APPENDIX II

Acoustic Bottom Classification Data Analysis and Results.

Table of Contents

List of Tables

List of Figures

Data Collection

The QTCView™ System

QTCView is a seabed classification system consisting of hardware and software developed by Quester Tangent Corp. (www.questercorp.com). The system uses acoustic information provided by a standard echo sounder to infer the properties of the seabed. Acoustic seabed classification involves the organization of sea floor echoes into "classes" based on a characteristic acoustic response. In the normal operation of a depth sounder, the acoustic pulse generated by an echo sounder travels through the water column, reflects from the seabed and returns to the transducer. There it is converted back into electrical energy and, after amplification and signal conditioning, recorded as a gray scale mark on paper or as colors of different hue and intensity on a video display. The data can also be stored as a digital time series (a set of numbers representing the amplitude of the signal in volts sampled at a regular time interval). QTCView "taps into" the electrical path between the transducer and the sounder. The detection of the transmitted signal going to the transducer is taken as the start of each record: the system then digitizes all of the data received by the transducer until such time as the signal associated with the sea bed has passed. This information constitutes a digital version of the echo trace.

Sophisticated signal processing algorithms are applied to the digitized echo, separating it into fundamental components (e.g., energy, frequency etc.). These components vary relative to each other as the signal reflects from differing sea beds. Sets of about 5 digitized traces are analyzed in this way, checked for consistency, averaged, and saved as a 166 element Full Feature Vector (FFV). These data are collected and saved during the data collection process, which occurs in tandem with grab sampling operations. The FFVs are input into a post-processing scheme that assumes echoes with similar component values come from sea beds with similar characteristics. Similar echoes are grouped into classes that may be related to the physical seabed characteristics by comparison with grab sampling results.

The Acoustics

The acoustic system used for this project is the ships echo sounder, a dual frequency SITEX CVS-108DF system. The transducer is a hull-mounted 50kHz/200kHz dual frequency unit. QTCView uses the 200 kHz signal: the half-power beam-width of the transducer at this frequency is 7°, and the beam is conical.

The Survey Area

The ABC survey covered the same area as the grab-sampling program, and in fact was carried out at the same time as the grab sampling.

The Survey Plan and Parameter Selection

The grab sampling plan covered areas with depths ranging from 5 meters to just over 100m. One of the parameters that must be defined for QTCView is the "reference depth", which is meant to be the average depth over a survey area. However, results are best if the maximum range of depths in a survey is less than 100m. Accordingly, although some of the expected depths at the outer edge of the sampling area were slightly greater than 100 meters, ABC data collection was restricted to areas with depths less than 100 meters, and used the parameters shown in Table A2-1.

Parameter	Values
	System
Base Gain (dB)	-5
Reference Depth (m)	45
Minimum Depth (m)	5
Maximum Depth (m)	100
	Sounder
Power	25W (RMS)
Pulse Length	648μs
Maximum Range(m)	80

Table A2-1: QTCVIEW parameters

Survey Operations

Details of the survey operations are detailed in Table A2-2.

Date	Km surveyed	No. of Records
23/8/2000	49.0	8,804
24/8/2000	141.4	22,789
25/8/2000	99.9	14,246
30/8/2000	56.6	7,653
31/8/2000	38.3	3,503
1/9/2000	38.6	4,958
2/9/2000	14.7	1,310
3/9/2000	73.6	9,849
5/9/2000	72.8	9,512
6/9/2000	78.7	9,718
7/9/2000	77.6	9,784
8/9/2000	26.5	2,307
10/9/2000	40.4	3,281
11/9/2000	39.8	3,964
Total	847.9	111,678

Table A2-2: Survey details.

Data Analysis

Data quality assessment and filtering

There are literally thousands of ABC records collected during a typical survey, and not all records are suitable for analysis. Records that are outside of the depth range specified by the system parameter settings must be removed, as well as "garbage" data. The most common problem is faulty depth picks. These occur when the QTCView system loses track of the bottom and then enters a search mode to find the bottom. Faulty depth picks are obvious in a plot of depth *vs.* time, and are easily removed. Faulty depth picks are rare overall, but are more common in deeper water, because of the attenuation of the acoustic signal with depth.

There were unique data quality problems with this particular survey. During the first few days of the survey in August, there were problems with the QTCVIEW "Blue Box" ceasing operations intermittently and having to be restarted during surveys. This behavior was unusual, and after discussions with Quester Tangent the problem was traced to some electrical equipment on the vessel. After the use of this equipment was discontinued the problems ceased. However, on analyzing data collected during this period, several data sets were found to be too noisy to be used. As a result, 23,949 records covering 89.5 km (10.6% of the total survey track) were discarded from the final data set. Fortunately most of these survey lines were in the northwest portion of the study area and away from the areas of greatest interest – *e.g.* Peacock Spit, Clatsop Spit, dredge disposal sites and the proposed deep water site. Details of record removal during quality assessment are shown in Table A2-3.

