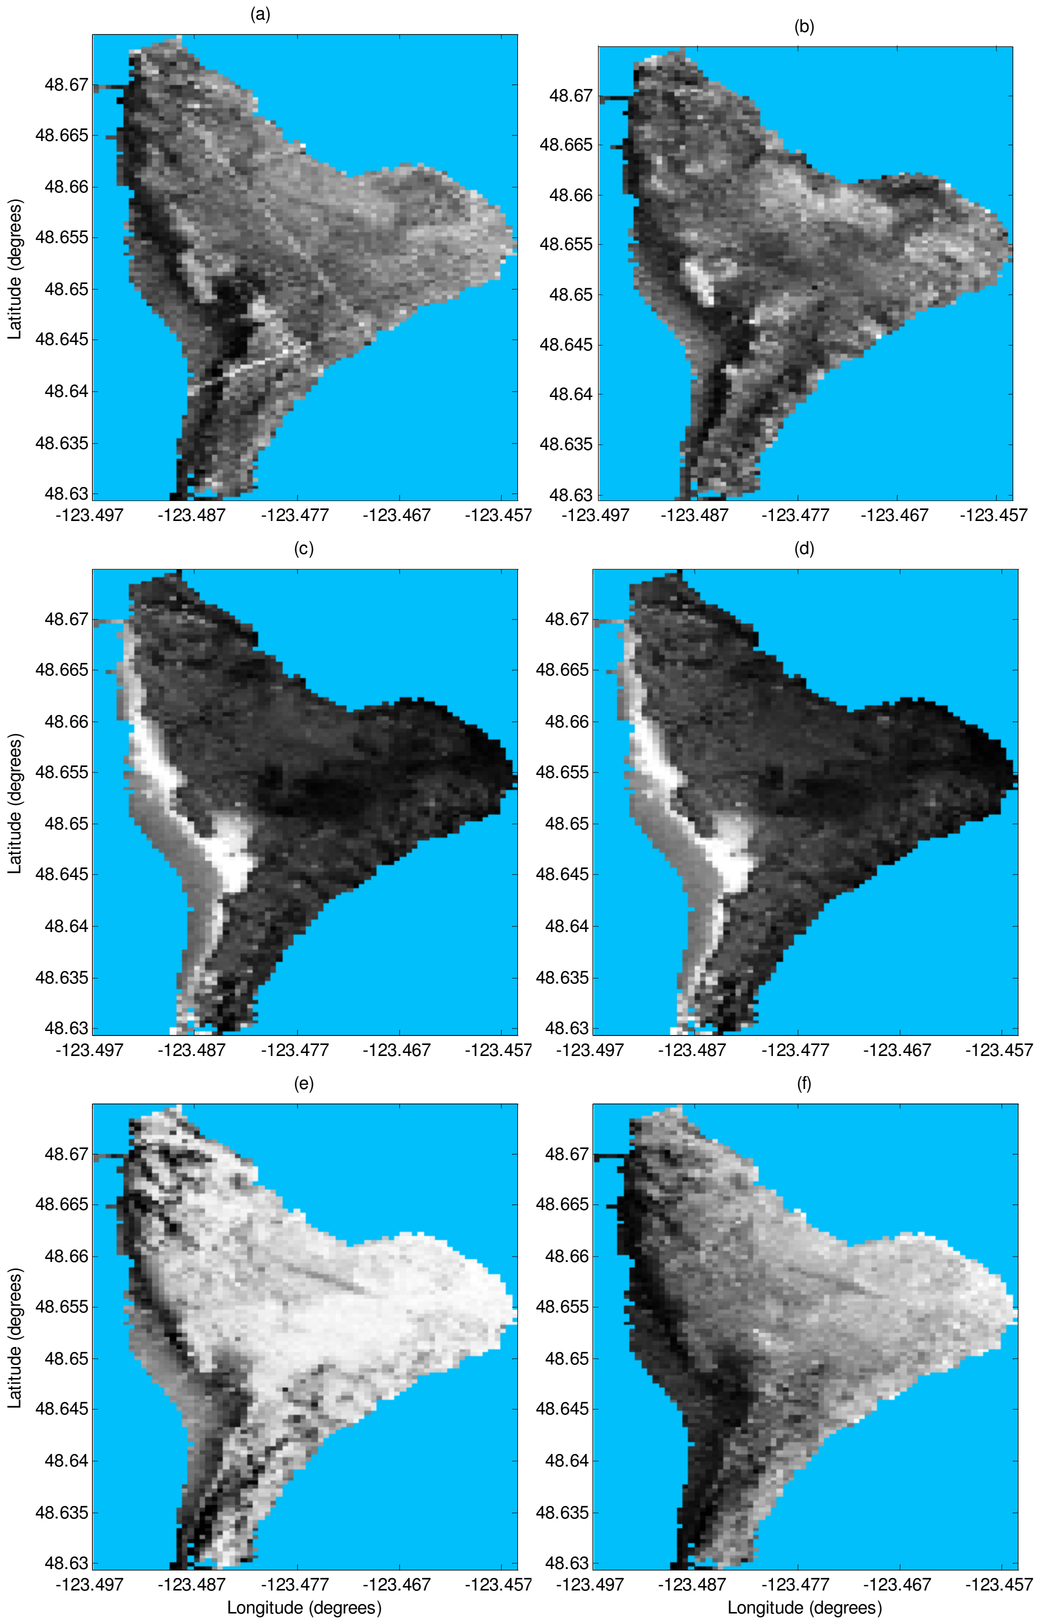

Figure 7-13. Feature maps for the Airmar38: (*a*) peak amp., (*c*) centre of gravity, (*d*) cum. eng. thres. 99%, (*e*) AEM x-cor. m. sand, (*f*) c.s. bandwidth. For Simrad200: (*b*) peak amp.

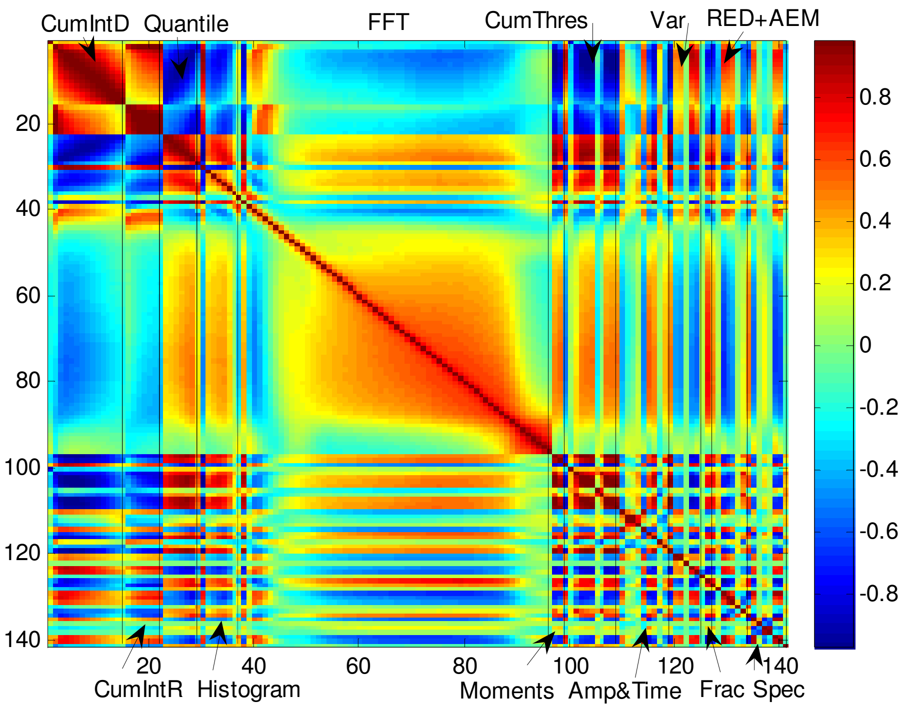

Figure 7-14. The correlation matrix for all available features from the Airmar38 survey as processed by SELwNF. The colourbar shows the correlation value. The numbers on the axes are sequential feature numbers that correspond to the features as listed in Table 7-17.

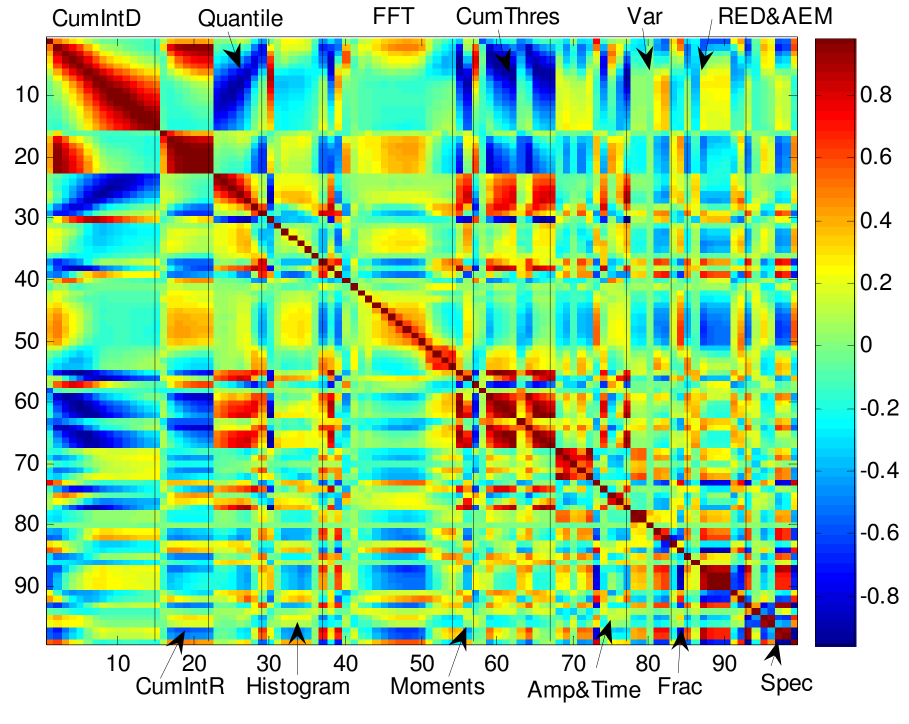

Figure 7-15. The correlation matrix for all available features from the Simrad200 survey as processed by SELwNF. The colourbar and axes annotations are as before.

After combining all of measures listed in Table 7-17 and Table 7-18, the reduced feature are comprised of 39 and 15 features for the low and high frequency cases, respectively. Most of these features are new features. The high-frequency reduced feature set contains fewer features than the low-frequency reduced feature set. Narrow-beam high frequency SBES are almost always dominated by incoherent surface backscatter so these types of SBES offer less discriminatory ability compared to lower-frequency moderate-beam SBES that experience a wider range of surface and volume scattering domains spread over a greater footprint. This is evident in comparisons of the features maps. However, the high frequency Simrad200 does exhibit beam-widening and impedance variation which are represented by the echo duration, time-spread and amplitude features (the peak amplitude is a particularly well-performing feature, see Figure 7-13). The amplitude features appear to be the best features for the high-frequency Simrad200 SBES; however, these features are not compensated for slope and are sensitive to slopes greater than the ½ beamwidth. This slope bias of the peak amplitude, in particular, will manifest in the classification shown in the next sub-section.

Possible future work will be to continue to improve the features. The trim-TVG technique can be expanded to include compensation for slope. The variability and spectral moments may be improved by geographic search, which will allow larger stacks. The spectral features will certainly be improved when applied to FWF data, and will likely perform as well as they did in testing with BORIS. Ideas for new features could be developed and evaluated in the same way as the new features listed here. The testbed framework is an excellent resource for feature development or research. New multibeam features may also be adapted from these features or be conceived from the physical basis of echosounding established here.

7.5.3 Classification Results and Discussion

In spite of how features perform on their own, they are only part of the improved classification methodology (see Figure 7-1). In this sub-section, the final classification results are presented and compared to the baseline TNORM results. The final classification method is dubbed *SELwNF*, which is a contracted acronym for **s**tandard **e**cho **l**ength variable **w**indow with effective depth and slope compensation with **n**ew **f**eatures processed with iterative outlier removal correlation PCA. The results are presented here following the same sequence as in section 7.3, starting with a qualitative analysis of the classification maps and ending with quantitative analysis.

Overall, the Airmar38 SELwNF result in Figure 7-16 is very good in comparison to the ground-truth map and is much improved from the TNORM result. It separates the three major classes of Patricia Bay very well and is insensitive to non-seabed influences. It does not discern the fine and medium sands, although that is not unexpected considering how similar those two seabeds are compared to differences between the mud and sand seabeds (especially in terms of echo duration, amplitude, etc.). The amplitude artefact that was visible as bright lines in the peak amplitude feature map, panel (*a*) of Figure 7-13, is apparent as class 8 in Figure 7-16 (*lower*). This minor artefact may be an operator error; such as change in transmit power setting during the survey. The class does not appear in the Simrad200 data. The Airmar38 classification result is slope and depth insensitive. There is no class confusion in the yellow, green and red boxes of Figure 7-17. The purple box of Figure 7-17 is an area where the multibeam and the Airmar38 disagree due to their different frequencies. This effect was noted earlier in SELw compensation and is apparent here. The 300 kHz EM3000 multibeam is strongly affected by relatively thin mud and silt veneers that the Airmar38 will penetrate. For a comparison of penetration depths, see Figure D-1. Thin veneers may also be responsible for slight disagreements on the boundaries of the gravel and sand regions.

The Simrad200 SELwNF result is insensitive to non-seabed influences, but suffers from the limited discriminatory ability of the high-frequency narrow-beam echosounder. In Figure 7-16 (*upper*), Classes 4 and 6 represent the mud seabed, while the low sloping gravels are found primarily in classes 3, 9 and 10. Class 11 corresponds to gravel – mud boundary, which appears to be an erosional moat in the multibeam relief

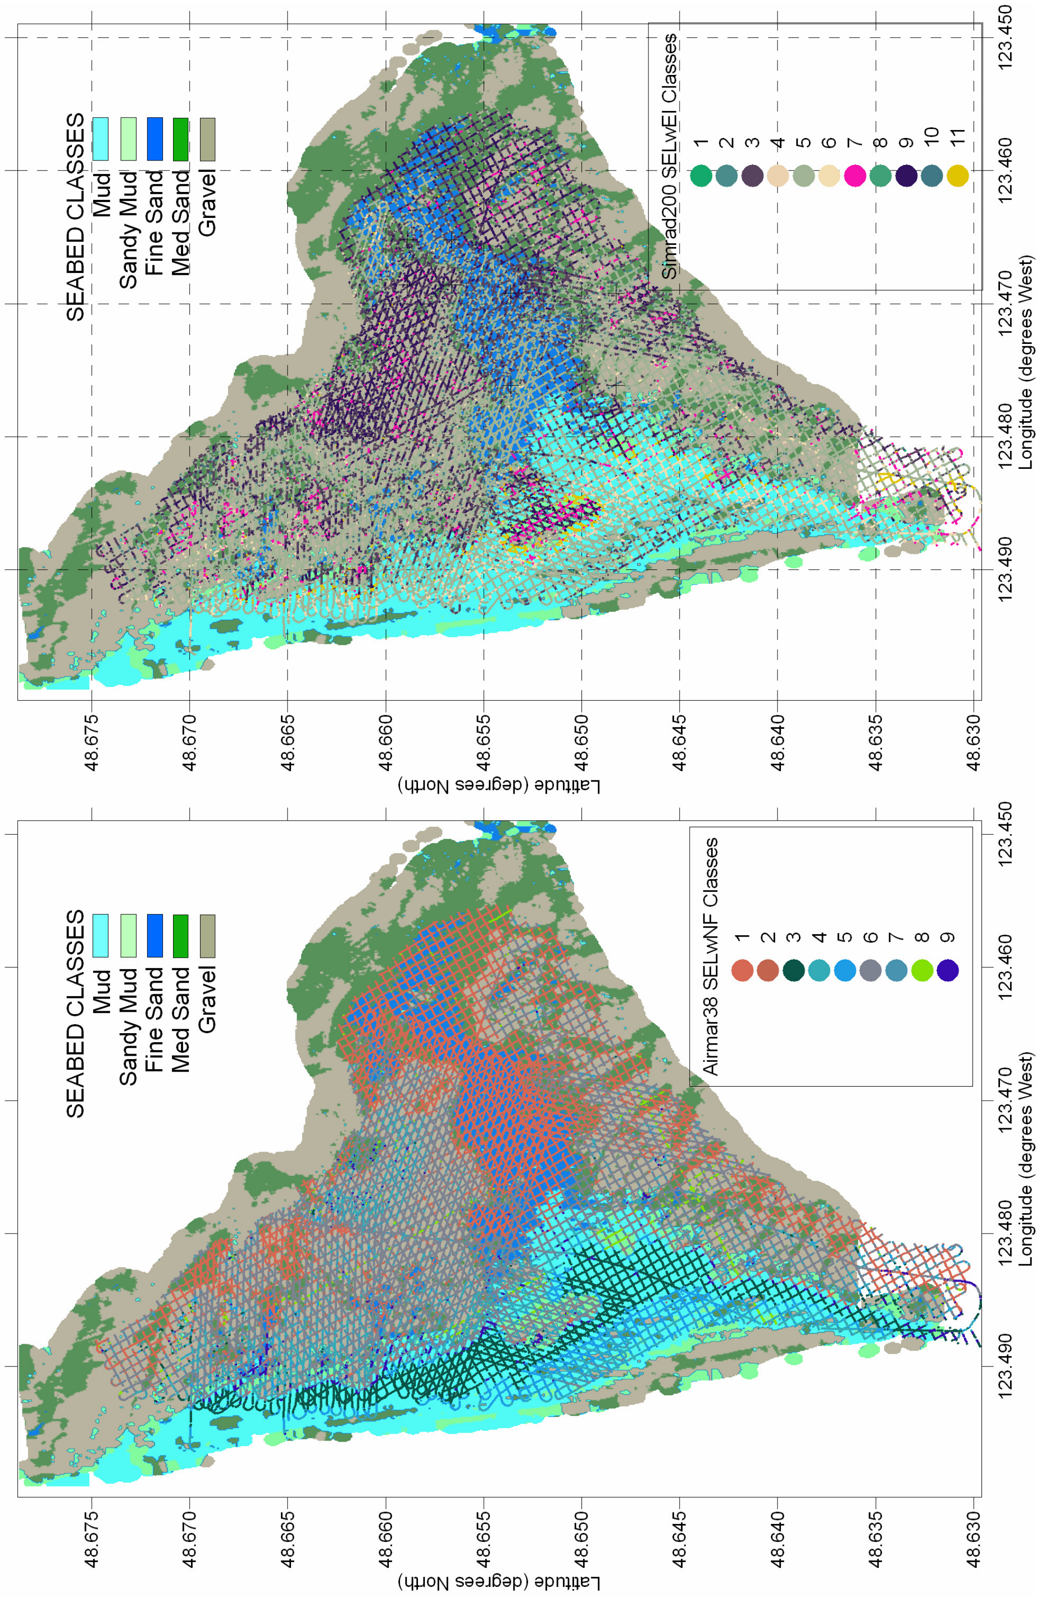

Figure 7-16. Airmar38 (*lower*) and Simrad200 (*upper*) SELwNF seabed classification overlaid on the multibeam seabed map.

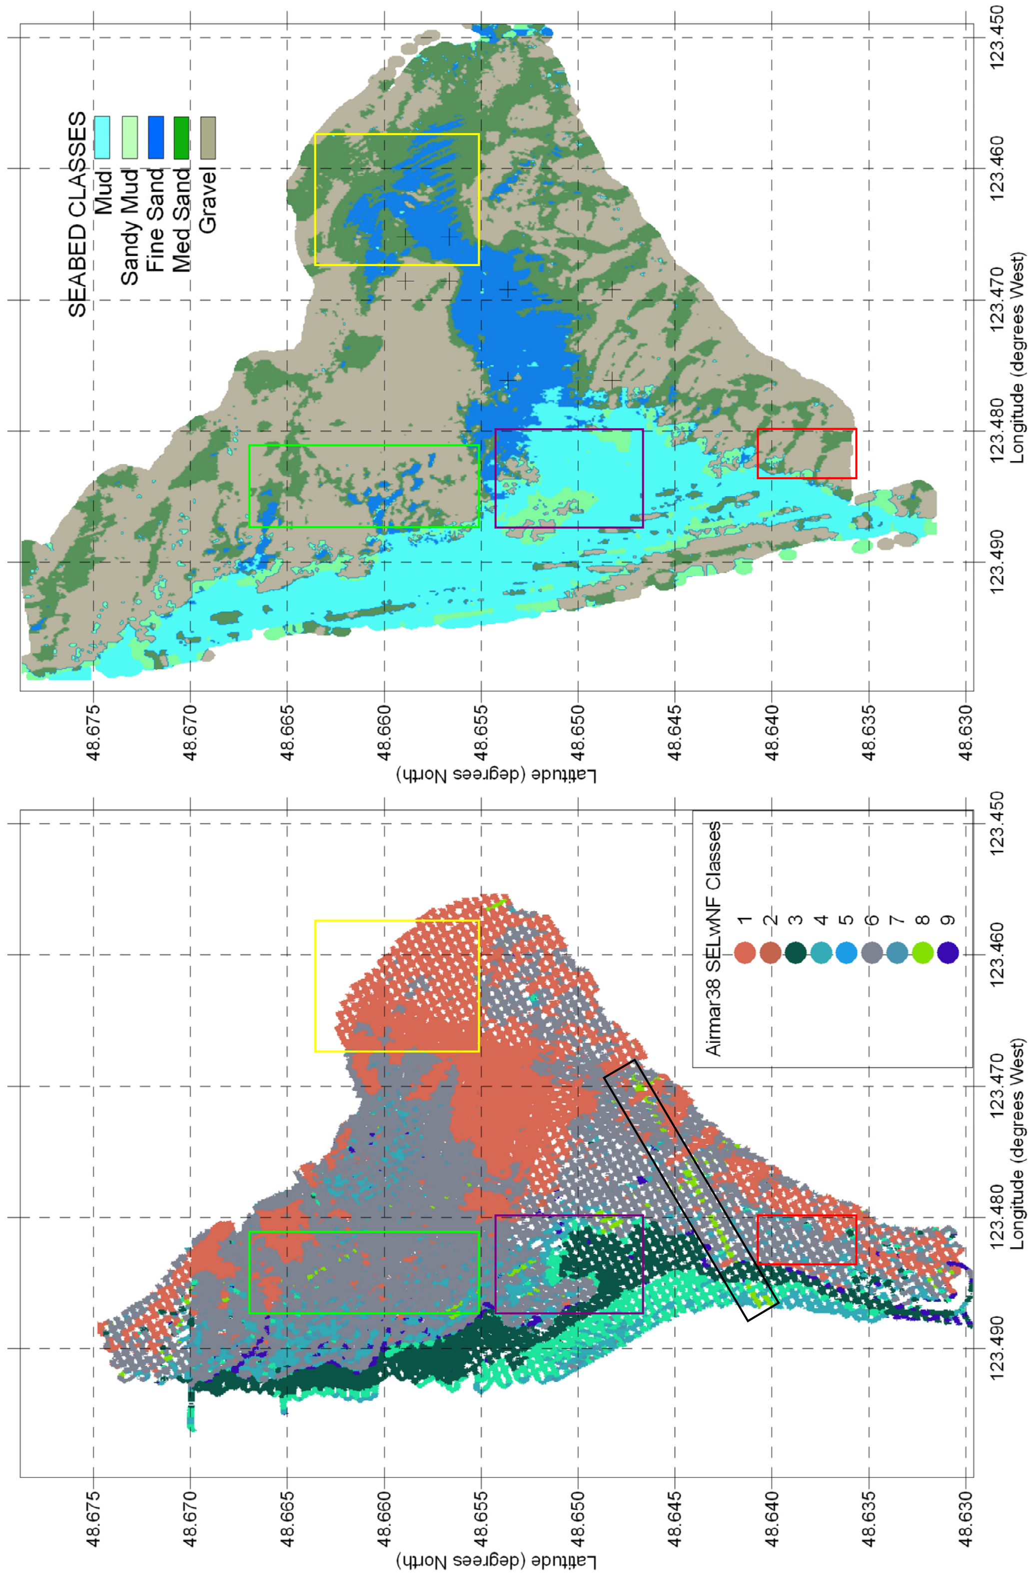

Figure 7-17. A comparison of the SELwNF classification map for the Airmar38 testbed survey (*lower*) to the multibeam seabed map (*upper*). The boxes are highlight areas to be discussed. Both maps were generated with an interpolation search radius of 30 m.

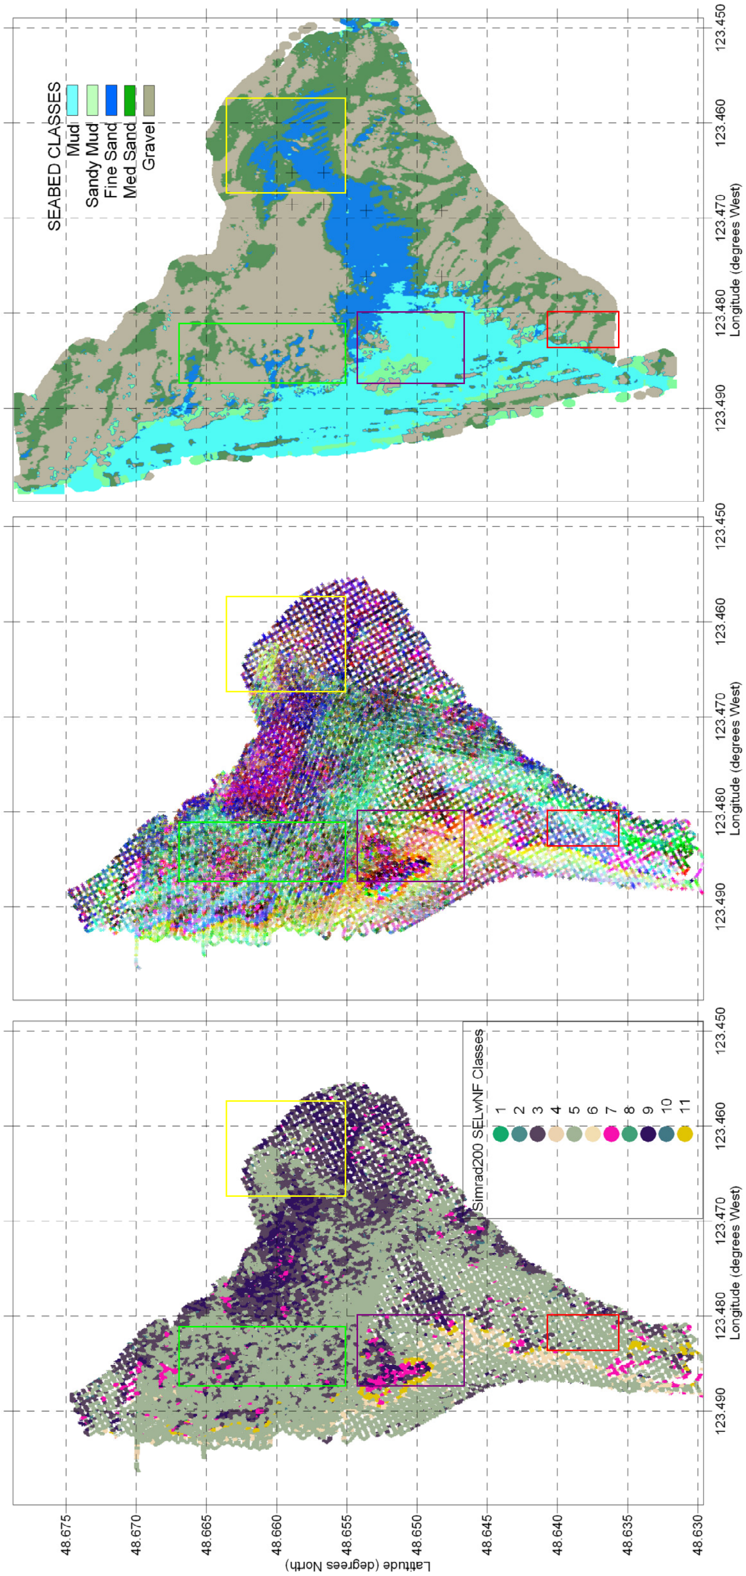

Figure 7-18. A comparison of the SELwNF classification (*lower*) and OCPR (*middle*) maps to the multibeam ground-truth class map (*upper*) for the Simrad200 testbed survey. The boxes are highlight areas to be discussed.

map (Figure 6-3). Long duration echoes are also observed in the vicinity of the moat. Class 5 is a catch all class with 51% membership. A dominant class like class 5 is an indication that cluster analysis may have under-split the distribution of points in Q-space. This is confirmed as one can better delineate seabed boundaries in the OCPR[57] map, Figure 7-18 (*middle*). A slope bias is also seen in the green and red boxes in Figure 7-18. This bias comes from the amplitude features[58], which are prominent in the reduced feature set. The EDM predicts, and section 6.3 confirms that narrow beams are more sensitive to small tilts than wider beams. Minor striping in the feature maps are due to a slight tilt of the transducer (1.5 °). Because some survey lines ran upslope and downslope, a small tilt was added or subtracted from the slope with each anti-parallel pass; as was seen in the slope experiment (Figure 6-18). These tilt artefacts are like the ensonification regime artefact in that they are not blatantly detrimental until the stronger non-seabed influences of depth and slope are removed. Wider-beam SBES are less sensitive to effects of tilt, seabed and vessel motion.

The quantitative measures of classification performance validate the above discussion. The Airmar38 shows marked improvement as seen by comparing the OA and AMI statistics for the TNORM and SELwNF results (Table 7-19), while the Simrad200 shows only a marginal difference. These statistics were calculated as described in section 7.2.

---

[57] OCPR – defined in sub-section 2.2.3

[58] Amplitude features are not compensated for seabed slope. However, these features show little sensitivity until the seabed slope is roughly one-half the beamwidth, at which point the specular component is lost. Shape features are sensitive at slopes as low as 2°. trim-TVG may be adapted to compensate amplitudes for slope using the compensation table approach used with QTC MULTIVIEW [Preston and Christney, 2005].

Table 7-19. AMI and OA statistics comparing the baseline classification result (TNORM) to the final classification result (SELwNF). Note that best-match OA is biased upwards for higher numbers of classes in the merge, which is why these numbers are provided.

TNORM	Airmar38	Simrad200
Best-match OA (%)	48	57
No. classes:	8	17
Merge-QC OA (%)	53	-
AMI (%)	15	13

SELwNF	Airmar38	Simrad200
Best-match OA (%)	70	56
No. classes:	9	11
Merge-QC OA (%)	78	-
AMI (%)	41	19

The merge-QC OA between the SELwNF classification for the Airmar38 and the multibeam ground-truth map was the best result overall at 78%; a result that is quite satisfactory considering the differences between ground-truth and the classification explained by the difference in centre frequency. For comparison, an S-QC OA of 90% was determined for multibeam data compared to itself in a similar scenario [Wilmut et al., 2009]; this is a estimated upper limited for OA values found using QC. The lower limit for OA values is 33%, which is the OA for random classification of three classes. The AMI statistic is in agreement with OA in showing improvement in classification performance from the TNORM method to the SELwNF method.

The affect of including features with poor discrimination ability from the reduced feature set was studied in a variant of SELwNF, dubbed *SELwNF+*, which was defined as follows: the features that were rejected due to poor discriminatory ability only (as determined by percent misassignment and S-QC percent correct assignment) were added to the reduced feature set to make a 'less' reduced feature set. For the low-frequency set, 21 features were added (to make 60 total), while 23 were added (to make 38 total) in the high-frequency less reduced set. Ideally, the result should not be affected by features with poor discrimination, as PCA should ignore them unless these features have correlated noise or are correlated to artefacts. Table 7-20 summarizes the quantitative evaluation of the SELwNF+ classification result compared to that of the SELwNF result. By including features with poor discrimination, the result is degraded from the SELwNF result (by 9% in AMI, 6% in merge-QC OA and 7% in best-match OA), but is still better than the result of SELw classification (see Table 7-8; improvements of 2% AMI, 7% merge-QC OA, and 0% best-match OA). Note that Figure 2-9 is the Q-space plot of the Airmar38 SELwNF+ Q-values used to select classes 6, 10 and 15 for training.

Table 7-20. AMI and OA statistics for the Airmar38 testbed survey comparing the reduced feature set to the less-reduced feature set.

SELwNF	Airmar38	Simrad200
Best-match OA (%)	70	56
No. classes:	9	11
Merge-QC OA (%)	78	-
AMI (%)	41	19

SELwNF+	Airmar38	Simrad200
Best-match OA (%)	63	44
No. classes:	16	9
Merge-QC OA (%)	73	-
AMI (%)	31	7

This comparison between the SELwNF and SELwNF+ results shows that poor features can degrade the classification performance. It is advantageous to reduce the feature set to the best features so that the performance of PCA is improved. This finding is consistent with the fact that several other seabed classification methods that have been successful with small feature sets, such as RoxAnn and the method presented of [van Walree et al., 2005], as described in section 2.4. The improvement of SELwNF+ over the SELw result is difficult to evaluate as not only have the features changed but also biased covariance PCA has been replaced with unbiased correlation PCA. Regardless of whether or not there is any improvement in classification, the addition of the new features is advantageous as these features are based in physics and have known relationships to geoacoustic properties of the seabed. The performance of the new features and small feature sets will be explored further in sections 7.6 and 7.8.

7.6 Supervised Classification

It is not entirely sufficient to evaluate classification methods based solely on the unsupervised classification outcome. Cluster analysis can over or under-split classes and PCA removes any connections to the features as discussed in chapter 2. To look more directly at the discriminatory power of the features by skipping the PCA and cluster steps, supervised quadratic classification (S-QC) is used. S-QC creates a seabed catalogue or trains on spatially defined areas.

For both echosounders, S-QC was trained on four geographical sites: the familiar homogenous experiment sites (see sub-section 6.1.1) and a flat area of the fine sand seabed, see Figure 6-1. After S-QC was run on the basis of four training sets, the fine sand and medium sand QC classes were merged since they had significant overlap in feature space. To create a square three by three confusion matrix, the five ground-truth classes were also merged to three classes (the muds and the sands were combined). In training the Airmar38 TNORM-processed data, the 1st cumulative integral and 1st wavelet (essentially amplitude) features were selected on the basis of the feature selection criteria as discussed above, resulting in a combined percent correct assignment of 91%. However, in training the Simrad200 TNORM-processed data, no combination of features could be found to achieve a value of percent correct assignment greater than 75%. Then, the Simrad200 TNORM-processed data could not be trained reliably, so its analysis was not taken further. For the Airmar38 TNORM-processed data, classification of the data set, based on the training above, resulted in an S-QC OA with the ground-truth of 53%; the same as the merge-QC result.

In training on the Airmar38 SELwNF feature data, two features were selected on the basis of their performance in discriminating the training sites and their correlations. These features were the centre of gravity and the time spread features (167 and 168, described in sub-section 7.5.1). The percent correct assignment obtained in training with these two features was 93%. Classification of the data set resulted in an S-QC OA of 77%. The resulting class and misclassification maps are shown in Figure 7-19. These maps show the same patterns as noted earlier. In particular, the largest disagreement between the classification and the ground-truth occurs for the purple box of Figure 7-17, which is due to the different centre frequencies of the systems used as explained earlier in this sub-section. In all, S-QC further confirms the satisfactory performance of the new features and SELw depth and slope compensation.

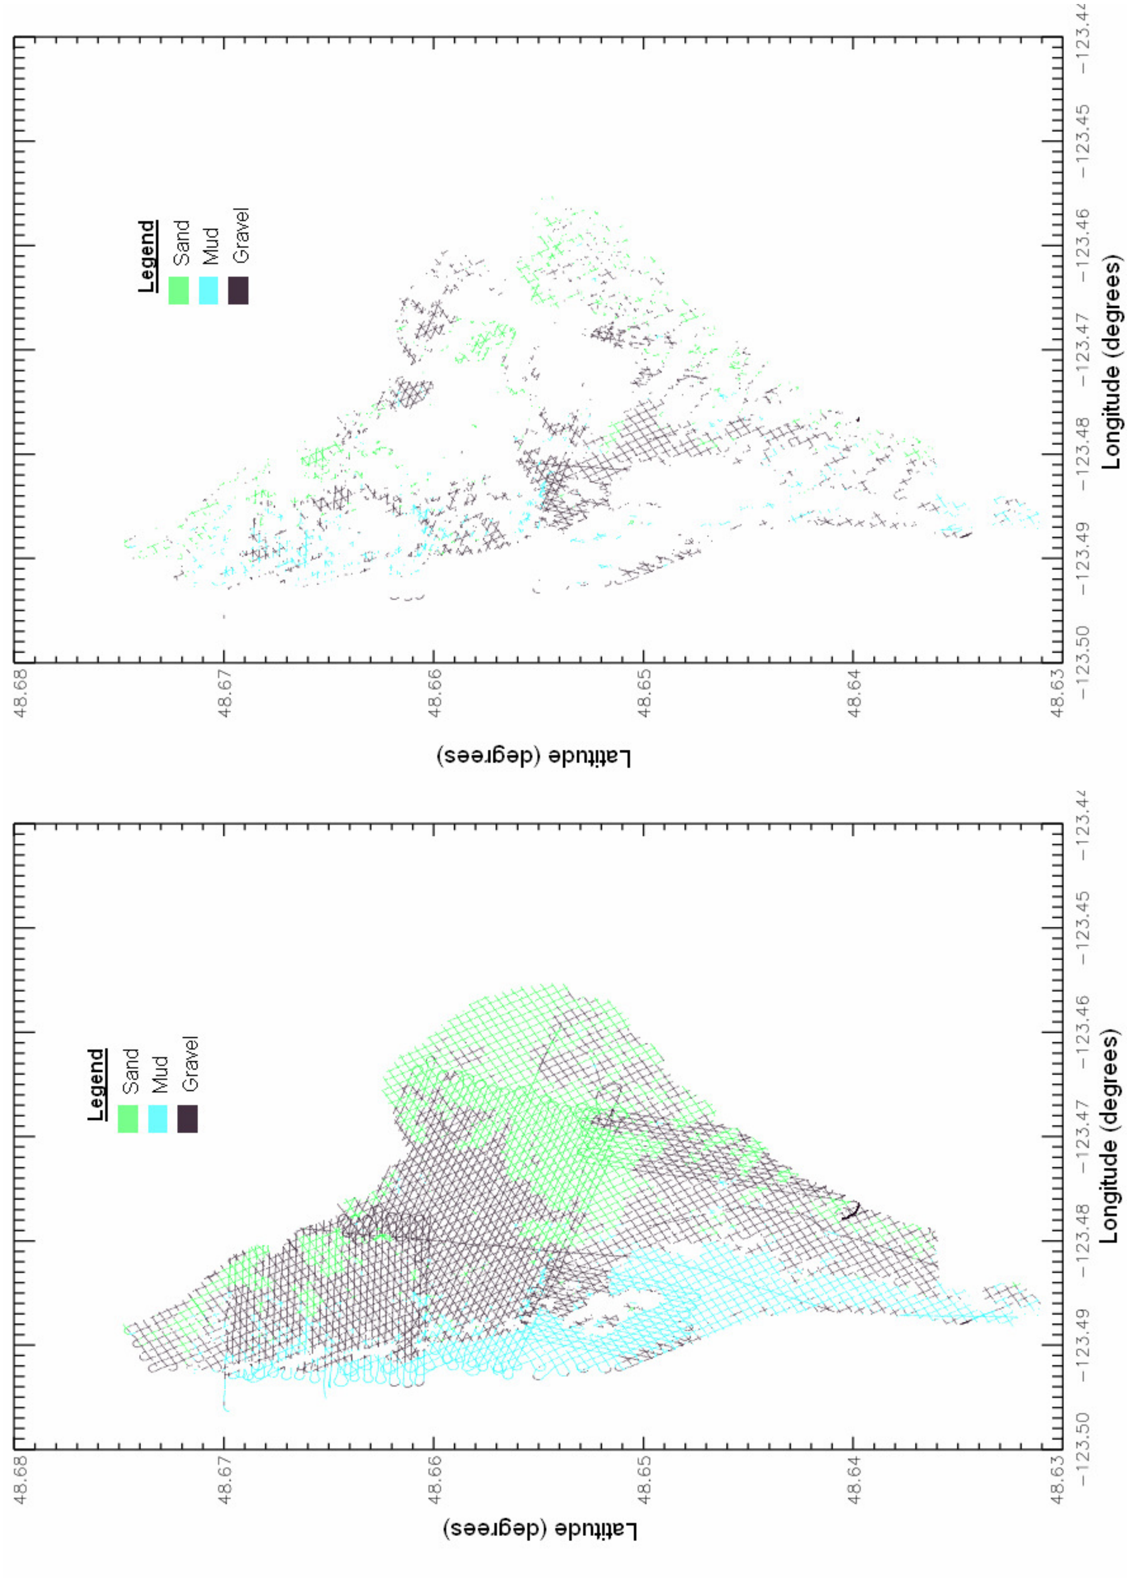

Figure 7-19. S-QC classification map (*lower*) for the Airmar38 survey, based on analysis of all features produced by the SELwNF method. A map of the misclassified points (upper) was derived relative to the ground-truth class map using a spatial search.

