Proceedings of the Institute of Acoustics

THE ISIS INTERFEROMETRIC SEABED INSPECTION SONAR

M. F. Geen, P. D. Hewitt, A. R. Adams

Submetrix Ltd., 76 Lorne Road, Bath, BA2 3BZ, U.K.

1. INTRODUCTION

This paper describes the ISIS Interferometric Sea-bed Inspection Sonar, showing how it provides close-coupled bathymetry and sonar amplitude data, and thus provides a new tool for scientists, hydrographers, and engineers.

2. OVERVIEW OF ISIS

2.1 History of ISIS.

During the 1960s it was found that, under certain conditions, sidescan sonars exhibited interference fringes, similar to those found in optics, and that these fringes were affected by the sea-bed geometry. Research into these effects was carried out during the 1970s. In the 1980s, this research resulted in several working systems, notably those of Denbigh [1] at Birmingham, and Cloet and Edwards [2] at Bath. In 1991 Submetrix was formed, to design and produce an interferometric swath-sounding system, using this work, but using up-to-date electronics, computing hardware, and design methods. This system is called ISIS, and will be commercially available in early 1993.

2.2 Functional Description of ISIS.

The basic function of ISIS is the measurement of depths in a "swath" either side of the surveying vessel. ISIS is capable of greater angular range, resolution and accuracy than is possible with other swath systems, and fully integrates amplitude "sidescan" data with the "bathymetry" depth measurements. ISIS consists of sonar electronics and computing hardware (Fig. 1), with real-time and post-processing software (Fig. 2). The sonar platform includes sonar transducers, processing electronics, and attitude measuring system. This platform is a tow-fish in the "standard" configuration, but the system has been designed in a modular manner, so as to allow easy mounting in hull-mounted pods, ROVs, etc. The sonar platform is connected to a computer system by a power and communications cable. Real-time data acquisition, control and processing software provides graphical seabed information and quality control, and records data for the post-processing software. The postprocessing software combines the sonar, attitude, position and tide data to build a depth model of the seabed. This depth model may be integrated with the sonar amplitude data, a process which is subject to continuing research.

The system transmits sonar pulses, or "pings", and uses the returned signal to measure depths along a "profile" at right angles to the track of the vessel. A series of these profiles forms a "swath", and a set of overlapping swaths is merged together to form a "surface model".

2.3 Market Aims.

ISIS is designed primarily as a hydrographic surveying tool, for use within the continental shelf. The main users are expected to be hydrographers, oil engineering companies, and port authorities. However, the addition of the new feature of coupled amplitude data may well result in new applications.