Description	Number
Starting number	111,678
Removed due to noise problem	23,949
Faulty depth picks, too shallow or deep	493
Final number of records	87,236

Table A2-3: Records removed during data assessment.

Unsupervised Classification

Classification may be either supervised or unsupervised. In supervised classification, local knowledge of the available bottom types specifies the classifications that will exist. The acoustic properties of these known bottom types are measured and used to form a catalogue that is employed in subsequent survey operations to classify the area. Unsupervised classification was applied to this project. An unsupervised classification is one in which no *a priori* judgments are made about the diversity of bottom types present in the survey area. FFVs are collected and analyzed after the fact to determine a reasonable division of the survey area into bottom type classes. Because a large number of bottom grab samples are collected as part of the Sediment Trend Analysis work, unsupervised classification is ideal for such projects. Nevertheless, some bottom types will classify out as different from others, although the associated grab samples may appear almost identical. This is because the properties of the acoustic return from the bottom depend on many factors, not all of which may be apparent from grab samples. The gross morphology of the bottom is a good example of how this can occur. Given a beam angle of 7° and a depth of 30m, the "acoustic footprint" on the bottom is roughly 3.7m in diameter. Two bottom types composed of exactly the same sediments, one perfectly flat and one with small sand waves due to, say, bottom currents, will have different acoustic returns. Although the grab samples may appear identical, the two regions may separate into distinct classes. Another example might be the presence of biota on the bottom. Two regions on the bottom composed of identical sediments may differ in that one is empty of biota and the other may have starfish or some other invertebrates scattered about. These invertebrates may not be evident in the grab samples, but will show up in the acoustic return.

Principal Component Analysis

The first stage in data analysis is to reduce the dimensionality of the FFV data. Principal Components Analysis (PCA), a mathematical procedure (see Murtagh & Heck 1987) that transforms a number of (possibly) correlated variables into a (smaller) number of uncorrelated variables called *principal components* is used to do this. The first principal component accounts for as much of the variability in the data as possible, and each succeeding component accounts for as much of the remaining variability as possible. Usually, principal component analysis is performed on a square symmetric correlation matrix (sums of squares and cross products from standardized data). The data are standardized because the elements of any FFV can differ by several orders of magnitude. Standardizing the data ensures that all element of the FFV are equally important in the PCA procedure. The objectives of principal component analysis are:

- To discover or to reduce the dimensionality of the data set.
- To identify new meaningful underlying variables.

The mathematical technique used in PCA is called eigen-analysis: one solves for the eigenvalues and eigenvectors of a square symmetric matrix with sums of squares and cross products. The eigenvector associated with the largest eigenvalue has the same direction as the first principal component; the eigenvector associated with the second largest eigenvalue determines the direction of the second principal component; and so on. The sum of the eigenvalues equals the trace of the square matrix (which is the number of variables; in our case 166) and the maximum number of eigenvectors equals the number of rows (or columns) of this matrix. Using the first three principal components has been found to be adequate for the purposes of ABC. The eigenvalues are examined to determine the effectiveness of the first three principal components in accounting for the variance in the data.

Using the results of the PCA the 166-element FFV for each data point in the cluster analysis can be replaced by the three-element PCA vector which approximates the FFV.

The Clustering Algorithm

K-means clustering (see Hartigan 1975) is used to partition the data into several classes or “clusters”. There are several variants of the k-means clustering algorithm, but most variants involve an iterative scheme that operates over a fixed number of clusters, while attempting to satisfy the following properties:

1. Each cluster has a center that is the mean position of all the samples in that cluster.
2. Each sample is in the cluster to whose center it is closest.

The algorithm works by first selecting N samples (where N is the chosen number of clusters) randomly as cluster centers. It then moves samples into the closest cluster, meanwhile recalculating the mean center of the cluster. This partition of samples into new clusters is repeated until any further movement of samples does not improve the mean square error of the partition. The space in which the classification takes place is that spanned by the three principal components, and the distance measure is Euclidean.

Determining the number of clusters to use is somewhat of an art. There is of course a practical limit to the number of clusters that can be reasonably represented in a region, based on the number of ABC records, the area covered, the assumed diversity of the environment, and the number of "ground-truth" records available. One method to determine the number of clusters is to keep track of the mean square error of the partition as the number of clusters is increased, and stop when it is judged that the mean square error of the partition does not decrease significantly with the addition of another cluster. Another method, the one used in this report, is to keep track of the Clustering Performance Index Rate (CPIR – see Kirlin and Desaji, 2000) and look for peaks. The results of this approach are shown in Figure A2-1, and show a small peak at 8 classes. The 8-cluster classification is the solution presented in this document.