7.7 Conclusions for Improved Classification

The new SELwNF classification methodology follows a similar general processing flow to the original methodology. However, the two are very different. The new physics-based features were designed for single-beam echosounders for improved discrimination, and the removal of non-seabed influences has been advanced considerably. These advances are combined with iterative outlier removal correlation PCA, a reduced feature set and evolved mapping and interpretation strategies. Strategies for reliable estimates of methods' parameters have been developed, (effective beamwidth and attenuation in particular). The methods are designed so that expert users can modify the general processing strategy with the tools provided for evaluating the choices of n, m, values of θ_{eff}, κ_{eff} and feature selection so that methodology is adaptable to any SBES. The methodology should be successful and robust in most survey scenarios, the reasons why are numerous:

- the methodology is flexible and adaptable;
- the methodology uses multivariate statistical segmentation which is a proven and robust technique;
- the methodology has been developed with thorough testing, modelling and with a basis in the physical understanding of echosounding;
- the methodology makes use of effective depth and slope compensation;
- the methodology makes use of new features that have improved discrimination and have physical meaning, which may be interpreted or combined with characterization methods to label classes without the need for ground-truth;
- the methodology makes use of iterative outlier removal correlation PCA and reduced feature sets to improve the performance of PCA;
- the method was shown, through several analysis methods, to be successful in Patricia Bay, which is a diverse and difficult survey scenario.

Future work should include applying the SELwNF method to survey unknown areas.

The new methodology improves the results in Patricia Bay in several ways. First, classes are not correlated to depth and slope. Second, the merge-QC OA for the Airmar38 survey improved from 53% to 78%, indicating a significant improvement from a marginal result to a result with good agreement with the ground-truth. The only major inconsistency is likely due to the different frequencies of the Airmar38 and EM3000 MBES that collected the ground-truth data. This satisfies the goal of improving classification in terms of the removal of non-seabed influences and improved discrimination. The other way in which the new methodology is an improvement is that the new features have a basis in physics and a physical interpretation. That is, the relative values of these features indicate an attribute of the seabed such as roughness or impedance. This can be used to identify the seabed type of the classes produced without the need for ground-truth data. This property of the new features is exploited in feature interpretation tables which are combined with other post-class characterization methods in chapter 8. This is step (D1) in Figure 7-1. Similar to post-class characterization is the idea of three-feature classification, which is presented next as it is a hybrid between classification and characterization.

7.8 Three-feature Classification and Relative Characterization

In this section, an alternative classification scheme is presented that has an advantage over SELwNF classification in that it is not phenomenological. This is so because the PCA step is skipped (the Q-values are replaced by the feature values) so the final classification has a connection to the physical interpretation of the features used. The classes produced can be labelled with relative seabed information generated by way of inspections of three-space plots of the feature data. Without the aid of ground-truth information, unsupervised SELwNF classes may be interpreted with seabed type labels by comparison with the three-feature seabed class map. One can use three-feature classification as a seabed map on its own, however, unsupervised SELwNF classification does perform better, and so the two are best used together. The process of identifying unsupervised classes with three-feature classification is shown in Figure 7-1 under steps

(D2) and (D3). Three-feature classification is a relative characterization scheme and is similar to the post-class characterization methods to be presented in chapter 8.

7.8.1 Relative Characterization by the Interpretation of One-Colour-Per-Record Plots

Theoretical models of echosounding indicate echoes only vary in certain ways, as was discussed in section 4.7, illustrated in Figure 4-5 and summarized in Table 4-1. The predicted behaviour of echoes in response to the seabed is supported by observations of models such as BORIS and the AEM (chapter 5) and observations of real echoes from the testbed (chapters 6 and 7). The three major ways in which echoes respond to changing seabeds are in their amplitude, duration and variability. Amplitude increases with impedance, and decreases with surface roughness. Secondary effects of attenuation and volume scattering inhomogeneities affect the amplitude of the volume echo. Echo duration, as well documented in chapters 5, 6 and 7, increases with roughness, via beam-widening, and increases as seabed attenuation decreases. Variability is a direct consequence of the incoherent component of echoes. It increases with surface roughness and with volume inhomogeneities.

Now, if one were to plot three features that are measures of amplitude, duration and variability in a three dimensional (3-D) scatter plot one might be able to separate the data in groups using cluster analysis, just as is done in Q-space in the standard QTC classification method. If the features are good measures of the three major attributes of echoes, then three-feature classification will produce a good result. However, classification by simply using cluster analysis on three-feature space would be ignoring the advantage of using three features directly. That is, the PCA step has been skipped and the positions of echoes in 3-D feature space have relationships to the geoacoustic parameters as mentioned above. However, the three major attributes do not have one to one relationships with the geoacoustic parameters, so instead it is *combinations* of the attributes that will be indicative of seabed type. And as will be shown, the combinations tend to be unique and there is little class confusion.

If echo positions in the 3-D feature space are painted by colours determined by their feature values, certain colour combinations will be indicative of seabed type. This is just like what is done with similarity colours for Q-space, except that the similarity colours class relationships indicated by the colours correspond to seabed types instead of acoustic diversity. In general, duration will be assigned red, amplitude green and variability blue. So every echo will have its own colour determined by the values of its three features. This is analogous to the one-colour-per-record (OCPR) maps shown in chapters 3 and 7, but in this case, the colours are associated with seabed characteristics instead of nebulous Q-values from PCA. Given the known relationships and colour assignments, one can predict the colours of the seabed as shown in Table 7-21.

Table 7-21. Relative strengths of the three major echo attributes and the expected three-space colour for the three major seabed types.

Seabed Type:	Echo Attribute: Duration (red)	Ampltude (green)	Variability (blue)	Expected Colour:
Gravel	moderate - high	low-high	moderate-high	grey-brown to white
Sand	low	high	low-moderate	green, blue-green
Mud	high	low	moderate-high	red, hint of purple

The features selected to be applied with this concept are the best performing features for each echo attribute of amplitude, variability and duration for each frequency as evaluated in Table 7-17 and Table 7-18. The features used are listed in Table 7-22. The selection criteria heavily weight the features' performance in terms of depth and slope bias and the ability to discriminate the major seabed types (see sub-section 7.3.2, Table 7-17 and Table 7-18). This is a method design process which does not aim to select features that best match the ground-truth data. There are other issues to note here: the amplitude features are biased for slopes greater than the half-beamwidth, the stack mean coherence is the negative of variability and the amplitude features for the Airmar38 had minor base gain artefacts as seen in Figure 7-13 (a). The feature values are plotted by colour in 3-D in Figure 7-20, while the resulting OCPR maps are shown as Figure 7-21 and Figure 7-22. Note that the features are scaled by their maximum values and the colour scale is linear spanning the range of 99% of the data.

Table 7-22. The features used for three-feature classification.

SBES:		Echo Attribute:		
		Duration	Ampltude	Variability
Airmar38	Feature Name:	Residual echo duration	Peak amplitude	Stack mean coherence
	Feature Number:	198	180	193
Simrad200	Feature Name:	Residual echo duration	Mean amp. pulse to tail	Amp. Stack standard deviation
	Feature Number:	198	182	191

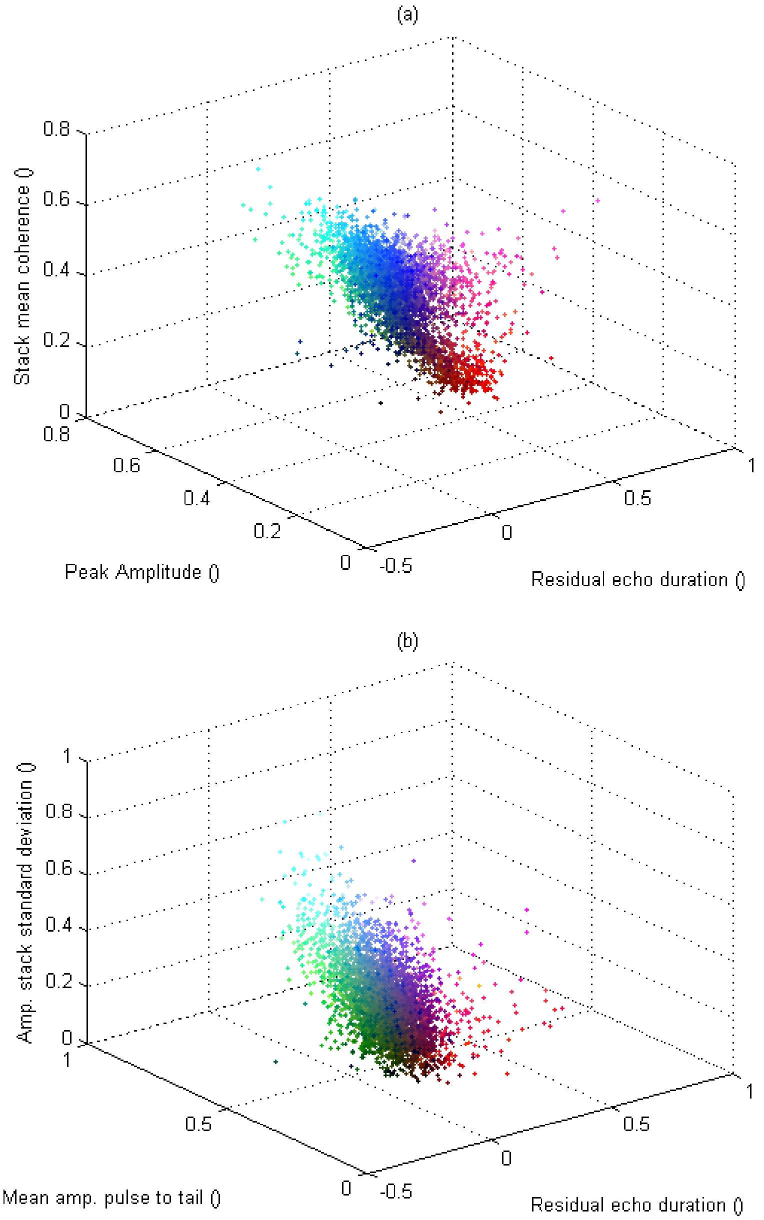

Figure 7-20. Coloured three dimensional feature space plots for the Airmar38 (*a*) and Simrad200 (*b*) testbed surveys. For clarity, only one out of every ten points is shown.

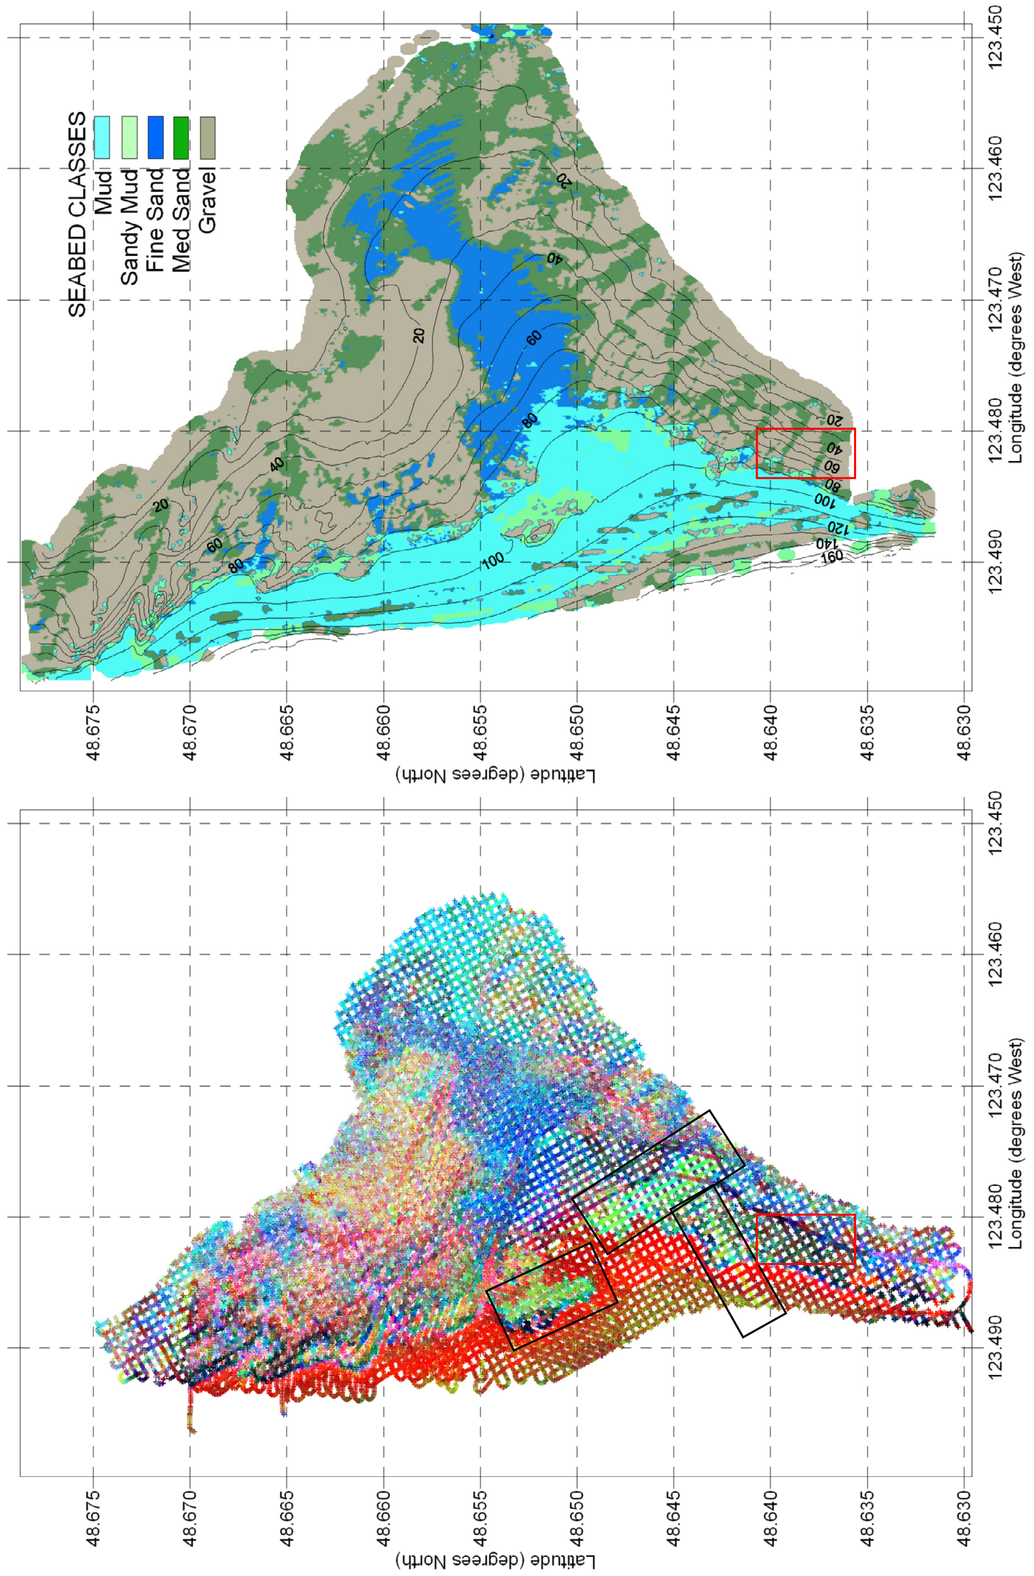

Figure 7-21. The OCPR three-feature map for the Airmar38 testbed survey (*lower*) compared to the multibeam seabed map (*upper*). The boxes denote areas of interest.

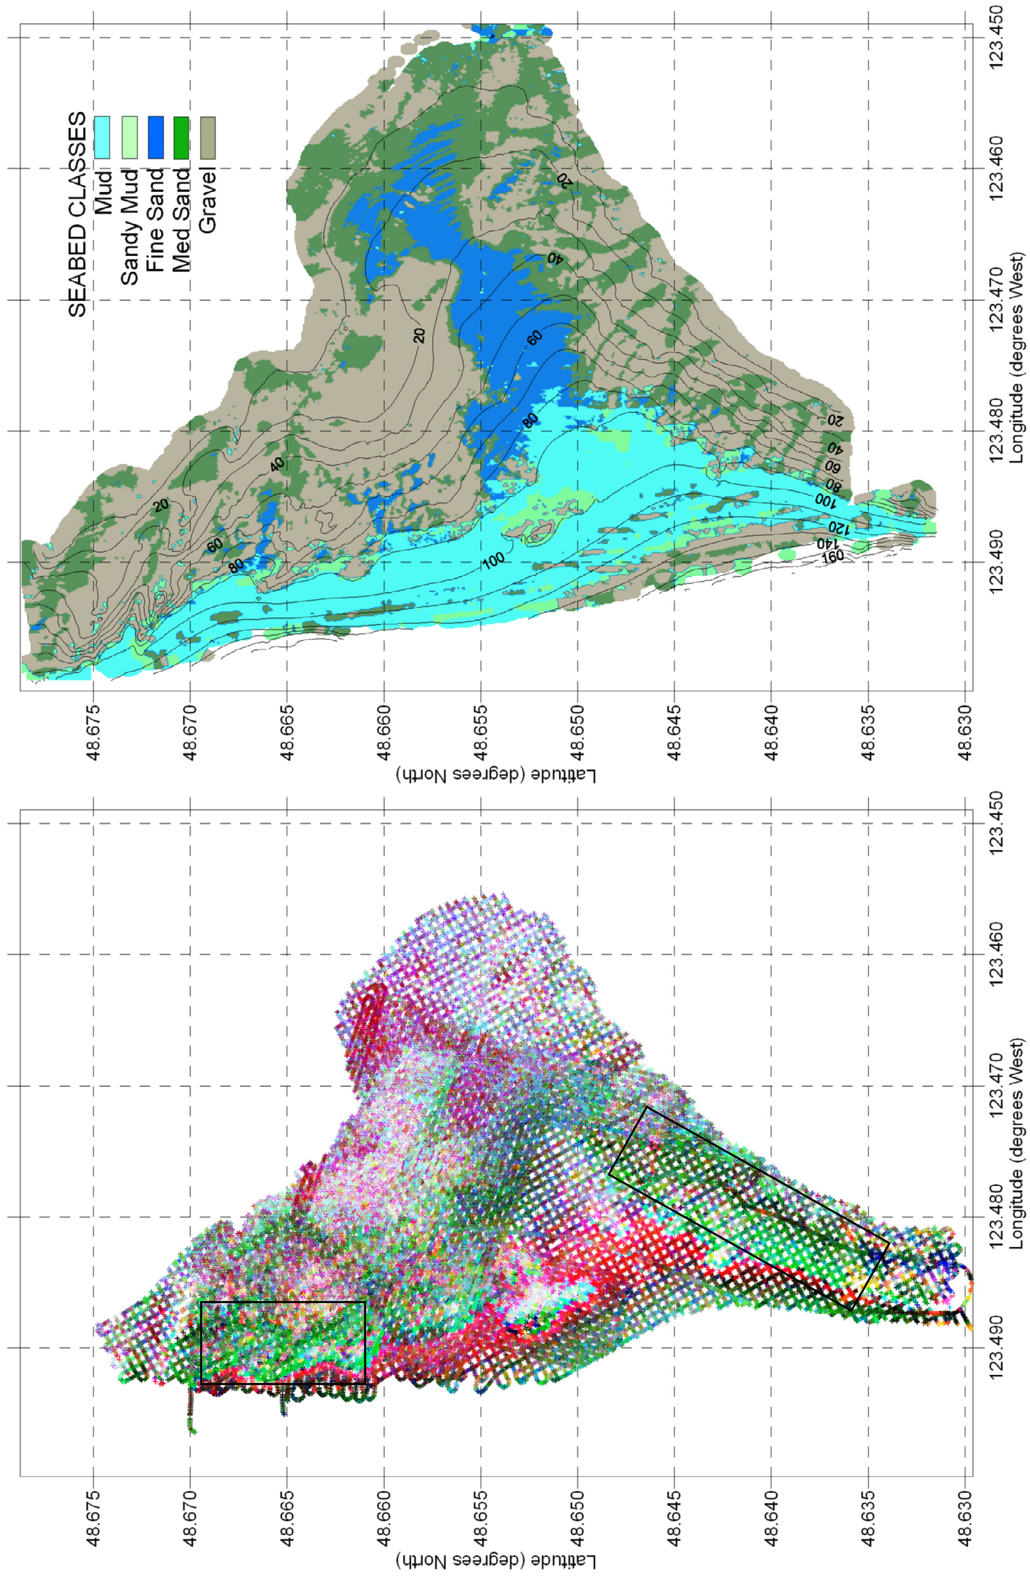

Figure 7-22. The OCPR three-feature map for the Simrad200 testbed survey (*lower*) compared to the multibeam seabed map (*upper*). The boxes denoted areas of interest.

The correspondence of the three-feature OCPR maps to the multibeam seabed map is clear, especially for the Airmar38 result. The Airmar38 map shows all the little patches of gravel, the major seabed boundaries and even the boundary between the medium and fine sand, which has not been seen in any result to this point. The black boxes on the Airmar38 map (Figure 7-21) show areas where a gain setting artefact (possibly due to operator error) boosts the peak amplitude, turning the map green. The red box shows an area of high slope that biases the peak amplitude feature. The expected colours of the seabed were not observed here simply because the stack mean coherence increases with decreasing echo-to-echo variability. Updating the expected colours table (Table 7-21) with this information and extra classes, results in a new colour interpretation table (Table 7-23).

Table 7-23. Colour interpretation table for the Airmar38 three-feature classification of Figure 7-21.

Seabed Type:	Relative values of the three features: Residual echo duration (red)	peak ampltude (green)	Stack mean coherence (blue)	Map Colour:
Gravel	moderate to high	low to high	moderate	brown-purple
Medium sand	low	high	high	blue-green
Fine sand	low	moderate	high	blue
Mud	very high	low	low	red
Mud Crust	high	low-moderate	low	red-brown
Extreme slope	-	-	-	black
Amp. artefact	-	-	-	green

The three-feature OCPR map for the Simrad200 is harder to explain. There are two slope artefacts at work here. First, echoes from narrow-beam SBES in areas of high slope tend to be shorter in duration than they should as predicted by the EDM. In particular, there are echo duration residual outliers that are large, which compress the red colour scale, so everything is red except for slope areas. The amplitude stack standard deviation also tends to go to low values in areas of slope, so with low levels of blue and red, the slope areas are coloured green. Sands have high amplitude so red plus blue equals purple. Muds have a low stack mean relative deviation because the nature of the variability is of modest amplitude but high frequency, which is a different sort of variability than is seen gravel echoes. It would be preferable to replace the stack mean relative deviation with a better behaved measure of variability, but it is the best one available in this case. Even with this slope artefact, the boundaries of the non-sloping gravels, sands and muds are clear.

Although the colour scheme interpretation had to be adjusted to accommodate for artefacts and a reversed feature, the colours agreed with the physical understanding of echosounding. The Airmar38 result was particularly good in identifying the seabed classes. To see this more explicitly, the three plots of Figure 7-23 show the expected relationships between the features and the seabed types. In this figure, the physical interpretations of the features are used to identify seabed types, which can be related to their OCPR colours or the seabed classes shown in the next section.

The analysis done here could be applied to other survey areas which have no ground-truth information. The colours and distributions of the echoes within the three-dimensional feature-space will indicate the seabed type well enough so that seabed type labels could be applied to the regular seabed class maps. For example, if the multibeam map of Figure 7-21 had no seabed type labels, the adjacent three-feature map could be used to ground-truth it. Taking this a step further, the colour interpretation table, Table 7-23, for the Airmar38 could be applied as a standard catalogue for supervised classification in other areas if the colour coding and features definitions were maintained. It is important to note that this result is only possible after careful removal of non-seabed influences on the features, a procedure that was established in chapters 5 through 7.

7.8.2 Three-feature Classification Results

For quantitative analysis, discrete classes are much easier to handle than a continuum of colours, especially given that all the ground-truth information is discrete. Classification of the three-feature data set is straightforward as the features replace the principal components (i.e. the Q-values) that are input to ACE for cluster analysis. The resulting class maps are shown as Figure 7-24 and Figure 7-25 .

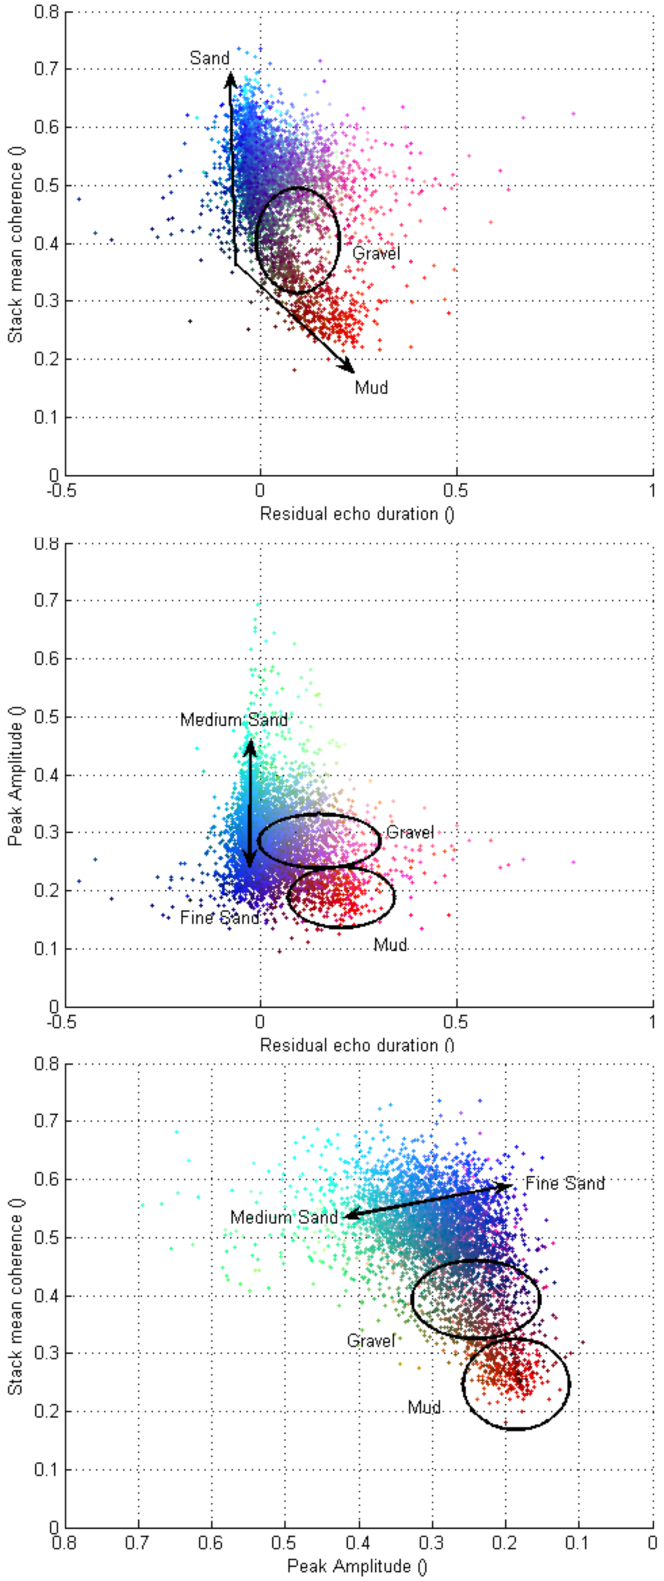

Figure 7-23. 2-D projections of the Airmar38 3-D feature space showing the relationships between the features and the seabed types based on theory, models and observations.

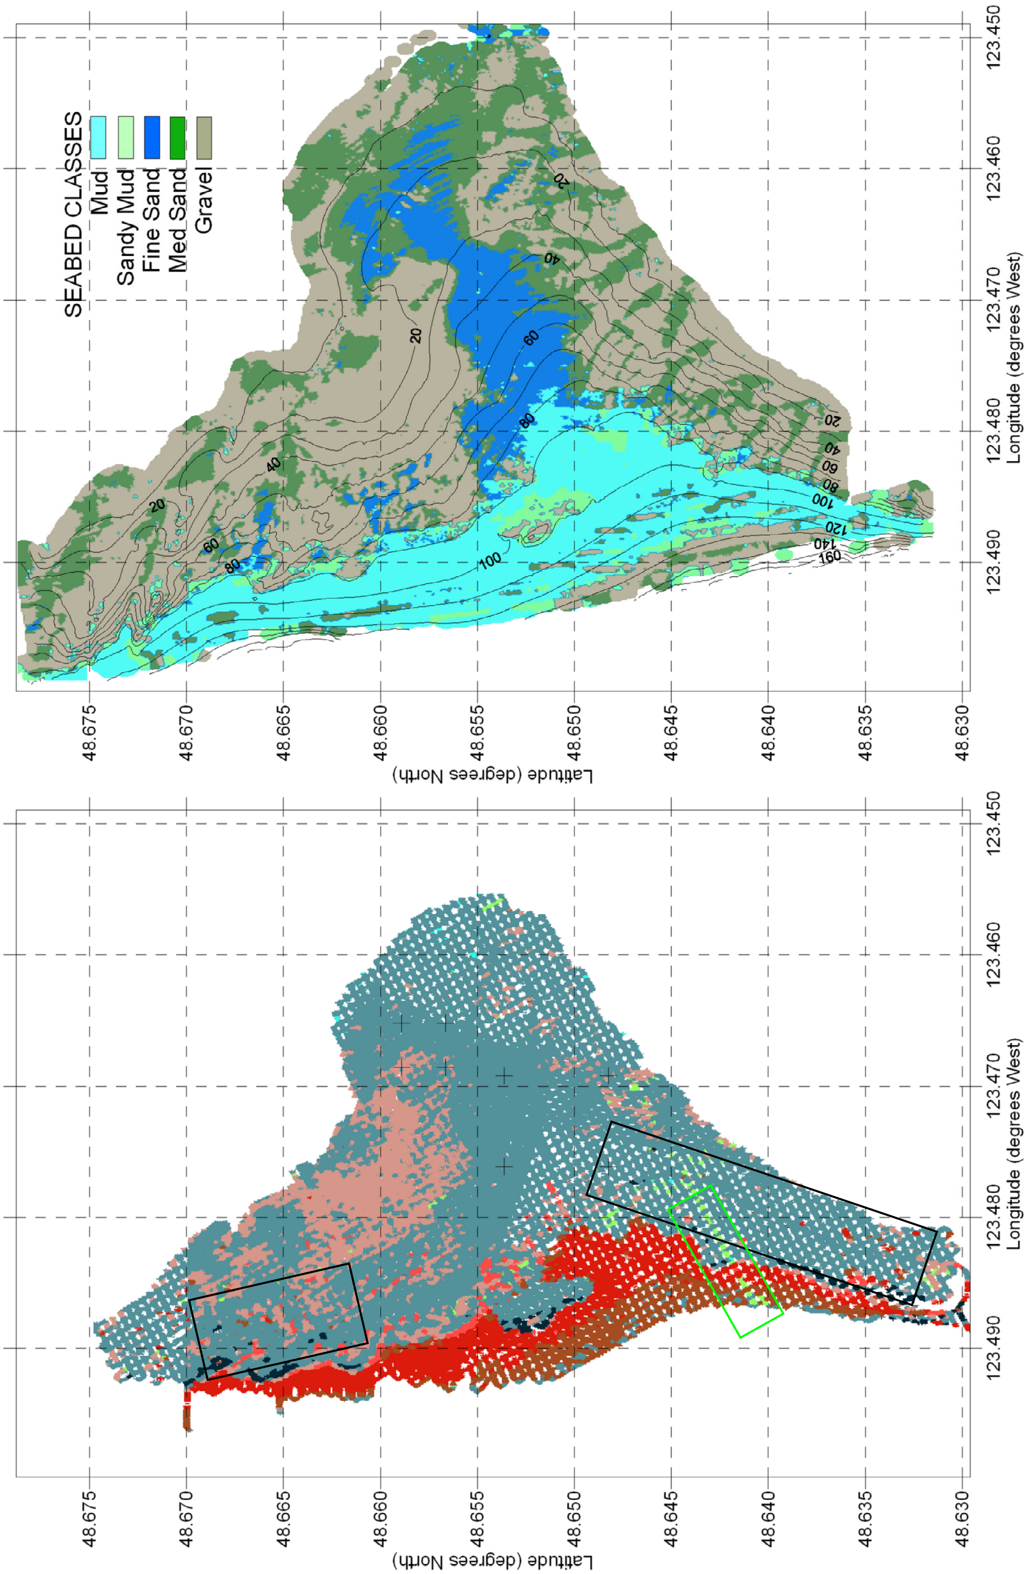

Figure 7-24. The three-feature classification map for the Airmar38 testbed survey (*lower*) to be compared to the multibeam seabed map (*upper*). The boxes are highlight areas to be discussed.

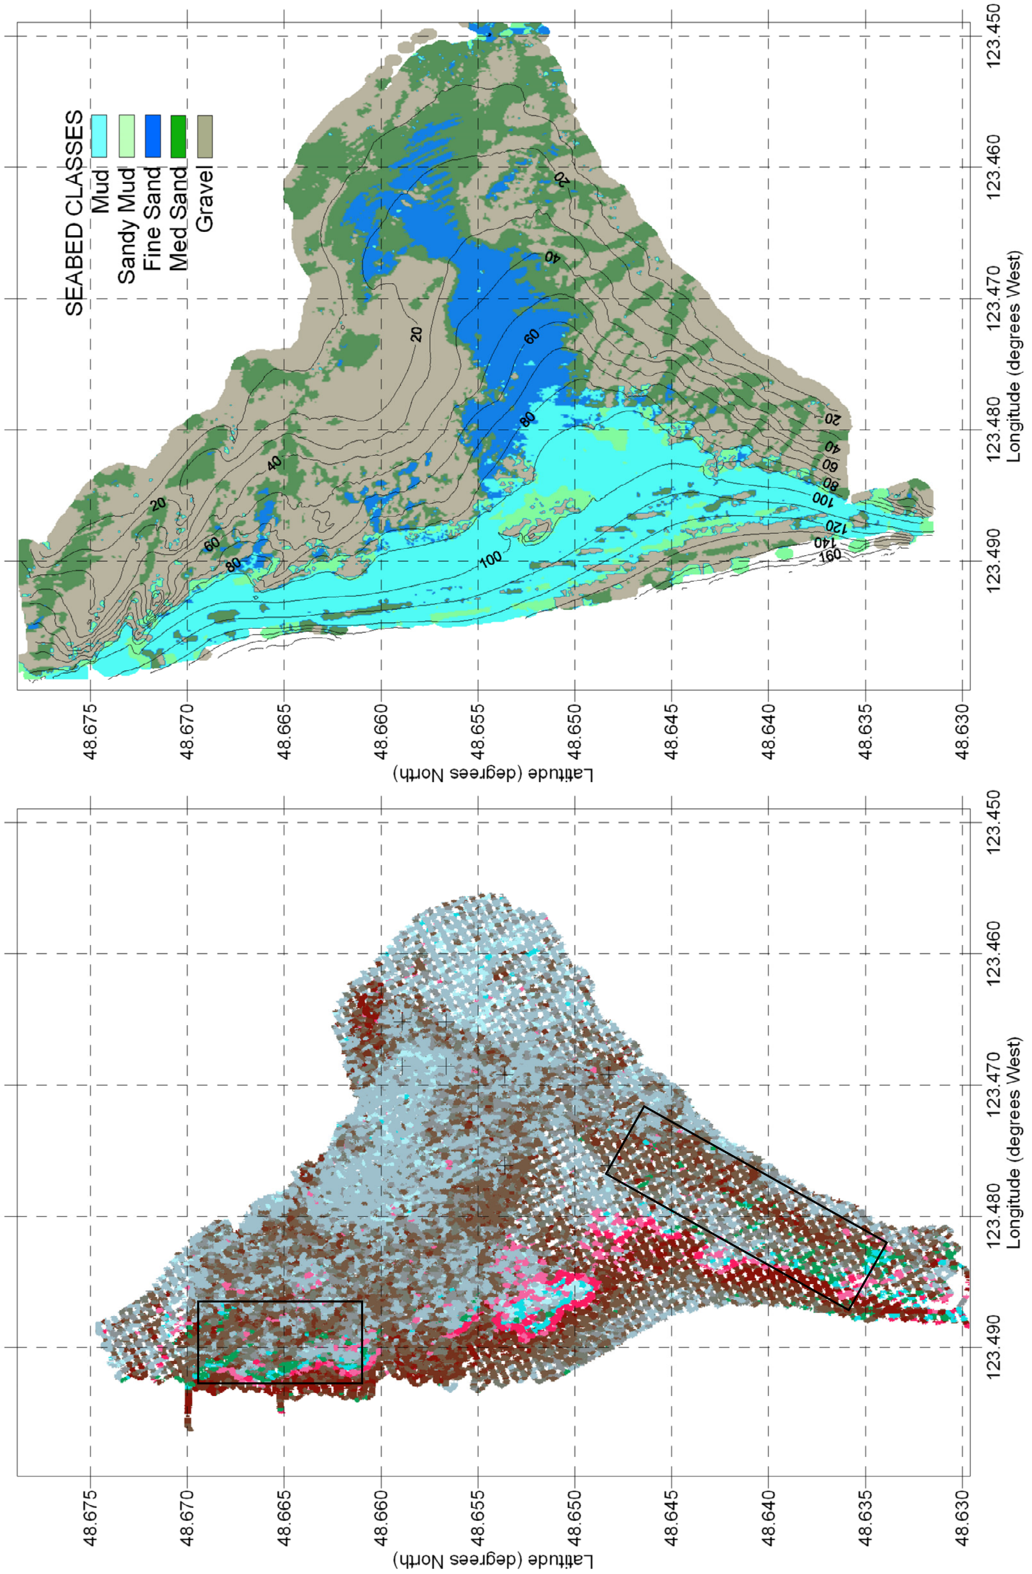

Figure 7-25. The three-feature classification map for the Simrad200 testbed survey (*lower*) to be compared to the multibeam seabed map (*upper*). The boxes are highlight areas to be discussed.

For all cluster analysis results, the class colours are determined by the location of the cluster centres. One can see how these colours relate to seabed type information in Figure 7-23 and Table 7-23. In the Airmar38 results, cluster analysis groups the sands and sloping gravels (shown in the black boxes) together. The Simrad200 classification does distinguish the muds from the combined sloping gravel and sand class, although the class colours are similar.

The quantitative measures of classification performance are compared to the baseline result in Table 7-24. For the Airmar38 survey, only the results with new features are better as indicated by AMI. The AMI for the Simrad200 three feature classification results is the best of any results with the Simrad200. These results further confirm the concept that echoes vary primarily in three dimensions.

Table 7-24. AMI, Kappa and OA statistics comparing the baseline classification result (TNORM) to the three-feature classification result (SELw3feat).

TNORM	Airmar38	Simrad200
Best-match OA (%)	48	57
No. classes:	8	17
AMI (%)	15	13

SELw3feat	Airmar38	Simrad200
Best-match OA (%)	58	58
No. classes:	8	11
AMI (%)	24	19

Overall, three-feature seabed classification was found to be a feasible method for relative seabed characterization. It demonstrates the power of physics-based features for characterizing the seabed in relative terms, without ground-truth data. Future development may include real-time characterization when slope information is available (i.e. from marine charts). Three-feature classification may also be a means for real-time supervised classification. Three-feature relative characterization can also be combined with unsupervised SELwNF results as it can be used to identify the SELwNF classes that most closely resemble the three major seabed types (gravel, sand, mud). This is done by spatial comparison. Data from other classes is reassigned to the identified classes (or to an outlier class) by merge-QC. This most promising approach will be developed in a future paper. The implications of three-feature classification for characterization will be developed further into characterization methods in the next chapter.

Chapter 8. Post-Class Seabed Characterization

As documented in the introductory chapters, attempts to characterize the seabed with single-beam echosounders have not been entirely successful. Previous methods, as discussed in section 2.3, focused on characterization on the basis of individual echoes or individual geoacoustic parameters. However, such an approach is clearly insufficient as echoes are highly variable due to the strength of incoherent scatter and, more importantly, are complicated functions of many geoacoustic parameters. The solution proposed in this chapter makes use of several methods that are used together on ensembles of echoes. The only requirement for defining the ensembles is that they are of one homogenous seabed type as non-seabed influences are accounted for. Then, ensembles can be defined geographically (if ground-truth information is available), or they can be defined by acoustic classes. In this dissertation, the results will focus on using acoustic classes, which is consistent with the general philosophy of developing practical operational methods. Then, the aim of *post-class characterization* is to add information on seabed type and geoacoustic properties to each acoustic class. The homogeneity of the classes is improved by allowing the user to remove classes that do not represent a specific area of seabed (i.e. outlier classes), by removing classes with few members, and by filtering out class members by the likelihood of them being assigned to that class (as measured by Mahalanobis distance). The idea of identifying acoustic classes was initially presented in the previous section with three-feature classification; it will be extended here. The methods of this chapter are used together to apply seabed type labels to classes based on estimated values of RMS roughness height and attenuation.

The characterization methods applicable to non-calibrated SBES have different regions of applicability, which is why a combination of methods is used. The combined scheme is shown in Figure 8-1. The two main methods are mean envelope inversion and attenuation by echo length (ABEL). All methods except ABEL measure RMS roughness height. ABEL measures the time corresponding to the 10 dB penetration depth which is then related to seabed type and attenuation via the extinction model (part of Equation 5-3). The effects of beam-widening and macro-roughness can also be observed by ABEL, making this method important, as will be shown in the next paragraph, for distinguishing long echoes due to high roughness from long echoes due to low attenuation.

The APL-UW seabed parameters, Table C-2, are used to find the seabed type that best matches the RMS roughness height and penetration time measured by the characterization methods. Figure 8-2 shows a plot of RMS roughness height and penetration time as a function of mean grain size. When the curves shown in Figure 8-2 are flat, matches between mean grain size and penetration time or roughness height will not be unique. Fortunately, the penetration time and roughness are complementary, with the penetration time being unique for sediments greater than 4.1 Φ, and the roughness height being unique otherwise.