ISIS INTERFEROMETRIC SEABED INSPECTION SONAR

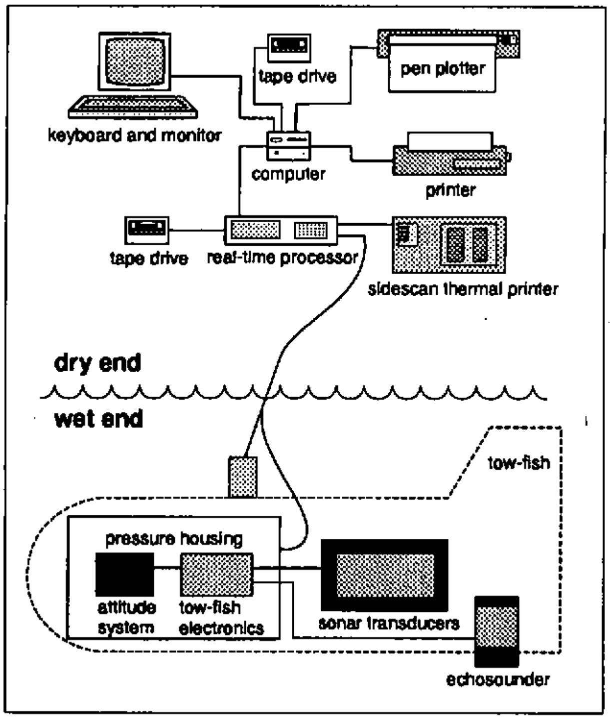

Fig 1 ISIS Hardware Block Diagram

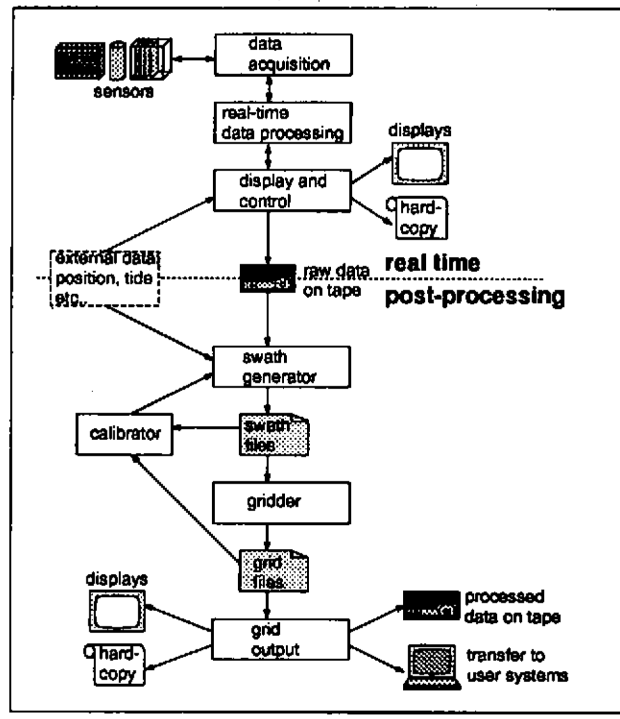

Fig. 2 ISIS Software Block Diagram

3. DATA ACQUISITION AND CONTROL

The real time survey system acquires, processes, displays and stores the survey data during the survey. The structure of this system is shown in Fig. 3. The system is divided into two main sections, wet and dry, separated by an umbilical carrying power and communications.

The wet unit is located inside the towed body. This section contains the main sonar, acquisition and processing electronics and various sensors.

The sonar section transmits a sonar pulse into the water under Transputer control. This insonifies the area of seabed to be analysed (profile). The return signal is then detected by the receive staves in the transducers and passed to two sonar electronics blocks. The first block isolates the return signal at the transmitted frequency for each stave channel, these signals are then passed to the phase meter. The second block feeds the signal of one of the staves through a time varying gain stage controlled by the Transputer network. The output of this stage is then filtered to produce an analogue signal related to the reflectivity of the seabed insonified. This signal is digitised and read by the Transputers to produce sidescan style data.

The phase meter produces a digital value proportional to the relative phase of the signals on the stave channels. The values are read by the Transputers along with sample time and status information.

According to the depth and seabed characteristics the Time Varying Gain (TVG) law applied to the analogue (sidescan) signal must be adjusted to maintain the signal integrity. This is done via the TVG interface. The Transputer downloads a number of digitally generated gain laws into the TVG hardware then selects the most appropriate law for the current depth and seabed conditions. If no stored gain laws are appropriate then the least appropriate gain law is over-written by the Transputers with a newly generated law.

ISIS INTERFEROMETRIC SEABED INSPECTION SONAR

A number of environmental sensors are required to enable the phase and amplitude data to be properly corrected and positioned. These sensors are connected to the Transputers via standard serial interfaces. The echosounder is used to measure the depth directly below the towed body. An integrated attitude package is included which contains various internal sensors that are compared and combined with the external pressure sensor and magnetic compass to produce comprehensive data on the position and orientation of the towed body.

The Transputers control the sonar transmit pulse and then asynchronously acquire the phase, amplitude and environmental data over a period determined by the range required. The raw phase samples are then submitted to a sequence of software filters. These filters convert the phase and timing values to elevation angles and ranges, then process these to remove noise, anomalies and to reduce the data size. When all the data is acquired and processed it is then packaged into a block representing a complete ping before transmission to the dry unit via the umbilical.