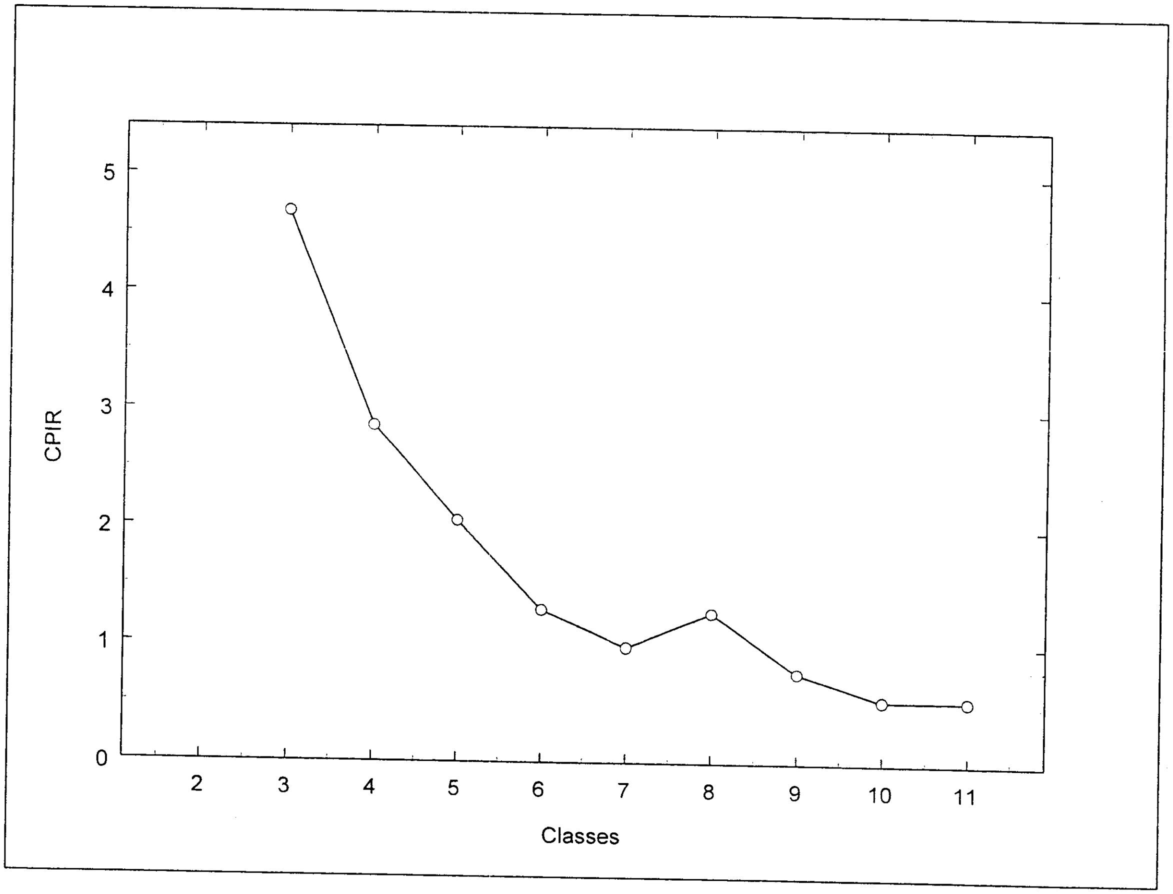

Figure A2-1: CPIR *vs.* number of classes.

Confidence Estimation

A means of assessing the confidence with which a datum is assigned to a cluster is required. Since all records are within a data space spanned by the three chosen principal components, and since the clustering algorithm uses a distance metric to assign data to clusters, a comparison of distances can be used to define confidence level. For each datum the distance D_c to its cluster center and its distance D_n to the center of the next closest cluster are calculated. Then the confidence level C, defined as a percentage, is given by:

$$C = \frac{D_n - D_c}{D_n + D_c} \cdot 100$$

Therefore, a datum that is at the center of its cluster is 100% confident, and a datum which is equally close to the next closest cluster center has a confidence level of 0%. This measure is intuitively conservative: a datum which is twice as close to its cluster center as to the next nearest cluster center has a confidence level of only 33%. To have a confidence level of 50%, a datum must be three times as close to its cluster center as to the center of the next nearest cluster.

Comparison with Collected Samples

A large number of grab samples were collected in the region surveyed for ABC, and these data can be used to help understand the meaning of the results of the unsupervised classification. To do this, the locations of all the ABC samples are compared with the locations of the grab samples and all of the ABC samples that are within a certain distance of a grab sample are selected. These ABC samples are then associated with that grab sample; this process is called *blending*. The radius within which the search for ABC samples is carried out is a variable that relates to the depth, being smaller in shallow water than in deep. In shallow water, the acoustic footprint is smaller than at deeper depths. In addition, pitch and roll of the vessel causes uncertainty in the location of the acoustic footprint on the bottom, and this uncertainty is greater in deep water than in shallow. In this project a radius of 20 meters was used.