All characterization methods are implemented to run autonomously after classification is complete using the classes as input data. Three additional methods, frequency-shift, the coherent stack and amplitude variability analysis, have also been implemented to the proof of concept stage, and these will be discussed briefly. These methods are presented as a demonstration of their feasibility only. The main source of data here is the Airmar38 SELwNF-processed echo time series and classes. This data set is suitable for demonstrating the methodology as Patricia Bay is a well-known environment that represents a broad range of seabed geoacoustic and geotechnical parameters (i.e. seabed types, depths and slopes). For this reason, it is reasonable to expect that if the methods work here, they will work in most locations. In the last section of this chapter, the results of the methods are combined to assign seabed types and other information to the Airmar38 classes.

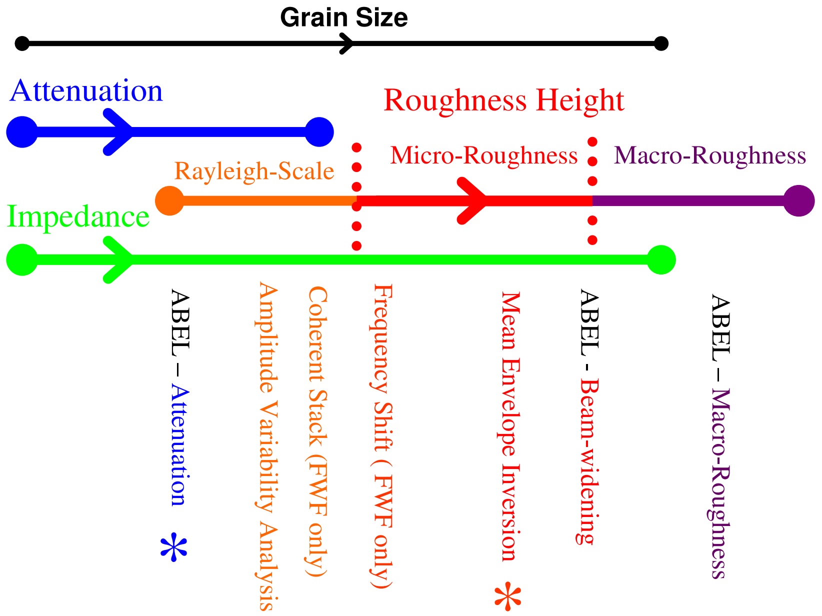

Figure 8-1. A depiction of the bounds of applicability of the characterization methods relative to the major seabed parameters which are empirical functions of grain size. All parameter values increase from left to right. The methods emphasized with asterisks are the main methods of post-class characterization.

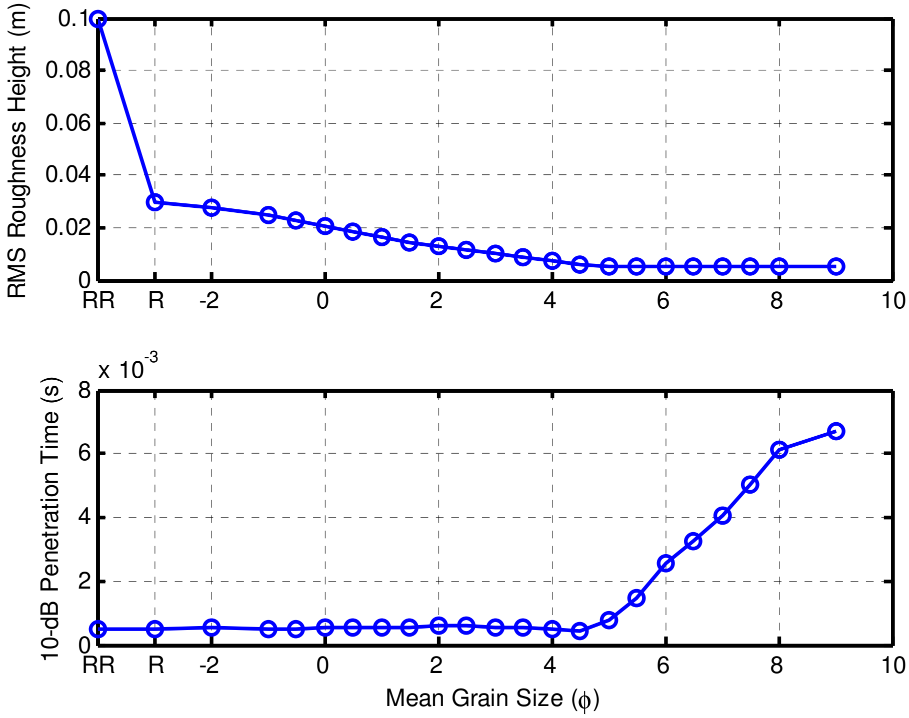

Figure 8-2. Geoacoustic parameters to be measured by the characterization methods of this chapter, plotted as functions of mean grain size calculated for the Airmar38 SBES. RMS roughness height (*upper*) is based on Table C-2, and the penetration time is based on the extinction model (see Equation 5-6 or Equation 8-5 later in this chapter).

8.1 Proof-of-Concept Methods

Proof-of-concept methods are not fully operational because either more development is required or the results obtained were not satisfactory. The frequency shift and coherent stack methods have not been further developed because more FWF data is required for testing purposes. These two methods are presented in [Biffard et al., 2007]. Amplitude variability analysis is a feasible method that suffers from a lack of information on a geoacoustic parameter, the correlation length, and may be improved easily with new information on that parameter in the near future. This method also includes pertinent theoretical background for the methods to be presented in sections 8.2 and 8.3. For these reasons, amplitude variability analysis appears below.

8.1.1 Amplitude Variability Analysis

This method is an adaptation of Stanton's method for measuring seabed roughness from distributions of echo peak amplitude – originally published in [Stanton, 1984] and further developed in [Stanton and Clay, 1986]. The method consists of two steps: fitting the Rice probability density function (PDF) to a histogram of echo peak amplitudes to measure the ratio of coherent to incoherent amplitude, γ, and then the RMS roughness height is calculated by applying scattering theory. There are two ways of relating γ to geoacoustic parameters: Stanton developed an expression based on KA theory, Equation 8-1, and this author adapted the Rayleigh reflection coefficient, Equation 8-2.

$$h_{RMS} = \sqrt{\frac{3\pi}{\gamma k^4 \theta^2 \left(\lambda_{0x}\lambda_{0y}\right)}} \tag{8-1}$$

$$h_{RMS} = \sqrt{\frac{\ln\left(\dfrac{\gamma}{1+\gamma}\right)}{-4k^2}} \tag{8-2}$$

In Equation 8-1, θ is the full beamwidth in radians, and $\lambda_{0x}, \lambda_{0y}$ are the surface correlation lengths taken in orthogonal directions. Stanton's model (Equation 8-1) makes use of a correlation function that can accommodate saw-tooth ripples, while Equation 8-2 assumes Gaussian roughness, which is applicable as long as the seabed surface does not have ripples or organized patterns of roughness. However, Equation 8-2 can be applied without assuming a value for the correlation length.

Through KA theory, Stanton established that the Rice PDF parameter γ is the ratio of coherent to incoherent energy [Stanton, 1984]. Stanton applied Equation 8-1 and an empirical relationship between the roughness height and the correlation lengths for seabeds with ripple fields (now referred to as sand waves). The empirical relationship is based on data published in [Heezen and Hollister, 1971]. Here, Equation 8-2 is used because there are no widely–published, reliable estimates for correlation lengths for non-rippled seabeds. The two models are roughly equivalent for a correlation of length of 7.5 cm for a 50 kHz centre frequency and a beamwidth of 20°.

For each seabed class, non-smoothed peak amplitudes are extracted for every echo in that class. Histograms of the raw, non-stacked, trim-TVG corrected peak amplitudes from each echo are fitted with the Rice PDF using a moment-matching approach[59] [Talukdar and Lawing, 1991]. The fit is scale invariant, so the raw values of the amplitudes are used. The Rice PDF changes from a Rayleigh distribution (for $\gamma \leq 0.3$) to a normal distribution (for $\gamma > 10$) [Stanton and Clay, 1986]. A Gaussian-shaped peak amplitude distribution is not expected as γ should range from zero to five for the range of frequencies and roughnesses here (based on Equation 8-2). The analysis is demonstrated on the three major Airmar38 SELwNF classes, which are: class 1 (sand), 3 (mud) and 6 (gravel) – see Figure 7-17 for the class map. These three classes comprise 76% of the survey data. The histograms and Rice PDF fits are shown in Figure 8-3. The resulting RMS roughness heights are shown below:

Table 8-1. RMS roughness height estimated by the amplitude variability analysis (AVA) method for the three major classes of the Airmar38 SELwNF processed testbed survey.

Class ▼	AVA Gamma - γ	AVA Estimated RMS Roughness Height (cm)	Expected RMS Roughness Height (cm)
1 - Sand	5.35	0.13	1.1 - 1.3
3 - Mud	1.13	0.25	0.5
6 - Gravel	0.91	0.27	2.5 - 3.0

---

[59] The raw peak amplitudes are required. For instance, running this analysis with the smoothed and stacked peak amplitudes (feature 180) results in reduced variability that leads to histograms that are all Gaussian.

The measured and expected values are not consistent, but at least they are relatively consistent in that the gravel exhibited more roughness than the sand. The apparent roughness of the mud is related to the variability caused by volume scattering – this method is not applicable to volume scattering and hence silt/clay/mud seabeds. The ABEL method will be used later to distinguish volume scattering from surface scattering. The maximum roughness measurable is determined by values of γ. For $\gamma<1$, the fits begin to lose sensitivity and for $0<\gamma<0.3$, the Rice PDF hardly changes. For $\gamma=0.3$, the roughness would be 0.38 cm as calculated by Equation 8-2 (at 38 kHz) – this is taken as the maximum roughness limit when using Equation 8-2. The Rayleigh criterion occurs at $\gamma=1.25$, which is of order of the γ values found here for the gravel and mud echoes[60]. These statements explain the low values – the sand and gravel seabeds have roughness higher than the upper sensitivity limit when using Equation 8.2. The gravel class also likely exhibits significant non-specular coherent components that are not accounted for. The results may be improved with a more sophisticated model.

Equation 8-1 requires a correlation length estimate to relate γ to roughness height. A value of 0.8 cm for the correlation length was found by fitting Equation 8-1 to the expected roughness height and the γ values from Table 8-1. This single value for the correlation length results in RMS roughness heights that are consistent for *both* the sand and the gravel seabeds. This assumed value for the correlation length is midway between the values for gravel (0.2 cm) and pebbles (1.2 cm) listed in [Heald, 2000]. However, the three correlation length values for sand, gravel or pebbles listed in [Heald, 2000] do not produce consistent results. These three values were measured from plaster molds of test tank sediments, not field data [Al-Hamadani, 1984]. However, these are the only three measurements available besides the sediment wave data of [Heezen and Hollister, 1971]. More testing with a fixed-value for the correlation length could be done to support the case for using Equation 8-1. This testing would involve empirically calibrating Stanton's model, Equation 8-1, to the measured γ values with known roughness heights, but in order to do this, more data are required. This type of calibration is done in the next

---

[60] A plot relating roughness height and γ can be found in [Biffard et al., 2007].

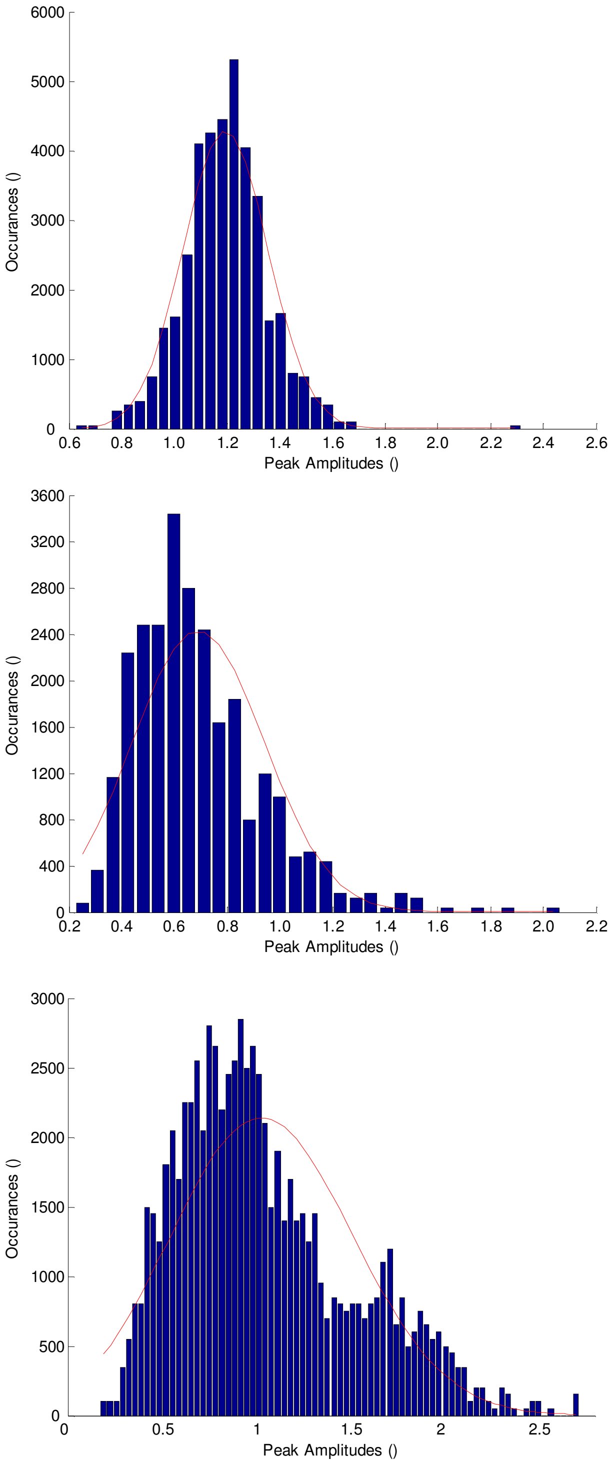

Figure 8-3. Echo peak amplitude histograms for the three major Airmar38 SELwNF classes as follows: (*top*) sand – class 1, (*middle*) mud – class 3, (*bottom*) gravel – class 6. The Rice PDF fits are the red curves. Peak amplitudes are in consistent relative units resulting from trim-TVG processing.

section, using the BORIS and AEM models but the values found are not applicable here for reasons that will be discussed. Without field measurements of correlation lengths or empirical calibration, Equation 8-1 cannot be used, so Equation 8-2 will continue to be used although the results were not satisfactory. Further testing using Equation 8-2 may also aid in resolving the issues with this method. The results using Equation 8-2 would likely improve with lower frequency SBES.

One aspect of this method that did turn out completely as expected is the relative amplitudes: the mud amplitudes were much lower than the sand, and the sand and gravel amplitudes were about equal with the gravel having a much larger spread of values. Because of this methods' potential, it is included when the results of the methods are combined in section 8.4 – not to demonstrate results, but to demonstrate the procedure of integrating such information.

8.2 Mean Envelope Inversion

Mean envelope inversion (MEI) estimates seabed and echosounder parameters by inverting the AEM forward model to mean echo envelopes for each seabed class. MEI estimates the mean RMS roughness for each class. These results can then be matched to a seabed type using Table 8-1 as described in section 8.4.

8.2.1 Description of the Method

The mean envelope inversion (MEI) method for post-class characterization is described as follows. The first step is to pre-process the data to find the mean echo envelope time series for each seabed class. A mean envelope is required as the forward model (the AEM, see section 5.2) is derived from mean seabed statistics predicting a mean echo intensity time series (which is easily converted to amplitude). A stable mean requires approximately 30 echoes [Sternlicht, 1999]. Echo envelopes of each class are aligned by their bottom picks and variable length windows are resampled to match the length of the echo at the mean depth of the class. This method will upsample to the longest echo if there is any danger of the Nyquist sampling frequency being violated. Some echoes are rejected – the top and bottom 5% by depth, in addition to rejecting echoes that have large Mahalanobis distances from the cluster centre (rejects an additional 35% of echoes). The stacks, or sample-wise means, of the interpolated echo envelopes are the data for the inversion, while the standard deviations for each sample are applied as uncorrelated data errors.

After calculation of class mean envelopes and their standard deviations, inversion begins with **a**daptive **s**implex **s**imulated **a**nnealing (ASSA) optimization [Dosso et al., 2001]. Optimization requires model parameter bounds, data and data error(s) as input. The forward model, the AEM, requires five parameters – these are listed in Table 8-2 along with the parameter bounds.

Table 8-2. Model parameters and bounds for mean envelope inversion. The first three bounds are relative to the parameter in the far right column, while the last two bounds consist of two setsas follows: for MEI / for BORIS testing.

Model Parameter	Parameter Bounds	Based on:
Depth	greater of ±1% or ± num. samples before pick	Mean depth
Beamwidth	-40% to +50% up to a maximum of 60 degrees	Eff. beamwidth
Pulse Duration	-0.1 to +0.3 ms	Given value
Roughness Ratio	3.5 to 150 / 0 to 500	-
Scale Factor / Ref. Coeff.	0.00005 to 10 / 0 to 1	-

ASSA optimization efficiently finds the model parameters that minimize the model to data misfit (otherwise known as the objective function). In this case, with uncorrelated data errors, σ_t, the misfit is given by:

$$E(\mathbf{m}) = \sum_t \frac{\left(p_t^{model}(\mathbf{m}) - p_t^{obs}\right)^2}{2\sigma_t^2} \tag{8-3}$$

where p_t^{model} and p_t^{obs} are the model data mean echo envelope time series, **m** is the model parameter vector and the sum is over all time samples, represented by the index *t*. Optimization only provides the model parameters that best fit the data. However, the quality of the fit is largely unknown, as optimization does not provide information on the uncertainties and marginal probability distributions of the model parameter estimates.

The second step in the process is to gather that information by searching the model space with Metropolis-Hastings sampling. This is a Markov-chain Monte Carlo method. It is similar to that described in [Dosso, 2001] and is combined with ASSA optimization in a later paper: [Dosso and Wilmut, 2006]. The version of Metropolis-Hastings sampling applied here samples a rotated model space, which greatly improves efficiency. The rotation is found by a process which is essentially PCA of the model parameter covariance matrix. The model parameter covariance matrix is initially found by calculating linear gradients, which leads to an initial rotated model space. This rotated model space is then sampled to find an accurate covariance matrix used for the actual sampling. In all, there are two initialization steps, followed by a cooldown, following by the actual sampling.

8.2.2 Calibration and Testing with BORIS Data

The testing of MEI with BORIS data has two purposes: to explore the performance of the method using data of known characteristics, and to create a means to relate the AEM roughness ratio parameter to seabed roughness heights. As discussed in section 5.2 and earlier in this chapter, the correlation lengths of seabed surfaces are not known. A relationship between RMS roughness height and correlation length is known for sand/gravel waves [Stanton, 1984; Stanton and Clay, 1986], and correlation functions are known for the air-sea interface [Medwin and Clay, 1998]. Neither is applicable here, so instead an empirical relationship is sought so that roughness ratios found by MEI can be equated to RMS roughness heights for characterization. This is essentially a calibration of the AEM to match the BORIS model and should not be used to infer information about correlation length.

One would expect that as roughness increases, the slope of the scattering facets (i.e. the RMS micro-slope of the seabed) would increase. Then, Equation 4-13 implies that the ratio of correlation length to roughness height should decrease as roughness increases. This is what was found in MEI testing with BORIS data. Mean envelopes were generated from simulated surveys of four seabeds (250 echoes each) with the simulated

Airmar38 SBES. The seabed types used span the range of roughness of most sediments (as listed in the APL-UW sediment parameters table, Table C-2), but exclude seabeds that tend to generate volume backscatter dominated echoes, (i.e. SiltyClay). MEI was run on the ensemble of echoes from each seabed, and the results for the roughness ratio are summarized in Table 8-3. The resulting values of the correlation lengths in Table 8-3 were calculated from the RMS roughness height multiplied by the MAP roughness ratio. The RMS roughness height is a known parameter as it was used as input to the BORIS model. The RMS micro-slope was calculated from Equation 4-13. An example set of marginal probability distributions is shown in Figure 8-4.

Table 8-3. The calculation of calibration correlation lengths from *M*aximum *a p*riori (MAP) roughness ratio estimates from MEI on BORIS simulated data. Note that the correlation and RMS micro-slope values listed are for calibration of the AEM only and need not correspond to other estimates.

BORIS Seabed Name	RMS Roughness Height (m)	MAP Roughness Ratio	Rough. Ratio 95% Credibility	Correlation Length (m)	RMS Micro-slope (degrees)
SiltySand	0.007	41.37	5.34	0.290	1.96
SandAG	0.015	16.01	2.64	0.240	5.06
Sand	0.020	7.57	1.29	0.151	10.71
Gravel	0.028	4.55	1.03	0.128	17.79

Mean envelope inversions of BORIS data found depth, beamwidth, pulse duration and reflection coefficient estimates that agreed with the input parameters to BORIS. Correlations and two-dimensional marginal probability distributions (not shown) were observed to show relationships between the parameters that were consistent with those established in chapters 4 and 5. The reflection coefficient is used here because the initial intensity of the transmit pulse, I_0, is known. Equation 8-4 is a linear empirical relationship between RMS roughness height and correlation length for the AEM model, as found by ordinary linear regression of the data in Table 8-3:

$$\lambda_0 = 0.347 - 8.258 h_{RMS} \tag{8-4}$$

The relationship between roughness ratio and roughness height is not linear, but it can be inferred from Equation 8.4:

$$h_{RMS} = \frac{0.347}{h_{ratio} + 8.258} \quad \text{where } h_{ratio} = {}^{\lambda_0}\!/_{h_{RMS}} \tag{8-5}$$

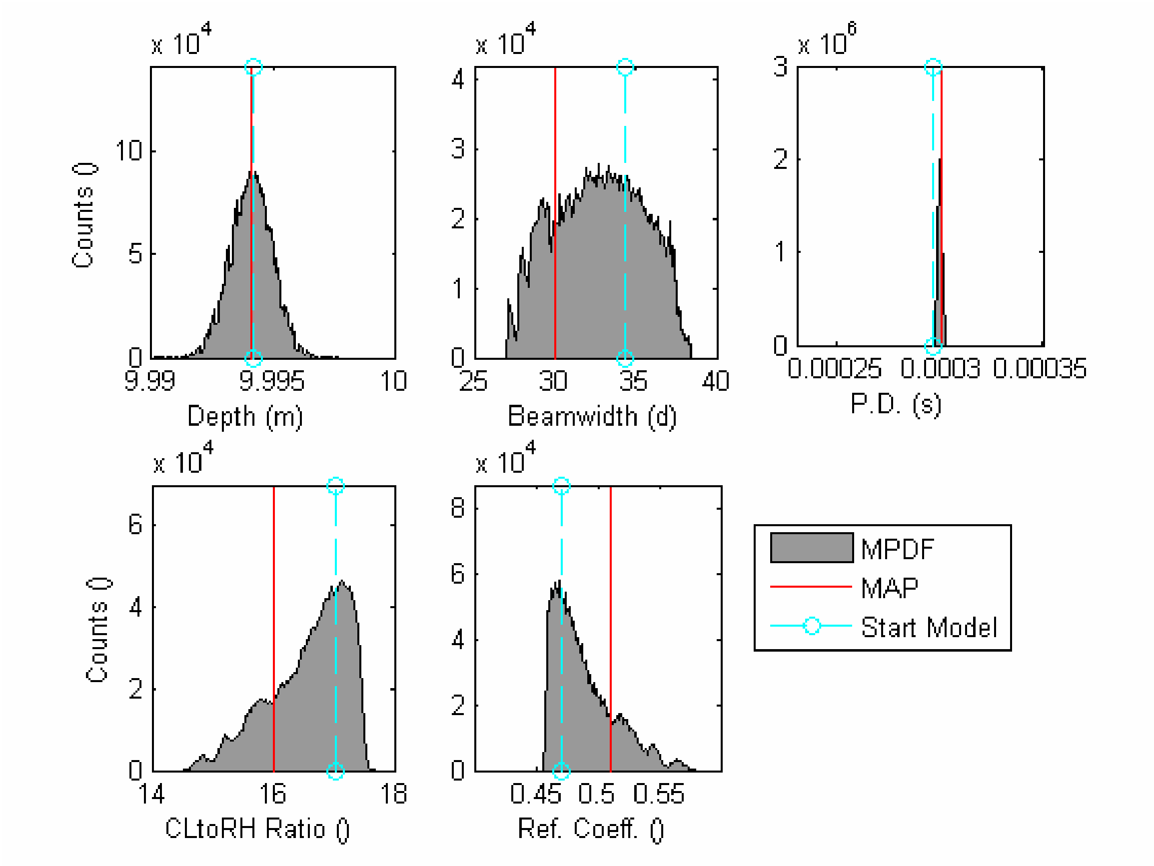

Figure 8-4. Marginal probability distributions of the AEM model parameters as found by Metropolis-Hastings sampling (part of MEI). The data is the mean of 250 echoes from a BORIS-simulated Aimar38 SBES survey on the SandAG seabed at 10 m water depth. Other BORIS parameters: -4.34 dB beamwidth: 22.8° (the effective beamwidth found to be 33° in chapter 7), pulse duration 0.0003 s, plane wave reflection coefficient 0.556. The sampling start model was found using ASSA optimization, while the MAP model is the lowest misfit model from Metropolis-Hastings sampling. The MAP does not necessarily coincide with the peak of the distributions. CLtoRH Ratio is an older label for the roughness ratio. The bounds of the optimization and sampling routines are listed in Table 8-2.

This relationship is to be applied to convert roughness ratios from the MEI to values of RMS roughness heights that can be matched to the seabed types via the APL-UW table (Table C-2). Equation 8-4 and Equation 8-5 are definitely not to be taken as a general result, but are only useful as a calibration of the AEM based on BORIS data. Further, these results are not consistent with the correlation lengths used by Heald, as discussed in reference to Equation 8-1 of the AVA method (section 8.1). The uncertainty of the roughness ratio from MEI can be estimated by the 95% credibility intervals. It was fairly consistent – the average percent uncertainty was 17.2%. The uncertainty in Equation 8-4 /

Equation 8-5 is taken to be 20%, which will be propagated into the RMS roughness height estimates for MEI in the next sub-section.

8.2.3 Mean Envelope Inversion Results

The mean envelopes from mean envelope inversions for the three major classes in the Airmar38 SELwNF classification result are shown in Figure 8-5. The selected classes were sand (class 1); mud (class 3) and gravel (class 6) as indicated in the figure. These mean envelopes are very revealing about the seabed just on the basis of their relative shapes. The gravel and the sand mean echoes reach their maximum amplitudes approximately one pulse duration after onset. This is what one would expect, as peak amplitude usually occurs when the ensonified area (or volume) of the seabed is at its maximum. In this case, the ensonification regime is annular so the ensonified area reaches a maximum at one pulse duration after onset. The exception may be when the roughness is extreme, i.e. macro-roughness. For echoes dominated by volume scattering, the maximum ensonified volume occurs sometime after one pulse duration after onset. This distinguishes the mud mean echo from the sand and gravel mean echoes, while the sand and gravel mean echoes can be distinguished from each other by the effect of beam-widening increasing the duration of the gravel mean echo. The ensonified volume is assumed to be a function of the penetration and sediment sound speed, and only approximately determines when maximum volume scattering amplitude occurs. It would usually occur after maximum surface ensonification, as can be inferred in this case.

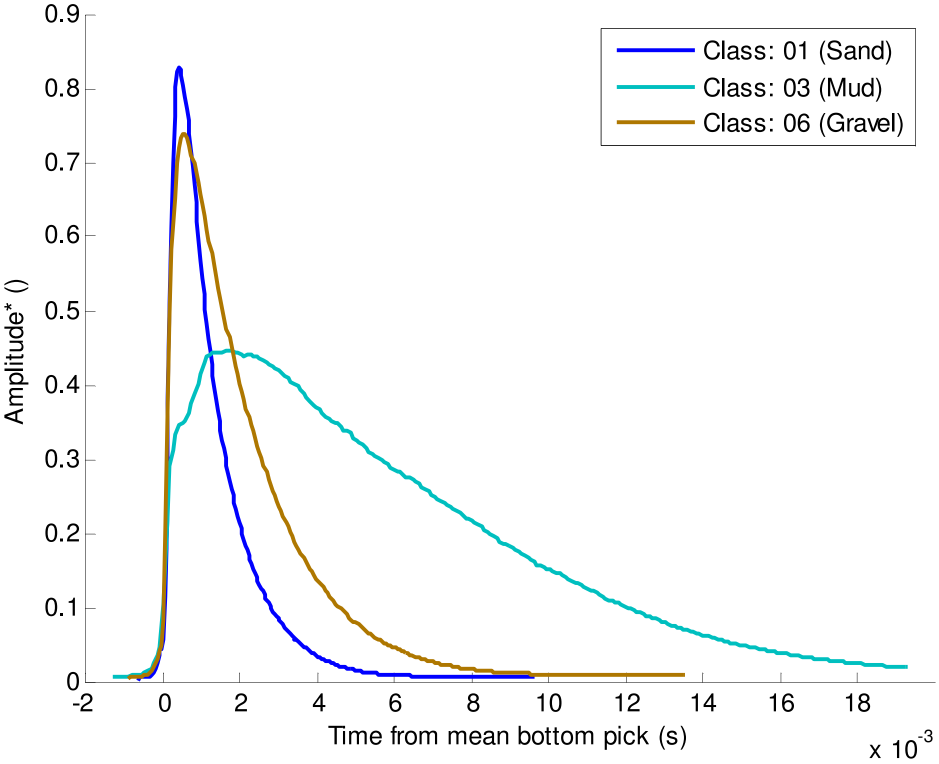

Figure 8-5. Mean envelopes for the three main classes of the Aimar38 SELwNF classification result. The amplitudes result from the sum of echoes individually normalized to an amplitude of unity[61].

The MEI results for the mean echo envelopes shown in Figure 8-5 are largely as expected. A comparison of the resulting best-fit AEM echoes and the mean echo envelope for the gravel class is shown in Figure 8-6. The marginal distributions for the same class are shown in Figure 8-7. For all three classes, the depth, beamwidth and pulse duration found by MEI are generally within their credibility limits of the known values. There are some notable exceptions. The beamwidth for the gravel class is wider than the effective beamwidth of 33°, as expected due to high roughness. The pulse duration and beamwidth for the mud class are at the upper bounds set in Table 8-2, as the AEM forward model attempts to fit an echo that is much longer than that generated at the upper bound of roughness (the AEM simulates surface backscattering only). In a survey where there is no information on what the classes are, this type of result may be a clue that

---

[61] Because of this, peak amplitudes in Figure 8-5 will respond positively to echo-to-echo coherency, which is related to the relative strength fo the coherent component. This is an interesting effect which could be used for seabed characterization.

either the seabed is fine sediment (silt, clay or mud), or that the roughness heights exceed the upper limit, which is 0.0284 m. This limit is based on the AEM minimum roughness ratio of 4 from section 5.2, converted to a roughness value using Equation 8-5). The results of the MEI method on the Airmar38 survey are shown in Table 8-4.

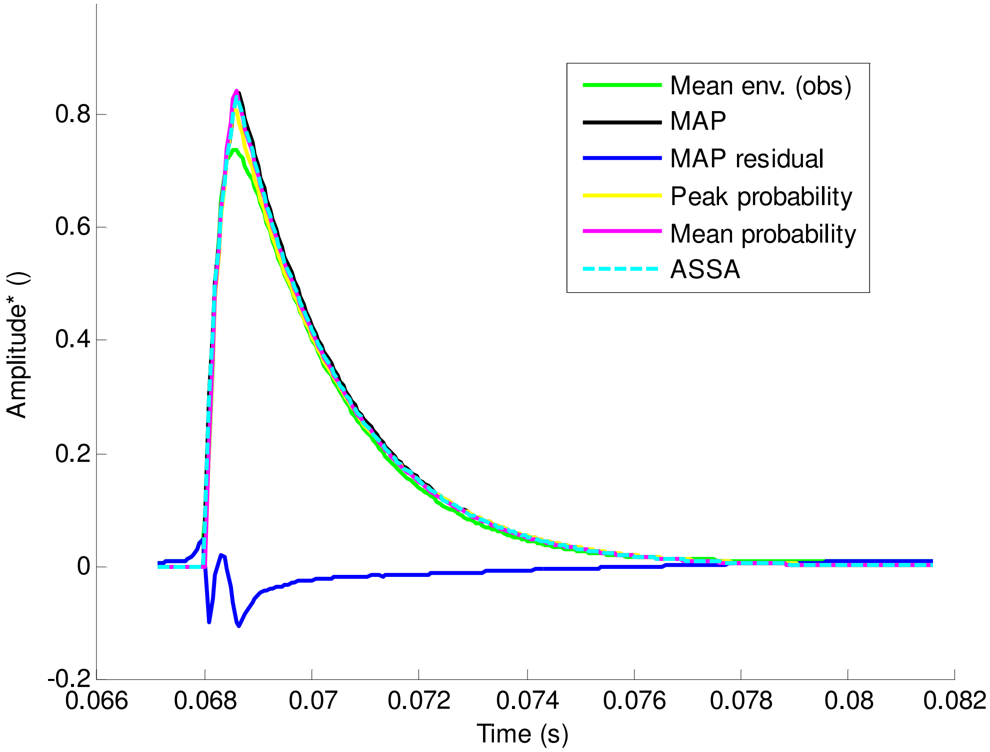

Figure 8-6. Mean envelope for class 6 (gravel) as seen in Figure 8-5, overlaid by best-fit AEM echoes. The MAP, peak and mean prob. model parameters can be seen in Figure 8-7.

Table 8-4. MEI class characterization results for the Airmar38 testbed survey. Note that the results for the mud class are discounted and colour yellow, as the AEM does not simulate volume scattering. The mud results (coloured grey) are only shown to demonstrate the values MEI will get with such seabeds.

Class Number	MEI estimated RMS Roughness Height (m)	Uncertainty in h_{RMS} (m)	Best-Match APL-UW Seabed Type	Corresponding Grain Size (Φ)	Ground-Truth Label	Experiment Site Mean Grain Size (Φ)	Expected RMS Roughness Height (m)
1	0.0111	0.0037	Fine Sand, Silty Sand	2.5	Sand	1.89	0.011-0.013
3	0.0292	0.0126	Rock	-	Mud	-	0.05
6	0.0243	0.0068	Sandy Gravel	-1	Gravel	-1.92	0.025-0.030

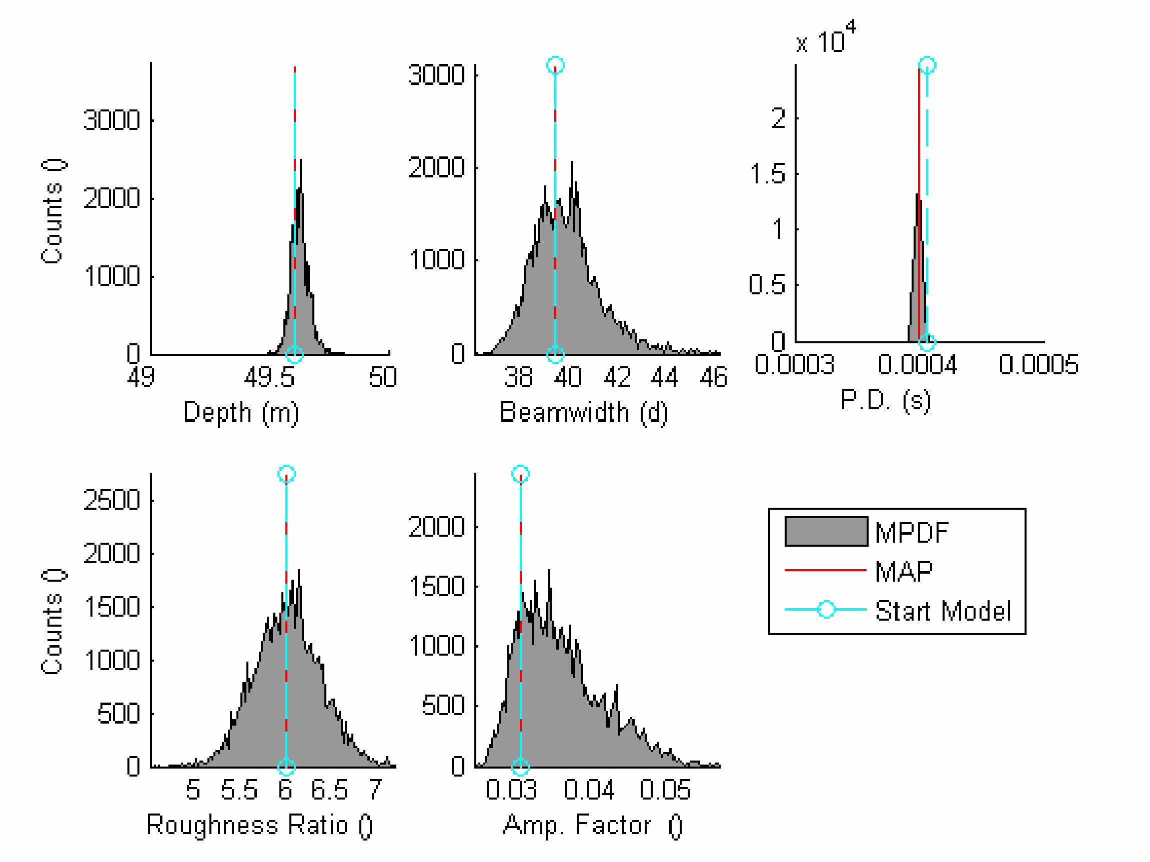

Figure 8-7. Marginal probability distributions for MEI on the mean envelope of class 6 (gravel). The bounds of the optimization and sampling routines are listed in Table 8-2.

8.2.4 Mean Envelope Inversion Conclusion

For the case of the Airmar38 full bay survey processed with the SELwNF classification method, the MEI method is able to identify seabeds from fine sands up to cobbles, with a maximum roughness of 0.0284 m. This is a higher upper bound than the characterization methods presented in the previous sections. This method does require calibration, although such a calibration may not be necessary if there was more information available about seabed correlation lengths. The method is insensitive to the ensonification regime depth effects, while moderate slopes (up to ½ the beamwidth) are handled well by SELw compensation. MEI does not characterize seabeds that produce volume scattering dominated echoes because the AEM does not include volume scattering. However, volume scattering dominated echoes can be identified as mentioned above, by their low mean amplitude seen in AVA, or by the next characterization method to be presented: ABEL. Further research in this area would be to test the method with different echosounders. Also beneficial here, and for the whole dissertation, would be in situ measurements of roughness height and correlation length by stereo photography or laser-line scan.

8.3 Attenuation by Echo Length (ABEL)

8.3.1 Description of the ABEL Method

The ABEL method is the application of the EDM for characterization. The idea is similar to that applied in sections 6.2 and 7.4: linear regression of echo duration to water depth is performed to find the effective beamwidth and attenuation. In this case, the effective attenuation value is not used. Instead, the penetration time, which is simply the regression or *y* intercept minus the pulse duration, is matched to modelled penetration time values for the 23 seabeds of the APL-UW table (Table C-2). This is done because an assumed value for the sediment sound speed is not satisfactory here. The best match between the measured and modelled penetration times is determined by the minimum absolute difference. The set of modelled penetration times is calculated for a 10 dB extinction of the signal, based on the 23 sets of parameter values from Table C-2:

$$t_p = \frac{2 \times 10dB}{c_s f_c \kappa} \tag{8-5}$$

Confidence intervals for the observed penetration time are also used to find the range of seabeds that could also match the penetration time.

The effect of slope is accounted for by the methods of sub-section 6.2.3, which are also applied in section 7.4, to find effective beamwidth and attenuation for the whole survey. The idea is the same here, except that the effective beamwidth and penetration time (substituted for effective attenuation) is found for subsets of the survey (the specific classes). The echo duration data can be further filtered by seabed slope. The class ensemble and its low-slope subset are processed by both the iterative slope removal method and with straight-forward linear regression. As before, the most advanced method that converges is used.

To use ABEL on post-classification ensembles, two conditions must be met. The depth spread throughout the class has to be sufficient to achieve accurate linear fits, and the seabed diversity within the class has to be low such that the echo duration versus depth plots are linear. If either or both conditions are not met, it will be reflected in statistical tests of significance and linearity, the data residuals, and, in particular, in the regression slope and intercept values and confidence intervals. If either of the two conditions is not met, the user can defer to a fit with the regression slope fixed to the overall survey effective beamwidth as established in section 7.4. The mean seabed slope of the class is also included when calculating the slope of the best fit line. Both fixed and free regression ABEL analyses are produced by ABEL and are demonstrated below.