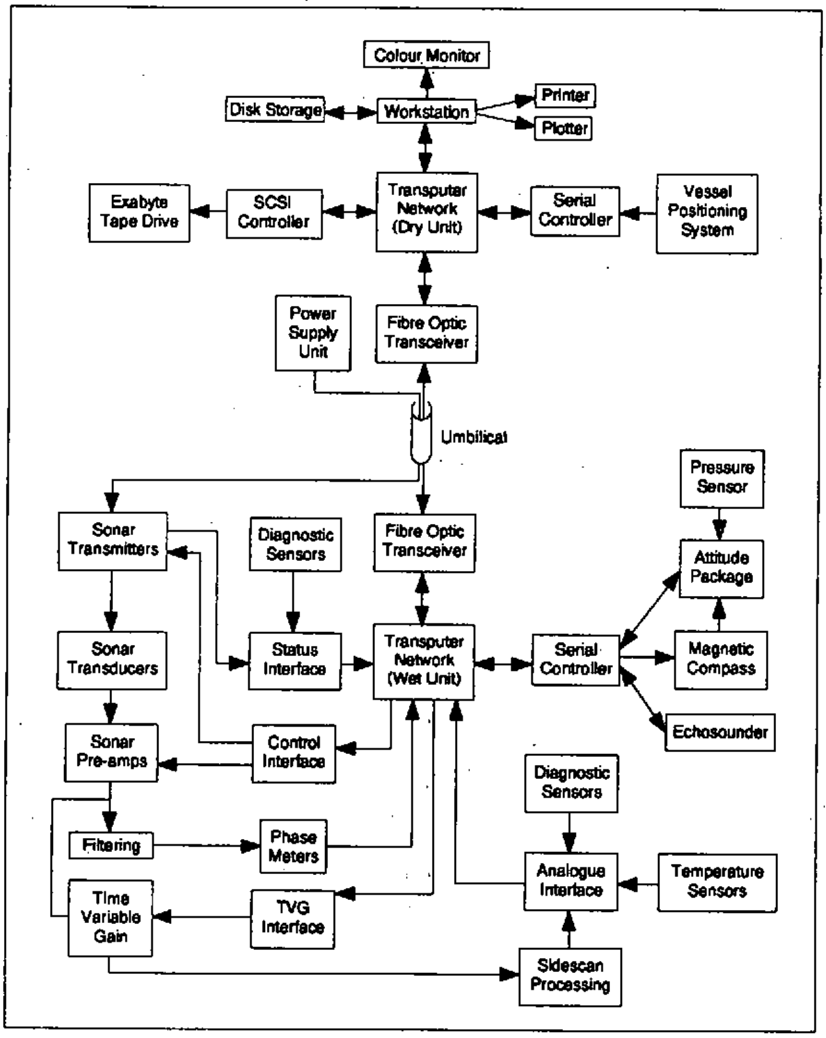

Fig. 3 Real-time System Block Diagram

The umbilical contains two power lines, one D.C. supply for the system electronics and one A.C. H.T. supply for the sonar transmitters. The umbilical also contains optical fibres for bidirectional digital control and data communications.

ISIS INTERFEROMETRIC SEABED INSPECTION SONAR

The dry unit is located on board the survey vessel. This section contains the power, control, display and storage facilities.

The Transputer network in the dry unit receives data packets transmitted from the wet unit. Added to each set of ping data is the information received from the vessel mounted positioning system, (this may also include vessel heading), via a standard serial interface. This complete block containing ping data and various additional information selected by the user is then stored on magnetic tape (Exabyte 8 mm helical scan format), for later post-processing.

A selected subset of the ping data is then passed through a sequence of display preprocessing filters on the Transputer network this display data is then passed to the Sun Workstation.

The workstation can display various diagnostic information for calibration and research purposes. The usual quality assurance display used during a survey would output a colour-encoded depth waterfall, showing the depths in the port and starboard profiles over the last few minutes. Any environment values (such as attitude, roll, pitch, etc.) can also displayed with their recent history alongside the depth display to cross-reference these parameters with the quality of the ping data at any one time. A further option is to emulate a sidescan recorder on the screen utilising the ping amplitude data.

Control options, such as profile range, survey line delimiters, etc., can be selected from the workstation and this information is sent down to the Transputer network for processing. Textual information relating to surveys, vessels, survey lines or any condition required to be reported may be entered by the user and recorded along with the ping data on the survey tape. This, and other control or status information can also be logged on a line printer for hardcopy output.