Once the list of all the ABC samples that are close enough to the location of a grab sample to be considered is found, the statistics of that set of data can be examined. Are they all classified the same? How do the confidence levels of the classifications compare? Then some statistics of the grab samples that are in regions that are classified identically can be calculated: for example, the mean and standard deviations of the grain-size parameters (mean, sorting and skewness) of the samples. The average percentage of mud, gravel and sand, the average confidence level of the classification, and the average depth of the samples are also calculated. In addition, the anecdotal descriptions of the samples in the field records can be examined to see if there is anything consistently different among the clusters.

Results

Statistics

After noisy data removal and filtering there were 87,729, or 78.1% of the original number of FFVs suitable for input into the classification process. The results of the principal component analysis are shown in Table A2-4: nearly 60% of the variance of the data set is included in the first three eigenvectors.

Eigenvalue	1	2	3	Sum
% Variance	43.5	10.9	3.7	58.1

Table A2-4: Results of PCA of ABC data.

The data were then analyzed using k-means clustering: the results are shown in Table A2-5. Based on practical considerations, examination of the cluster sizes and the CPIR changes, a classification using 8 clusters was chosen.

Cluster	1	2	3	4	5	6	7	8	9	Sum of Squares
2	58198	29038								11,392
3	36787	27078	23371							8,119
4	30745	19348	18637	18506						4,643
5	28333	19068	14960	14949	9926					3,550
6	24549	18657	13237	12028	11765	7000				2,972
7	19367	15466	13068	11653	11331	9565	6786			2,184
8	18803	15369	10861	10799	10564	9544	9374	2422		1,876
9	15504	15244	10144	9652	9319	8979	8901	7795	1698	1,641

Table A2-5: Numbers of samples in each cluster, and sum of squares for shallow data. The cluster numbers in this table are not the actual class numbers used in the following descriptions; they are arranged in decreasing order by number of records.

The number of records in each class and the cluster centers in eigenvector space are shown in Table A2-6, where the columns refer to the eigenvector numbers (column 1 is the cluster center location on the eigenvector 1 axis). Table A2-7 shows the normalized (largest inter-cluster difference is set equal to 1) inter-cluster separations in eigenvector space. Note that class 6 has both the largest positive and the largest negative value, and appears to be a nodal point for the second and third eigenvector.

Class	Records	Vector 1	Vector 2	Vector 3
1	2422	+0.808	+0.067	-0.440
2	9374	+0.821	+0.296	-0.459
3	10799	+0.796	+0.467	-0.517
4	10861	+0.751	+0.653	-0.592
5	10564	+0.695	+0.835	-0.662
6	18303	+0.619	+1.000	-0.704
7	15369	+0.490	+0.903	-0.594
8	9544	+0.419	+0.729	-0.502

Table A2-6: Number of records in each class and cluster center locations in eigenvector space.

Class	1	2	3	4	5	6	7	8
1	0	23	41	62	82	100	92	78
2	23	0	18	39	60	78	71	60
3	41	18	0	21	41	60	54	46
4	62	39	21	0	21	39	37	36
5	82	60	41	21	0	19	23	34
6	100	78	60	39	19	0	20	40
7	92	71	54	37	23	20	0	21
8	78	60	46	36	34	40	21	0

Table A2-7: Normalized inter-cluster distance in eigenvector space, shown as a percentage.

The largest separations are between clusters 1 and 6, and clusters 1 and 7. Clusters 2 and 3 and clusters 5 and 6 are the closest together in eigenvector space, although there are several other pairs almost equally close together. Cluster 6, which is the cluster with the most records, is the furthest away from another cluster (Cluster 1) and also almost the closest to another (Cluster 5).

The space spanned by the first three eigenvectors is an artificial space: in order to get a feeling for what these classifications mean, it is necessary to look at some measured or observed characteristics of the sediments in areas classified differently. The results of *blending* can be used to do this.

The Data File

Data are provided in an ASCII format, as a comma-separated variable (CSV) file that can be read by any text editor or spreadsheet program, such as Microsoft Excel. The first line in the file is a header line that describes the content of each record. The first two records on each line are the position of the ABC record in meters of Easting and Northing in UTM Zone 10 co-ordinates (WGS84 datum). The next record is the depth of the ABC record in meters, and the next is the classification of the point, a number from 1 to 8. The final record in each line is the confidence of the classification in percent.

Blending

There were 1,238 grab sample grain-size results and 586 anecdotal descriptions available to match up with the ABC data. Using a radius of 20m and a minimum acceptable confidence level of 50% for the ABC classification, 6,506 ABC records were connected to 525 grab samples. The results of that analysis are shown in Table A2-8, Table A2-9, Table A2-10 and Table A2-11.

Class	#of records	%Sand	%Mud	Depth(m)
1	10	99.7 ± 1.0	0.3 ± 1.0	7.0 ± 0.9
2	49	99.2 ± 2.8	0.8 ± 2.8	9.8 ± 0.9
3	32	99.6 ± 2.5	0.4 ± 2.5	12.9 ± 0.7
4	24	96.7 ± 7.9	3.3 ± 7.9	19.3 ± 5.1
5	20	96.4 ± 6.9	3.6 ± 6.9	20.9 ± 2.0
6	60	96.8 ± 5.2	3.2 ± 5.2	35.3 ± 5.2
7	54	86.6 ± 10.6	13.4 ± 10.6	67.6 ± 7.5
8	45	79.1 ± 13.7	20.8 ± 13.7	90.4 ± 9.1

Table A2-8: Mean textural properties and depth of each class, shown with the 95% confidence level of the mean.