The advantage of the free regression ABEL is that the effective beamwidth it finds responds positively to increased scattering by way of beam-widening. Effective beamwidth and penetration time should be anti-correlated, as low attenuation usually precludes roughness as seen in Figure 8-2. High values of both are only found in the case of seabeds with macro-roughness and weak volume backscatter, i.e. gravels and rough rock. In this case, the regression intercept is sensitive to the lengthening of the echo caused by the height of scatterers above the mean seabed level, instead of lengthening due to volume backscatter. The case of micro-roughness becoming macro-roughness was explored in section 5.4.1. In Figure 5-18, for a simulated 50 kHz SBES, the RMS roughness height at which beam-widening saturated and micro-roughness became macro-roughness occurred between two and five centimetres. This is convenient as this is also the upper limit of roughness measurable with MEI at these frequencies (2.84 cm as discussed in sub-section 8.2.4). This convenient co-occurrence is not a coincidence. This critical level of roughness is the point at which scattering occurs throughout the footprint, saturating the area of the seabed from which backscatter can be generated. This is when beam-widening reaches its maximum as can be seen in Figure 5-18 (and Figure 5-26 to a lesser extent), and is why the AEM (the forward model for MEI) saturates for a roughness ratio of four. For real SBES, the effective beamwidths can become large enough to include the first sidelobe (section 6.2) for RMS roughness heights of order 3 cm (gravel/cobbles). Simulations indicate that for roughness larger than a critical level of roughly 3 cm, effective beamwidths should saturate at widths in the first sidelobe (as observed in sections 5.4.3 and 6.2) and the penetration time can be related to the macro-roughness height, l_H, by the following, extracted from Equation 5-3 (which was discussed in detail in sub-section 5.4.1, in reference to Figure 5-18):

$$l_H = \frac{2t_p}{c_w} \tag{8-6}$$

Macro-roughness is usually zero. Indeed, roughness greater than 3 cm is only observable in Patricia Bay in small areas of high slope, i.e. rocky outcrops. Class 9 of Figure 7-17 corresponds to these areas. However, the ABEL results for class 9 are not shown as the class is very limited in number and also suffers from the tendency of tail picking to underestimate the duration of echoes from highly sloping seabeds (slope greater than the effective beamwidth). Measurements of macro-roughness should be much more useful and applicable to low to moderate sloping roughness seabeds such as the coral reefs surveyed in [Gleason, 2009]. ABEL results for attenuation are shown in the next section and discussed with consideration of the relative levels of scattering indicated by the effective beamwidth.

8.3.2 ABEL Results

The ABEL results for the Airmar38 testbed SELwNF classification illustrate the utility of this method. For the classes associated with the sand and gravel seabeds, classes 1 and 6, the ABEL analysis with the iterative LS slope-removal method provides good results with effective beamwidths and penetration times indicative of the seabed type as expected. An example plot of echo duration as a function of depth is shown as Figure 8-8, while the results are summarized in Table 8-5. A large effective beamwidth was found for the gravel class indicating beam-widening due to scattering, while for the sand class the effective beamwidth was slightly less than the overall effective beamwidth of 33°. The measured penetration times matched a range of modelled penetration times; with the best matches finding seabed types that are consistent with the ground-truth data (see section 6.1). The gravel class covers areas of gravel, sandy gravel, sandy muddy gravel and rocky outcrops (see Table 6-1), while the sand class combines the medium and fine sand multibeam classes (Figure 6-1). The resolution of ABEL is somewhat low for sands and gravels as these sediments have similar and overlapping penetration times, as discussed earlier in reference to Figure 8-2. MEI is well suited to characterizing these seabeds, while ABEL is proficient at characterizing fine sediments (muds, clays, silts), so that ABEL and MEI are complementary. For the mud class, the ABEL result is relatively close to the ground-truth data, however, the confidence intervals are large due to the limited depth range of this class (only 15 m).

The fixed-regression ABEL analysis provides a better result for the mud class, with narrower confidence intervals and fewer matching seabed types, see Table 8-6. However, the fixed- regression ABEL results are not suited for the sand and gravel classes as the overall effective beamwidth is not appropriate for those seabeds due to the effects of beam-widening. The gravel class exhibits a much larger effective beamwidth due to strong beam-widening, while the sand class exhibits slightly smaller beamwidths due to weak beam-widening (as seen Table 8-5). ABEL analysis is very sensitive to the slope of the regression line (i.e. the effective beamwidth), so it is very important to account for the effects of beam-widening properly. Under the circumstances of the survey, the effective beamwidths found were quite reliable. The fixed-beamwidth ABEL analysis is seen as a reliable backup to free regression ABEL for the case of little depth variation or a small survey. In all, the three final results are shown in black in Table 8-5 and Table 8-6, while the results that are not used, for the aforementioned reasons, are shown in grey.

8.3.3 ABEL Conclusion

The characterization information offered by macro-roughness and the effective beamwidth are the two additional ABEL methods that were shown in Figure 8-1. This information complements the characterization information provided by measured attenuation. By understanding the effects of both surface and volume scattering on echo duration, ABEL analysis is useful for the full range of seabed types.

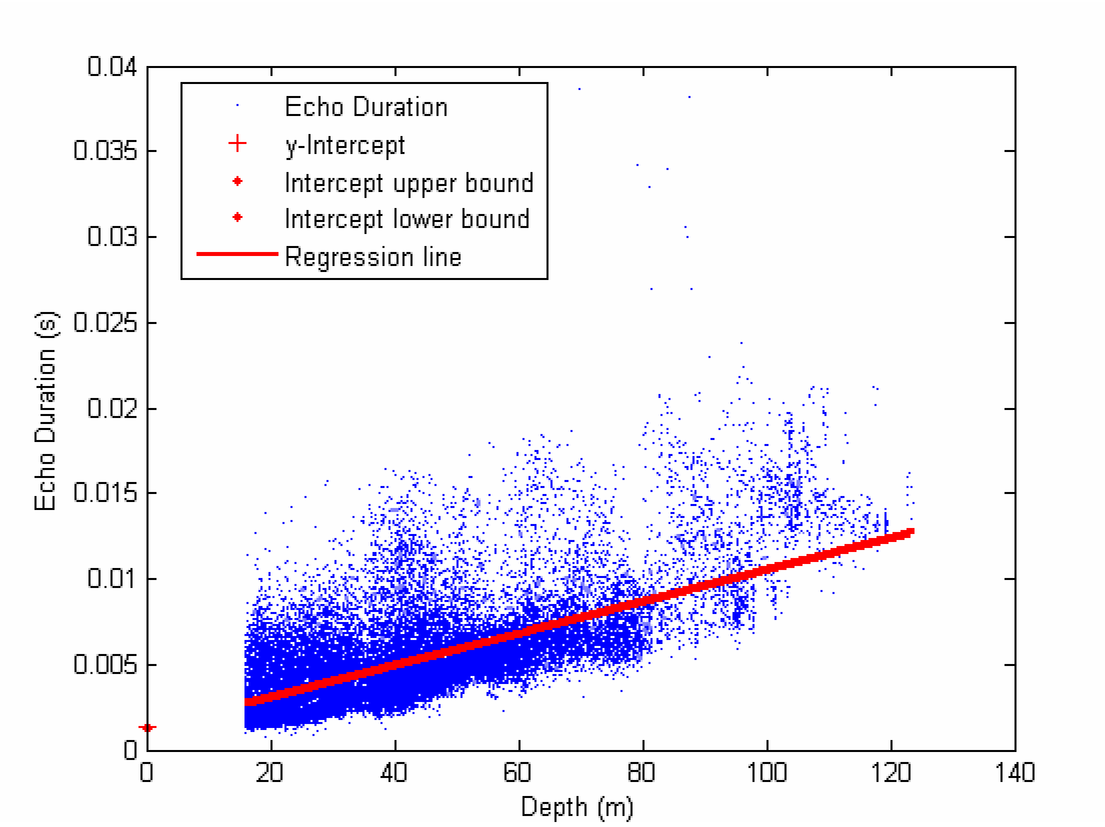

Figure 8-8. Linear regression of slope-removed echo duration to depth for the echoes of class 1 (Airmar38 SELwNF).

Table 8-5. Free regression ABEL results for the three main classes of the Airmar38 SELwNF classification. Iterative low-slope algorithm converged for all classes. All regressions show a clear a linear trend in the data. Note that overall effective beamwidth was 33°. Results in grey are not used.

Class ▼	Effective Beamwidth (degrees)	Penetration Time (s)	Grain Sizes in Range (Φ)	ABEL-free Best Match Grain Size (Φ) and Type	Expected, (Experiment Site Measured) Mean Grain Size (Φ)
1 - Sand	29.7+/-0.56	0.645+/-0.087	1.5 2 2.5 3	2.5 (Fine Sand)	1.5 - 2.5, (1.89)
3 - Mud	42.9+/-5.6	4.358+/-2.561	6 6.5 7 7.5 8 9	7 (Sandy Clay)	8 - 9, (-)
6 - Gravel	52.2+/-2.6	0.557+/-0.102	R -2 -1 -0.5 0 0.5 1 1.5 2 2.5 3 3.5 4	-2 (Cobble, Gravel, Pebble)	-2 to -1, (-1.92)

Table 8-6. Fixed effective beamwidth ABEL results for the three main classes of the Airmar38 SELwNF classification. The regression slope was fixed to a value calculated from the overall effective beamwidth of 33° and the mean slope for each class. All regressions show a clear linear trend in the data. Results in grey are not used.

Class ▼	Penetration Time (s)	Grain Sizes in Range	ABEL-fixed Best Match Grain Size (Φ) and Type	Expected, (Experiment Site Measured) Mean Grain Size (Φ)
1 - Sand	-0.273+/- 0.103	-	-	1.5 - 2.5, (1.89)
3 - Mud	6.105+/- 1.151	7.5 8 8.5	8 (Silty Clay)	8 - 9, (-)
6 - Gravel	1.537+/- 0.148	5.5	5.5 (Medium Silt)	-2 to -1, (-1.92)

Overall, the ABEL method is capable of accurately characterizing fine sediments while discerning long echoes due to roughness from those due to volume scattering. This is beneficial to MEI inversion in particular, which does not model volume scattering, but does characterize coarse sediments (sand, gravels) well. The two methods are complementary, while the AVA method is potentially capable of characterizing sediments that lie in between – fine sand for example. The next section discusses combining the characterization methods.

8.4 Combining the Characterization Methods for a Single Result

As discussed above, the three main characterization methods are complementary in their ranges of applicability. In particular, the ABEL method is most useful for determining whether the classes in question are fine sediments or rough sediments based on penetration time and beam-widening.

The first step in combining the results is to reject and accept each method's result on the basis of the limits of applicability and other considerations as was discussed for each method. For example, the sand and gravel results of the free regression ABEL method (Table 8-5) are accepted, while the mud class result of the fixed regression ABEL method (Table 8-6) is accepted. The MEI result for mud is not-applicable, while the MEI results for sand and gravel are compareable with the ABEL results. AVA results are largely out of bounds. In all, the MEI results for coarse sediments are taken as the final result, while the ABEL results for fine sediments are taken as the final result. This is consistent with the ranges of applicability discussed in reference to Figure 8-1.

Further clues to the nature of the acoustic classes may be obtained by interpreting feature values and the mean echo envelopes, (which were shown earlier in Figure 8-5). This information provides relative characterization similar to that provided by interpretation of three-feature classification results of section 7.8 and is useful for evaluating and confirming the results of the characterization methods. Many of the physics-based features have reliable relationships with geoacoustic parameters, which were discussed with the introduction of each feature family in sub-section 7.5.1, and summarized in Table 7-10. The suite of characterization algorithms produces a large table of mean and standard deviations of these features for each class. A small version of this table is shown below as Table 8-7. For brevity, only the three best features of the amplitude, variability and duration categories are shown, along with the best AEM cross-correlation feature. The first three features vary as expected as was discussed in section 7.8[62], while the AEM cross-correlation medium sand feature will increase with echoes that are most similar to the AEM for medium sand.

Table 8-7. Feature interpretation table for the three main classes of the Airmar38 SELwNF classification result.

Class:	Peak Amplitude ( )	Stack Mean Coherence ( )	Residual Echo Duration (ms)	x-cor Med. Sand ( )
1 - Sand	1.188+/- 0.223 (1)	0.540+/- 0.087 (1)	-0.296+/- 0.370 (3)	0.947+/- 0.027 (1)
3 - Mud	0.720+/- 0.316 (3)	0.411+/- 0.080 (2)	5.987+/- 2.426 (1)	0.821+/- 0.062 (3)
6 - Gravel	1.019+/- 0.684 (2)	0.282+/- 0.041 (3)	1.364+/- 0.627 (2)	0.918+/- 0.042 (2)

Note that the fixed-beamwidth ABEL penetration times are quite similar to the mean residual echo duration values; the difference is in how seabed slope is handled. The mean slope is applied to the fixed-beamwidth ABEL analysis, while the residual echo duration feature is calculated with the seabed slope for each echo. In all, the benefits of physics-bases features are further demonstrated here by the ability for relative characterization and interpretation.

---

[62] As summarized in Table 4-1 and the discussion of section 7.4, the three features will respond to the seabed parameters as follows: the peak amplitude will increase with impedance and decrease with roughness; the stack mean coherence will decrease with variability, which is related to the strength of incoherent surface and volume backscatter; the residual echo duration will increase with beam-widening and decrease with attenuation.

8.5 Conclusion for Characterization Methods

The end result of post-classification characterization is a seabed type label for the acoustic classes produced by the SELwNF classification method. The seabed types are drawn from the APL-UW table, but other labels can be applied, such as 'sloping rock', 'coral heads of relief greater than 1 m', etc. Some bedforms such as sand waves or bioturbation may confuse the characterization scheme. In evaluating the performance of post-class characterization, a pragmatic approach must be taken. No reliable characterization methods exist for uncalibrated SBES. Classification methods provide class boundaries only, with absolutely no clues to what the classes may represent. The body of information provided here will, at the very least, facilitate much easier and accurate interpretation of seabed classification information; the difficulties of which were discussed in chapters 2 and 3.

Overall, the two main methods work well for the 38 kHz echosounder tested here. Additional methods may contribute in the near future with more development. The performance of the characterization methods improve with lower frequencies and wider beams, as do the classification methods. All of the methods of chapters 7 and 8 share a synergy in that the characterization methods have complementary limits of applicability, are much more reliable when combined, and require homogenous acoustic classes which are produced by the classification methods. In turn, the classification methods are built on effective slope and depth compensation (chapters 6 and 7) and the physical basis in chapters 4 and 5. The various problems with current methods were identified in chapter 1 through 3, with a research plan established therein. In all, characterization is not possible unless built upon the foundation of all of the preceding chapters.

Chapter 9. Conclusion

The design of this dissertation is such that each chapter has a conclusion section; then, the aim of this chapter is to summarize the major achievements, to summarize the contributions of this dissertation relative to the current knowledge in the field and to summarize possible applications and future directions of this research.

The major achievements are the seabed classification and post-classification characterization methods as presented in chapters 7 and 8. A flowchart of the combined methodology is shown as Figure 7-1, a list of novel contributions and publications appears as appendix H. The concepts behind the methods are established in chapters 4, 5 and 6, with the need to improve the current methods demonstrated in chapters 2 and 3.

The few other researchers in this field have, in the past, established new seabed classification methods [Burns et al., 1989; Caughey, 1996; van Walree et al, 2005], while this dissertation builds upon their work with a substantive physical basis. Although phenomenological statistical methods may be complex, the process of seabed classification is fundamentally incomplete until the physics are understood. Many statistical methods blindly classify the seabed on the basis of the dominate artefacts instead of the seabed response that is contained within the signal. Heald provided a basis in physics for the RoxAnn method [Heald, 2001]. However, the RoxAnn method has been shown to suffer from depth and slope artefacts, requires a second echo and is limited in its ability to discriminate the full diversity of seabeds. Sternlicht presented an excellent inversion method based on calibrated sonar – the sonar and the method are not generally applicable for classification or characterization [Sternlicht, 1999]. The research of this dissertation is intended to improve the statistical methods of Caughey, Burns and alike, with physics-based methods similar to those of Sternlicht and Heald. The new physics-based methods presented here are designed to be generally applicable and robust to a wide range of conditions, with the aim of replacing the QTC single-beam classification method as the preferred method in the field. The combined classification and characterization methods of this dissertation will spur on new research and exploration of the seabed.

Areas of further research were identified throughout the text. However, the main avenue of continuing the research is to apply the methods generally. That is to test the methods in various survey conditions and then to put the methods to use in mapping habitat, mapping sediment type for underwater installations or resource management, to aid fishers and to apply the methods to the numerous other applications. The methods and the physical basis can also be applied to multibeam echosounders. Some of the new features could be applied to multibeam data (amplitudes, rasters or per beam time series), while adaptation of the characterization methods to multibeam echosounders could prove to be the most utilized method from this dissertation, especially for hydrographic applications. For practical applications, it is envisioned that the methods could be applied to fish-finding and depth echosounders found on virtually every vessel. At the time of this writing, QTC is developing a real-time seabed classification software and hardware framework that could apply a suitable method from this dissertation. The leading candidate is the reduced feature set combined with a standard catalogue (trained and calibrated using the methods of chapters 6 and 7) with slope compensation based on data from hydrographic charts. Quadratic classification could be used effectively here. Three-feature classification is another method that may not offer as good discrimination as the full SELwNF method, but does offer relative characterization which may be combined with the SELwNF method. Characterization methods could be applied in real-time by accumulating ensembles of the last 20 to 50 echoes. More echoes are required for AVA, while a library of AEM data could be generated pre-survey to speed up the process of finding the optimal model. Real-time classification and characterization is a technically difficult prospect, as all the methods have to be robust to varying conditions, from picking to characterization. The methods here were designed with this level of reliability in mind to enable real-time classification. The exact details of the implementation are left to the engineers, as the basis has been established in this dissertation.

The wide-spread availability and use of real-time seabed classification and characterization would be of great practical benefit for people working on the sea. Echosounders with the methods of this dissertation would also be an invaluable tool to explore and understand the oceans.

Appendix A. Picking – the Detection of the Onset and Termination of Echoes

Establishing the onset of the echo, a process known as bottom picking, is essential for the seabed classification and characterization methods of this thesis. Not only is the bottom pick used to measure the water depth (Equation 2-1), but it is also used to frame a section of time series data that contains the echo (i.e. the analysis window). For classification and characterization, highly accurate and consistent bottom picks are required to align echoes for echo-to-echo averaging or 'stacking' (sub-section 2.2.2). Some methods also require measurements of the duration of echoes. Then the end of the echo must also be picked; this pick is termed the tail pick. Measurements of echo duration depend on the accuracy of the picking routine, so finding the best possible methods is important. Tail picks are useful for evaluating the performance of the analysis window; if the tail pick occurs outside of the analysis window, something is amiss.

A.1 Bottom Picking

The usual purpose of a bottom pick is to detect the onset of the echo indicating the two-way travel time to the seabed, which leads to the measure of water depth (e.g. see Equation 2.1). The process for bottom picking is well established and there are numerous standards applied by hydrographic agencies and companies. The usual method for single-beam sonar is threshold detection. In his thesis, Sternlicht found that the optimal method of picking and aligning echoes for normal incidence canonically distributed echoes (i.e. for a normal seabed like medium sand) is threshold detection [Sternlicht, 1999]. Sternlicht found that his group delay alignment technique gave the best results for atypical echoes from high seabed slope, high incidence angle or a combination of high roughness and high frequency [Sternlicht, 1999]. Bottom picking becomes especially difficult when there abundant reverberation in the water column just above the seabed [Preston et al., 2006]. van Walree managed to avoid the issue of picking by applying a static 10 ms analysis window around the peak intensity of the entire trace and then using a median filter to effectively reduce the variability of feature values [van Walree et al., 2005]. The feature extraction algorithms used by [van Walree et al., 2005] are also mostly insensitive to picking as long as the echo is contained within the analysis window. Other approaches exist such as matched-filter processing; however, the vast majority of published studies use threshold detection.

The QTC threshold detection algorithm for single-beam systems is as follows: A search ‘gate’ is established based on the previous pick (if none exists, the entire trace is used), see Figure A-1. The peak amplitude within the gate is found. The threshold for detection is determined as a percent of the peak amplitude of the last several echoes. It is usually 35% as set by the user, ranging from 5 to 50% depending on the data conditions. The gate is then searched in one of two directions: from the peak amplitude to the start of the gate, or from the start of the gate to the peak amplitude. The search is stopped and the pick established when two consecutive samples are above (or below) the threshold. The algorithm has other options for rejecting picks due to factors such as noise level, signal to noise ratio, clipping and low amplitude. Echoes without bottom picks are excluded from further processing in all methods presented. In all the surveys presented in chapters 3, 6 and 7, the bottom picking parameters were tailored specifically to each data set to maximize the performance. The performance was judged by consistency of the position of the picks relative to the echoes and maximising the number of echoes picked while still rejecting all bad traces.

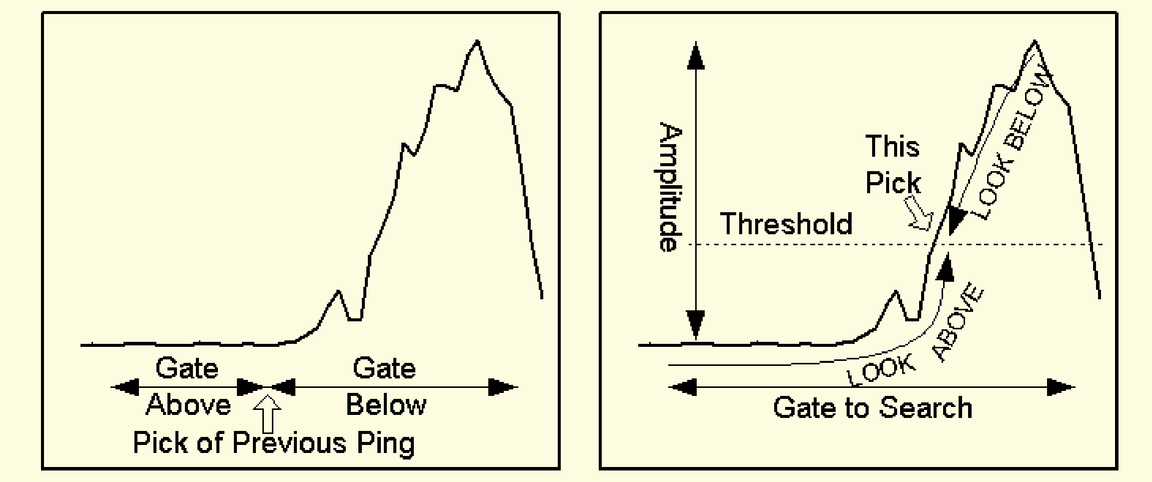

Figure A-1: The QTC IMPACT bottom picking algorithm. (*left*) the search gate is established based on the previous pick. (*right*) the gate is searched for samples above or below the threshold. Courtesy of QTC, from [QTC, 2004a].

In the processing of all data presented, details of the bottom picking are as follows unless otherwise noted. The bottom pick algorithm is from QTC IMPACT version 3.5. The bottom pick thresholds are 10% of peak amplitude for low frequencies (24, 38 and 50 kHz) and 15% for high frequency (200 kHz). The gates are specified in two-way travel time converted from metres: 3 m above / 5 m below (low frequency) and 1.5 m above / 3 m below (high frequency). The minimum signal to noise ratio was 1.5 (low) and 3 (high). Clipped echoes and echoes below 5% signal threshold were unpicked.

A.2 Tail Picking

Tail picking is critical for the echo duration methods of chapters 7 and 8, and the importance of tail picking cannot be understated.

Tail picking with a threshold method similar to the bottom picking method described above yields marginally useful results. The problem is that while the onset of an echo is abrupt, the termination or tail is not. As echoes decay in amplitude most will have temporary decreases in amplitude, best described as ‘valleys’, due to the sporadic effects of reverberation. A threshold pick algorithm often falsely picks a ‘valley’ as the tail. Figure A-2 has an example of a 'simple' threshold tail pick stuck in a valley. In work with BORIS, a 'smart' generalized threshold picking scheme was developed with some intelligence built-in to recognize valleys. It looked in the vicinity around the pick for the mean amplitude and trend, if either was indicative of a valley, the algorithm would move on. However, even the smart threshold pick would still get caught in larger valleys. Its performance was observed to be not as consistent at tail picking as the cumulative methods.

Cumulative methods take advantage of the fact that there is a bottom pick already established. Cumulative methods sum the trace from the bottom pick to the earliest of the following: the end of the trace, the onset of second echo (which is the time of the bottom pick times two plus a small safety margin) and a predefined maximum (defined by time or depth, usually 0.010 s for high frequency (200 kHz) and 0.050 s for low frequency (24-38 kHz)). There are two variations: cumulative energy (CE) and cumulative amplitude (CA). The algorithm for both is simple: Calculate the cumulative sum of amplitude or energy (amplitude squared) for each sample in the direction of increasing time (deeper). Normalize to one by dividing by the total amplitude or energy. Find the earliest sample with a cumulative amplitude or energy exceeding a pre-selected threshold.

The above methods, particularly the CA method, work well in cases of no noise, but for real data there is always some noise. Real traces usually have a non-zero minimum signal level or DC-offset, this value is the minimum of the entire trace. It is removed by subtracting it from the entire trace, so that the minimum trace value is zero. Two extensions to compensate for noise are applied, prior to the calculation of the cumulative amplitude/energy curves. One is to use a symmetric running average to smooth out the noise level (denoted RA), the other is to find the mean noise level per sample and subtract that from each sample (denoted NLS). For the NLS method, time-varying gain acts to boost the noise level (either electronic or acoustic) with time. Therefore, the noise level must be measured at the end of the current trace. The noise level is then measured as the mean amplitude from a small region at the end of the trace (a minimum of 5 samples to a maximum of 1/10th of the entire trace). The noise level used on each trace can be the average over the data set or that for the specific trace – the latter was found to be more effective because of gain variations within the data sets. NLS works because of the summation of noise in the cumulative sum – NLS subtracts the mean noise which is accumulated by the cumulative sum. For cumulative amplitude picking, the trace is allowed to have negative values as this acts to remove the noise from the cumulative sum accurately. For cumulative energy picking, NLS does not work because negative samples are squared to positive values. The best noise compensation solution is to combine the two methods to smooth the noise to get a better measurement of the mean offset amplitude at the end of the trace. The combined noise compensation method (RANLS) works with either picking method, and is effective and unique.

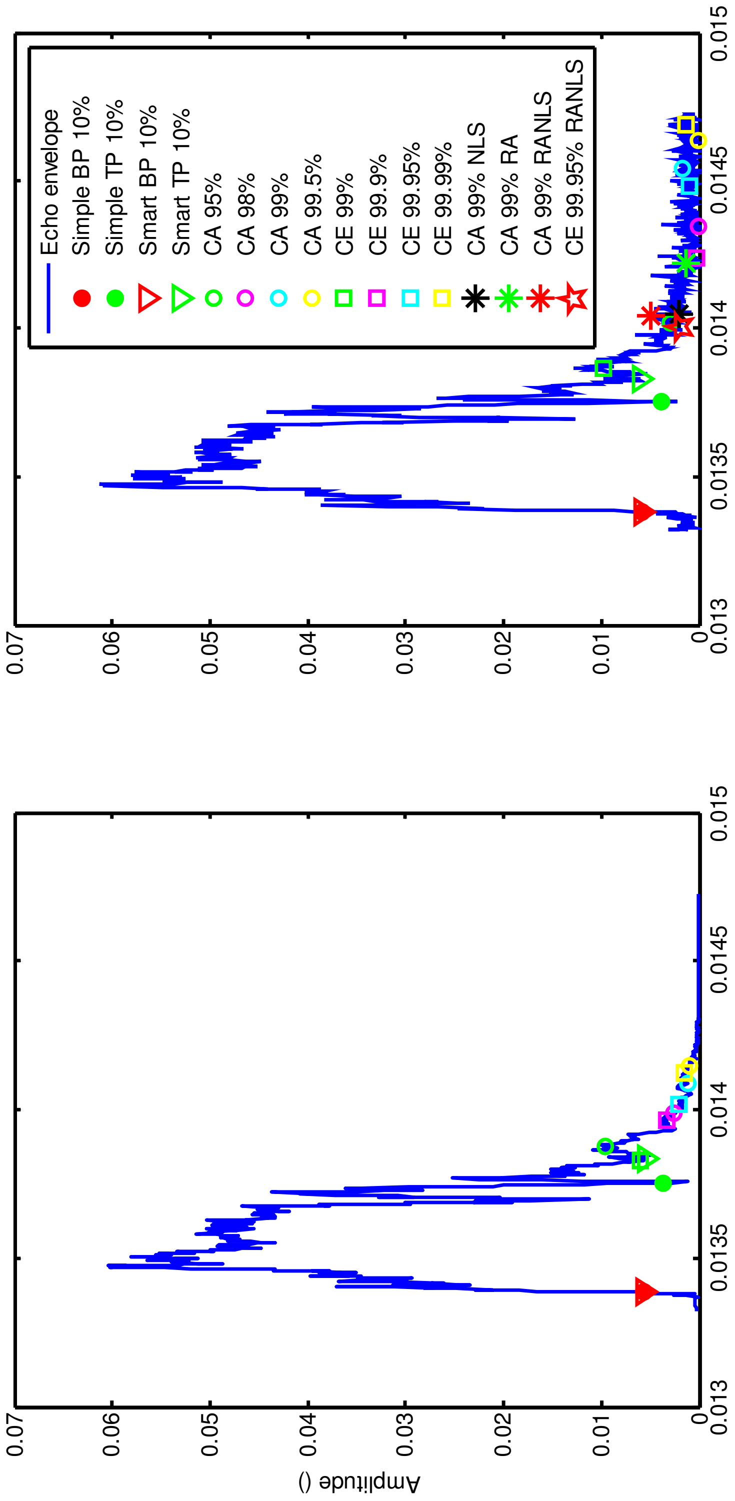

Figure A-2. A BORIS-generated echo amplitude trace to illustrate the results of the various picking methods; the left trace has no noise, while the right trace is the left trace plus 2% one-sided Gaussian noise. Model setup: 10 m water depth, Suzuki50 SBES, sandy gravel seabed (see appendix C for specific parameters).

Figure A-2 is an example of BORIS data with picks from the various available methods with or without noise at several threshold levels. Comparing the picks between noiseless and noisy cases, one sees the threshold picks are unaffected, while the low threshold (less than or equal to 99%) CE tail picks are less sensitive to noise than are all CA tail picks. However, CE tail picks are sensitive to the distribution of amplitudes; they are affected by echo shape, as a 'pointy' echo of the same duration as a 'flat' echo would be picked erroneously shorter. This effect can be seen in Figure 5-16. In that figure, the CE tail picks occur erroneously early for echoes that have RMS roughness heights around 2 cm and 12.5 to 20 cm. Because of this bias, CA picks are the preferred picks.

In general, the 99% CA pick is the most sensible tail pick as seen by inspecting BORIS data (such as the left side of Figure A-2) and analysis of the low-noise data from the Patricia Bay depth dependence experiment (see chapter 6.2). However, because CA tail picks weight each sample evenly, they are sensitive to any noise, which infers that they are very sensitive to any multi-path echoes or stray signals occurring from the end of the echo to the end of the tail picking analysis window. Without noise compensation, CE picks are more robust to noise and variability than CA picks, which is why CE picks were used (on BORIS results in particular) until the noise compensation algorithms were developed. CA tail picks for all the data of chapters 6 and 7 were initially attempted without noise compensation and the results tracked the duration of the picking analysis window and not the echo. Therefore, noise compensation is required for tail picking.

As demonstrated in Figure A-2 the noise compensation methods are quite effective at reducing the biases of cumulative picking for noise. The 99% CA tail pick is consistent with the noiseless 99% CA tail pick when the NLS or RANLS method is used. The stability of the combined RANLS noise compensation method is preferred. Both NLS and RANLS methods tend to cause the picks to occur a minuscule amount of time earlier, however the effect is far less than the natural variability of the data. This is related to the nearly imperceptible decay of echo amplitude after the 99% CA pick, which can actually persist for a surprisingly long duration, up to 50 ms for an echo that appears to be 3 ms. This was observed in the low-noise data from the depth dependence experiment –

particularly the 24 kHz data from gravel site. This effect is explained by the fact that real beampatterns do not cut off to zero, nor does volume scattering cut off to zero – echoes decay asymptotically. If the picking analysis window is limited to a short duration that is not clear of residual signal from the primary echo, the noise level measurement (by the RANLS or NLS methods) will be artificially high and the tail pick will be erroneously short. In the other extreme, one would expect that exceedingly long picking analysis windows would also cause artificially high noise measurement because of amplification by TVG. The effect of TVG is not an issue with the RANLS method as long as the duration of the picking analysis window is constant throughout the data set, as a varying picking window can bias the picking (this is important for data with multiple echoes, such the data presented in section 6.2).

In general, stable and accurate picking is achieved with picking analysis windows that are consistent and of sufficiently long duration. The RANLS 99% cumulative amplitude tail pick is the best overall tail pick.

Appendix B. Data Reduction and Principal Component Analysis

An important step in the processing of echo time series for seabed classification by phenomenological segmentation is data reduction. This step occurs after calculation of echo statistics (known as features) and prior to cluster analysis. The number of features varies, depending on the processing scheme and the data available, but is always a number greater than 20. This is too many dimensions considering many features are redundant, correlated or even collinear and that cluster analysis[63] is best viewed in three-space. Features must be combined or removed to make a total of *three*. Anderson et al. state that up to 80 parameters are required to describe the seabed fully [Anderson et al., 2007]. From chapter 4, the physics of seabed scattering utilizes at least 10 parameters (depending on the theory). However, echoes primarily respond to three main seabed factors: seabed roughness, interface impedance and volume inhomogeneity. Ideally, three combinations of features would be found that map echoes onto a basis corresponding to the three major factors. Data points would then be plotted in this three-space and could be grouped spatially into seabed classes using cluster analysis. However, each data set is unique with different features being important each time. So instead, the approach is to find three orthogonal linear combinations of the features that embody the maximum amount of information possible. Principal Component Analysis (PCA) is a multivariate statistical technique that meets this objective well. General PCA is described in part 1 this appendix, followed by a discussion and evaluation of four specific PCA methods. Part 2 describes a feature filtering method that is paired with PCA as the data reduction processing applied in chapter 5.

B.1 PCA

The approach of PCA to maximize the amount of variance represented by linear combinations of the data variables is similar to approaches of other multivariate

---

[63] Cluster analysis is discussed in section 2.2

techniques. All use linear combinations of variables: multiple regression uses linear combinations of independent variables to best predict dependent variable(s), canonical correlation analysis uses linear combinations of a subset of variables to correlate with linear combinations of another subset of variables, and discriminate analysis seeks linear combinations to separate known groups within the data [Rencher, 2001]. If separate groups within the data are defined prior to analysis, discriminate analysis may be a better approach than PCA. In seabed classification, these separate groups are usually defined geographically and called training sites [Wilmut et al., 2009]. The problems with this technique include failing to represent a group within a survey, and having a non-homogenous training site. Both problems are related to not having good prior knowledge of the area to be surveyed. However, a lack of knowledge of an area is why an area is usually surveyed. Therefore, PCA is the most robust and generally applicable technique for seabed classification. Discriminate analysis is briefly discussed in chapter 5.

The orthogonal linear combinations of the variables defined by PCA are called the principal components. The principal components are contained in an orthogonal rotation matrix, **V**, also called the projection matrix as it projects the data on the principal components. **V** is also said to contain the principal component *loadings*, which are the coefficients of each variables in the linear combinations that make up the principal components. The top three principal components (as determined by the fraction of variance explained) are contained in the columns of **V**, and are extracted to a *reduction matrix* **V*.** The reduction matrix is then used to project the echo feature data, **X,** into the vector space defined by the reduced principal components. QTC calls this three-dimensional space Q-space and the location of each data point is called its Q-value. The Q-values (normally known as factor scores) are contained in the matrix **F**, where **F = XV***. There are several variations of PCA methods, the following paragraphs will explain each and then the variations will be evaluated.

In the standard PCA that appears in textbooks such as [Rencher, 2001], there are two similar methods: *covariance* PCA and *correlation* PCA. The difference between the two is in the data pre-processing steps. Consider a data matrix **X** composed of *n* rows of observations and *m* columns of variables (or echo features). In both methods, the mean (of each column) is subtracted (for each variable) and each entry in the data matrix is divided by $\sqrt{n-1}$. In correlation PCA, the additional pre-processing step is to divide each column of $\mathbf{X}$ by the column's Euclidean norm, which is the square root of the sum of all the squared data values for each column/variable/feature. In the former case of not dividing by the norm, the matrix product $\mathbf{X}^{\mathbf{T}}\mathbf{X}$ is equal to the *m* by *m* covariance matrix. In the latter case when the data are divided by the norm, the matrix product $\mathbf{X}^{\mathbf{T}}\mathbf{X}$ is equal to the *m* by *m* correlation matrix. For the next step in both cases, the matrix $\mathbf{X}^{\mathbf{T}}\mathbf{X}$ undergoes eigenvalue decomposition to produce a diagonal matrix $\mathbf{D}$ of generalized eigenvalues and a full matrix $\mathbf{V}$ whose columns are the corresponding eigenvectors. The eigenvalues correspond to the amount of covariance or correlation explained by the corresponding eigenvector. The reduction matrix $\mathbf{V^*}$ is formed by the top three eigenvectors as determined by the largest eigenvalues and is an *m* by *3* matrix. The Q-values / factor scores matrix $\mathbf{F}$ is an *n* by 3 matrix, with each row being the stacked echoes' position in Q-space. Groupings of points in Q-space are then found by cluster analysis to define seabed classes.

An important last step is to store the following information for later use: the data mean and mean norm values that are used to pre-process the data for PCA, the reduction matrix and the cluster definitions (provided by cluster analysis). This information comprises a *seabed classification catalogue*. It is used to classify supplemental data by projecting the supplemental data into the original Q-space and then assigning each datum to the nearest cluster centre (as measured by Bhattacharya distance). In this case, the original data set can be thought of as a training survey. Classifying acoustic data in this way is called supervised classification, as described in chapter 2.

There are two variations on standard PCA employed by QTC in their IMPACT and MULTIVIEW processing suites. In IMPACT processing, a modified covariance PCA is used. The feature algorithms are designed to output values with a strict range from zero to one, but no attempt is made to normalize the expected distributions; i.e. features should have normal distributions with the same mean and standard deviation in order for unbiased processing with covariance PCA. Since no normalization is done, the covariance matrix is calculated directly from the feature data. The unmodified feature data are then multiplied by the reduction matrix to project them into Q-space. However, good results can be obtained by this method, as has been shown by many successful surveys through a diversity of of conditions (as discussed in chapter 2, see [Hamilton, 2005; Gleason, 2009] for lists of sucessful applications of the method). The reason this application of covariance PCA does not hinder the results is the design of the features: whether it be by design or is simply fortunate, the best features tend to have the largest means and covariances. This is how it works: a feature with a larger covariance will have higher (on average) component loadings in the reduction matrix. Then the loadings multiply the features' values; if those values are high, this will result in a much larger contribution to Q-values (and the overall result) than features with low mean and/or low covariance. Conversely, if a number of IMPACT's original features had high mean values but were basically noise, the results would have been poor. As an example of this, it is known that virtually all FFT features and most wavelet features have little discriminatory information in all survey scenarios (see chapters 2 and 7). This is especially true of the higher-order FFT and wavelets that measure higher frequency content in the signal. This is due to under-sampling by TNORM depth compensation and suppression of high frequency content in the signal by stacking. Fortunately, these poor features do not disrupt the results as their means and covariances are low. Figure B-1 summarizes an example of IMPACT-style covariance PCA. The scree plot indicates how the variance of the data set is concentrated in the first few principal components, while the component loading plots show each feature's contribution to the respective principal components. Features that are furthest from the centre of the component loading plots contribute the most, while features that are near to the centre contribute little and are useless[64]. Figure B-1 shows that all but the lowest order wavelet and FFT features do not contribute to the Q-values (the reason these features contribute is that they are effectively uncompensated relative amplitude features as was discussed further in section 7.5).

---

[64] Conventional interpretation of component loadings with loading plots, such as Figure B-1, is done on the basis of rotated principal components. In this case, the aim is to maximize covariance or correlation. Applying a rotation to maximize some other criterion to aid interpretation would result in less than optimal principal components for the purpose of data reduction. Therefore rotations are are not applied. The interpretation of unrotated loadings is still valid since it is these unrotationed values that are applied to the features to produce scores or Q-values, Q-space and subsequent classifications. As such, component loading plots indicate the relative contributions of each feature to the final result of classification. This is similar to the direct interpretation of loadings advocated in [Rencher, 2001], pp. 403-404.

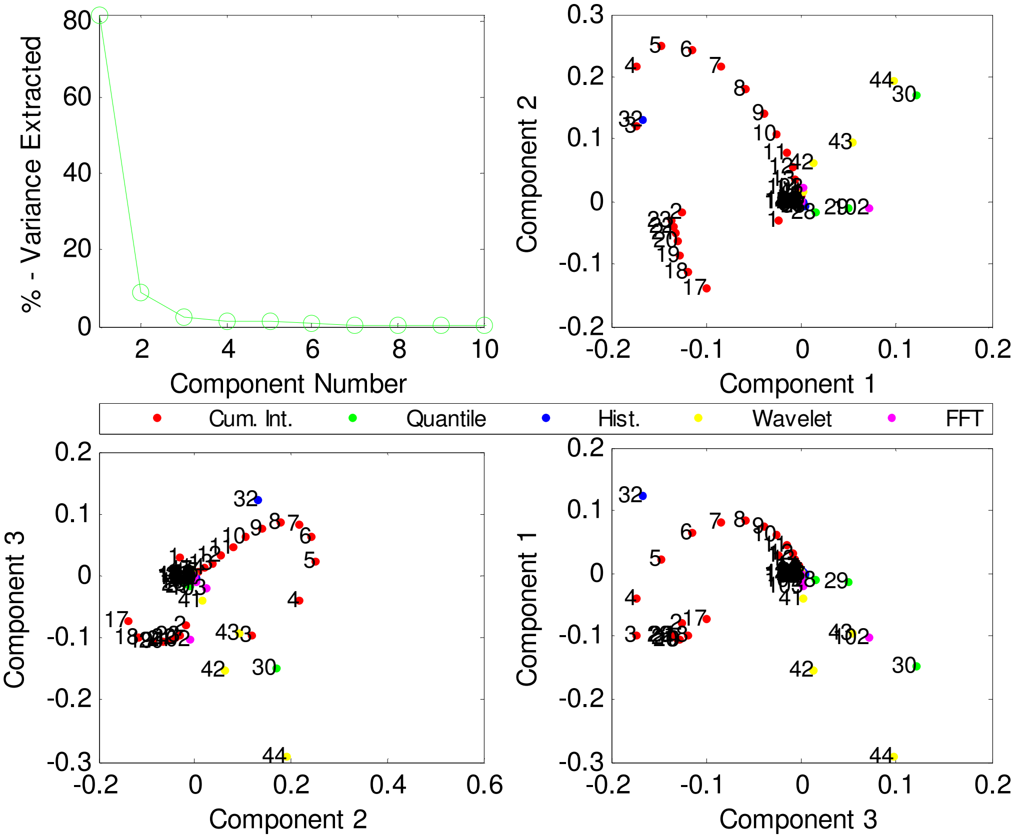

Figure B-1. An example result of IMPACT PCA on Patricia Bay TNORM-processed 38kHz data. *top left*: a scree plot of the percent variance explained by each principal component. *top right:* weighted component loading plot for principal component 1 versus principal component 2. *bottom left/right:* weighted component loading plots of the remaining components. 'Weighted' refers to scaling by each feature's mean value to make the values plotted representative for each features' contribution in determining the scores/Q-values. This means that the farther the features are from the centre (0,0) of the component loading plots the more they contribute to the respective component scores. Feature locations are coloured by feature algorithm.

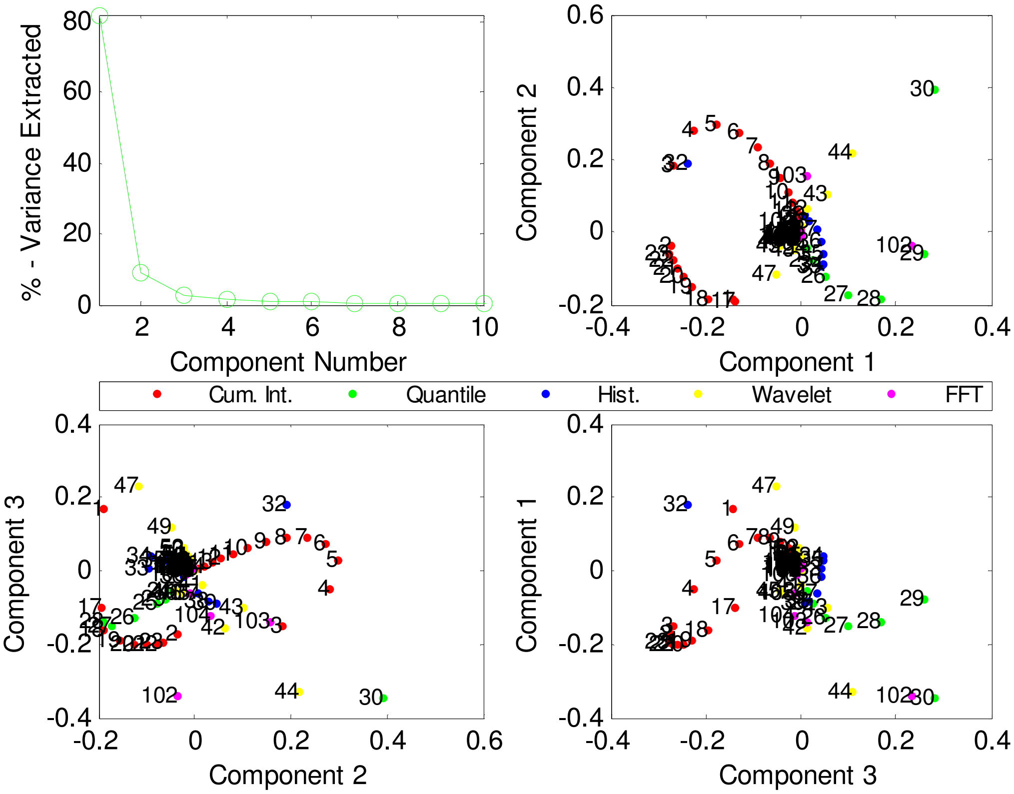

Figure B-2. An example scree plot and component loadings for standard covariance PCA as calculated from TNORM-compensated 38 kHz Patricia Bay data. The component loadings are not modified, as the data set is zero-mean and unity norm.

The PCA processing in QTC MULTIVIEW is a variation on correlation PCA. MULTIVIEW PCA skips the pre-processing step and calculates the reduction matrix from a correlation matrix that is calculated without modifying the data. The reduction matrix is applied to data that are not zero-mean and unity norm. However, the component loadings are scaled by the standard deviation of each feature's values, which is effectively similar to scaling the data by its Euclidean norm.

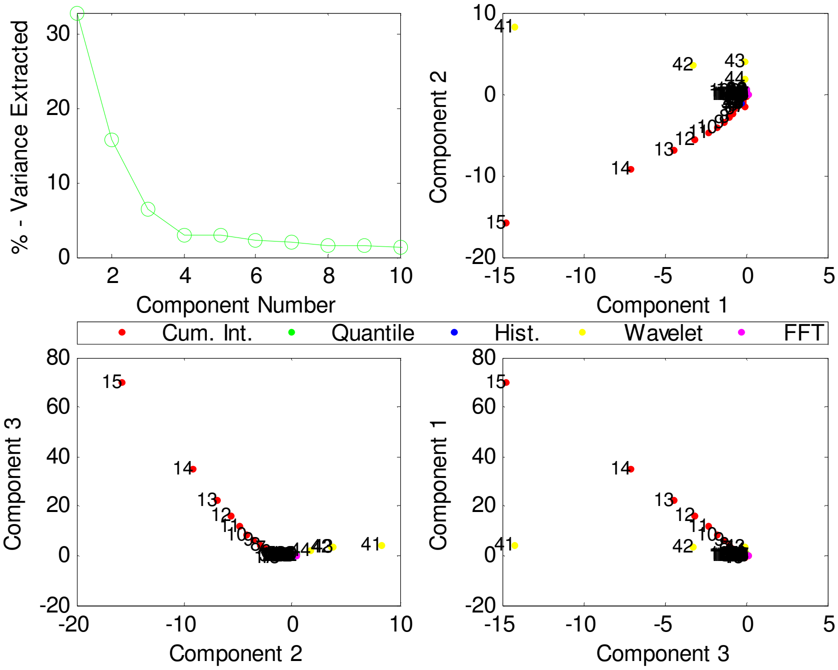

Figure B-3. Example scree plot and weighted component loadings for MULTIVIEW type correlation PCA as calculated from TNORM-compensated 38 kHz data from Patricia Bay.

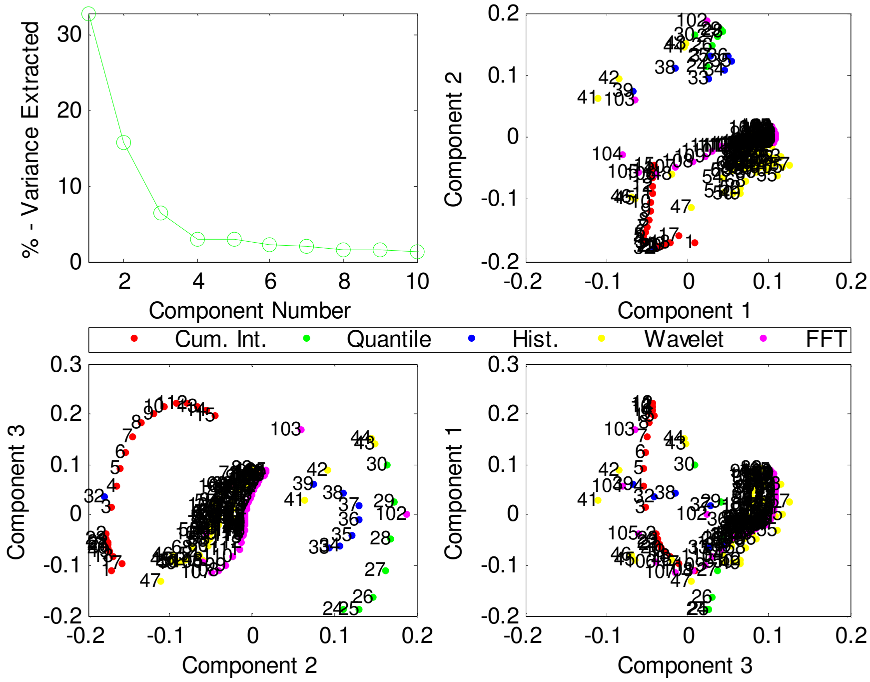

Figure B-4. Example scree plot and component loadings for standard correlation PCA. As calculated from TNORM-compensated 38 kHz data from Patricia Bay.

It is difficult to compare the performance of the various PCA methods. One can rely on theory, which states that correlation PCA is the most robust PCA, while covariance PCA will perform well with data that are all of the same type, i.e. ocean temperature data measured at numerous locations. One can also evaluate the results on real data. Figures Figure B-1 through Figure B-4 show typical results of IMPACT PCA, standard covariance PCA, MULTIVIEW PCA and standard correlation PCA. The IMPACT PCA result is dominated by relatively few features, even compared to the standard covariance PCA which has a similar result. It is difficult to compare these two PCA by the loading plots since there are different weightings involved. The correlation methods (Figure B-3 and Figure B-4) result in principal components that explain a significantly lower fraction of the data correlation than the covariance method. This is likely caused by noisy low-mean features contributing to the correlation; the covariance method ignores them due to their low values. However, there could have been useful information in the low-mean features that would be ignored by the covariance PCA. The correlation PCA methods, especially the standard correlation PCA, weight a greater diversity of features than the covariance methods. To compare the PCA methods further, cluster analysis results were observed for the same Patricia Bay data processed with different PCA procedures. In the case of the original IMPACT feature set, the results of the covariance PCAs were very similar to each other, as were the correlation PCAs. All methods performed well, with IMPACT PCA and standard correlation PCA being the best. In the case of new features, the QTC PCA methods were completely disrupted by the new features' un-bounded values. The standard covariance PCA emphasizes only a few of the new features, with most features not contributing. The standard correlation PCA provided a good result as viewed in the loadings plots, Q-space and final class map.

In all, the choice of PCA method depends on the feature data provided. For the original IMPACT 166 features that had fortunate weightings, IMPACT PCA performs better than the standard correlation PCA. IMPACT PCA is used for this scenario in chapter 3 and in parts of chapter 7. Standard correlation PCA is the best method for the new features introduced in section 7.5.

B.2 PCA Iterative Outlier Removal

The introduction of new features with unbounded range in section 7.5 created a new problem of extreme Q-values coming out of the PCA processing step. The PCA result also appeared somewhat biased toward producing these outliers instead of capturing the maximal correlation of the data - one can see a few outliers well separated from the main distribution of Q-values in Figure B-5. Both PCA and cluster analysis are easily biased by outliers as these analyses rely on L2 norms that are heavily influenced by extreme values. Solutions to this problem were investigated. Possible methods of mitigation were found to include outlier removal from the features via standard methods such as Grubb's test [Rencher, 2001], Thompson's Tau, etc., outlier removal from the reduction matrix, and outlier removal from the Q-values. The difficulty with removing outliers from the features is, first, the features may not be normally distributed (a condition for most outlier tests) and, second, feature outliers may not be outliers once transformed to Q-space. Methods were investigated for making PCA more robust to outliers. One method in particular is promising, shown in [Garbys, 2006]; however this solution does not fully prevent passing outliers on to cluster analysis. A simple, novel method was developed that both removes outlier Q-values prior to cluster analysis and prevents outlier feature-values from influencing PCA.

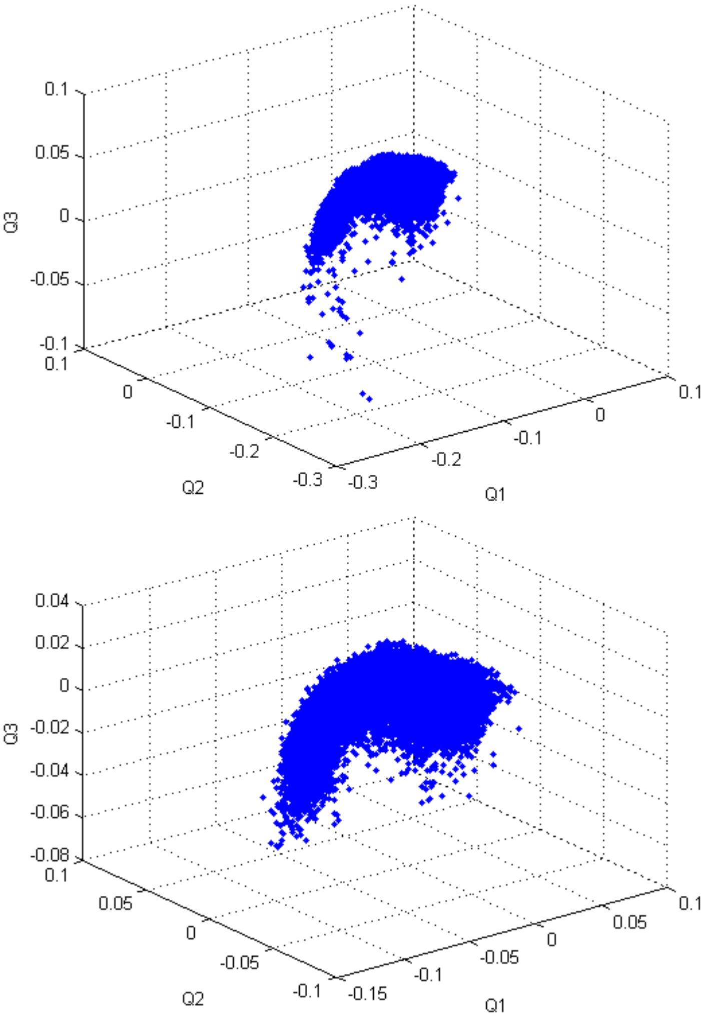

Figure B-5. Q-space plot of the Q-values of the Patricia Bay 38 kHz survey processed with the SELwNF method (see section 7.5) with correlation PCA before (*upper*) and after iterative outlier removal (*lower*).

The idea behind the new PCA outlier removal method is the initial data (and associated feature vectors) associated with outlier Q-values drive PCA to produce those same outlier Q-values in the first place, whether or not those data points have feature values that are outliers themselves. A subsequent assertion is that removing these data points (that correspond to outlier Q-values) prior to PCA will shift the PCA result to produce fewer and less extreme outlier Q-values. Removing these new outliers will cause a smaller shift in the principal axes because this second group of outliers has less of an effect on PCA. The algorithm is then an iterative loop. For the first iteration, PCA is run normally, and the outliers in the Q-values are found. In subsequent iterations, data / feature-values corresponding to the Q-value outliers from the previous iteration are removed for the calculation of a new reduction matrix. Then, *all* Q-values are re-calculated from the new reduction matrix and new Q-value outliers are found. At the end of every iteration, the new reduction matrix is compared to the previous reduction matrix by a mean absolute difference calculation. If this number is not below a small value (1%), the loop is run again. If the change in reduction matrix is small, the PCA result has stabilized and the last set of outlier Q-values are removed so they do not bias the cluster analysis. The iterative loop is the important part of the process, the method of selecting outliers is less important. Outliers are found based on the Mahalanobis distance calculated from the entire set of Q-values. Q-values with the largest of Mahalanobis distances are the outliers. Removing a very small fraction of Q-values is found to be sufficient to shift the PCA result significantly. The optimum fraction of the Q-values is difficult to define. If the fraction is too large, the loop does not stabilize; the largest fraction at which any observed data set stabilized was 10%. If it is too small, no change is seen and many outliers remain. Based on the distribution of Mahalanobis distances, the fraction of Q-values to be removed was found to be 0.05%.

In the case of the 38 kHz Patricia Bay survey, (specifically the SELwNF results of section 7.5), the PCA iterative outlier removal loop stabilized after three iterations (about 5 seconds processing time), with a mean absolute difference in the reduction matrix of 7.1%. The reduction of outliers is observable by comparing the upper and lower plots of Figure B-5. The shift in the reduction matrix and hence the definition of Q-space is more subtle. In other test cases, larger shifts in the reduction matrix were seen (up to 65%), but the results were not used in this dissertation. Classification results did improve slightly as seen in qualitative testing. Often the 'outlier' class was eliminated, and the consistency of class maps improved. A more quantitative evaluation should be undertaken, but is left for future work. This technique is seen as successful in removing outliers, improving PCA and cluster performance, and it is unlikely to cause any detriment to the results, as the shift in the reduction matrix is relatively minor. Overall, iterative outlier PCA is a successful technique and is applied with the new features in section 7.5 to overcome the problem of outliers from unbounded feature algorithms.

Appendix C. Tabulations of Geoacoustic and Geotechnical Parameters

This appendix contains a summary of the geoacoutic and geotechnical parameters used throughout this dissertation. First, a table of the various SBES geotechnical parameters is shown as Table C-1. This table is also a complete list of BORIS simulated echosounders – every echosounder used in the field was also simulated with BORIS. Where parameters are indicated as simulated, the simulations are based on physical quantities, such as transducer diameter, as much as possible. Table C-2 is based on the APL-UW handbook [APL-UW, 1994]. Note that Sternlicht parameterizes the volume scattering cross-section for all seabeds of Table C-2 as equal to the attenuation times the frequency times 0.04, which is different from the values of the handbook that vary from 0.01 to 0.02 [Sternlicht, 1999]. The volume scattering cross-section is not of interest here as the BORIS model uses a measure of volume inhomogeneities instead. A list of BORIS simulated seabeds follows as Table C-3. As mentioned, there is no good single source of BORIS seabed parameters. This best is [Berntsen, 2001], but is not consistent with the values from [APL-UW, 1994]. Table C-3 is a compromise of values from [APL-UW, 1994; Berntsen, 2001; Bergem et al., 1999] with '96' tag referring to seabeds from [BORIS, 1996]. The letters 'AG' refer to seabeds created for A. Gleason's dissertation work, which were used for some tests here. The PSL habitat seabed of Figure 4-3 was never developed further – it is the 'Sand' seabed with the correlation lengths switched and roughness increased to 5 cm RMS.

Table C-1. SBES parameters used throughout this dissertation. Circ50_20 is a virtual SBES used for BORIS testing only. All of the other real SBES are also available as BORIS simulations, with BORIS 3-D beampatterns simulated or constructed from 2-sided measured profiles and BORIS transmit pulse files simulated to match usual settings or measured. The frequencies, durations and beamwidths are all measured from the BORIS numerical beampattern and transmit pulse files. For two -6 dB beamwidths there was not enough resolution to accurately measure the beamwidths (as indicated by the ~ symbol).

System:	Frequency (kHz)	Pulse Duration (μs)	Beamwidth (degrees) -3 dB	-4.34 dB	-6 dB	Beampattern type	Pulse type
Simrad12	11.9	106	15.8	18.8	21.8	2-p meas.	sim.
Odom24	24.7	311	19.3	23.8	27.7	2-p meas.	sim.
Simrad38	37.8	296	8.2	13.3	11.2	2-p meas.	sim.
Airmar38	37.8	296	16.1	22.8	~24	2-p meas.	sim.
Circ50_20	50.0	285	19.7	23.4	27.1	sim.	sim.
Deep50	50.0	106	19.2	23.2	25.6	2-p meas.	sim.
Suzuki50	48.2	422	28.6	36.2	44.2	sim.	meas.
Suzuki200	202.0	294	12.0	14.2	16.4	sim.	meas.
Odom200	200.0	305	4.3	4.9	5.1	2-p meas.	sim.
Simrad200	200.2	299	7.0	9.1	~11	2-p meas.	sim.

Table C-2. Sediment parameters from the APL-UW models, [APL-UW, 1994]. The most closely related BORIS seabed is listed for cross-referencing purposes.

Sediment Name / Folk Class	Grain Size (Φ)	Sound Speed (m/s)	Density (kg/m^3)	RMS Roughness Height (m)	Attenuation (dB/kHz/m)	Related BORIS Seabed Name
Rough Rock	-	3750	2500	0.100	0.196	HRHB
Rock	-	3750	2500	0.030	0.196	LRHB, HardFlat
Cobble,Gravel,Pebble	<-2	2700	2500	0.028	0.273	-
Sandy Gravel	-1	2006	2492	0.025	0.456	Gravel, Gravel96
Very Coarse Sand	-0.5	1961	2401	0.023	0.456	
Muddy Sandy Gravel	0	1917	2314	0.020	0.456	
Coarse Sand, Gravelly Sand	0.5	1875	2231	0.018	0.468	
Gravelly Muddy Sand	1	1836	2151	0.016	0.480	
Medium Sand	1.5	1767	1845	0.015	0.492	Sand, Sand99, SandAG
Muddy Gravel	2	1710	1615	0.013	0.505	
Fine Sand, Silty Sand	2.5	1661	1451	0.011	0.517	Sand96
Muddy Sand	3	1620	1339	0.010	0.571	
Very Fine Sand	3.5	1586	1267	0.009	0.634	
Clayey Sand	4	1554	1223	0.007	0.696	SiltySand
Coarse Silt	4.5	1527	1194	0.006	0.757	
Sandy Silt, Gravelly Mud	5	1500	1168	0.005	0.450	
Medium Silt, Sand-Silt-Clay	5.5	1482	1149	0.005	0.244	
Sandy Mud	6	1481	1149	0.005	0.140	MudShells
Fine Silt, Clayey Silt	6.5	1479	1148	0.005	0.111	
Sandy Clay	7	1478	1147	0.005	0.088	
Very Fine Silt	7.5	1476	1147	0.005	0.071	
Silty Clay	8	1473	1146	0.005	0.059	SiltyClay, Clay96
Clay (all grades)	9	1470	1145	0.005	0.054	

Table C-3. A list of the parameter values used to simulate seabeds for BORIS simulations. The five representative seabeds used for effective beamwidth and attenuation testing are shown in bold. The HardFlat seabed is a construct for testing the effects of roughness, seen in Figure 5-16.

BORIS Seabed Name	Sound Speed (m/s)	Shear Speed (m/s)	Density (kg/m^3)	RMS Roughness Height (m)	Exponent ( )	Attenuation (dB/kHz/m)	Volume RMS Height (m)	Horizontal Corr. Length (m)	Vertical Corr. Length (m)	Approximate APL Sediment Name
HRHB_AG	3750	100	2500	0.250	4.00	0.100	0.010	0.30	0.06	Rough Rock
LRHB_AG	3750	100	2500	0.050	3.75	0.100	0.010	0.30	0.06	Rock
HardFlat	4500	0	2500	variable	3.25	No Volume	-	-	-	Rock
Gravel	**2006**	**0**	**2492**	**0.028**	**3.00**	**0.456**	**0.040**	**0.30**	**0.10**	Sandy Gravel
Gravel96	2000	0	2000	0.080	3.00	0.063	0.040	0.40	0.15	Sandy Gravel
Sand	**1767**	**0**	**1845**	**0.020**	**3.25**	**0.492**	**0.020**	**0.15**	**0.03**	Medium Sand
Sand99	1767	0	1845	0.010	3.25	0.492	0.020	0.10	0.02	Medium Sand
Sand96	1720	0	1900	0.020	3.00	0.113	0.004	0.50	0.20	Fine Sand
SandAG	1780	100	1900	0.015	3.30	0.492	0.020	0.15	0.03	Medium Sand
SiltySand	**1554**	**0**	**1223**	**0.007**	**3.25**	**0.696**	**0.015**	**0.10**	**0.02**	Clayey Sand
MudShells	**1515**	**0**	**1225**	**0.020**	**3.25**	**0.140**	**0.020**	**0.20**	**0.05**	Sandy Mud
SiltyClay	**1473**	**0**	**1146**	**0.010**	**3.25**	**0.059**	**0.010**	**0.05**	**0.01**	Silty Clay
Clay96	1496	0	1550	0.020	3.00	0.063	0.003	0.40	0.08	Silty Clay

Appendix D. Scattering Theory for Single-Beam Echosounders

D.1 Volume Scattering Theory

The contribution of scattering within the seabed volume to seabed echoes is usually less than 10% (by total amplitude) for SBES[65]. Because of this, volume scattering is usually ignored. However, a significant fraction of the transmit pulse always passes into the seabed volume, as the normal incidence reflection coefficient rarely approaches a value of unity. Normal incidence reflection coefficients range between 0.4 and 0.8 for virtually all seabeds, including all the seabeds of appendix C. The reflection coefficient, sediment sound speed and attenuation together control how deep into the seabed the transmit pulse will travel as shown in Equation D-1 [Sternlicht, 1999], and plotted in Figure D-1.

$$l_p = \frac{10 + 10\log_{10}\left(1 - \Re_{\min}{}^2\right)}{\kappa f_c} \tag{D-1}$$

For Equation D-1, l_p is the penetration depth into the seabed at which the strength of the transmit pulse is reduced by 10 dB relative to its strength before interacting with the seabed (assuming accurate TVG is applied). $\Re_{\min}$ is the minimum value of the reflection coefficient within the beam of the SBES. Volume scattering will be significant when the penetration depth is sufficient to allow interaction with volume scatterers.

---

[65] As observed generally and in simulated echo times series generated by the BORIS software, which will be described in detail later in this chapter

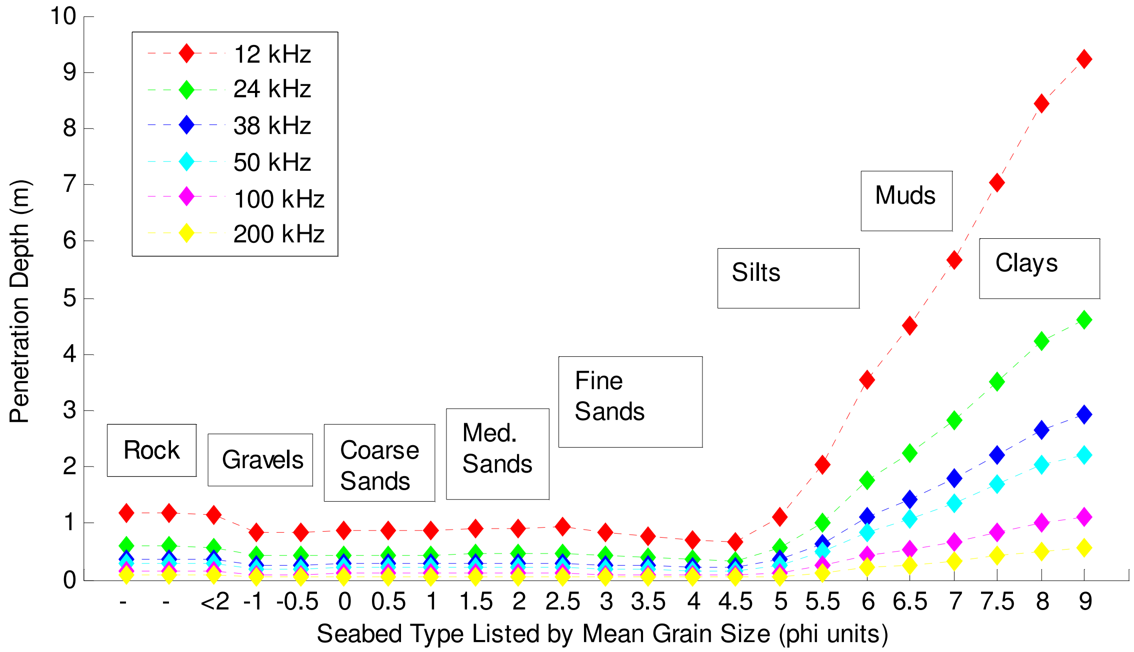

Figure D-1. 10 dB skin or penetration depths for the seabeds of the APL-UW handbook (Table C-2), based on the skin depth model (Equation D-1), for typical SBES frequencies.

A problem with attenuation is that field measurements cannot separate intrinsic attenuation (due to absorption of sound by the medium) from losses due to scattering by volume inhomogeneities. For this reason, attenuation is usually measured in seabeds that are homogenous and are not layered or bioturbated. The term homogenous refers to a consistent grain size distribution throughout. Mixing in pockets of different size sediment grains, or even large individual grains (cobbles, boulders), would add volume inhomogeneities and increase volume scattering above that included in the measurement of attenuation. However, even in the most homogenous sediments, there is still some scattering – measured attenuation will always have some component of scattering. For volume scattering models, it is important to evaluate the relative strength of the scattering included in the attenuation; this has been done indirectly using models of sound propagation.

Sound propagation within the seabed is a topic of active research. Some models predict either a linear or non-linear relationship between sound absorption and frequency: visco-elastic theory (as mentioned in chapter 1) and Buckingham's grain-shearing theory

[Buckingham, 2005] predict a strictly linear dependence, while Biot theory [Biot, 1956b] predicts a non-linear dependence. Buckingham's newer pore fluid viscosity grain-shearing model [Buckingham, 2007] predicts a non-linear dependence for frequencies of 1 to 10 kHz, and then approaches the grain-shearing model asymptotically for frequencies greater than 10 kHz. Both of Buckingham's models indicate a linear dependence of absorption on frequency for frequencies ranging from 10 to 400 kHz, which is almost precisely the range of frequencies applicable for this dissertation. These models are supported by data from the SAX99 experiment [Buckingham, 2005; Buckingham, 2007; Thorsos et al., 2001].

Since field measurements of attenuation contain both absorption and losses due to scattering, the above models of absorption should correspond to lower bounds of the measured attenuation values. In [Buckingham, 2005], the grain-shearing model fits one set of measurements very well indicating that scattering is a small, almost negligible, component of attenuation. Another result in the same paper, but in a completely different experimental setup, shows the model matching the lower bound of the data, with some substantially larger values caused by scattering. In the volume scattering model applied in [Jackson et al., 1986], the scattering component of attenuation is assumed to be a small proportion of the overall attenuation – the ratio of volume scattering cross-section to the attenuation is set low enough to avoid multiple scattering (0.004 $c\kappa$). Judging by the fits of the model to the data, this was found to be on the upper limit of the data, while modifying the ratio to optimize the results generated ratios of the volume scattering cross that ranged from 0.002 $c\kappa$ for bioturbated seabeds down to 10^{-4} $c\kappa$ for uniform sediments. Jackson et al. also found that the scattering coefficient was proportional to frequency implying varying sizes of scatters, and the amount of scattering was largely independent of the overall attenuation, i.e. variables such as bioturbation or the presence of volume inhomogeneities (layers, pebbles, etc.) could increase the volume scatter independent of seabed type [Jackson et al., 1986].

Even though volume scattering loss is generally small compared to absorption loss, this does not mean the fraction of the transmit pulse backscattered is insignificant. As mentioned, there are several mechanisms that enhance backscattering over the intrinsic scattering of the uniform sediment grain matrix, such as layering (i.e. sediment laminae) and individual scatters. Volume scattering theories are less developed than surface scattering theories, and have only recently begun to account for the effects of layers and scatterers within the seabed. There are two main approaches. The older empirical models calculate the reverberation from units of volume based on a volume scattering cross-section that is due to volume in homogeneities [Jackson et al., 1986]. In [Sternlicht, 1999], the volume scattering cross-section is a fixed ratio of the attenuation – 0.004 as discussed above. This model also assumes isotropic (non-directional) scattering. A newer, more complex, model considers scattering by individual scatters which are modelled as impedance differences within the volume, [Pouliquen et al., 1999]. The implementation is the same as surface scattering; it uses the tangent-plane approximation which is explored in the surface scattering sections later in this chapter.

D.2 The Helmholtz-Kirchhoff Integral Equation

The Helmholtz-Kirchhoff (HK) integral equation is a formulation of the boundary-value problem that addresses the problem of wave scattering from random rough surfaces. Therefore, understanding the HK equation is essential. General solutions to the three-dimensional acoustic wave-equation (Equation D-2) are constrained on the boundaries by the HK equation. The method of solution is to build solutions of the wave equation into the HK equation so that solutions of the HK equation are also solutions of the general wave equation. The general acoustic wave equation is as follows, with p as the acoustic pressure field:

$$\nabla^2 p(\mathbf{r},t) = \frac{1}{c^2}\frac{\partial^2 p(\mathbf{r},t)}{\partial t^2} \tag{D-2}$$

The derivation of the HK equation presented here is based on the material found in chapter 7.2 of [Medwin and Clay, 1998] and [Ogilvy, 1992]. The development starts by deriving Green's Theorem from Gauss's Theorem, which is an expression of conservation that the acoustic field produced inside a region is equal to the total acoustic field leaving a surface containing the region. The HK equation is arrived at through a number of steps, beginning by defining the region as not having any sources or sinks of acoustic waves. Depending on the definition, one can arrive at one of two expressions, based on whether the scattering surface is interior to the region, or exterior to the region. The different geometries are for reflection and transmission, respectively. The interior case is shown as Equation D-3; however, both expressions have the same basic form, although there are some differences[66]. Following Medwin and Clay, the Helmholtz-Kirchhoff integral equation for the scattered field, p_{sc}, is:

$$p_{sc}(Q) = \frac{1}{4\pi}\int_S \left[p_{sc}\frac{\partial}{\partial n}\left(\frac{e^{-ikR}}{R}\right) - \left(\frac{e^{-ikR}}{R}\right)\frac{\partial p_{sc}}{\partial n}\right]_S dS \tag{D-3}$$

where R is the distance between the point Q and the surface S with partial derivative with respect to the normal to the surface. The surface is closed and is partially made up of the scattering surface. Equation D-3 is the starting point for virtually all models of seabed echoes, including the two most relevant to SBES: [Sternlicht, 1999] and [Pouliquen, 1999].

The HK equation is basically the combination of Green's theorem, the acoustic wave equation and the assumption of no sources or sinks. The common thread between many theories of wave scattering is they are various approaches to solving the boundary-value problem posed by the HK integral equation – that is, one needs information about the boundary to solve the HK integral equation. Exact solutions to the HK equation are often categorized as integral equation theory, which includes limited special cases and numerical results. All other theories are approximations and may be best thought of as theoretical scattering models. Theories based on the HK equation that are not covered in this chapter are listed here with references:

- [Ogilvy, 1992] : integral equation theory, the extinction theorem, some extensions of Perturbation and Kirchhoff Approximation theories, multiple scattering theories

---

[66] Here there are some differences between the two references used. In the case of [Medwin and Clay, 1998], one of the acoustic fields is defined as the solution of the spherical wave equation, while [Ogilvy, 1992] leaves the field undefined. Medwin and Clay also assume the incoming and scattered fields are CW, but are separated in time so as not to occur at the same space and time. This allows them to separate the total field or incoming field from the scattered field leaving an equation for the scattered field only.

- [Medwin and Clay, 1998] : various extensions of Perturbation and Kirchhoff Approximation theories
- [Jackson et al., 1986; Sternlicht, 1999] : composite roughness model

General Perturbation theory (including Rayleigh theory) and general Kirchhoff Approximation (KA) theory will be summarized in the following three sections. These models are the basis of the BORIS echo simulating model and software that appears in section 4.5. A summary of the implications of scattering theory follows in section 4.7. chapter 5 presents an analytic echo time series model based on the KA theory.

D.3 Rayleigh Theory

Rayleigh originally developed his theory of scattering from periodic slightly rough surfaces in 1907, [Rayleigh, 1907]. The theory is presented in terms of acoustics in Rayleigh's famous book “The theory of Sound,” [Rayleigh, 1945]. The theory was later extended to arbitrary angles of incidence by LaCasce and Tamarkin, [LaCasce and Tamarkin, 1956], and subsequently extended to the more seabed-representative case of a random rough surface by Rice, [Rice, 1951]. There are also numerous contemporary papers discussing the topic.

For the case of a periodic rough surface, Rayleigh was able to derive an expression for the scattered field, p^{sc}. This theory for periodic rough surfaces converges for ‘slightly’ (to be defined later) rough surfaces. Multiple scattering effects are not included due to the assumption that the only outgoing waves are those that are singularly scattered. In this simplistic case, the ‘corrugated’ surface is defined by:

$$h(x) = a_0 \cos Kx \tag{D-4}$$

where h is the surface height along the horizontal co-ordinate x, with amplitude of a_0 and a spatial wavenumber K which has a corresponding wavelength of Λ. This wavelength is the Rayleigh wavelength that also occurs in perturbation theory as will be shown in the next section.

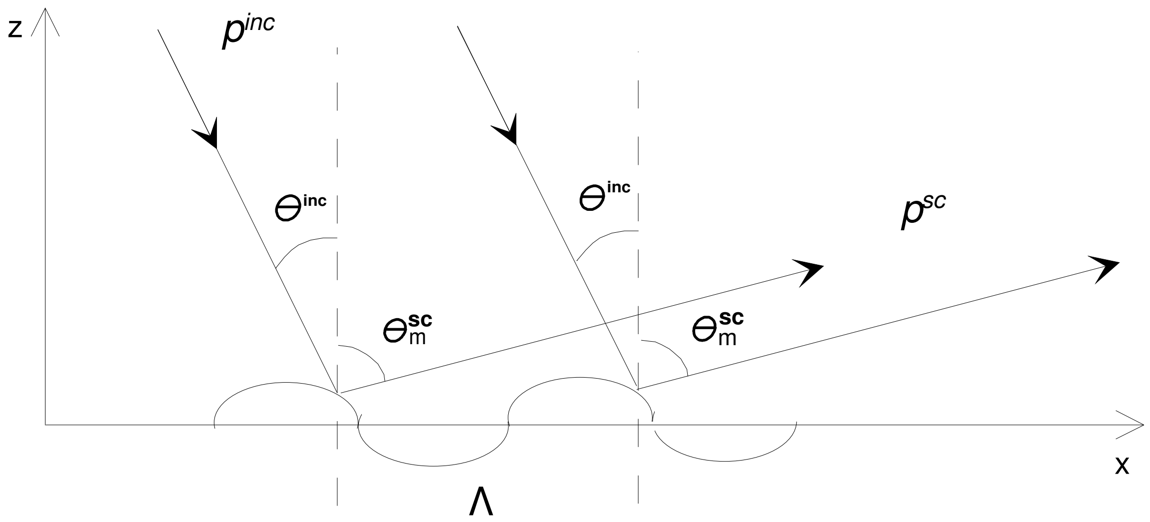

Figure D-2. Diagram for the derivation of Rayleigh scattering from a periodic rough surface.

To find the scattered field, p^{sc}, directly, consider a plane wave incident on the periodic surface of Equation 4-11 as shown in Figure D-2. Pairs of rays from the incident wave that are a Rayleigh wavelength apart horizontally will constructively interfere for a specific combination of incident and reflected angles. Non-Rayleigh wavelength ray pairs will have various small heights that introduce phase differences so will be ignored. However, for a regular surface, all rays have Rayleigh pairs. For constructive interference, the standard criterion is used: the difference of the path length between two waves/rays is an integer number of acoustic wavelengths. The relationship between the incident angle, the wavelengths and the scattering angle that results in constructive interference is found by geometry (based on Figure D-2). The condition for constructive interference is:

$$\sin\theta_m^{sc} = \sin\theta^{inc} + m\frac{\lambda}{\Lambda} \quad (\mathrm{m} = 0, \pm 1, \pm 2, \ldots) \tag{D-5}$$

By summing over all grating indices, m, all the reflected rays are included, including those from the various surface heights. This assertion holds if the surface heights are small compared to the path length differences, which are of a similar scale to the Rayleigh wavelength. The path length difference is then due to the various small slopes (or incident angles) of the periodic surface, instead of the heights so that the effect of the surface height variation is ignored. The scattered wave field is then:

$$p^{sc} = -\sum_{m=-\infty}^{\infty} A_m e^{i\mathbf{k}_m^{sc} \bullet \mathbf{r}} \tag{D-6}$$

where:

$$\mathbf{k}_m^{sc} = k\left(\hat{\mathbf{i}} \sin\theta_m^{sc} + \hat{\mathbf{k}} \cos\theta_m^{sc}\right) \tag{D-7}$$

The scattered field is defined in terms of the scattering directions $\mathbf{k}_m^{sc}$ (derived from Equation D-5) and coefficients A_m. To find A_m, boundary conditions are required. A_m corresponds to plane wave scattering coefficients for diffraction gratings. A variety of boundary conditions and approaches has been used to find the unknown coefficients in Equation D-6. As a demonstration of the most common approach, known as the Rayleigh-Fourier method, the Dirichlet boundary condition for the total field is assumed: $p(\mathbf{r})|_{z=h} = 0$. Since the total field is the sum of the incident and scattered fields and the incident field is just a plane wave: $\exp(i\mathbf{k}^{inc} \bullet \mathbf{r})$, the Dirichlet boundary condition is used to equate the scattered field to the incident field:

$$e^{-ikh(x)\cos\theta^{sc}} = \sum_{m=-\infty}^{\infty} A_m e^{i\left(mKz - kh(x)\cos\theta_m^{sc}\right)} \tag{D-8}$$

The resulting expression is then solved for A_m in one of two ways. One is to expand both sides in a Taylor series in terms of h and equate terms of the same order of h – this is analogous to Perturbation theory that will be discussed later. The other method is known as the Rayleigh-Fourier method. It involves expanding both sides as a Fourier series and equating coefficients of equal components. This derivation is the same as for optical diffraction gratings: this is why A_m is called the diffraction grating scattering coefficient. The result is a series of Bessel functions that quickly converge for the condition of a ‘slightly’ rough surface: $ka_0 << 1$. This series then results in a set of simultaneous linear equations that may be solved for A_m, which ends up having no angular dependence. The physics expressed by Equation D-5 contains the angular dependence. The scattered field becomes a function of the *difference* in wavenumber of the surface undulations (K) and incident plane wave (k); see the right side of Equation D-8. This result is a frequency-shift of the scattered field.