The data in Table A2-8 show that the classes are numbered by mean depth, Class 1 being the shallowest and Class 8 the deepest. The sediments are sandy: Classes 1,2 and 3 are nearly pure sands, Classes 4,5, and 6 have a trace of mud, and only classes 7 and 8 have any appreciable amount of mud.

Class	# of records	Mean	Sorting	Skewness
1	10	1.98 ± 0.22	0.54 ± 0.04	0.00 ± 0.04
2	49	2.09 ± 0.22	0.60 ± 0.15	0.11 ± 0.38
3	32	2.18 ± 0.25	0.59 ± 0.14	0.04 ± 0.27
4	24	2.34 ± 0.39	0.72 ± 0.38	0.11 ± 0.32
5	20	2.48 ± 0.36	0.71 ± 0.27	0.14 ± 0.33
6	60	2.52 ± 0.41	0.64 ± 0.15	0.28 ± 0.53
7	54	2.71 ± 0.52	1.13 ± 0.48	1.69 ± 0.59
8	45	3.08 ± 0.65	1.49 ± 0.57	1.33 ± 0.62

Table A2-9: Mean grain-size parameters (in Phi units) for each class, with standard deviations.

Class	2	3	4	5	6	7	8
1	0/0/0	*/0/0	*/0/0	**/0/0	**/*/0	**/**/**	**/**/**
2	-	0/0/0	**/0/0	**/*/0	**/0/0	**/**/**	**/**/**
3		-	0/0/0	**/*/0	**/0/*	**/**/**	**/**/**
4			-	0/0/0	0/0/0	**/**/**	**/**/**
5				-	0/0/0	0/**/**	**/**/**
6					-	*/**/**	**/**/**
7						-	**/**/**

Table A2-10: Significance ("0" = not significant, "*" = 95%, "" = 99%) from t-tests of differences in mean/sorting/skewness of each class. The table is symmetrical, so the bottom half is not filled in.**

Some comments can be made about the classes based on the textural properties. The mean grain size becomes uniformly finer (mean phi becomes bigger) with class number (and mean depth). Sorting becomes generally poorer with class number, and skewness becomes generally more positive. The tests of significance of differences in textural properties show that adjacent class numbers are not significantly different in terms of their textural properties. Of the adjacent pairs only 6/7 and 7/8 are significantly different. Class 8 is the only one that is significantly different at the 99% confidence level from all other classes for all three textural properties. Classes 1 through 6 have very similar textural properties: for example the sorting and skewness of classes 1 through 5 are statistically identical.

In order to try and differentiate these classes using some objective measure, the anecdotal reports must be used. A summary of the analysis of these data is shown in Table A2-11.

Class	N	% Firm	% Biota	% Worms	% Shell	% Molluscs
1	10	100	0	0	0	0
2	48	90	35	53	6	47
3	27	100	44	83	0	17
4	19	95	68	92	0	15
5	20	95	90	78	6	17
6	59	98	88	87	17	6
7	54	50	76	90	22	2
8	45	42	69	87	32	0

Table A2-11: Descriptive properties of the classes. The numbers are percentages, and 'N' is the number of anecdotal descriptions available for each class.

Table A2-11 shows the utility of adding the descriptive properties to the interpretation of the classes. The results in Table A2-8 indicate that Classes 1 through 6 are virtually identical insofar as some of their textural properties are concerned, but the descriptive data show that these classes differ in other ways. For example, Class 1 is the only class in which no biota were found. Class 2 is the only one of the shallowest 4 classes to have any shells or shell debris present, and had a very high incidence of live mollusks. Some general trends can be seen in the data: Classes 1 through 6 are firm in texture, likely because they contain very little mud, and Classes 7 and 8 are commonly loose in texture. Biota tend to be most common at intermediate depths; worms or worm tubes were the most common biota to be noted; mollusks are most common in shallower depths, and shells and shell debris is most common at depth.

Maps

Figure A2-2 is a map showing the classification of areas along the vessel track over the entire study area, and Figure A2-3 is a closer view of the river mouth. The colors of the classes are given in Table A2-12. The points plotted in the map are those for which the confidence of the classification is greater than 50%. Regions in which classification points are sparse are usually found at the boundary between two classes where confidence levels are low. Such regions are more common between deeper classes (*e.g.* 6 and 7 and 7 and 8).

The pattern of depth-dependence of the classes is clear in Figure A2-2. Note how the classification bands follow the bulge of the contour lines around the river mouth. Some of the deeper classes (4, 5 and 6) are found in the dredged channel and south of Jetty 'A' where depths are deeper due to dredging and scour.