Rice extended Rayleigh theory to random rough surfaces, but did so for polarized electromagnetic waves on a two-dimensional surface [Rice, 1951]. The theory is readily adaptable to acoustics. The result obtained for the average field intensity is actually the same result that one can derive with second order Perturbation theory (shown later), which is also very similar to Bragg scattering theory. All of these similar theories address the same problem – that of scattering from surfaces with very small scale roughness heights such as diffraction gratings and crystal lattices. Rice's results are valid for surfaces where the correlation length, λ_o, is much greater than the incident wavelength: $\lambda_o >> \lambda$. The conditions of applicability basically ensure that the heights are small enough to avoid shadowing and multiple scattering. Also of note, these theories are for general scattering, not exclusively *back*scattering.

D.4 Perturbation Theory

Perturbation theory, like almost all scattering theories, aims to solve the boundary value problem posed by the Helmholtz-Kirchhoff integral equation (Equation 4-10). To solve the HK equation, boundary conditions and definitions of the surface/boundary are needed. In most situations, the exact surface geometry is not known, only its bulk statistics such as RMS height, h_{rms}. Then, an appropriate approximation to make is to define the surface as the mean plane, $z=0$, and find a new boundary condition to account for the height deviations from the mean surface. To accomplish this, the acoustic field is perturbed by the height deviations, h, a known boundary condition is applied and then a new boundary condition is found that can be used to solve the Helmholtz-Kirchhoff integral equation for the total or scattered field. That is the overall approach of Perturbation theory. It was initially applied to acoustics by Brekhovskikh in 1952. An overview is presented here, mainly based on [Ogilvy, 1992].

If the surface height function, *h(x,y)*, about a mean surface plane of *z=0* is small relative to the acoustic wavelength and has small gradients or local slopes, then the total acoustic field near *z=0* may be expanded as a Taylor series:

$$p(x,y,z) = p(x,y,0) + h(x,y)\frac{\partial}{\partial z}p(x,y,0) + \frac{h(x,y)^2}{2}\frac{\partial^2}{\partial z^2}p(x,y,0) + \ldots \qquad \text{(D-9)}$$

The solution for the HK equation is written as sum of terms of order h^n. The explicit conditions for the validity of such an expansion are given by:

$$\begin{aligned} k\left|h(x,y)\right| &<< 1 \\ \left|\nabla h(x,y)\right| &<< 1 \end{aligned} \qquad \text{(D-10)}$$

Initial boundary conditions generally fall under two categories, Dirichlet (known field values) and Neumann (known field derivative), for example:

$$\begin{aligned} p(\mathbf{r})\,|_{z=h} &= 0 \\ \frac{\partial p}{\partial n_0}\,|_{z=h} &= 0 \end{aligned} \qquad \text{(D-11)}$$

where n_0 is the unit surface normal. Both have physical interpretations: Equation D-11a is the boundary condition for a free, soft or pressure release surface such as the ocean-atmosphere surface (a 'hard' surface such as the seabed would have high pressure), while Equation D-11b is the boundary condition for a rigid boundary such as the seabed surface (a soft boundary such as the ocean-atmosphere surface allows particle motion).

D.4.1 First-Order Perturbation Theory

The first-order solutions for the Dirichlet boundary condition of Equation D-11a and for the Neumann boundary condition of Equation D-11b are very similar. The incident and scattered fields are expanded as a series and the boundary conditions are applied, then terms of equal order of h^n are equated essentially defining new boundary conditions relating the incident and scattered fields. The zeroth-order scattered field for both turns out to be the reflection of the incident field, or the solution to the smooth half-space scattering problem satisfying either condition. By conservation of energy, the first-order correction to the scattered field should then be zero. To find the first-order correction, a solution to the HK equation is sought.

For the Dirichlet boundary condition, the solution for the first-order correction is a function of the surface heights. Since the mean value of surface heights over the surface is zero, the mean first-order correction for all **r** is zero. Appendix E shows that such a mean is the coherent field - there is no effect of roughness on the coherent field, and the coherent field conserves energy. However, there is a first-order contribution to the diffuse field (the zeroth-order diffuse field is zero). Equation D-12 is the total first-order correction intensity, which is only the diffuse field intensity as the first-order coherent field and its intensity are zero:

$$\begin{gathered}\left\langle I_1^{sc}(\mathbf{r})\right\rangle = \frac{4k^4\cos^2\theta_1\cos^2\theta_2}{r^2}A_m P(kA,kB) \\ A = \sin\theta_1 - \sin\theta_2\cos\theta_3 \text{ and } B = -\sin\theta_2\sin\theta_3\end{gathered} \tag{D-12}$$

where θ_1 is incident angle of the incident field, while θ_2 and θ_3 specify the scattered field directions (in a 3-D sense so θ_2 is the scattering elevation and θ_3 is the azimuth). A_m is the area of the mean plane $z=0$ and $P(s_1,s_2)$ is the surface height power spectrum.

The first order perturbation theory result for the Dirichlet boundary condition does not conserve energy, but does again demonstrate a frequency dependence on the scattered intensity. The arguments of the power spectrum of Equation D-14 are the components of the change of the wave vectors due to change in direction. The strength of surface roughness at these frequencies contributes to the directionality of the scattered intensity. In a sense, this is an expression of the ability of the surface to ‘perturb’ the reflected field. This is similar to the selective scattering process which will be seen in Kirchhoff Approximation theory in the next section.

Evaluation of perturbation theory for the Neumann boundary condition follows the same approach with marginally different results. The zeroth-order solution is again the reflection of the incident field, or the solution to the smooth half-space scattering problem satisfying Neumann boundary conditions. The HK equation for the first-order correction contains gradients of the surface height function unlike the Dirichlet boundary condition. However, it is directly proportional to the surface height distribution, like the

Dirichlet problem, which again leads to the first order correction to the coherent scattered field being found to be zero. The Neumann boundary condition first-order correction to the total field intensity is:

$$\left\langle I_1^{sc}(\mathbf{r})\right\rangle = \frac{4k^4(1-\sin\theta_1 \sin\theta_2 \cos\theta_3)^2}{r^2} A_m P(kA, kB) \qquad \text{(D-13)}$$

with terms defined as in Equation D-12. Although the formulation and derivation is more complicated for the Neumann boundary condition than for the Dirichlet case, the result is very similar. The physical interpretation is also the same: the coherent scattered field is unchanged due to roughness, while the first-order correction adds a diffuse field due to roughness that does not satisfy conservation of energy. In Equation D-13, one can see the 'selective scattering' effect in the surface height power spectrum, indicating that scattering is a function of incident pulse frequency.

D.4.2 Second-Order Perturbation Theory

Second-order Perturbation theory is difficult to solve, but there are some very useful results. A paper by Kuperman is a good example [Kuperman, 1975]. The most important result is that the second-order coherent field is generally not zero. There is also an alternative to the approach presented in the previous paragraphs. An approach that yields solutions is to express the total field as the sum of the coherent and diffuse fields, instead of a sum of the incident and scattered fields:

$$p(\mathbf{r}) = \langle p(\mathbf{r})\rangle + p_d(\mathbf{r}) \qquad \text{(D-14)}$$

Solving the HK equation for a first-order Dirichlet boundary condition for the average total field intensity, based on Equation D-14, actually leads to a second-order (h^2) integral equation for the coherent field. The solution for the coherent field is the perfectly reflected incident field reduced by a coefficient:

$$V = 1 - 2k^2 h_{RMS}^2 \cos^2\theta_1 \quad \text{for} \quad k\lambda_0 >> 1 \qquad \text{(D-15)}$$

The diffuse field can also be found with this approach. The result for the diffuse field intensity is also a function of the coefficient, *V*, and the surface power spectrum as in Equation 4-19. The coherent field is no longer a complete reflection of the incident field, while the diffuse field combines the aforementioned ‘selective scattering’ process with the same coefficient that responds to roughness.

This important result for second-order perturbation theory, which is consistent with Rayleigh theory, shows that as roughness increases, the coherent field decreases and the diffuse field increases. The same is true for increasing centre frequency. This implies that as roughness and frequency increase, the echo variability will increase. For transmit pulses of finite bandwidth, there will also be a frequency shift within the spectral band. The frequency of the echo would likely be lower than that of the transmit pulse, but not necessarily so as selective scattering will affect the echo as well as the 'reflection' coefficient of Equation D-15. These relationships between the seabed surface and the echo may be exploited for seabed classification and characterization.

One of the issues of perturbation theory is the limited range of applicability. Although higher-order theory handles shadowing and multiple scattering, the applicability is limited to the scales of roughness (h_{RMS}) less than the 0.3 of the wavelength. The criteria on the correlation length are somewhat inconsistent. Because of the ability to handle shadowing and multiple scattering, Perturbation theory is most useful as applied to high or grazing angles of incidence. Recent theoretical developments, [Thorsos and Borschat, 1995], have developed perturbation theory based corrections for Kirchhoff Approximation theory with the benefits of wide applicability and the ability to handle grazing incidence. This theory is applied in the BORIS model presented in section 4.5.

D.5 Kirchhoff Approximation Theory

Like Perturbation theory, Kirchhoff Approximation (KA) theory aims to solve the HK equation (Equation D-3) for the acoustic field by applying a simplifying boundary condition. The simplification is that the reflection coefficient can be applied at every point along the surface to relate the incident field to the scattered field:

$$\begin{aligned} p^{tot}(\mathbf{r}) &= [1+\Re(\mathbf{r})]p^{inc}(\mathbf{r}) \\ p^{sc} &= \Re(\mathbf{r})p^{inc}(\mathbf{r}) \end{aligned} \tag{D-16}$$

This is the tangent plane approximation. The scattered field arises from reflections from small planes, or *facets*, tangent to the surface. As the surface becomes rougher, the tangent planes become steeper, deviating more from horizontal, directing more rays in wider directions. The vertical offsets of the planes also increases causing phase shifts that lower the coherent component while raising the incoherent component. A simple analogy is a disco ball with its array of small mirrors.

In Equation D-16, the reflection coefficients are dependent on position and, more importantly, angle. In some cases, the HK equation can still be solved, however, it is much easier to solve if one approximates the problem by assuming a constant reflection coefficient. This additional simplification is valid when the distances from the source and receiver are large compared to the wavelength (for plane waves), and when the angle of incidence is not near critical so that the reflection coefficient is not varying widely. This approximation is best for the source and receiver in specular orientation, such as a normally incident monostatic SBES. In that case, the sections of the surface that will backscatter are facets oriented normal to the incoming rays, while sound scattered from non-normal aspects of the surface does not contribute to the echo. This means one can apply the perpendicular reflection coefficient, Equation 4-1, as a constant, thereby removing this complication from the problem. With this approximation, the problem becomes one of handling the geometry and the specifics of the seabed surface.

The assumption of a constant reflection coefficient is generally valid for the case that concerns this dissertation: nadir incidence and small surface gradients (as described below). In cases where the assumption is not valid, a variable reflection coefficient can be evaluated analytically for certain cases (not shown) and numerically for any others. A good example of a numerical approach that applies a variable reflection coefficient is the BORIS model [Pouliquen et al., 1999]. However, some of the criteria of validity are the same as those for the full normal facet assumption, (i.e. those concerned with shadowing and multiple-scattering).

The validity of the Kirchhoff Approximation is dependent on how planar each surface facet is. The plane on which the reflection coefficient is applied must be longer than the incident wavelength (Equation D-17a), and must not deviate from planar by much more than a fraction of the wavelength (Equation D-17b). The first criterion is stricter than the second criterion:

$$\begin{aligned} kr_c \cos^3 \theta_{inc} &>> 1 \\ kr_c \cos \theta_{inc} &>> 1 \end{aligned} \qquad \text{(D-17a,b)}$$

where r_c is the local radius of curvature and θ_{inc} is the angle of incidence. The radius of curvature is:

$$\begin{aligned} r_c^{\min} &\sim 0.1 \frac{\lambda_0^2}{h_{RMS}} \\ r_c^{median} &\sim 0.3 \frac{\lambda_0^2}{h_{RMS}} \end{aligned} \qquad \text{(D-18)}$$

From Equation D-17, the validity of KA theory improves as wavelengths shorten with higher frequencies.

For normal incidence, the angle of incidence should be taken as the SBES half-beamwidth, $\theta/2$, plus the foot-print scale seabed slope, ϕ. The derivation of Equation D-17 is based on the geometry shown in Ogilvy's book page 103 [Ogilvy, 1992]. Equation D-17 effectively limits the surface gradients and the angles of incidence for applicability.

There are numerous other conditions of validity, based on various arguments mostly due to shadowing and multiple-scattering effects. The best overall validity conditions are the ratios of the correlation length to wavelength and to RMS height (Equation D-19). These ratios express the radius of curvature conditions and the multiple-scattering / shadowing limits:

$$\begin{gathered}\frac{\lambda_0}{\lambda} \geq 1 \\ \frac{\lambda_0}{h_{RMS}} > 1\end{gathered} \tag{D-19}$$

To further simplify the problem and make it tractable, the far-field or Fresnel approximation is applied. Ogilvy uses the term far-field, although this approximation is not related to the normal far-field criteria [Ogilvy, 1992]. Instead, this approximation is that all distances and heights along the seabed are small compared to the distance, r, from the surface at which the fields are computed. For a surface of infinite dimensions, this would not be possible. However, virtually all incident waves have an angular extent such that their footprint on the scattering surface or seabed is limited to an approximate diameter $d_{footprint}$, which may be calculated by Equation 2-2. The far-field approximation is only valid if the following conditions hold:

$$\begin{gathered}r >> d_{fp} \\ r >> \frac{d_{fp}^{\;2}}{\lambda_0} \\ d_{fp}^{\;2} >> \lambda_0^2\end{gathered} \tag{D-20}$$

The third condition is necessary to ensure the applicability of the stochastic model of the seabed heights, especially if one were generating seabed heights numerically.

The general far-field KA theory result is found by combining the aforementioned approximations and further manipulations, which in limited circumstances can make the problem solvable. The full form is quite involved. One may find it in [Medwin and Clay, 1998], page 609 and in [Oligvy, 1992], page 80. For the purposes of this dissertation, the conditions of validity and an understanding of the approximations made are sufficient. Many researchers use the KA theory result to find a scattering cross-section to which they compare measured values, or they use the KA theory result to derive an apparent reflection coefficient for calculating bottom loss, etc. The KA theory result is also useful for scattering from the sea-air interface (this problem is much better understood than that of the seabed) and from a number of targets particularly cylinders and wedges. For this dissertation, the general KA result will be used as a basis for the BORIS model and a new analytical echo time domain model.

D.6 Special Tractable Cases of KA Theory

Although the general KA theory result is very difficult to solve analytically, there are several special cases that either have solutions or imply some useful results. First is the case of zero roughness. Next, the implications on the nature of the coherent and diffuse fields can be found without assuming anything about the scattering surface. Then, the surface is assumed to have a Gaussian probability height distribution, which allows solutions to be found for the average intensity of the scattering.

D.6.1 KA Theory with Zero Roughness

In the case of zero roughness, KA theory produces the specular reflection – exactly as one would expect. All of the approximations made are perfectly valid in this case. The scattered field is directly related to the incident field via the reflection coefficient, and the Snell-Descartes law which implies a change in direction. This reflection can be thought of as a virtual image behind the surface – the beampattern of the transmit pulse continues undisrupted.

D.6.2 The Coherent Field from General and Gaussian Roughness

Evaluating general KA theory for the mean or coherent scattered field results in the following equation:

$$\left\langle p^{sc} \right\rangle = \chi(kC) p_0^{sc} \tag{D-21}$$

where p_0^{sc} is the field reflected from a smooth surface of the same extent as the mean rough surface (i.e. the specular reflection) and $\chi(kC)$ is the one dimensional characteristic function of the surface (see section 4.2, evaluated for the change in wavenumber between incident and scattered directions in the z-directions). Equation D-21 is a general result that holds as long as KA theory is valid. In the case of a Gaussian height probability distribution, the characteristic function is defined as a Gaussian and Equation D-21 becomes:

$$\begin{aligned} \left\langle p^{sc} \right\rangle &= p_0^{sc} e^{-g/2} \\ g &= k^2 h_{RMS}{}^2 \left(\cos\theta_1 + \cos\theta_2\right)^2 \end{aligned} \tag{D-22}$$

where θ_1 and θ_2 are the incident angles of the incident field and reflected ray. $\sqrt{g}$ is known as the Rayleigh parameter such that the Rayleigh criterion defines the condition in which the coherent field is dominant over the incoherent field:

$$\sqrt{g} < \frac{\pi}{2} \tag{D-23}$$

When the Rayleigh criterion is valid, the surface is considered to be 'smooth', and the term $\Re e^{-g/2}$ can be applied as the coherent reflection coefficient [Lurton, 2002]. The definition of Equation D-23 stems from a geometrical derivation, which is described in the next section. The limit of applicability of Equation D-21 is the same as that for KA theory. The assumption of a Gaussian roughness probability distribution is also generally valid. One problem is that the coherent field can have contributions from large, normally orientated facets that are not considered in this model. These facets appear for significant roughness: $g > 1$. This non-specular coherent field can be seen to boost the coherent field (as seen in Figure 13.2.2 of [Medwin and Clay, 1998]) and is explored further in sub-section 5.4.1. Complications of shadowing and multiple scattering also affect the accuracy of this model. The literature considers that Equation D-22 breaks down when Equation D-23 no longer holds. However, it does continue to apply as the *specular reflection coefficient* to much higher levels of roughness, representing the specular part only, instead of the total coherent field. At these levels of roughness, non-specular coherent contributions become the dominate part of the total coherent field (this is explored further in section 5.4.1).

There are numerous papers that document the use of Equation D-22 to characterize the RMS roughness height. First, Medwin and Clay discuss applying the ratio of the coherent to diffuse energies (total intensity) found by fitting theoretical distributions to measured amplitude distributions [Medwin and Clay, 1998]. They apply the approach for g up to 25, which is substantially larger than the Rayleigh criterion. However, there is limited sensitivity for $g > 1$, and the approach quickly breaks down in the case of non-Gaussian roughness or shadowing. Second, the relative strength of the diffuse component also determines the ping-to-ping variability. Then, from Equation D-22, variability is directly related to roughness. Stanton and others apply a result similar to Equation D-22 that is also valid in the case of low roughness: $g < 1$ [Stanton, 1984]. The third approach is the frequency shift method of Berntsen et al. [Berntsen et al., 1999]. For an incidence pulse of a finite bandwidth, the coherent scatter will show a centre-of-mass frequency-shift related to the frequency dependence in Equation D-22.

As seen in sections 4-5 and 4-6, Rayleigh and Perturbation theory also show a frequency dependence of the scatter from a rough surface. However, Brekhovskikh and Lysanov derive a scattering coefficient (from KA theory) that is independent of frequency [Brekhovskikh and Lysanov, 1988]. Their result is based on Gaussian-distributed heights and Gaussian correlation functions. Heald used this result to provide a theoretical background for the RoxAnn seabed classification method [Heald, 2000]. Using the same surface definition, Thorne and Pace derived a similar frequency independent backscatter coefficient [Thorne and Pace, 1985]. This frequency independence is contrary to results and other theories as noted in [Lurton, 2002]. The reason for the discrepancy is that the theory does not separate coherent and incoherent scattering and assumes a monochromatic incident field (i.e. no bandwidth to shift). The frequency shift method of Berntsen et al. works because the coherent and incoherent fields are separated by direction – the incoherent field scatters away from the receiver leaving a frequency-shifted coherent field [Berntsen et al., 1999]. This assertion works better for low frequencies because incoherent scatter is less likely to dominate, whereas at high roughness relative to the acoustic wavelength, incoherent backscatter dominates.

D.6.3 The Geometrical Rayleigh Criterion, Bragg and Lambertian Scattering

In [Lurton, 2002] and [Rayleigh, 1945], one may find the original geometric derivation of the Rayleigh criterion of Equation D-23, and a similar solution for the coherent specular reflection given by Equation D-22. The major difference in this geometric Rayleigh theory result is that scattered angle is assumed to be equal to the incident angle, which removes the second cosine term and adds a factor of two. The derivation is based on the phase change between two parallel rays being less than $\pi/2$ when scattering from a random roughness surface, as depicted in Figure D-2 except that $\theta_{sc} = \theta_{inc}$. It is remarkable that the KA theory result for low for Gaussian-distributed roughness is exactly the same as the result from this geometric theory. The second-order perturbation result of appendix D.4.2 for the coherent field is also very similar.

Although this theory is a special case of Rayleigh theory presented in appendix D.4 it is presented here as it is part of a group of simple scattering models that include Bragg and Lambertian scattering. Lambertian scattering is an empirical model. It is based on the following premise: the incident intensity of a plane wave is proportional to the cosine of the angle of incidence and the radiated intensity is proportional to the cosine of the scattering angle (as can be shown by dividing the incident or scattering power by the ensonified area). Then, Lambert's law is that the backscattering cross-section is proportional to $\cos^2\theta_{inc}$. Many scattering theory results, including Rayleigh, Perturbation and KA theories, have an angular dependence that is Lambert's law or is at least a function of a power of a cosine term. For this reason, empirical descriptions of scattering are often based on cosine angular dependence. Bragg scattering is a geometric model that is most often applied in scattering from regular lattices. Familiar examples are X-Ray crystal diffraction (which led to the discovery of DNA for instance). The theory is that scatterers have to be located with a spacing d such that the backscattered field constructively interferes with itself; this requires a path-length difference between neighbouring rays of multiples of the wavelength. The rule is: $2d\sin\theta_{inc} = n\lambda$. This is effectively a selective scattering regime: certain wavelengths will be observed at certain angles depending on the nature of the surface roughness. More specifically, the spatial wavenumber responsible for scattering at an angle θ_{inc} is $2k\sin\theta_{inc}$. If one assumes a power-law surface height spectrum with an exponent of three, the backscattering strength will be a function of $\sin^{-3}\theta_{inc}$. For a fluid-fluid interface, the result from [Lurton, 2002], for Bragg scattering theory is:

$$\sigma_{Bragg}(\theta_{inc}) \propto k\frac{\cos^4\theta_{inc}}{\sin^3\theta_{inc}}\left|U(\theta_{inc})\right|^2 \tag{D-24}$$

where $U(\theta_{inc})$ is an intensity reflection coefficient that is also a function of the impedances of the two media. $\sigma_{Bragg}(\theta_{inc})$ is a scattering cross-section which can be taken as the ratio of the scattered to incident intensities. Bragg scattering has a relatively weak negative frequency dependence which is characteristic of incoherent fields. This is consistent with the frequency trends already seen since the Bragg scattering regime applies in conditions where the incoherent field dominates. Indeed, the Bragg scattering result is similar to Perturbation theory results for high relative roughness.

In all, Perturbation theory, Rayleigh theory, Bragg and Lambertian scattering apply to high angles of incidence where the seabed 'looks' like a diffraction grating. KA theory and the Rayleigh specular reflection criterion apply at near normal incidence where the seabed looks like broken mirror or a flat disco ball. The remainder of appendix D.5 discusses the remaining solutions of KA theory and the details of applying scattering theory to model seabed echoes from SBES.

D.6.4 The Incoherent Field from General Roughness

It is not possible to define an average incoherent field – such a field would be zero (see appendix E). However, the difference between the average total intensity and the coherent intensity can be called the average incoherent intensity. The form of the solution for KA theory for general roughness can be found in [Ogilvy, 1992], which is not duplicated here as only the implications are important. The major implication is that the frequency dependence of the incoherent field is positive – higher frequencies contribute more to the incoherent field than do lower frequencies. The exact dependence depends on solving an integral involving the characteristic function.

D.6.5 The Total Mean Intensity in The Case of Gaussian Roughness

The total mean intensity for a plane wave scattered from a surface with a Gaussian height probability distribution and a Gaussian characteristic function is:

$$\langle I \rangle = \langle p^{sc} p^{*sc} \rangle = I_0 e^{-g} + \langle I^{dif} \rangle \qquad \text{(D-25)}$$

where *I₀* is the intensity reflected from a smooth surface (all the results in this sub-section are from [Ogilvy, 1992]). The coherent intensity is well defined; however, the mean diffuse intensity is difficult to evaluate because the integral of the surface characteristic function is improper in general. Nevertheless, solutions can be found with additional assumptions based on the three levels of roughness: slight ($g << 1$), moderate ($g \approx 1$) and high ($g >> 1$).

For slightly rough surfaces, the total mean intensity is:

$$\langle I \rangle = I_0 e^{-g} + \frac{k^4 C^2 F^2}{r^2} A_m P(kA, kB) \qquad \text{(D-26)}$$

where *A,B,C* are wavevector components and *F* is an angular component related to *A,B,C*. The term $P(kA, kB)$ is the selective scattering term that was presented and discussed in Appendix D.4 with Perturbation theory. Equation D-26 is also similar to the results of first-order Perturbation theory. Note that Equation D-26 is also frequency dependent in both the coherent and incoherent parts because of terms of *k*. In addition to the selective scattering process, the incoherent field undergoes a positive frequency shift, while the coherent component undergoes a negative shift relative to the transmit pulse: the bandwidth of the transmit pulse is increasingly divided between the coherent (lower band) and diffuse (upper band) components with increasing roughness. This division of the energy in the band is logical as it complies with the conservation of energy. This rule holds while there is a coherent component of the echo: slight to moderate roughness.

For moderately rough surfaces, the total mean intensity is as in Equation D-25; however, the mean diffuse intensity cannot be found explicitly but is approximated by:

$$\frac{k^2 F^2 \lambda_0^2 g e^{-g}}{4\pi r^2} A_m \exp\left(-\frac{k^2\left(A^2+B^2\right)\lambda_0^2}{4}\right) < \left\langle I^{dif} \right\rangle \le \frac{k^2 F^2 \lambda_0^2}{4\pi r^2} A_m \qquad \text{(D-27)}$$

Equation D-27 is an expression for the bounds on the mean diffuse intensity and is valid for g of order unity: $0.5 < g < 5$. This scattering regime is common for typical echosounders of interest such that the diffuse component of the echo dominates. However, there is still a marginally significant coherent component. The frequency dependence of the diffuse field is not as strong as it is for the slight roughness regime.

For highly rough surfaces, the coherent field is vanishingly small and is taken to be zero. Then the total mean intensity is:

$$\langle I \rangle = \left\langle I^{dif} \right\rangle = \frac{k^2 h_{RMS}{}^2 F^2}{r^2}\left(\frac{1}{g}\right) A_m P_{12}\left(\frac{A}{C}, \frac{B}{C}\right) \qquad \text{(D-28)}$$

where P_{12} is the single point joint probably distribution for the two-dimensional surface gradients. The term P_{12} is a selective scattering regime that is based on the surface gradients and not on the surface heights, as was the case in slight to moderate roughness. Selective scattering does occur for the moderate roughness diffuse component although it is not apparent from Equation D-27 because only the lower limit of Equation D-27 has a spectral term that resembles selective scattering. The most important complication of the high roughness case is that there is no overall frequency dependence (k^2 is cancelled out by ${}^{1}\!/_{g}$). Since there is no coherent component to partition lower frequencies into, there can be no centre-of-mass frequency shift. Ogilvy reports that many papers call this high roughness regime the ‘deep-phase modulation’ regime [Ogilvy, 1992]. Large phase shifts occur because of the large height variation of the surface relative to the wavelength. Large height variations also cause scattering is dominated by large facets. This is seen in Figure 4 of [Preston and Biffard, 2007] as strong speckle-like features in the echogram.

D.6.6 KA Theory Extensions Required to Model SBES

The KA and Perturbation theory presented in previous sections 4.7 and 4.8 are for planar incident fields. The major difference between this and the real fields produced by SBES are that real fields have a time dependence (a transmit pulse duration) and a finite aperture (a beamwidth).

The problem with directivity is most easily handled by inserting a Gaussian directivity function into the field definitions when solving the HK equation. A Gaussian directivity function is the most tractable modification. As noted in chapter 2, beampatterns are Bessel functions, and they have sidelobes. A Gaussian function is an adequate approximation for the main lobe of the beampattern.

There are a number of approaches to handle the time dependence. For specific time dependent solutions, the surface heights must be known perfectly. Such a solution, if possible, would give an analytic representation of a single echo. Then, many iterations would be required to find the coherent and incoherent fields. Surface height fields can be generated stochastically to match given bulk parameters. Given a high enough resolution, the scattered field can be calculated at each surface element, propagated to the receiver and summed – this is the basis of the BORIS numerical model. When it is preferable to work with the bulk characteristics of the seabed to produce analytic results, KA theory solutions are possible – they are outlined in [Ogilvy, 1992].

An alternate approach is to view scattering as a convolution of the transmit pulse with the impulse response of the surface. The impulse response is the Fourier transform of the frequency response of the surface. For the coherent field, the frequency response is a function of $e^{-g/2}$. Obtaining the analytic result for the echo time series is then straightforward, but is only a simplistic model of the coherent component of the echo. In his thesis, Park used this approach to attempt isolate the impulse response of the seabed for the purpose of improving QTC seabed classification [Park, 2003]. It was later demonstrated that this convolution approach is comparable to simulation results from the original version of the BORIS model [Zielinsky et al., 2004].

A more conventional approach is to apply the backscatter coefficient determined by KA or Perturbation theory within a propagation model. The propagation model accounts for the time dependence by modelling the movement of the incident beam's leading and trailing edges sweeping outward from the centre of the footprint. This requires defining the sweep or elevation angle as a function of time. The resulting ensonified area, defined by the area between the leading and trailing edges, is the scattering surface (and volume) over which the backscatter is summed by integration. Integration leads to a function of time. In special cases, this is solvable analytically as will be seen in sections 4.9 and 5.3.

A more general approach is to do the integration numerically. The results by Sternlicht are the best recent example [Sternlicht, 1999]. The propagation model for the intensity is:

$$I(t) = \int_{\varphi=0}^{2\pi} \int_{\theta=0}^{\theta(t)} d^2 I(\theta,\varphi) \qquad \text{(D-29)}$$

In the co-ordinate system of Sternlicht [1999], θ and φ are elevation and azimuth angles, while the surface scattering intensity contribution, $d^2 I(\theta,\varphi)$, is:

$$d^2 I(\theta,t) = I_{transmit}\left(t - \frac{2R}{c}\right) \frac{S(\theta)\beta^4(\theta,\varphi)}{R^4 10^{\alpha_w \frac{R}{5}}} dA \qquad \text{(D-30)}$$

where $S(\theta)$ is the angular dependent backscatter coefficient found by KA theory, R is the range as before, α_m is the water column attenuation, and $\beta(\theta,\varphi)$ is the directivity function (assumed the same for transmit/receive). The angular time dependence $\theta(t)$ is found by a geometric argument. The directivity function accounts for the beampattern and for the incident angle due to slope or tilt by numerical projection of the beam. The integral of Equation D-29 is discretized and computed numerically.

Appendix E. Field Definitions

First, some terminology: the terms field, amplitude and pressure are used interchangeably in some instances, but are distinct. An acoustic field includes information about the direction, frequency, phase and amplitude of an acoustic wave. Pressure is the value of the field at a point, so a lowercase p represents the acoustic field. Intensity is the squared-pressure divided by the acoustic impedance (density multiplied by sound speed). Power is the intensity multiplied by the ensonified area, while energy is the power multiplied by the time. Energy should not be confused with the sum of digitally sampled amplitude over time; one may call that value the cumulative amplitude which is only used as an echo feature.

The *coherent field* is defined as the pressure amplitude field averaged over many instances of the scattering surface so that phase fluctuations in the received field average out. The coherent field is usually the specular reflection. However, at higher degrees of roughness (several times the Rayleigh criterion roughness), the coherent field may include some non-specular contributions. A 'stack' of full waveform echoes generated by CW (continuous wave) transmit pulses approximates the coherent field, especially for a large number of echoes. BORIS results indicate that a stack of 30 echoes is sufficient (a stack is the aligned time-average). In theoretical derivations, the averaging shown below is evaluated by applying the average values of the surface statistics. The coherent field and its intensity, as one might measure by a stack of echoes or derived theoretically, is defined by:

$$\left\langle p^{coh} \right\rangle = \left\langle p^{sc} \right\rangle \tag{E-1}$$

$$\left\langle I^{coh} \right\rangle \propto \left\langle p^{sc} \right\rangle^2 = \left\langle p^{sc} \right\rangle \left\langle p^{*sc} \right\rangle \tag{E-2}$$

where the angle brackets indicate a time-average, p^{sc} is the scattered field and the star indicates complex conjugate (as per convention). Conversely, the average total field intensity is:

$$\left\langle I^{tot} \right\rangle \propto \left\langle p^{*sc} p^{sc} \right\rangle \tag{E-3}$$

With real data, the total field intensity is found by the square of the Hilbert transform. The Hilbert transform essentially multiples a signal by its complex conjugate as per Equation E-3. The total echo amplitude is the absolute value of the Hilbert transform, and this is the standard echo envelope or analytic signal that is used throughout this dissertation (most echosounders de-modulate the signal with a rectify and filter approach as described in chapter 2). Stacking is the averaging process in Equation E-3. The total field exists for every echo, while the coherent field is only accessible as an average. The difference between the average total and coherent intensities is the average diffuse field intensity $\langle I^{dif} \rangle = \langle I^{tot} \rangle - \langle I^{coh} \rangle$, while the average diffuse field may be found by:

$$\langle p^{dif} \rangle = \sqrt{\langle p^{sc} p^{*sc} \rangle} - \langle p^{sc} \rangle \qquad \text{(E-4)}$$

In general, the diffuse field has a wide angular spread and lacks a phase relationship with the incident field. It is not possible to define the average diffuse field directly, as the mean of the diffuse field is identically zero. It is only possible to discuss the diffuse intensity or the diffuse field as being what is leftover after the coherent part is removed from the total field. The coherent field is concentrated around the specular direction and usually has a constant phase relationship to the incident field. Theoretical coherent fields are usually reflections, which do have a constant phase, although this is not necessarily true. In some cases, especially for high roughness, there is no specular reflection and the coherent field may be the sum of coherent selective scattering which could have varying phase and direction – as observed in BORIS results from chapter 5 (Figure 5-17).

Appendix F. The Setup and Calibration of the BORIS Model for the Simulation of SBES (From Section 5.1)

The BORIS model has 81 input parameters that are specified via three text files it reads upon its initiation. The simulation is executed via command line syntax that makes use of another 13 tags to specify the various options including the type of file (binary or ASCII) used for time series output, beampattern specification, etc. Because of all this complexity, a MATLAB toolbox was developed to automate the specification of BORIS parameters, to run BORIS and to process its output[67]. This makes running BORIS much easier. The advantages of using MATLAB in this way are numerous. First, this allows batch runs of BORIS, which is necessary, as BORIS is extremely slow (it takes up to 1.5 hours simulate a single echo, with an average of 8 minutes). To generate data for the numerous experiments and simulated surveys of this dissertation, BORIS was run on two high-end Windows PCs (circa 2006, 3 GHz clock speed, 1 GB RAM) for a total computation time of about *five years*. The MATLAB toolbox also facilitates the easy simulation of echoes with a simple command:

```
>> initializeborisini('ODOM24','AGSAND')
runboris('show',1)
```

This would initialize the input files to simulate the Odom 24 kHz SBES on a medium sand and then run the simulation plotting the echo time series upon completion. The parameters that describe the seabed and the echosounder are discussed in sub-section 5.1.1. Sub-section 5.1.2 shows how the specific seabed and echosounder parameters determine a large number of model parameters. The model parameters are largely established by a number of test runs, with the aim of generating realistic seabed echoes and seabed surveys while minimising computation time. Some example results are shown in section 5.1.3.

---

[67] This MATLAB toolbox is not to be confused with limited set of functions contained in the BORISAUX toolbox delivered with the BORIS software.

F.1 Seabed and Echosounder Parameters for the BORIS Model

The geotechnical parameters required to describe the seabed and echosounder comprise the larger proportion of the BORIS model parameters, while the number of geoacoustic parameters pertaining to the seabed is relatively few. For instance, BORIS describes the location and orientation of the transmitter and receiver separately in 3-D Cartesian co-ordinates measured in metres, relative to a point of origin at the centre of the first seabed tile (the seabed is composed of tiles of stochastically generated seabed). This allows bistatic scattering simulations, which are not explored here. As SBES are monostatic, the transmitter and receiver are co-located in all simulations. Varying the vertical distance to the seabed simulates varying water depth, while varying the tilt of the transducer simulates seabed slope. The transmit pulse and beampattern are supplied to BORIS numerically via separate binary files, instead of by beamwidth and pulse duration parameters. However, the MATLAB toolbox can also generate the beampattern from either a given beamwidth value, a measurement of a transducer diameter, or a one or two dimensional beampattern. Example beampatterns are shown as Figure F-2 and Figure F-3. The transmit pulse time series file can be simulated knowing a centre frequency, pulse duration and sampling frequency (which is set high enough to avoid aliasing). An example of this is shown in Figure F-1. The toolbox also has the capability to extract the above parameters from the beamwidth and pulse files: this is important as these parameters affect the various model parameters discussed in the next section. Table C-1 of appendix C summarizes all of the echosounder parameters used throughout this dissertation. The parameters listed in Table C-1 were measured by the transmit pulse and beampattern simulation and analysis tools of the MATLAB toolbox.

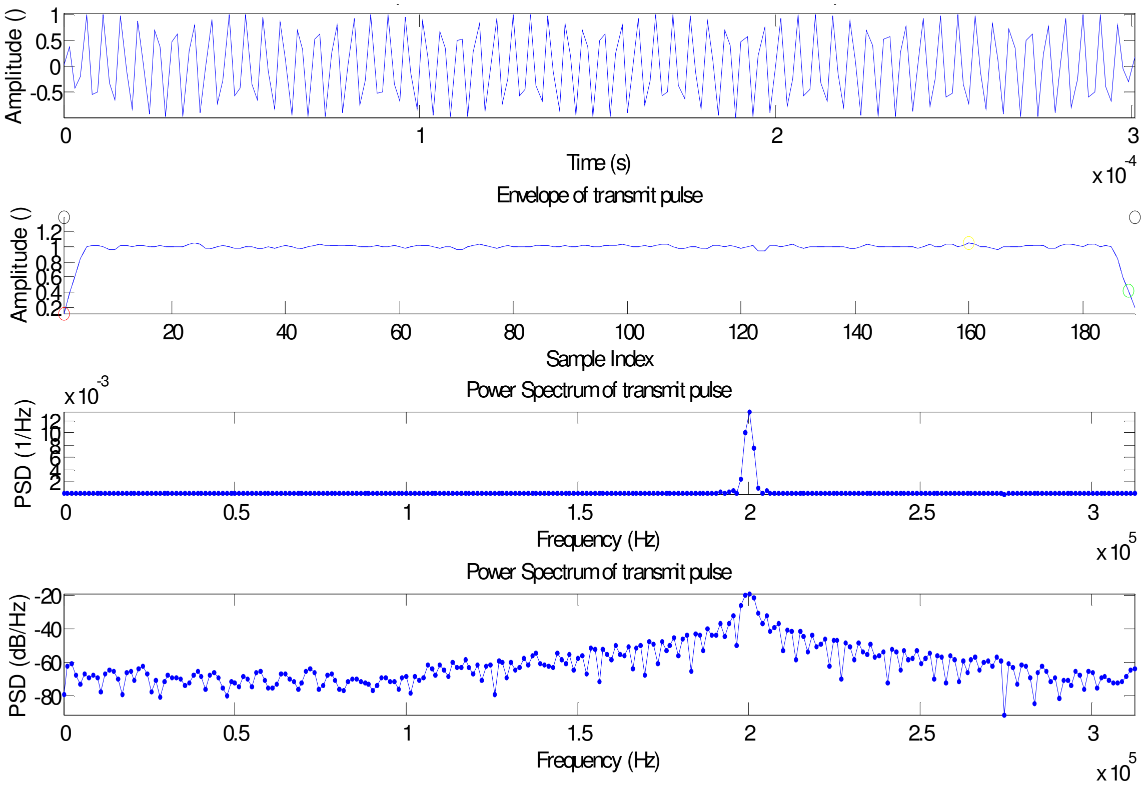
Envelope of transmit pulse

Power Spectrum of transmit pulse

Power Spectrum of transmit pulse

Figure F-1. Transmit pulse and analysis for the Odom200 SBES. Bottom and tail picks are indicated by the red and green circles on the plot of the envelope.

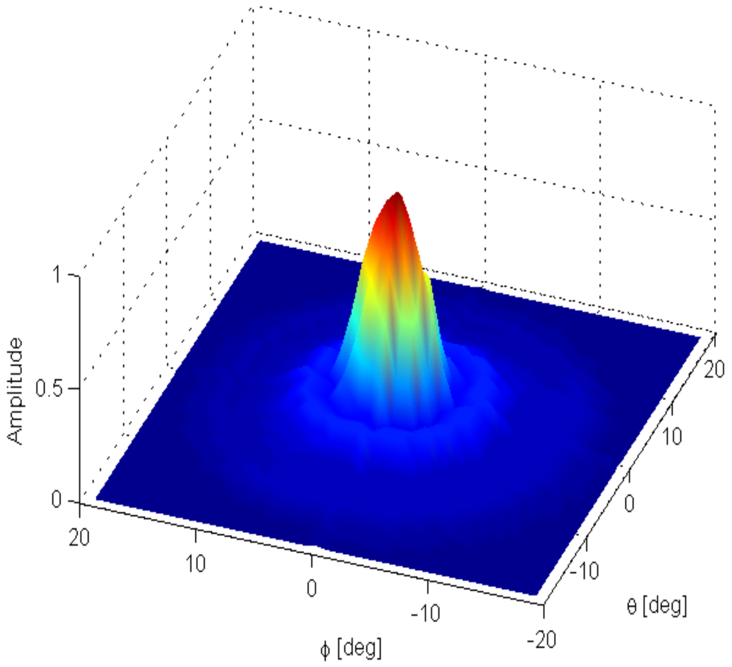

Figure F-2. The beampattern for the Odom200 SBES as constructed from two-sided beampattern profiles provided by the manufacturer.

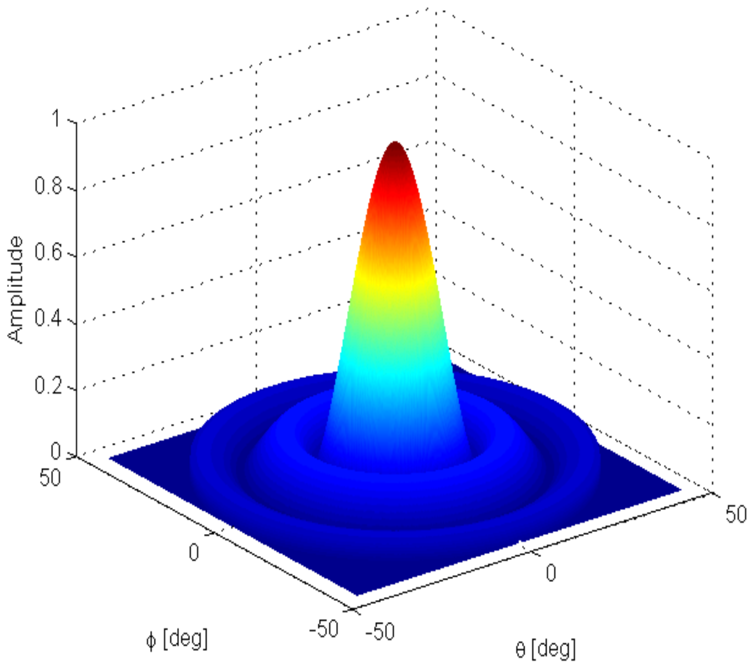

Figure F-3. Simulated beampattern for the virtual Circ50_20 SBES. This beampattern was designed to have a moderate beamwidth at a moderate frequency for use as a general case to study the effects of roughness.

The process of describing the seabed was briefly discussed in chapter 4. The end results are the numerical stochastically-generated seabed and volume shown Figure 4-3. This data set is generated at run-time and can be stored in two binary files for use as input for additional simulations (this option saves computation time). BORIS offers a number of options for seabed roughness height models: Gaussian, power-law, power-law saturated, Pierson-Moskowitz [Canepa et al., 2004]. The Pierson-Moskowitz model is used to simulate ocean-air surfaces (to study scattering from ocean-air interfaces with upward-looking sonars). BORIS has additional capabilities of adding sediment waves to the above roughness surfaces. This would be useful for simulating PSL burying habitat for instance. Based on recent research, in part summarized by Jackson and Richardson [Jackson and Richardson, 2007], the best model for roughness on scales less than one to ten metres is the power-law option. The values of the seabed surface parameters (RMS roughness height, power-law exponent) are shown in appendix C. The saturated power-law is most appropriate for simulating seabeds that lack bedforms or 'macro-roughness' such as sediment waves. The saturation is such that the amplitude of the roughness of spatial wavenumbers less than 7 radians per metre is held to a constant value instead of becoming infinitly large for small wavenumbers. Another safeguard against unrealistic seabed heights are a number of spatial filters (a 'hard' filter was used here). For instance, the high-pass and low-pass frequencies of the hard spatial filter were set to 5 and 400 rad/m, respectively. These filters eliminate extreme values of roughness that are either too small to affect acoustic waves or approach the domain of macro-roughness. Both values were established empirically by a process dubbed 'parameter testing' with BORIS. The procedure is simple: vary each parameter independently to establish the values for which no variation in the resulting echoes are observed. For example, Figure F-4 is the parameter test for the high-pass spatial filter frequency for the seabed surface.

The model of volume inhomogeneities is similar, but with the added complexity of an extra dimension. The volume can be assumed to be isotropic, so the inhomogeneities that cause the scattering are uniformly distributed. However, that is not usually the case unless the sediments are bioturbated, as sediments are layered. This is expressed by the horizontal correlation length being much larger than the vertical correlation length; the BORIS manual cites a finding that the horizontal correlation length is approximately five times the vertical correlation length [Canepa et al, 2004]. The volume model used by the APL-UW handbook [APL-UW, 1994] is not compatible with that of the BORIS model and there are few measurements of volume scattering parameters (aside from attenuation). Therefore, the volume parameters used here are estimates based on the few sources available: [Berntsen, 1999; Bergem et al., 1999; Canepa et al., 1997]. The volume parameters are tabulated along with the surface parameters in appendix C. The seabed volume also has filter parameters similar to the surface filter parameters, and their values were also established by parameter testing.

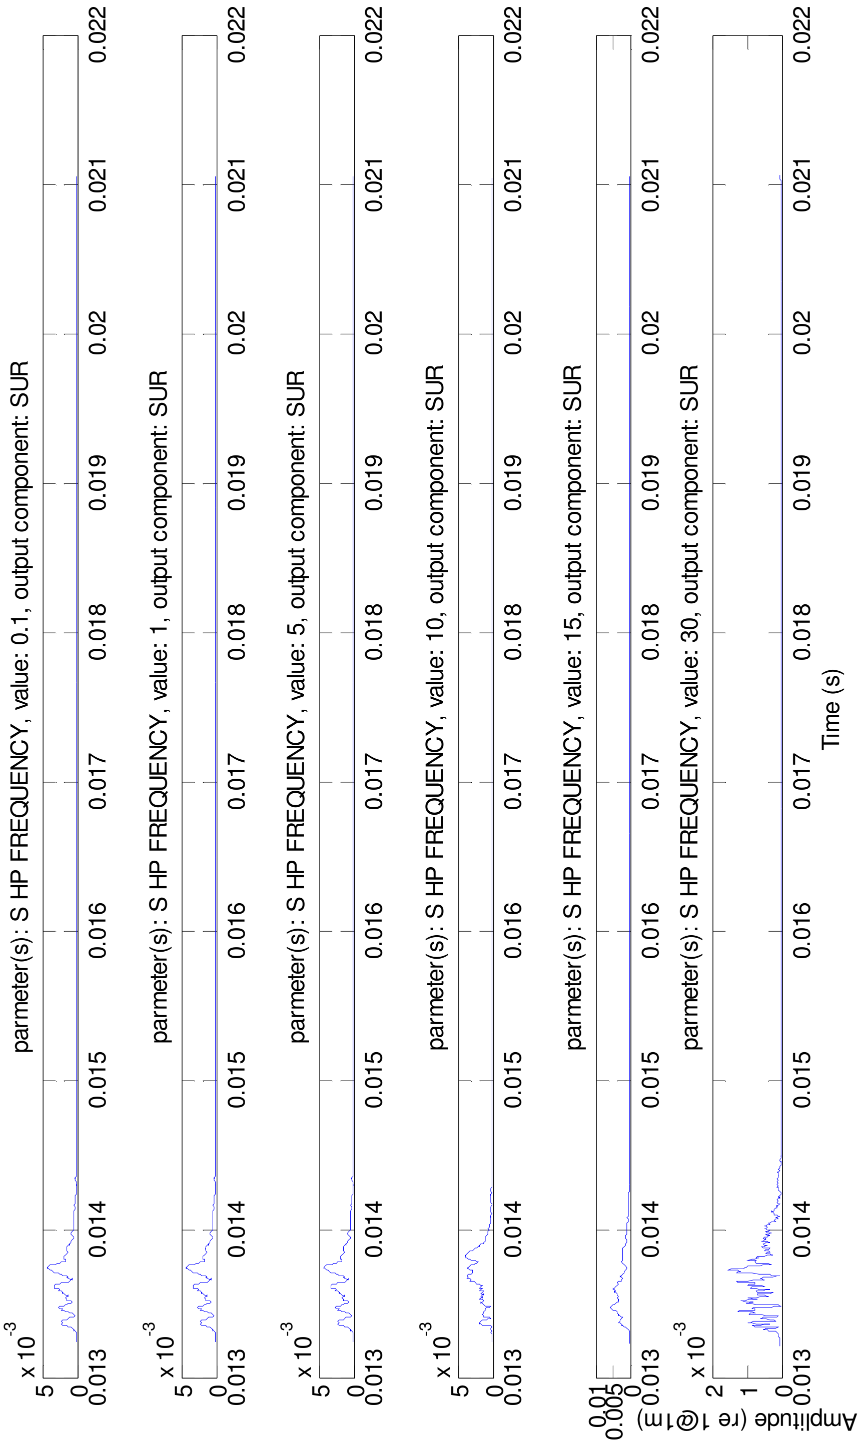

Figure F-4. BORIS parameter test results for the surface roughness high-pass frequency. One can see that values less than 5 rad/m do not affect the echo, as shown in envelopes generated by interface or surface scattering only (output component: SUR).

Another important parameter is the penetration depth, which is the maximum depth into the seabed to which sound penetrates. It controls the extent of the volume echo calculation, and is determined by the skin depth model, Equation D-1. If it is set to zero, BORIS does not compute the volume echo. Through parameter testing, it was found that a 20-dB skin depth was necessary to not affect the volume echo.

F.2 Other BORIS Model Parameters

The penetration depth limits the ensonification volume vertically, however the horizontal extent of the simulation along and within the seabed must also be defined and then minimized to speed up the computation. Generally, this is determined by the extent of the beampattern, which is usually large enough to include at least one sidelobe. Parameter testing has found that excluding the second sidelobe does slightly modify the echo, and the effect is greater when beampatterns are limited to exclude the first sidelobe. Therefore, the beampatterns used, as seen in Figure F-2 and Figure F-3, are bounded at the edge of the second sidelobe. This determines the beampatterns' angular extent, which is generally taken as the angular extent of the simulation.

For virtually all simulations run, the memory requirements to store the simulated seabed and volume for the full spatial extent of the model would exceed the computers' memory capacities (only 1 gigabyte in this case). To overcome this deficiency, the seabed and volume are specified in *tiles* such that the tiles are self-continuous on their edges so that many can be put together to create the required simulation seabed and volume. The dimensions of the tiles are calculated so that they occupy a fixed amount of memory – 128 megabytes for the seabed surface and 256 megabytes for the volume. The other factor that determines these dimensions is the spatial sampling frequency. The BORIS manual recommends a spatial sampling frequency of 2 to 20 times that of the acoustic centre frequency [Canepa et al., 2004]. Parameter testing found the spatial sampling frequency for the surface should be equal to or greater than 10 times the acoustic centre frequency. For the volume, the result was slightly less stringent but the rule of 10 times stands. This agrees with Rayleigh's finding that roughness on scales less than one-tenth the acoustic wavelength has no effect [Rayleigh, 1945]. The spatial sampling frequency is extremely important for realistic echoes, as BORIS works by summing the acoustic response at every sample point. It also exerts a strong influence on the computation time. The temporal sampling frequency is also important to preserve the bandwidth of the transmit pulse. The same considerations of Nyquist frequency as discussed in chapter 2 apply here. To approximate the data from a QTC5 acquisition system, the sample frequency was taken to be 625 kHz for all simulations regardless of echosounder centre frequency. This is the upper limit as established in section 2.1 (Equation 2.4) and is acceptable as the computation time is only linearly dependent on sample frequency and not much would be gained by making it variable. The computation time also depends:

- linearly on the spatial extent, or integration area, which implies it is quadratic with depth;
- quadratically with spatial sampling frequency for the surface and the SSA order;
- linearly with penetration depth and the spatial sampling frequency in the volume.

Since the beamwidth is a moderate function of the acoustic centre frequency and the spatial frequencies are linear functions of the centre frequency, the frequency of the SBES to be simulated is the largest single factor in the computation speed. Water depth is also important; the default value for simulations is a resonably shallow 10 m. The SSA order is the order of the perturbation theory correction (small slope approximation). For normal incidence situations, an SSA correction is not expected to affect the results. However, parameter testing showed differences the between fourth, second and zeroth order results, indicating the most accurate simulation (fourth order) should be used. The differencs are due to the affects of off-normal rays at the edges of the beams. Parameter testing also established the size of the SSA fourth order integration size to be 21 as recommended in the BORIS manual [Canepa et al, 2004].

In general, the BORIS simulations of this work were run with the parameters as established above.

Appendix G. Observations of Echo Characteristics and Potential New Features in a BORIS Simulated Survey (From sub-section 5.4.4)

In this appendix, a modelling study of a set of trail features is presented. The goal is to develop new physics-based features, by examining their relationship to geoacoustic and geotechnical parameters. In field experiments, such parameters are unknown or uncertain. With BORIS simulated surveys, these parameters are known exactly. The results here are based on an example simulated survey that was shown in section 5.1.2 (Figure 5-3). The simulated survey had realistic survey conditions, including varying depth and slope. The echogram shows a number of effects that can also been seen in features. The current QTC phenomenological-type features only capture an echo's overall shape, while proposed physics-based features would measure an attribute of the echo that is known to be influenced by some characteristic of the seabed such as roughness. A number of examples will be shown and discussed here that will lead to the new physics-based feature set presented in chapter 7.

This initial set of new features consists of several algorithms or feature families. First are the measures of time shown in Figure G-1. Rise time is the duration from the bottom pick to the peak amplitude, while fall time is the duration from the peak to the tail pick. The 'spec.', 'scat.' and 'vol.' durations are the durations of three regions within the echo where one may expect, based on theory, that the specular reflection, the surface scatter or the volume scatter will dominate. For reference, see Figure 4-5 for an illustration of the echo components. The specular region ('spec.') is defined as the first two-thirds of the echo between the bottom pick and the peak, while the volume region ('vol.') is defined as the last two-thirds of the echo from the peak to the tail pick. The region remaining in between is the surface ('scat.') scattering region. These divisions are somewhat arbitrary, but are based on the typical divisions in an echo. The two major influences on echoes are evident in these features: depth and slope. In this scenario, the features are all correlated to the overall echo duration. Because of this correlation, the echo duration, corrected for depth and slope, can represent all the features. However, the other features may be useful in other situations. An interesting side-track is that the ratios of some of these features are largely depth invariant as seen in Figure G-2. Some systems do use feature ratios, however not by design (i.e. RoxAnn and Biosonics VBT), and this finding explains why such features are effective. However, feature ratios are still affected by seabed slope.

Energy (mean amplitude squared) is often used as a feature, see Figure G-3. The energy and amplitude of echoes are functions of the impedance and roughness (recall the Rayleigh reflection coefficient, Equation D-22). In cases where the impedance is low, the energy and amplitude will respond to variations in the volume backscattering cross-section. Like other ratio-type features, the ratios of the energy from the three regions of the echo to the total energy are also mostly invariant with depth as shown in Figure G-4. Also interesting is that duration is largely dominated by the echo tail or volume region, while the energy is concentrated around the surface scattering region – this implies energy and amplitude features may be complementary with duration features. Standard measures of amplitude and energy are strongly influenced by seabed slope, see Figure G-3. These measures should be depth invariant as TVG has been applied. For small slopes of 3 and 6 °, (traces 0 to 141) the measures appear to increase slightly, possibly due to improved backscattering from the surface. For seabed slopes of 10 and 22 ° (traces 142 to 300), the measures decrease substantially inferring that the specular reflection and s ome backscattering are lost. One would expect this to occur for slopes greater than one-half of the beamwidth. The Suzuki50 SBES has a rectangular transducer that is wider across track (42 ° re -3 dB, roll direction) than along track (16 ° re -3 dB, pitch direction, direction of seabed gradient). For such a SBES, the wider beamwidth is usually used for SEL depth compensation. However, it is apparent that the narrow part of the beam is interacting with the seabed slope. Cross-slope and along-slope survey lines will interact with the seabed slope differently, causing classification error. It is clear that non-symmetrical beampatterns are undesirable for seabed classification because of this directional bias and because of the difficulty determining beamwidth. It is possible to address these issues with the methods of chapters 6 and 7, but to do so is difficult and the results are not as reliable as those with an axial-symmetric beampattern.

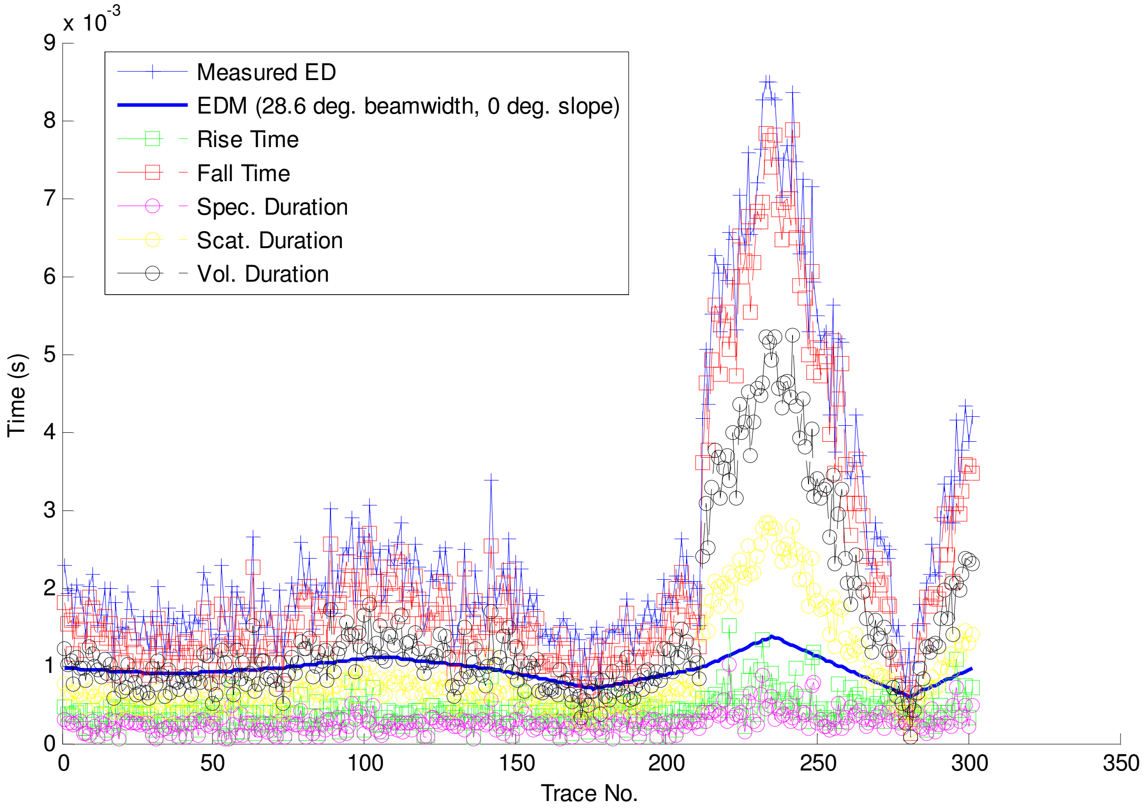

Figure G-1. Measures of time for the example simulated survey with the Suzuki50, SandAG seabed (see Figure 5-3 for complete parameters and other details). The seabed slope increases from 3° to 6° at trace no. 71, to 10 deg. at trace n. 141 and to 22 deg. at trace no. 211. Measured echo duration (ED) is measured by the 99% cumulative energy tail pick, while the model ED is based on the EDM with the – 3 dB beamwidth and zero seabed slope. The input values to the EDM are the usual values one would use for SEL depth compensation – the poor performance of these values indicates that the effective beamwidth and seabed slope need to be considered for accurate modelling of echo duration and for effective depth compensation.

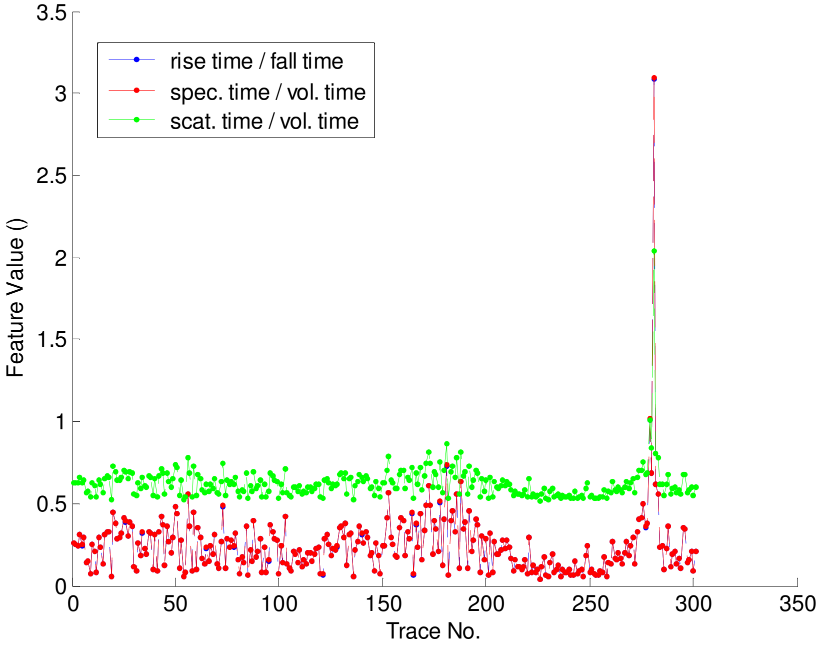

Figure G-2. Ratios of the measures of time for the example simulated survey (see Figure 5-3 for parameters). The extreme value around index 275 is a consequence of a short echo in sloping shallow water.

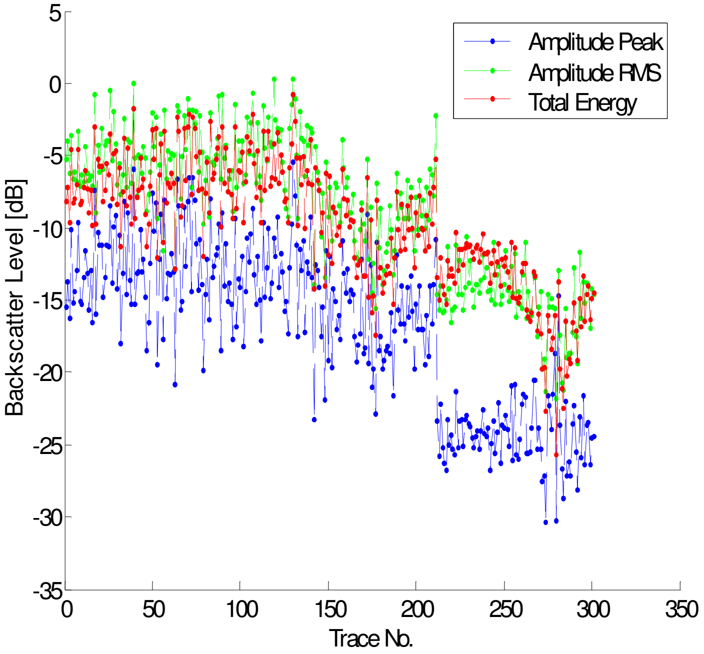

Figure G-3. Measures of amplitude and energy for the example simulated survey (see Figure 5-3 for parameters).

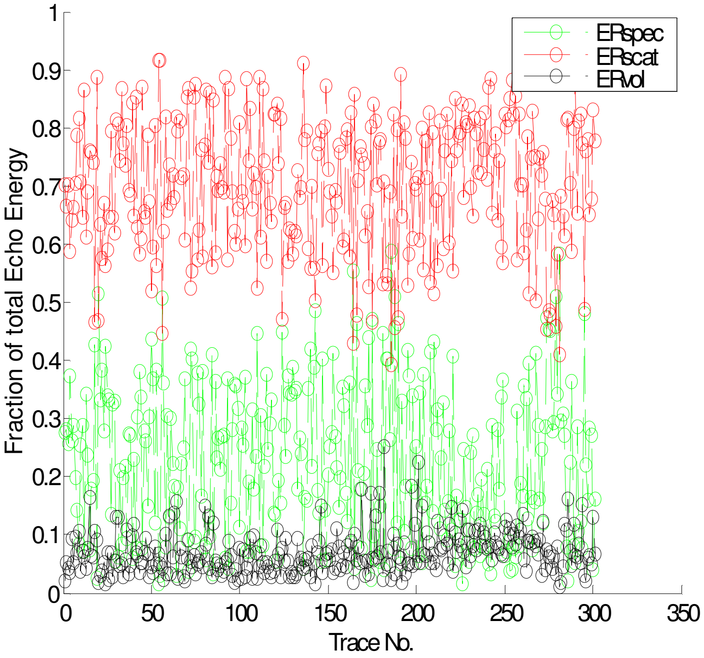

Figure G-4. Energy ratios for the example simulated survey (see Figure 5-3 for parameters). The energy ratios are calculated as the 'spec', 'scat' and 'vol' energy divided by the total energy.

An interesting set of features, shown in Figure G-5, are, in theory, invariant of both depth and slope if the effects of both are purely in stretching the echo in time. The features are the cumulative amplitude thresholds within the echo. This feature family is based on the cumulative sum of echo amplitude from the bottom pick to the tail pick resampled to 128 samples and normalized to unity. Three thresholds are picked: 0.2, 0.5 and 0.8. The feature value is the sample number at which the cumulative sum exceeds this threshold. This is another way one might avoid the depth compensation issue, which is by looking at the relative times of echo characteristics within the echo. However, these features do appear to be correlated with depth and slope. The cumulative energy picking used here is likely biased as mentioned in the previous section. When the normalization to the echo length is removed, this cumulative amplitude threshold algorithm is similar to the best current QTC feature algorithm, the cumulative integral. It shows promise for expressing the shape of the echo if the depth and slope effects can be mitigated. Without the effects of depth and slope, the shape of an echo is determined by the relative contributions of the backscattering components, which respond to the geoacoustic properties of the seabed (mainly roughness, impedance and attenuation).

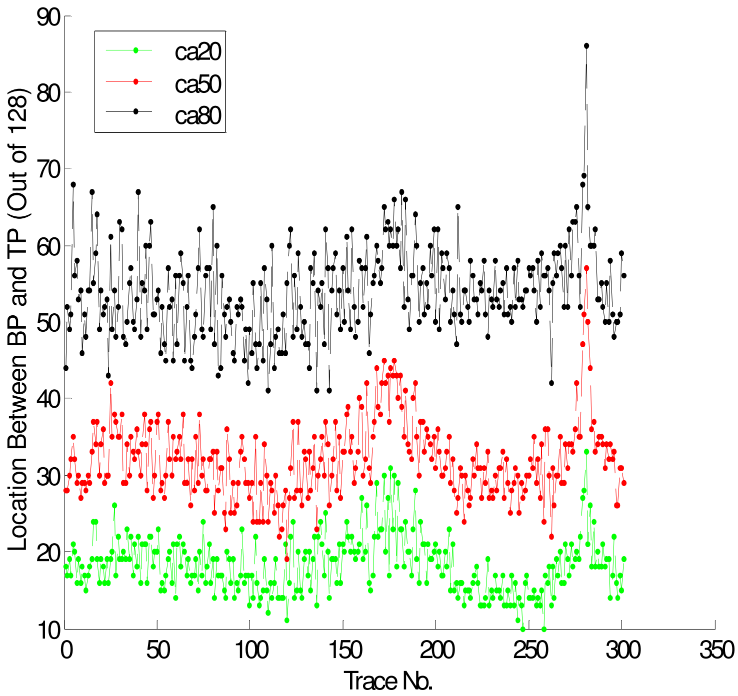

Figure G-5. The within-echo cumulative amplitude threshold features for the example simulated survey (see Figure 5-3 for parameters).

In the current QTC single beam classification method, there are two feature algorithms that capture spectral information – FFTs and wavelets. As noted in chapter 2 and appendix A, these algorithms are not useful. Only the first two FFT and the first wavelet features are weighted in the PCA results observed. Since the echo time series are normalized by their maxima as part of these features, these features measure the mean amplitude relative to maximum amplitude, not any spectral information.

As discussed in chapters 2 and 4, the spectra of echoes should be similar to that of the transmit pulse but with less backscattering of higher frequencies, resulting in an apparent frequency shift lower (as seen in the Rayleigh reflection coefficient) and modified by selective scattering processes. Since the transmit pulse used for the simulations is a gated CW pulse, the spectra should resemble a sinc function, which is what is seen in the results: Figure G-6.

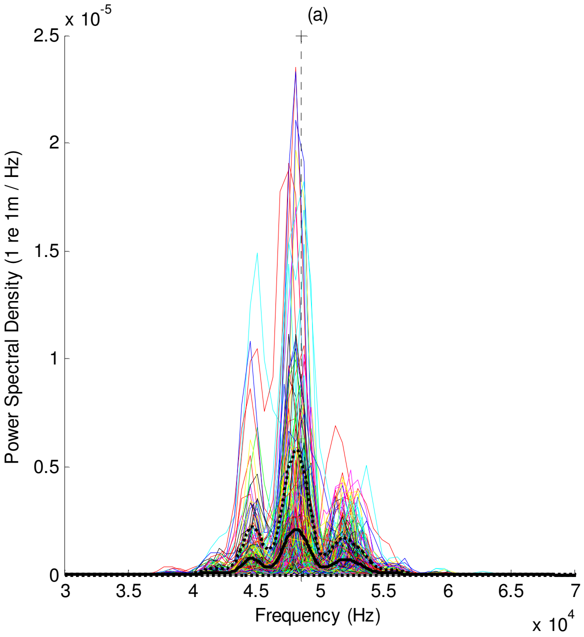

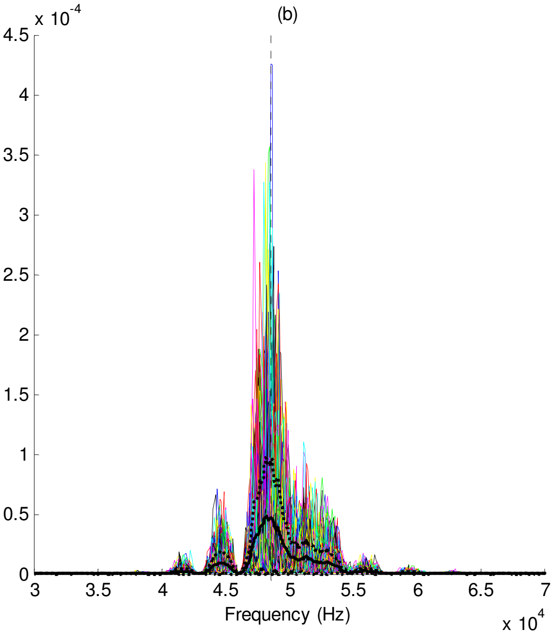

Figure G-6. Power spectral density for the example simulated survey (Suzuki50 echosounder, SandAG seabed) (*a*) and a similar survey with the HRHB seabed (i.e. rough rock) (*b*). See Table C-3 for seabed parameters. Each simulation is a coloured line, while the thick black line is the mean spectrum and the thick dotted line is the mean plus or minus the standard deviation. The dotted line is the transmit pulse centre of gravity frequency.

Instead of using the FFT components as features it is apparent that it would be far more effective to use descriptive statistics of the spectra. Measures of the centre frequency are shown in Figure G-7. van Walree applied spectral moments as features [van Walree et al., 2005]; the first moment is the centre of gravity frequency, the second moment is essentially expressed by the bandwidth statistics (Figure G-9), and the third and fourth moments are the skewness and kurtosis in Figure G-8. All of these spectral measures show a decrease in centre frequency and bandwidth, and modification of the echo spectra due to increased roughness as shown by comparing the simulated survey of the sand and the high relief hard bottom (HRHB) seabed. The frequency shifts are quantifiable with KA theory and can be used for characterization as presented in chapter 8. The effects of roughness (and to a lesser extent volume scattering) will, in theory, always decrease the higher end of the spectrum, lowering the bandwidth and increasing the skewness and kurtosis; these features are currently relative measures suitable for physics-based classification. van Walree applied spectral moments as features without demonstrating how and why they respond to the seabed [van Walree et al., 2005], while the results and theory of this dissertation do explain the behaviour of echo spectra.

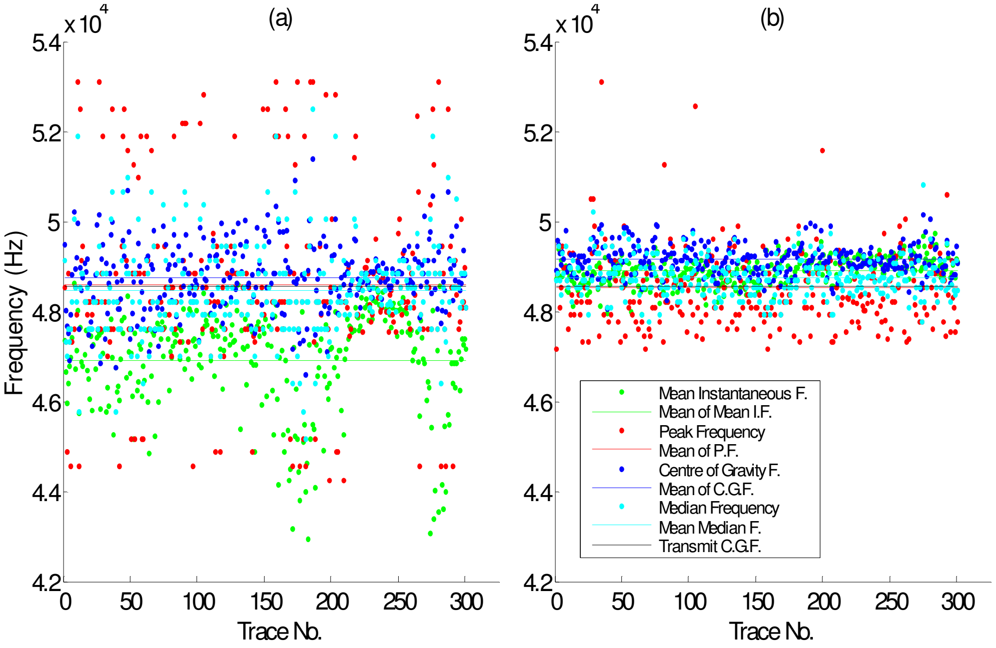

Figure G-7. Measures of the centre frequency for the example simulated survey with the SandAG seabed (*a*) replaced by the HRHB seabed (*b*). The lines are mean values, not linear fits. The instantaneous frequency is derived from the Hilbert transform. The mean of the peak frequencies is hidden under the transmit pulse centre of gravity frequency.

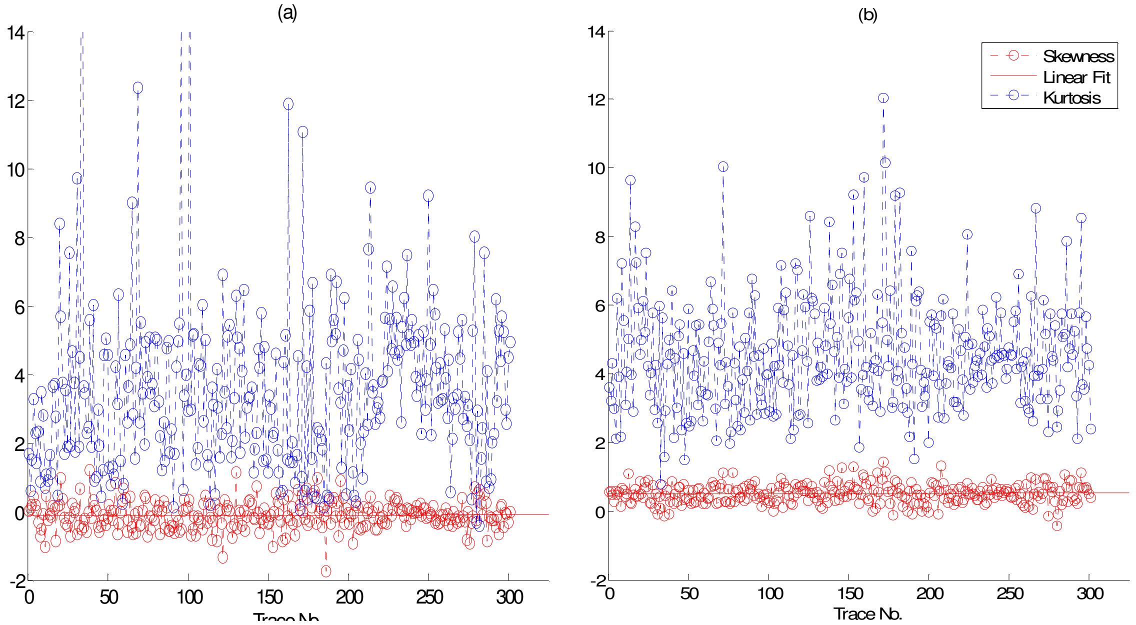

Figure G-8. Spectral moments for the SandAG (*a*) and HRHB (*b*) simulated surveys. (The y-axis is unit-less).

Spectral measures shown above should be applied to improve seabed classification, since they have the advantage of being based on individual FWF (full wave form) time series, instead of stacked envelope echo time series, as is the case for the current QTC single-beam method. The current method suffers from each echo's spectra being 'folded' on itself during envelope formation, so the asymmetries in the spectra that are characteristic of roughness are wiped out. Also, the stacking processing acts as a low-pass filter eliminating even more information useful for classification. It is unavoidable that virtually all field data, except that from the QTC5 system, will be envelope data; however, it is possible that the spectral measures will work with such data as long as these features are calculated from individual echoes and then averaged over the echoes in the stack (in order to be compatible with the features derived from the stacked echoes).

The prospective features developed in this appendix will be applied to new classification methods in chapter 7 and to new characterization methods in chapter 8. Some of the new features presented in this sub-section were designed to be largely invariant of the influences of depth and slope on their own; however, observations of depth and slope correlations indicate that feature design cannot eliminate the adverse effects of depth and slope. Feature ratios such as those used in the RoxAnn style methods are not completely effective and they lack the flexibility to separate the effects of the seabed from the effects of depth and slope. Depth and slope compensation should be handled separately. The final set of new features then contains a redesign of the above features, based on dedicated depth and slope compensation.

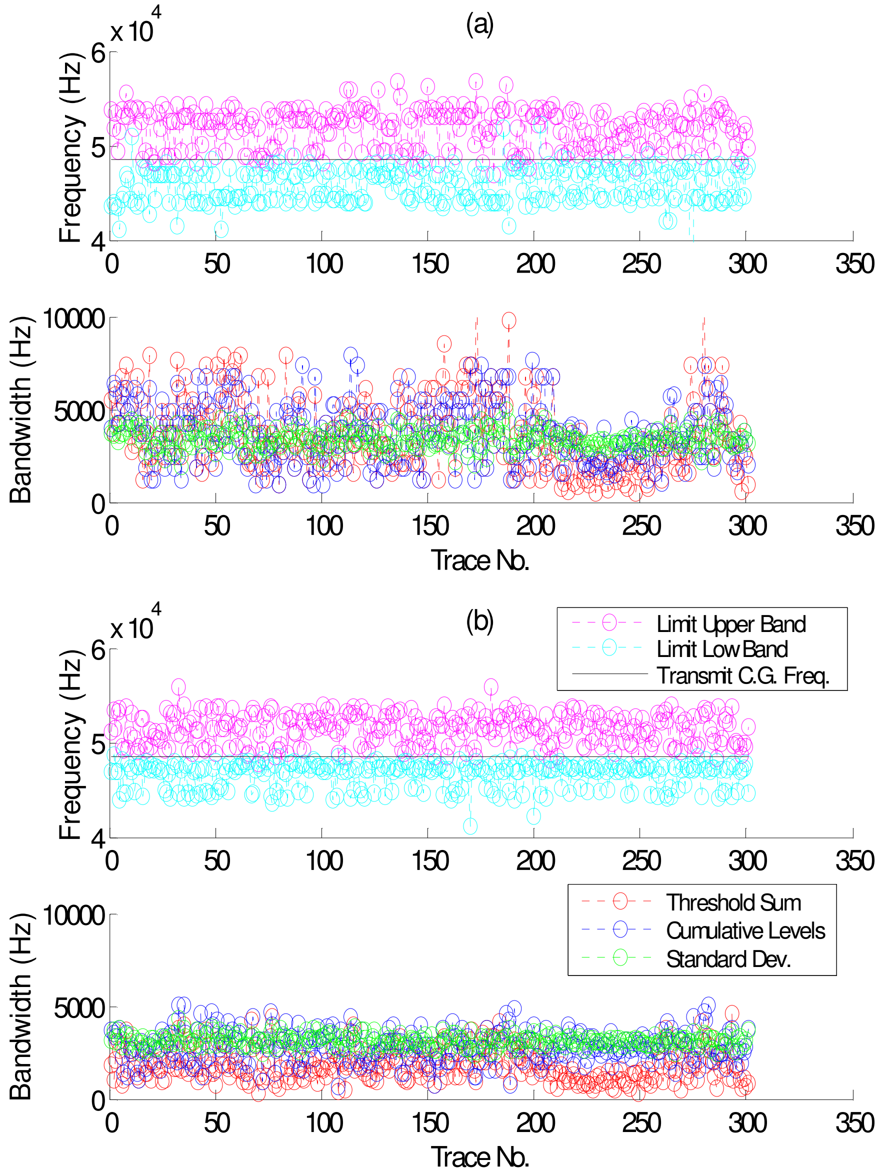

Figure G-9. Measures of bandwidth as applied to FWF data from the SandAG (*a*) seabed and the HRHB (*b*) seabed. The threshold sum is the sum of the number of frequency bins containing power spectral density above 10% of the peak power spectral density, while the cumulative levels are similar to the cumulative tail picks described in appendix A. The upper panels in (*a*) and (*b*) are upper and lower bandwidths found from the cumulative level algorithm.