There does not seem to be any ABC "signature" associated with material in the dredge disposal sites: no anomalous patterns are seen associated with these sites, and there is no evidence in the ABC of any difference between sediment at those sites and the surrounding sediments.

Class Number	Color
1	Red
2	Orange
3	Yellow
4	Green
5	Cyan
6	Blue
7	Violet
8	Purple

Table A2-12: Colors used to identify classes in the maps.

Figure A2-2: (Overleaf)Classification map of the study area.
Figure A2-3: (Overleaf)Classification map of the area around the river mouth.

Figure A2-2: Classification map of the study area.

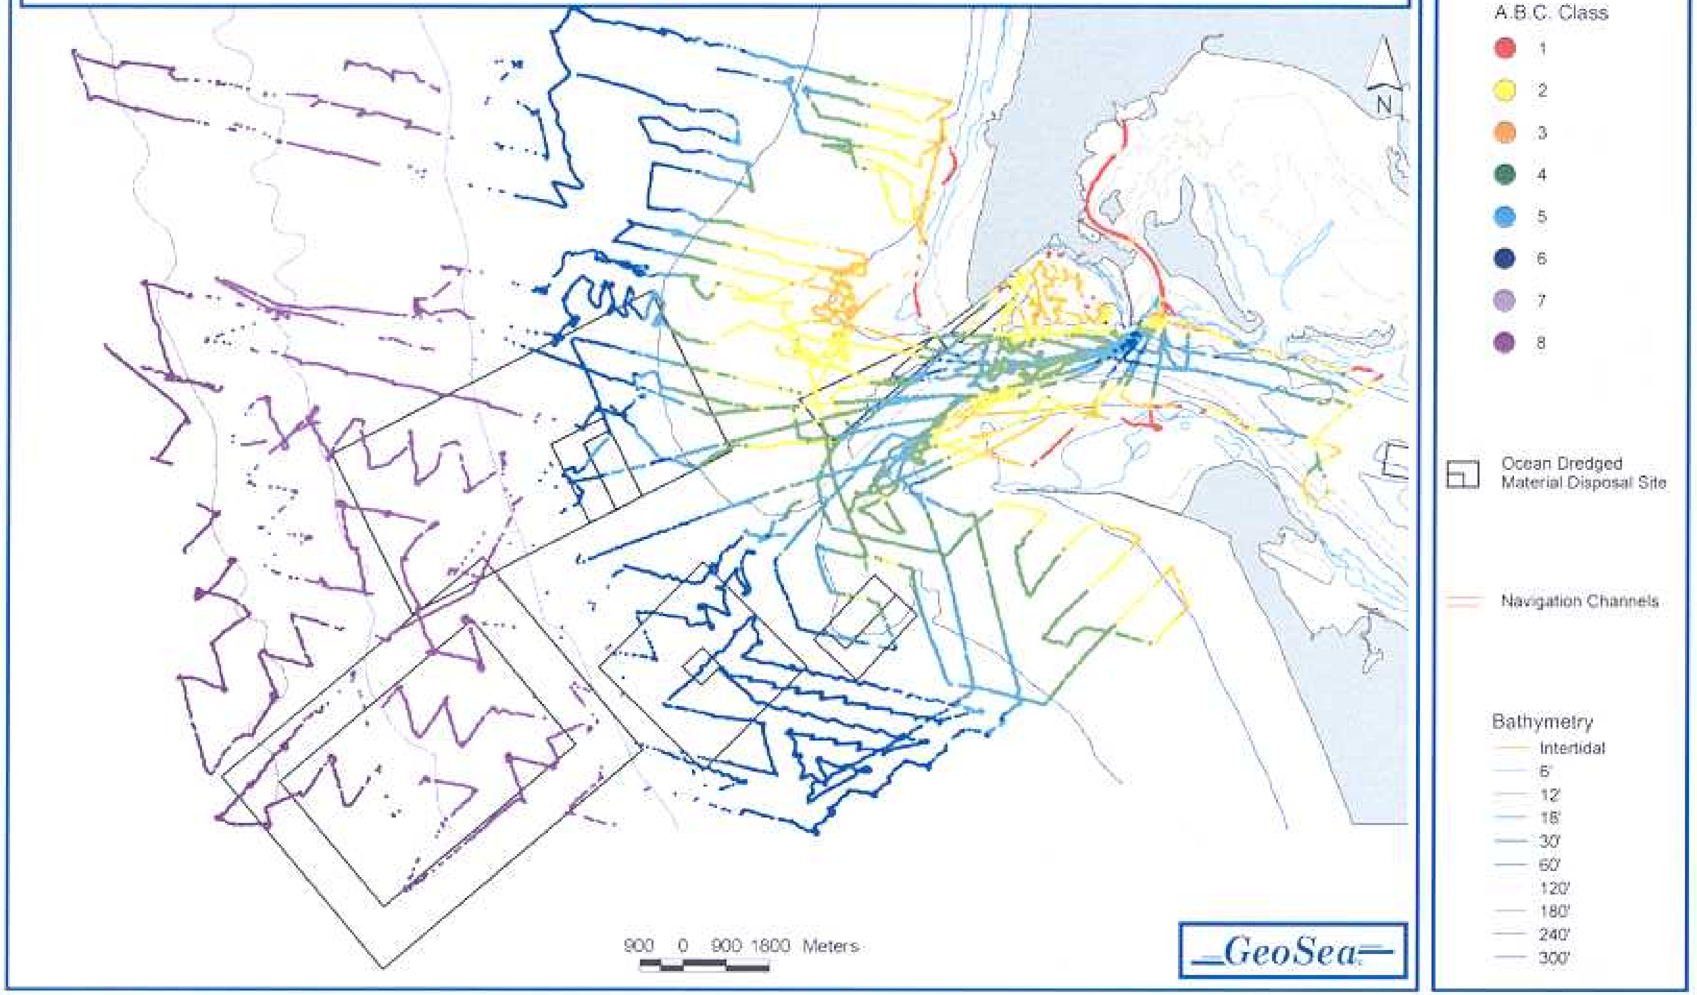

Figure A2-3: Classification map of the area around the river mouth.