Appendix H. Dissertation Outline, Summary of Novel Contributions and Publications (From Chapter 9)

In order to find improvements to current classification methods and to discover functional characterization methods, many novel contributions were made through the course of this work. Overall, the novel contributions can be summed up as the application of a basis in physics, including new physical models, new methods and novel applications of the current methods. To aid in the introduction of this work and to highlight its contribution to the science, the following is a point-form outline of the dissertation.

Chapter 2. Current Seabed Classification and Characterization Methods using Single-Beam Echosounders

- A comprehensive overview of SBES basics, the QTC seabed classification methodology and other SBES seabed classification methods. (A compilation of information has never been published.)

Chapter 3. Applications of Current Methods in Single-Beam Echosounder Seabed Classification

- A number of novel applications of the standard QTC seabed classification method are presented aimed at evaluating its performance. All surveys throughout the dissertation were planned and carried out by this author unless otherwise noted.
- Application of current seabed classification methods in a novel integrated habitat mapping project for Pacific Sand Lance.
- Application of current seabed classification methods onboard a remotely-operated vehicle to map gas hydrates and glass sponges.
- Application of current seabed classification methods to map sunken logging debris, which is disrupted by high seabed slope.

Chapter 4. Existing Theory and Models

- An overview of: geoacoustic properties (especially those related to grain size), models of surface roughness, volume and surface scattering theory. The theory and implementation of the BORIS[68] model of echosounding is introduced.
- This overview offers a unique and comprehensive understanding of the physics behind SBES.

Chapter 5. New Theory and Models

- Application of the BORIS numerical time series model for simulating SBES surveys; establishment of appropriate BORIS parameter inputs and testing.
- Insights on backscattering garnered from experiments with BORIS.
- The derivation of the geometric echo duration model; implications of the echo duration model for bathymetry, depth and slope compensation are presented.
- The derivation of the analytic echo times series model; testing of the analytic echo time series model with implications for the ensonification regime.

Chapter 6. Controlled Experiments in the Patricia Bay Testbed

- There are very few controlled experiments in acoustic remote sensing research; controlled experiments require a thorough knowledge of the test area – a testbed. This dissertation establishes one of only a few such areas using MBES classification, video and sampling methods.
- Experiments with a dual frequency transducer simulate changing water depth and various seabed slopes. This is done over well characterized testbed

---

[68] the name of the echosounding model BORIS is an acronym for **bo**ttom **r**eflection from **i**nhomogeneities and **s**cattering

sites, with good control of positioning and attitude. The results are analyzed thoroughly to verify the echo duration model and to calibrate and evaluate depth and slope compensation routines. This fundamental experiment has never been done before.

- Experimental verification of new theory.

Chapter 7. New and Improved Seabed Classification Methods

- A comprehensive high-density SBES survey of Patricia Bay is analyzed to compare the performance of standard QTC classification against new improved methods.
- The current methodology is improved with new bottom and tail picking algorithms, new depth and slope compensation, new physics-based features (and removal of poor features), and outlier removal principal component analysis (PCA). Results at each stage are compared to those of the standard method and to ground-truth class map of chapter 6.
- The final improved methodology is novel and based on physics – it is no longer strictly phenomenological.
- The evaluation of improvements is done rigorously, including the use of confusion tables, statistics and quadratic classification (for comparison).

Chapter 8. Post-class Seabed Characterization

- The methodology combines new characterization methods with improved classification results and establishes an objective way of interpreting the results. The aim is to assign geoacoustic properties to seabed classes, thus the term post-class characterization. Some characterization methods require large ensembles of data, others can characterize a single echo but they are best when combined.

- The frequency shift, amplitude variability and coherent stack methods are drawn from the literature. The mean envelope inversion method combines a new echo time series model with Bayesian inversion techniques. The ABEL method is entirely novel and is the most effective method.
- No characterization result has ever been entirely successful, especially not on an un-calibrated system. The post-class characterization approach of combined methods working on classes should also be robust to varying conditions, composite seabeds, etc., and apply to any SBES.

H.1 Publications to Date

Some of the content of this dissertation has been published, see the following references: [Biffard et al., 2005a, Biffard et al. 2005b, Biffard et al., 2006; Biffard et al., 2007; Preston and Biffard, 2007; Bloomer et al., 2007; Biffard et al., 2008; Biffard et al., 2009; Biffard et al., 2010a; Biffard et al., 2010b].

References

[Al-Hamadani, 1984] Al-Hamadani, Z. K. S. Laboratory studies of underwater acoustic backscattering from rough surfaces. Ph.D. Thesis, University of Bath, 1984

[Anderson et al., 2007] Anderson, J. T., Holliday, V., Kloser, R., Reid, D., and Simard, Y. Acoustic seabed classification of marine physical and biological landscapes. *ICES Cooperative Research Report*, **286**, 2007

[Anderson et al., 2008] Anderson, J. T., Holliday, D. V., Kloser, R., Reid, D. G., and Simard, Y. Acoustic seabed classification: current practice and future directions. *ICES J. Mar. Sci.*, **65**, pp. 1004–1011, 2008

[APL-UW, 1994] APL-UW High-Frequency Ocean Environmental Acoustic Models Handbook, Applied Physics Laboratory Technical Report, APL-UW TR 9407, AEAS 9501, University of Washington, October 1994

[Barbu et al., 2006] Barbu, M., Kaminsky, E., Trahan, R.E.Jr. Acoustic seabed classification using fractional Fourier transform and time-frequency transform techniques. *Proceedings of IEEE OCEANS*, 6 p., 2006

[Barnes and Tunnicliffe, 2008] Barnes, C.R., Tunnicliffe, V. Building the world's first multi-node cabled ocean observatories (NEPTUNE Canada and VENUS, Canada): Science, realities, challenges and opportunities. *Proceedings of IEEE OCEANS*, 8 p., 2008

[Bergem et al., 1999] Bergem, O., Pouliquen, E., Canepa, G. and Pace, N. G. Time-evolution modelling of seafloor scatter. II. Numerical and experimental evaluation, *J. Acoust. Soc. Am.* **105**(6), pp. 3142-3150, 1999

[Berntsen et al., 1999] Berntsen, B., Hovem, J. M. and Bergem, O. Characterization of the seafloor using normal incidence acoustic backscattered time domin signals from a parametric sonar. *Proceedings of IEEE OCEANS*, pp. 30-36, 1999

[Berntsen, 2001] Berntsen, B. Model-based estimation of seafloor parameters by use of acoustic backscattering. Ph.D. thesis, Norwegian University of Science and Technology, Trondheim, Norway, 2001

[Biffard et al., 2010a] Biffard, B. R., Bloomer, S. F., Chapman, N. R. and Preston, J. M. The role of echo duration in acoustic seabed classification and characterization, *Proceedings of Acoustics Week in Canada*, *Canadian Acoustics,* **38**(30), pp. 48-49, 2010

[Biffard et al., 2010b] Biffard, B. R., Bloomer, S. F., Chapman, N. R. and Preston, J. M. The role of echo duration in acoustic seabed classification and characterization, *Proceedings of IEEE OCEANS*, 8 p., 2010

[Biffard et al., 2009] Biffard, B. R., Robinson, C. and Chapman, N. R. Benthic sampling of pacific sand lance (Ammodytes hexapterus) in the Gulf Islands of British Columbia, Canada - report on acoustic seabed classification surveys, October 2008 to January 2009, *C-MARS technical report*, 2009

[Biffard et al., 2008] Biffard, B. R., Robinson, C., Chapman, N. R., Haynes, T. B. and Bloomer, S. F. Benthic sampling of pacific sand lance (Ammodytes hexapterus) in the Gulf Islands of British Columbia, Canada - report on a trial acoustic seabed classification survey, October 29-30, 2007, *C-MARS technical report*, 2008

[Biffard et al., 2007] Biffard, B. R., Preston, J. M. and Chapman, N.R. Acoustic classification with single-beam echosounders: processing methods and theory for isolating effects of the seabed on echoes. *Proceedings of IEEE OCEANS*, 6 p., 2007

[Biffard et al., 2006] Biffard, B. R., Bloomer, S. F., Chapman, N. R., Preston J.M. and Galloway, J. L. Single-beam seabed characterization: a testbed for controlled experiments. *Proceedings of the Eighth European Conference on Underwater Acoustics*, Carvoeiro, Portugal, 2006

[Biffard et al., 2005a] Biffard, B. R., Bloomer, S. F., Chapman, N. R. and Preston J. M. Single-beam seabed classification: direct methods of classification and the problem of slope. *Boundary Influences in High Frequency, Shallow Water Acoustics*, pp. 227-232, 2005

[Biffard et al, 2005b] Biffard, B. R., Bloomer, S. F., Chapman, N. R. and Preston J. M. Single beam echosounding: Considerations of depth and seabed slope. *J. Acoust. Soc. Am.*, **117**, pp. 2442 (A), 2005

[Biot, 1956a] Biot, M. A. Theory of propagation of elastic waves in a fluid-saturated porous solid. I. Low-frequency range. *J. Acoust. Soc. Am.,* **28**(2), pp. 168-178, 1956

[Biot, 1956b] Biot, M. A. Theory of propagation of elastic waves in a fluid-saturated porous solid. II. High-frequency range. *J. Acoust. Soc. Am.,* **28**(2), pp. 179-191, 1956

[Blondel et al. 1998] Blondel, P., Parson, L., Robigou, V. TexAn: textural analysis of sidescan sonar imagery and generic seafloor characterisation. *Proceedings of IEEE OCEANS*, pp. 419-423, 1998

[Bloomer et al., 2007] Bloomer, S. F., Biffard, B. R., Chapman, N. R. QTC DEEP – A ROV-Mounted Single Beam Acoustic Seabed Classification System for High-Resolution Mapping. *Proceedings of IEEE OCEANS*, 6 p., 2007

[Bloomer et al., 2000] Bloomer, S. F., Mosher, D. C., Collins, W. T., Preston, J. M., and Rosenberger, A. Subsurface Classification of High-resolution Seismic Data with Multivariate Statistical Techniques: Case Study from the Strait of Georgia. *Proceedings of IEEE OCEANS*, pp. 1677-1684, 2000

[Bloomer et al., 1999] Bloomer, S. F., Chapman N. R., Collins W. T., and Simpkin P. G. Subsurface classification of shallow gassy river sediments using broadband seismic data, *137th Meeting of the Acoustical Society of America*, 4 p., 1999

[Brekhovskikh and Lysanov, 1991] Brekhovskikh, L.M. and Lysanov, Y. P. *Fundamentals of Ocean Acoustics*. Springer-Verlag, 2nd Ed, 1991

[Brekhovskikh and Lysanov, 1988] Brekhovskikh, L.M. and Lysanov, Y. P. *Fundamentals of Ocean Acoustics*. Springer-Verlag, 1st Ed, 1988

[Briggs, 1989] Briggs, K. B. Microtopographical roughness of shallow-water continental shelves. *IEEE J. Ocean. Eng.*, **14**(4), pp. 360-366, 1989

[Broschat and Thorsos, 1997] Broschat, S. and Thorsos, E. An investigation of the small slope approximation for scattering from rough surfaces. Part II. Numerical studies. *J. Acoust. Soc. Am.,* **101**(5), pp. 2615-2625, 1997

[Buckingham, 2007] Buckingham, M. J. On pore-fluid viscosity and the wave properties of saturated granular materials including marine sediments. *J. Acoust. Soc. Am.,* **122**(3), pp. 1486-1501, 2007

[Buckingham, 2005] Buckingham, M. J. Compressional and shear wave properties of marine sediments: comparisons between theory and data. *J. Acoust. Soc. Am.,* **117**(1), pp. 137-152, 2005

[Burczynski, 2001] Burczynski, J. A technical description of VBT-bottom classifer™ classification system. *Biosonics Inc. Technical Report*, 2001

[Burns et al., 1989] Burns, D. R., Queen, C. B., Sisk, H., Mullarkey, W. and Chivers, R. C. Rapid and convenient acoustic seabed discrimination. *Proceedings of the Institute of Acoustics,* **11**, pp. 169-178, 1989

[Caiti and Zoppoli, 1998] Caiti, A. and Zoppoli, R. Seafloor parameter identification from parametric sonar data. *Proceedings of the Fourth European Conference on Underwater Acoustics*, 5 p., 1998

[Canepa et al., 2008] Canepa, G., Sessarego, J., Tesei, A., Guillermin, R. and Soukup, R. A time domain model of scattering from small discret volume particles: tank validation. *Proceedings of Acoustics 08*, pp. 4119-4124, 2008

[Canepa et al., 2004] Canepa, G., Pautet, L. and Pouliquen, E. BORIS-SSA: BOttom Response from Inhomogeneities and Surface using Small Slope Approximation. *NATO URC Special Report*, M-152, La Spezia, Italy, 2004 (version 1.0)

[Canepa et al., 1997] Canepa, G., Bergem, O. and Pouliquen, E. The Implentation of BORIS-3D: BOttom Response from Inhomogeneities and Surface using Small Slope Approximation. *SACLANTCEN Special Report*, M-125, San Bartolomeo, Italy, 1997 (version 1.0)

[Caughey, 1996] Caughey, D.A. *Seabed classification from acoustic echosounder returns*. Ph.D. Thesis, University of Victoria, Victoria, B.C., Canada, 1996

[Chakraborty et al., 2001] Chakraborty, B., Kaustubha, R., Hegde, A., Pereira, A. Acoustic seafloor sediment classification using self-organizing feature maps. *IEEE Trans. Geosci. Remote Sens.,* **39**(12), pp. 2722-2725, 2001

[Chapman et el., 2009] Chapman, R., Jiang, Y. and Goff, J. Measurements of attenuation of low frequency sound in marine sediment. In proceedings of *Underwater Acoustics Measurements*, 2009

[Chivers et al., 1990] Chivers, R., Emerson, N. and Burns, D. R. New acoustic processing for underway surveying. *Hydrographic Journal,* **56**, pp. 9-17, 1990

[Clarke and Hamilton, 1999] Clarke, P. A., and Hamilton, L. J. The ABCS program for the analysis of echo sounder returns for acoustic bottom classification. Report DSTO-GD-0215, Maritime Operations Division, Aeronautical and Maritime Research Laboratory (Defence Science & Technology Organisation, Commonwealth of Australia), Melbourne, Australia, 1999

[Collins and Galloway, 1998] Collins, W. T. and Galloway, J. L. Seabed classification and multibeam bathymetry: tools for multidisciplinary mapping. *Sea Technology*, **39**(9), pp. 45-49, 1998

[Cooke et al., 2003] Cooke, K., Kieser, R. and Stanley, R. D. Acoustic observation and assessment of fish in high-relief habitats. *ICES J. Mar. Sci.,* **60**, pp. 658-661, 2003

[DREP, 2001] STING users manual. Originally produced by the Defence Research Establishment of the Pacific, Canadian Forces, modified by A.G.O. Environmental Electronics, Inc. and Jasco Research, Ltd, 2001

[Dommise et al. 2005] Dommisse, M., Urban, D., Finney, B., Hills, S. Potential depth biasing using the Biosonics VBT seabed classification software, *Marine Technology Society Journal*, **39**(2), pp. 90-93, 2005

[Dorgan et al., 2005] Dorgan, K. M., Jumars, P. A., Johnson, B., Boudreau, B. P. and Landis, E. Burrowing mechanics: Burrow extension by crack propagation, *Nature,* **433**, pp. 475, 2005

[Dosso, 2002] Dosso, S. E. Quantifying uncertainty in geoacoustic inversion I: A fast Gibbs sampler approach. *J. Acoust. Soc. Am.,* **111**(1), pp. 129-142, 2002

[Dosso et al., 2001] Dosso, S. E., Wilmut M. J. and Lapinski, A. L. An adaptive hybrid algorithm for geoacoustic inversion. *IEEE J. Ocean. Eng.*, **26**, pp. 324–336, 2001

[Dosso and Wilmut, 2006] Dosso, S. E. and Wilmut M. J. Data uncertainty estimation in matched-field geoacoustic inversion. *IEEE J. Ocean. Eng.*, **31**, pp. 470–479, 2006

[Dunsiger et al., 1981] Dunisger, A. D., Cochrane, N. A., Vetter, W. J. Seabed characterization from broad-band acoustic echosounding with scattering models. *IEEE J. Ocean. Eng.*, **6**, pp. 94-106, 1981

[Edwards, 2005] Edwards, N., Marine Controlled Source Electromagnetics: Principles, Methodologies, Future Commercial Applications, *Surveys in Geophysics* **26**, pp. 675–700, 2005

[Ellingsen et al., 2002] Ellingsen, K. E., Gray, J. S., and Bjornbom, E. Acoustic classification of seabed habitats using the QTC VIEW™ system. *ICES J. Mar. Sci.*, **59**, pp. 825-835, 2002

[Finn, 1993] Finn, J. T. Use of the average mutual information index in evaluating classification error and consistency. *International Journal of Geographical Information Systems,* **7**(4), pp. 349-366, 1993

[Folk, 1980] Folk, R. L. *Petrology of Sedimentary Rocks*, Hemphill Publishing Company, Austin, TX, 1980

[Fonseca and Mayer, 2007] Fonseca L., and Mayer, L. Remote estimation of surficial seafloor properties through the application Angular Range Analysis to multibeam sonar data. *Marine Geophysical Researches,* **28**, pp. 119–126, 2007

[Gabrys et al., 2006] Gabrys, B., Baruque, B. and Corchado, E. Outlier Resistant PCA Ensembles, *10th International Conference on Knowledge-Based Intelligent Information and Engineering Systems*, Howlett, R. J.and Jain, L. C. (Eds.), Part III, pp. 432 – 440, 2006

[Galloway, 2001] Galloway, J. Benthic habitat mapping with acoustic seabed classification. In *Proceedings of IEEE OCEANS*, pp. 2642-2644, 2001

[Galloway, 2004] Galloway, J. Acoustic seabed classification: Integration into hydrographic survey programmes. *Hydro International,* **8**(5), 2004

[Galloway, 2008] Galloway, J., Systematic acoustic seafloor habitat mapping of the British Columbia coast, in *Marine Habitat Mapping Technology for Alaska*, J.R. Reynolds and H.G. Greene (eds.), Alaska Sea Grant College Program, University of Alaska Fairbanks, pp. 198-202, 2008

[Gleason et al., 2006] Gleason, A. C. R., Eklund, A. M., Reid, R. P. and Koch, V. Acoustic seabed classification, acoustic variability, and grouper abundance in a forereef environment. *NOAA Professional Papers,* NMFS **5**, pp. 38-47, 2006

[Gleason et al., 2009] Gleason, A. C. R., Preston, J. M., and Bloomer S. F. Reproducibility of single-beam acoustic seabed classification. *Proceedings of IEEE OCEANS*, 8 p., 2009

[Gleason, 2009] Gleason, A.C.R. *Single-beam acoustic seabed classification in coral reef environments with application to the assessment of Grouper and Snapper habitat in the Upper Floride Keys, USA*. Ph.D. Thesis, University of Miami, 2009

[Greenstreet et al., 1997] Greenstreet, S. P. R., Tuck, I. D., Grewar, G. N., Armstrong, E., Reid, D. G. and Wirght, P. J. An assessment of the acoustic survey technique, RoxAnn, as a means of mapping seabed habitat. *ICES J. of Mar. Sci.*, **54**, pp. 939-959, 1997

[Gucler and Gross, 1964] Gucluer, S. and Gross, M. Recent marine sediments in Saanich Inlet, a stagnant marine basin. *Limnology and Oceanography*, **9**(3), pp. 359-376, 1964

[Hamilton and Bachman, 1982] Hamilton E. J. and Bachman R. T. Sound velocity and related properties of marine sediments. *J. Acoust. Soc. Am.,* **72**, 1891-1904, 1982

[Hamilton et al., 1999] Hamilton L. J., Mulhearn P. J. and Poeckert R. Comparison of RoxAnn and QTC-VIEW bottom classification system performance for the Cairns area, Great Barrier Reef, Australia. *Continental Shelf Research,* **19**, pp. 1577-1597, 1999

[Hamilton, 2005] Hamilton, L.J. A Bibliography of acoustic seabed classification. Technical Report 27, CRC for Coastal Zone Estuary & Waterway Management, Australia, 2005

[Heald and Pace, 1996] Heald, G. J. and Pace, N. G., Implications of a bi-static treatment for the second echo from a normal incidence sonar. *Proceedings of the Third European Conference on Underwater Acoustics*, pp. 649-654, 1996

[Heald, 2001] Heald, G.J. High frequency seabed scattering and sediment discrimination. *Proceedings of Institute of Acoustics,* **23**(2), pp. 258-267, 2001

[Heald, 2000] Heald, G.J. *An analysis of normal incidence acoustic backscatter for seabed discrimination*. Ph.D. thesis, University of Bath, Bath, UK, 2001

[Heezen and Hollister, 1971] Heezen E. C. and Hollister C. D. *The Face of the Deep.* New York: Oxford Univ. Press, 1971

[Hyndman et al., 1990] Hyndman, R. D., Yorath, C. J., Clowes, R. M. and Davis, E. E. The Northern Cascadia Subduction Zone at Vancouver Island – seismic structure and tectonic history. *Canadian Journal Earth of Science*, **27**, pp. 313-329, 1990

[Isakson et al., 2005] Isakson, M., Chotiros, J.P., Guild, M. A comparison of models for the interaction of acoustic waves and sediments with reflection coefficient measurements from a sand/water interface at the sediment acoustics experiment 2004 (SAX04). *Boundary Influences in High Frequency, Shallow Water Acoustics*, pp. 73-81, 2005

[Jackson et al., 1986] Jackson, D. R., Winebrenner, D. P., Ishimaru, A. Application of the composite roughness model to high-frequency bottom backscattering, *J. Acoust. Soc. Am.*, **76**(5), pp. 1410-1426, 1986

[Jackson and Briggs, 1992] Jackson, J. R. and Briggs, K. B. High frequency bottom scattering: roughness versus volume scattering. *J. Acoust. Soc. Am.*, **92**, pp. 962-977, 1992

[Jackson and Richardson, 2007] Jackson, D. R. and Richardson, M. D. *High-Frequency Seafloor Acoustics*, Springer Science, 2007

[Jiang and Chapman, 2008] Jiang, Y. and Chapman, N. R. Bayesian geoacoustic inversion in a dynamic shallow water environment. *J. Acoust. Soc. Am.*, **123**(6), pp. 155-161, 2008

[Kitami et al., 1997] Kitami, M., Ogawa, K., Takano, N. and Motooka, S. Experiments on discrimination between the same materials of different grain size. *Japanese Journal of Applied Physics,* **36**(5B), pp. 3350-3353, 1997

[Kenny et al., 2003] Kenny, A. J., Cato, I., Desprez, M., Fader, G., Schu¨ttenhelm, R. T. E., and Side, J. 2003. An overview of seabed-mapping technologies in the context of marine habitat classification. *ICES J. Mar. Sci.*, **60**, pp. 411–418, 2003

[Kuperman, 1975] Kuperman, W. Coherent component of specular reflection and transmission at a randomly rough two-fluid interface. *J. Acoust. Soc. Am.,* **58**, pp. 365-370, 1975

[LaCasce and Tamarkin, 1956] LaCasce and Tamarkin, Underwater sound reflection from a corrugated surface. *J. Appl. Phys.,* **27** 138-148, 1956

[Lambert et al., 1993] Lambert, D. N., Crawford, J. C. and Walter, D. J. Development of a high resolution acoustic seafloor classification survey system. *Proceedings of the Institute Acoustics*, **15**(2), 1993

[Lambert et al., 1999] Lambert, D. N., Walter, D. J., Young, D. C., Griffin, S. R. and Benjamin, K. C. Acoustic sediment classification developments. *Sea Technology*, pp. 35-41, September 1999

[Lee et al., 2009] Lee, G. H., Kim, H. J., Kim, D. C., Yi, B. Y., Nam, S. M., Khim, B. K., and Lim, M. S. The acoustic diversity of the seabed based on the similarity index computed from Chirp seismic data. *ICES J. Mar. Sci.*, **66**, pp. 227–236, 2009

[Legendre and Legendre, 1998] Legendre, P., and Legendre, L. *Numerical Ecology*. 2nd English ed. Elsevier Science BV, Amsterdam, 1998

[Legendre et al., 2002] Legendre, P., Ellingsen, K. E., Bjørnbom, E. and Casgrain, P. Acoustic seabed classification: improved statistical method. *Canadian Journal of Fisheries and Aquatic Science,* **59**, pp. 1085–1089, 2002

[Legendre, 2003] Legendre, P. Reply to the comment by Preston and Kirlin on ‘‘acoustic seabed classification: improved statistical method". *Canadian Journal of Fisheries and Aquatic Science,* **60**, 1301–1305, 2003

[Lied et al., 2004] Lied, T. T., Walday, M., Olsgard, F., Ellingsen, K. E., and Holm, S. SEABEC—a single beam echosounder, seabed classification system. *Proceedings of IEEE OCEANS*, pp. 2024–2028, 2004

[Lubniewski, and Pouliquen, 2004] Lubniewski, Z. and Pouliquen, E. Sensitivity of echo parameters to seafloor properties and depth variability. *Proceedings of the Seventh European Conference on Underwater Acoustics*, pp. 763-768, 2004

[Lurton and Pouliquen, 1992] Lurton, X. and Pouliquen, E. Automated sea-bed classification system for echo-sounders. *Proceedings of IEEE OCEANS*, pp. 317-321, 1992

[Lurton, 2003] Lurton, X. *An Introduction to Underwater Acoustics: Principles and Applications.* Springer Praxis Books, Springer, 2002

[Medwin and Clay, 1998] Clay, C. S. and Medwin, H. *Fundamentals of Acoustical Oceanography*, Academic Press, 1998

[Mantua et al., 1997] Mantua, N., Hare, S., Zhang, Y., Wallace, J. and Francis, R. A Pacific interdecadal climate oscillation with on salmon production. *Bulletin of the American Meteorological Society*, **78**, pp. 1069-1079, 1997

[Michaloppoulou et al., 1994] Michalopoulou, Z., Alexandrou, D. de Moustier, C. Application of a maximum likelihood processor to acoustic backscatter for the estimation of seafloor roughness parameters. *J. Acoust. Soc. Am.*, **95**(5), pp. 2467-2477, 1994

[Mulhearn, 2002] Mulhearn P. J. Influences of penetrometer probe tip geometry on bearing strength estimates for mine burial predictions. Technical Report, Defence Science and technology Organization, Austrailia, 2002

[Motooka et al., 1995] Motooka, S., Kitami, M., Ogawa, K., Takano, N. Ocean experiments to discriminate seabed materials by echo pulse and neural network at Uchiura Bay. *Japanese Journal of Applied Physics* **34**(5B), pp. 2930-2935, 1995

[Mourad and Jackson, 1989] Mourad, P. D. and Jackson, D. R. High-frequency sonar equation models for bottom backscatter and forward loss. *Proceedings of IEEE OCEANS*, pp. 1168-1175, 1989

[Ogilvy, 1991] Ogilvy, J. A. *Theory of wave scattering from random rough surfaces*, Adam Hilger, Bristol, U.K., 1991

[Orlowski, 1984] Orlowski, A. O. Application of multiple echoes energy measurement for evaluation of seabottom type. *Oceanlogica*, **19**, pp. 61-78, 1984

[Mandelbrot, 1982] Mandelbrot, B. B. *The Fractal Geometry of Nature.* W. H. Freeman, San Francisco, 1982

[Osler et al., 2005] Osler, J., Furlong, A. and Christian, H. A sediment probe for the rapid assessment of seabed characteristics, in: *Acoustic Inversion Methods and Experiments for Assessment of the Shallow Water Environment*, eds. A. Caiti, N.R. Chapman, J.-P. Hermand and S.M. Jesus (Eds.), Springer, New York, pp. 171-182, 2005

[Panda et al., 1994] Panda, S., LeBlanc, L. R., and Schook, S. G. Sediment classification based on impedance and attenuation estimation. *J. Acoust. Soc. Am.*, **96**, pp. 3022-3035, 1994

[Pace, 1990] Pace, N. G. Acoustic backscatter and seabed characteristics, *Proceedings of the Institude of Acoustics*, **12**(1), pp. 21-31, 1990

[Park, 2003] Park, J. *Distance measures among classes of deconvolved sea-bottom sonar responses*, M.A.Sc. thesis, University of Victoria, 2003

[Parnum et al., 2009] Parnum, P., Siwabessy, J., Gavrilova, A. and Parsons, M. A. Comparsion of single beam and multibeam sonar systems in seafloor habitat mapping. In proceedings of *Underwater Acoustic Measurement*, 8 p., 2009

[Pouliquen, 2004] Pouliquen, E. Depth dependence correction for normal incidence echosounding. *Proceedings of the Seventh European Conference on Underwater Acoustics*, 6 p., 2004

[Pouliquen, 2002] Pouliquen, E. Multi-sensor analysis of the seabed in shallow water areas: overview of the MAPLE'2001 experiment. *Proceedings of the European Conference on Underwater Acoustics*, pp. 21-28, 2002

[Pouliquen et al., 1999] Pouliquen E., Bergem O. and Pace N.G. Time-evolution modelling of seafloor scatter. I. Concept. *J. Acoust. Soc. Am.* **105**, 3136-3141, 1999

[Pouliquen and Lurton, 1992] Pouliquen, E. and Lurton, X. Sea-bed identification using echosounder signals. *Proceedings of the European Conference on Underwater Acoustics*, Luxembourg, pp. 535-538, 1992

[Prager et al., 1995] Prager, B. T., Caughey, D. A., and Poeckert, R. H. Bottom classification: operational results from QTC VIEW. *Proceedings of IEEE OCEANS*, pp. 1827-1835, 1995

[Preston, 2009] Preston, J. M. Automated acoustic seabed classification of multibeam images of Stanton Banks. *Applied Acoustics*,**70**, pp. 1277–1287, 2009

[Preston and Biffard, 2007] Preston, J. M. and Biffard, B. R. Acoustic classification with sounder echoes: isolating effects of the seabed on the echo. In proceedings of *Underwater Acoustic Measurements*, 8 p., 2007

[Preston et al., 2006] Preston, J. M., Inouchi, Y. and Shioya, F. Acoustic classification of submerged aquatic vegetation, *Proceedings of the Eighth European Conference on Underwater Acoustics*, 6 p., 2006

[Preston and Christney, 2005] Preston, J. M. and Christney, A. C. Compensation of sonar image data primarily for seabed classification. US Patent No.6868041, filed April 30, 2003, granted Mar. 15, 2005, UK Patent 2403013, 2005

[Preston et al., 2004a] Preston, J. M., Christney, A. C., Collins, W. T. and Bloomer, S. Automated acoustic classification of sidescan images. *Proceeding of IEEE OCEANS*, 6 p., 2004

[Preston et al., 2004b] Preston J. M., Christney A. C., Beran L. S. and Collins W. T. Statistical seabed segmentation - from images and echoes to objective clustering. *Proceedings of the Seventh European Conference on Underwater Acoustics*, pp. 813-816, 2004

[Preston et al., 2004c] Preston, J. M., Christney, A. C., Collins, W. T., McConnaughey, R. A. and Syrjala, S. E. Considerations in large-scale acoustic seabed characterization for mapping benthic habitats. *Proceedings of the ICES Annual Science Conference*, 2004

[Preston, 2004] Preston, J. M. Resampling sonar echo time series primarily for seabed sediment classification. US Patent No. 6801474, filed April 30, 2003, granted Oct. 5, 2004

[Preston et al., 2003] Preston, J. M., Christney, A. C. and Collins, W. T. Comparisons of acoustic classifications with multibeam, sidescan, and single-beam sonars. In proceedings of *Third Conference on High Resolution Surveys in Shallow Water*, 1 p., 2003

[Preston et al., 2002] Preston, J. M., Christney, A. C., Beran, L. S., Collins, W. T. and McConnaughey, R. A. Objective measures of acoustic diversity for benthic habitat characterization. In proceedings of the *Symposium on Effects of Fishing Activities on Benthic Habitats: Linking Geology, Biology, Socioeconomics, and Management*, Tampa, FL, November 12-14, 2002

[Preston et al., 2001] Preston, J. M., Christney, A.C., Bloomer, S.F., and Beaudet, I.L. Seabed classification of multibeam sonar images. *Proceedings of IEEE OCEANS*, pp. 2616-2623, 2001

[Preston and Collins, 2000] Preston, J. M. and Collins, W. T. Bottom classification in very shallow water by high-speed data acquisition. *Proceedings of IEEE OCEANS*, pp. 1277-1282, 2000

[Preston et al., 1999] Preston, J.M., Collins, W.T., Mosher, D.C., Poeckert, R.H. and Kuwahara, R.H. (1999) The strength of correlations between geotechnical variables and acoustic classifications. *Proceedings of IEEE OCEANS*, pp. 1123-1128, 2009

[QTC, 2004a] QTC IMPACT User's Manual and Reference, Quester Tangent Corp., Sidney, B.C., Canada, 2004

[QTC, 2004b] Confidence and probability in QTC acoustic classification, *Quester Tangent Corp. Technical Note*, Sidney, B.C., Canada, 2004

[QTC, 2004c] QTC CLAMS User's Manual and Reference, Quester Tangent Corp., Sidney, B.C., Canada, 2004

[Rayleigh, 1945] Rayleigh, Lord. *The Theory of Sound*. McMillan, 1945

[Rayleigh, 1907]. Rayleigh, Lord. On the dynamical theory of gratings. *Proceedings of Royal Society* A79, pp. 399-416, 1907

[Rice, 1951] Rice, S. O. Reflection of electromagnetic waves from slightly rough surfaces. *Comm. Pure Appl. Math.,* **4**, pp. 351-378, 1951

[Rencher, 2002] Rencher, A. *Methods of multivariate analysis*. 2nd ed. John Wiley and Sons, Inc., 2002

[Richardson and Briggs, 1993] Richardson, M. D. and Briggs, K. B. On the use of acoustic impedence values to determine sediment properties. *Proceedings of the Institute of Acoustics,* **15**, pp. 15-24, 1993

[Robidoux et al., 2008] Robidoux, L., Fonseca, L. and Wyatt, G. A qualitative assessment of two multibeam echosounder (MBES) backscatter analysis approaches. *Canadian Hydrographic Conference and National Surveyors Conference*, Victoria, BC, 2008

[Schlagintweit, 1993] Schlagintweit, G. E. O. Real-time acoustic bottom classification for hydrography: a field evaluation of RoxAnn. *Proceedings of IEEE OCEANS*, pp. 214-219, 1993

[Siemes et al., 2010] Siemes, K., Snellen, M., Amiri-Simkooei, A. R., Simons, D. G., Hermand, J. P. Predicting spatial variability of sediment properties from hydrographic data for geo-acoustic inversion. In press, *IEEE J. Ocean. Eng.*, 2010

[Simons et al., 2009] Simons, D. G., Snellen, M. and Siemes, K. The potential of inverting geo-technical and geoacoustic sediment parameters from single-beam echo sounder returns, In proceedings of *Underwater Acoustic Measurement*, 8 p., 2009

[Siwabessy, 2001] Siwabessy, P. J. W. *An investigation of the relationship between seabed type and benthic and bentho-pelagic biota using acoustic techniques*. Ph.D. Thesis, Curtin University of Technology, Perth, Western Australia, 2001

[Stanton and Clay, 1986] Stanton, T. K. and Clay, S. C. Sonar Echo Statistics as a Remote-Sensing Tool: Volume and Seafloor. *IEEE J. Ocean. Eng.*, **11**(1), pp. 79-96, 1986

[Stanton, 1982] Stanton, T. K. Effects of transducer motion on echo-integration techniques. *J. Acoust. Soc. Am.*, **72**(3), pp. 947-949, 1982

[Stanton, 1984] Stanton, T. K. Sonar estimates of seafloor microroughness. *J. Acoust. Soc. Am.*, **75**(3), pp. 947-949, 1984

[Stepnowski et al., 2003] Stepnowski, A, Moszyñski, M. and Dung, T. V. Adaptive neuro-fuzzy and fuzzy decision tree classifiers as applied to seafloor characterization. *Acoustical Physics,* **49**(2), pp. 193-202, 2003

[Stepnowski et al., 1996] Stepnowski A., Moszyñski M., Komendarczyk R. and Burczyñski J. (1996) Visual real-time Bottom Typing System (VBTS) and neural networks experiment for seabed classification, *Proceedings of the Third European Conference on Underwater Acoustics*, pp. 685-690, 1996

[Sternlicht and de Moustier, 2003a] Sternlicht D. D. and de Moustier C. P. Time-dependent seafloor acoustic backscatter (10–100 kHz) *J. Acoust. Soc. Am.*, **114**(5), pp. 2709-2726, 2003

[Sternlicht and de Moustier, 2003b] Sternlicht D. D. and de Moustier C. P. Remote Sensing of sediment characteristics by optimized echo-enveloped matching. *J. Acoust. Soc. Am.*, **114**(5), pp. 2727-2743, 2003

[Sternlicht, 1999] Sternlicht, D. D. *High frequency acoustic remote sensing of seafloor characteristics*. Ph.D. Thesis, University of California, San Diego, 1999

[Talukdar and Lawing, 1991] Talukdar, K. K., and Lawing, W. D. Estimation of the parameters of the Rice distribution, *J. Acoust. Soc. Am.*, **89**(3), pp. 1193–1197, 1991

[Tao and Motooka, 2006] Tao, L. and Motooka, S. Improvement in discrimination by eliciting properties of echo signals from different seabed materials. *Japanese Journal of Applied Physics,* **45**(5B), pp. 4829-4833, 2006

[Thorne and Pace, 1984] Thorne, P. D. and Pace, N. G. Acoustic studies of broadband scattering from a model rough surface. *J. Acoust. Soc. Am.*, **75**(1), pp. 133-144, 1984

[Thorsos et al., 2005] Thorsos, E. I., Williams, K. L., Tang, D. J. and Kargl, S. G. SAX04 overview, *Boundary Influences in High Frequency, Shallow Water Acoustics*, pp. 3-12, 2005

[Thorsos et al., 2001] Thorsos, E. I., Williams, K. L., Chotiros, N. P., Christoff, J. T., Commander, K. W., Greenlaw, C. F., Holliday, D. V., Jackson, D. R., Lopes, J. L., McGehee, D. E., Piper, J. E., Richardson, M. D. and Tang, D. J. An overview of SAX99: acoustic measurements. *IEEE J. Ocean. Eng.*, **26**(1), pp. 4-25, 2001

[Thorsos and Broschat, 1995] Thorsos, E. I. and Broschat, S. L. An investigation of the small slope approximation for scattering from rough surfaces. Part I. Theory. *J. Acoust. Soc. Am.*, **97**(4), pp. 2082–2093, 1995

[Tollefsen et al., 2006] Tollefsen, D., Dosso, S. E. and Wilmut, M. J.. Matched-field geoacoustic inversion with a horizontal array and a low-level source. *J. Acoust. Soc. Am.*, **120**(1), pp. 221-230, 2006

[Tsehmahman et al., 1997] Tsehmahman, A. S., Collins, W. T. and Pager, B. T. Acoustic seabed classification and correlation analysis of sediment properties by QTCView. *Proceedings of IEEE OCEANS*, pp. 921-926, 1997 [von Szalay and McConnaughey, 2002] von Szalay, P. G. and McConnaughey, R. A. The effect of slope and vessel speed on the performance of a single beam seabed classification system. *Fisheries Research,* **54**, pp. 191-194, 2002

[van Walree et al. 2006] van Walree, P., Ainslie, M., Simons, D. Mean grain size mapping with single-beam echo sounders. *J. Acoust. Soc. Am.*, **120**(5), pp. 2555-2566, 2006

[van Walree et al., 2005] van Walree, P., Tegowski, J., Laban, C. and Simons, D. Acoustic seabed discrimination with echo shape parameters: a comparison with the ground truth. *Continental Shelf Research*, **25**, pp. 2273-2293, 2005

[Walpole and Myers, 1985] Walpole, R. E. and Myers, R. H. *Probability and statistics for engineers and scientists*, 3rd Edition, Macmillan Publishing Inc., 1985

[Walter et al., 1997] Walter, D. J., Lambert, D. N., Young, D. C. and Stephens, K. P. Mapping sediment acoustic impedance using remote sensing acoustic techniques in a shallow-water carbonate environment. *Geo-Marine Letters,* **17**, pp. 260-267, 1997

[Wang and Philpot, 2007] Wang, C. and Philpot, W. Using airborne bathymetric lidar to detect bottom type variation in shallow waters. *Remote Sensing of Environment,* **106,** pp. 123–135, 2007

[Wentworth, 1922] Wentworth, C. K. A scale of grade and class terms for clastic sediments. *J. Geology,* **30**, pp. 377–392, 1922

[Wilmut et al., 2009] Wilmut, M., Bloomer, S., Preston, J., Discriminant analysis in image-based seabed classification. In proceedings of *Underwater Acoustic Measurement*, Crete, Greece, 2009

[Wood and Lindwall, 1996] Wood, W. T. and Lindwall, D. A. Full waveform inversion of field sonar returns for a visco-acoustic earth; a comparison of linearized and fully nonlinear methods. *IEEE J. Ocean. Eng.*, **21**(4), pp. 423-431, 1996

[Zelinsky et al., 2004] Zielinski, A., Lubniewski, Z. and Tegowski, J. Seafloor backscattering: numerical simulation vs. filter theory approach. *Proceedings of the Seventh European Conference on Underwater Acoustics*, Delft, the Netherlands, 6 p., 2004