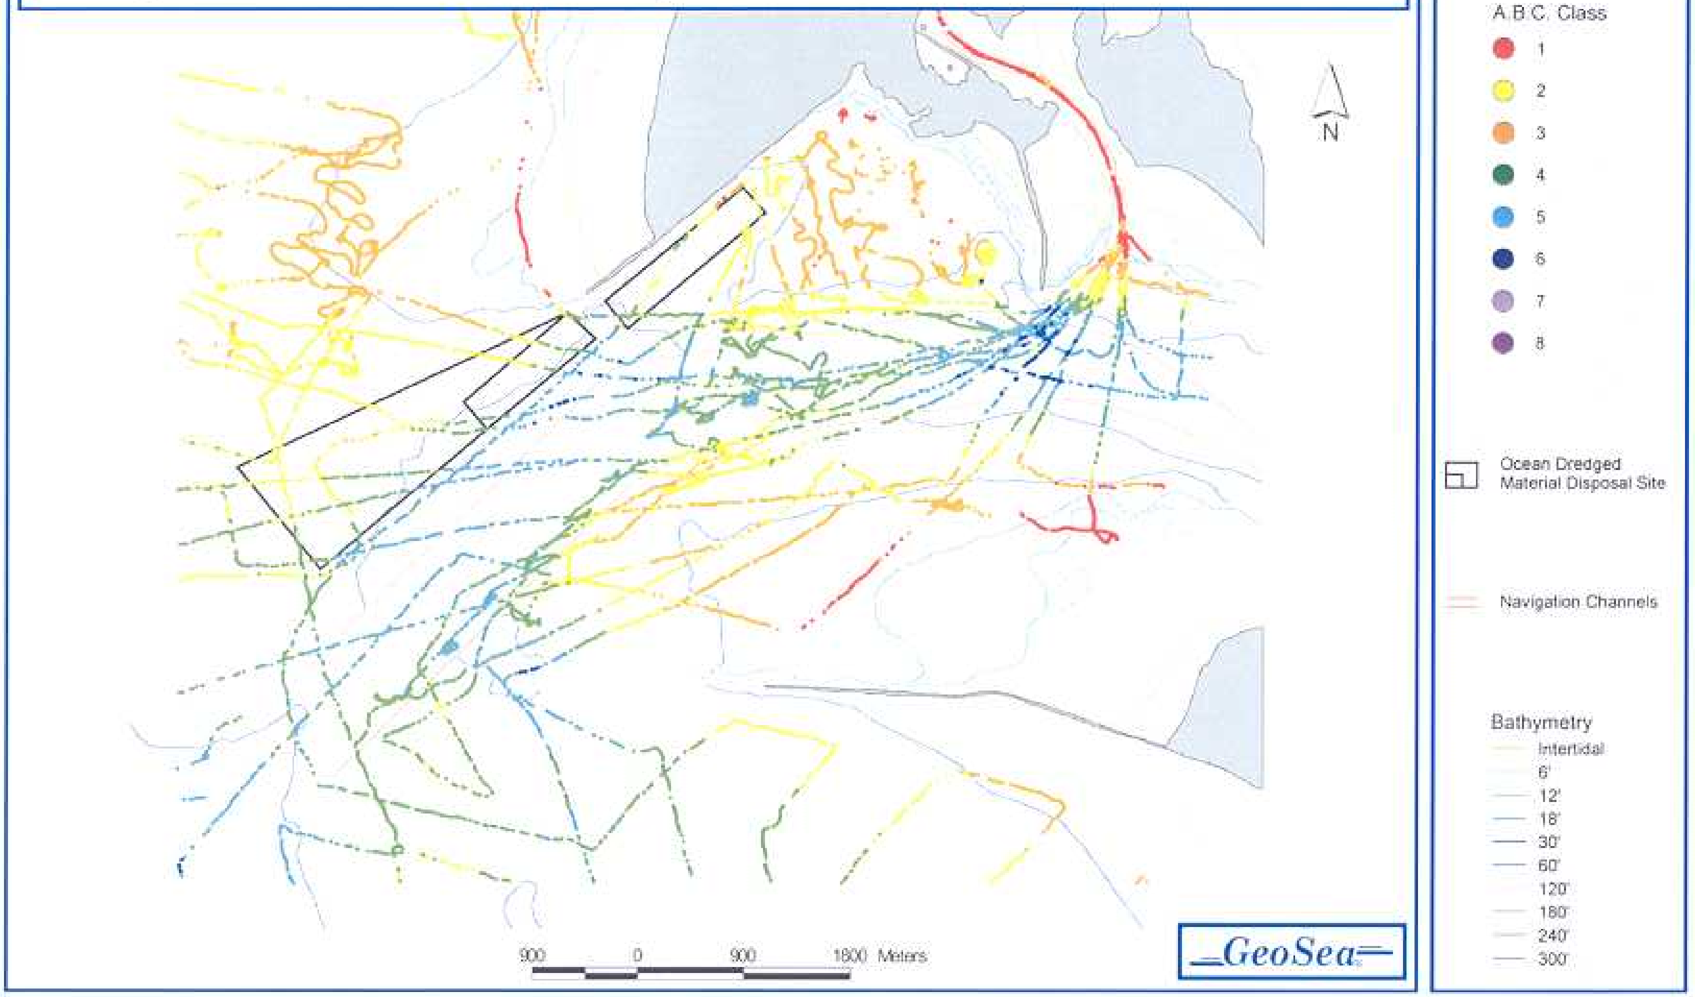

Summary

The following are descriptions of each of the eight classes.

Class 1: Medium pure sand, well sorted, un-skewed. The texture is firm, with no biota present. Found close inshore near Peacock Spit, on the northern edge of Clatsop Spit, and up the channel to Ilwaco.

Class 2: Medium to fine pure sand, well sorted and un-skewed: the grain-size parameters for this class are statistically identical to those of class 1. Generally firm in texture with biota sometimes found: worms/worm tubes and live mollusks are about equally common. Shell debris found occasionally. Found inshore between the North Jetty and Jetty 'A', along the western edge of Clatsop spit, and in shallow areas north and south of the South Jetty.

Class 3:. Medium to fine pure sand, well sorted and un-skewed: the grain-size parameters for this class are statistically identical to those of class 2. Biota slightly more common than class 2, but almost exclusively worms and/or worm casings; mollusks occasionally present. Found along the edges of the dredged channel between the North Jetty and Jetty 'A' and north of the South Jetty, and in a depth-related band along the western edge of Clatsop and Peacock Spits.

Class 4: Fine sand with a trace of mud, well sorted and un-skewed: the grain-size parameters for this class are statistically identical to those of class 3. Firm in texture, two-thirds of samples had biota present, almost entirely worms and/or worm casings with occasional mollusks.

Class 5: Fine sand with a trace of mud, well sorted and un-skewed: the grain-size parameters for this class are statistically identical to those of class 4. Generally firm in texture; 90 percent of the samples contained biota, with worms and/or worm casings common, and occasional mollusks and shell debris. Found in the dredged channel and in a band spanning the study area from north to south, roughly following the depth contours and just offshore (and therefore slightly deeper) of the locations of Class 4.

Class 6: Fine sand with a trace of mud, well sorted and un-skewed: the grain-size parameters for this class are statistically identical to those of class 5. Almost all samples were firm in texture, and biota was present in 88 percent of samples. Worms and worm casings were most common, followed by shells and shell debris and lastly live mollusks. Found in a small area just south of Jetty 'A' and in a band spanning the study area from north to south, following the depth contours and offshore and deeper than Class 5.

Class 7: Fine sand with more than 10% mud. Less well sorted than the previous six classes, and positively skewed. Half the samples were loose in texture and half firm, and biota were less common than in Classes 5 and 6. Almost all samples with biota contained worms and worm casings, with occasional shells and/or shell debris. Found in a depth-related band near the western margin of the study area.

Class 8: Very fine sand with more than 10% mud. Not well sorted, and positively skewed. More than half the samples were loose in texture and the rest firm, and biota were slightly less common than in Class 7. Almost all samples with biota contained worms and worm casings, with occasional shells and/or shell debris. Found in a depth-related band at the western margin of the study area.

References

F. Murtagh and A. Heck, 1987. *Multivariate Data Analysis*, Kluwer Academic, Dordrecht.

J. Hartigan, 1975. *Clustering Algorithms*, John Wiley & Sons, Toronto.

R.L. Kirlin and R.M. Dizaji, 2000. *Cluster Order Using Clustering Performance Index Rate, CPIR*, Proceedings of the IEEE Nordic Signal Processing Symposium (NORSIG 2000)