Using a 200 millisecond ping period a 150 metre profile range (each side) can be captured. This also allows 5 pings to be produced each second. These survey parameters produce around 1000 to 20000 sets of raw phase samples (around 10 to 200 kbytes) per ping. These are then filtered to produce around 1000 to 3000 angle/range samples (around 10 to 30 kbytes), which are stored on tape. Diagnostic options are available to store raw data values but only at a much reduced rate due to the bandwidth limitations of the tape drive. A range of ping periods can used from 100 ms (75 m range, 10 pings per second) to 400 ms (300 m range, 2.5 pings per second).

4. REAL-TIME DATA PROCESSING AND INTERFEROMETRY

The ISIS system currently has a pair of sonar transducer arrays ; one each side of the towfish. Each array has a transmitter which produces a beam narrow in azimuth but wide in elevation. In addition to the transmitter each array has multiple receive staves. Each receiver measures the phase of the returned signal around one hundred thousand times per second. The system simultaneously samples the amplitude of the signal tens of thousands of times per second. By comparing the phases measured on spatially separated receivers a measure of the difference in path length covered by the signals can be calculated. See figure 4.

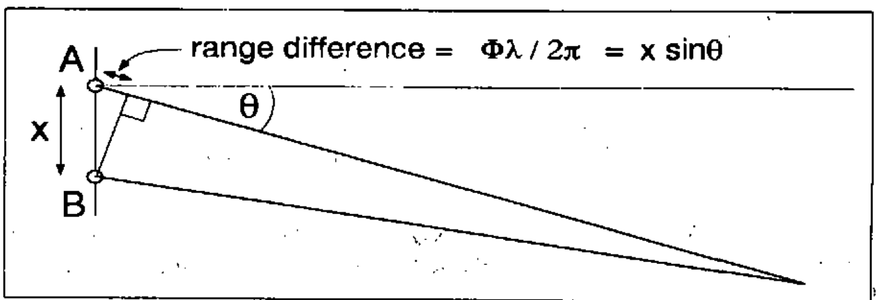

Fig. 4: Geometry of Phase Measurement

ISIS INTERFEROMETRIC SEABED INSPECTION SONAR

The hardware cannot measure phase shifts of +/- $2n\pi$ (n any integer), as far as it is concerned a phase shift of 3.5π is indistinguishable from a 1.5π phase shift and will be measured as 1.5π. This hardware limitation means that the phase, and therefore the path difference, cannot be unambiguously measured unless it is known to actually lie within the range 0 to 2π radians. For receivers spaced a half wavelength apart the path difference must lie between $+\lambda / 2$ and $-\lambda / 2$, expressed in phase terms this is $+\pi$ to $-\pi$ radians. This means that the phase measured is unambiguous so θ can be uniquely determined using the following formula ;

$$\theta = \arcsin(\phi \lambda / 2\pi x) \qquad (1)$$

where ϕ is the actual phase difference between staves A and B. λ is the wavelength of the sonar signal.

Limitations to the accuracy with which the phase can be measured and the effects of noise on the signals, however, mean that this value for θ is too poor to use on its own. Looking at equation (1) we can see that the effect of phase errors and resolution limits on θ can be reduced by increasing x, the receiver separation. By using receivers 2λ apart the effect of phase errors on $\sin(\theta)$ is quartered, but this is where hardware limitations mentioned above take a hand. The measured phase ϕ_m can only lie between 0 and 2π, whereas the actual phase ϕ can lie anywhere between -4π and 4π. Physically possible solutions for ϕ, the actual phase difference are given by equation (2).

$$\phi = \phi_m + 2n\pi \qquad -2 <= n < 2 \qquad (2)$$

where ϕ_m is the phase measured by the hardware.

This means that there are four possible solutions to equation (1).

So we have a trade-off; by increasing receiver spacing we gain resolution but also increase the number of solutions for θ. Fig. 5 plots the phase of the incoming signal on the Y axis against the solution(s) for θ, the incoming signal angle, on the X axis, for each of the receiver spacings. On each of the plots there are two horizontal lines plotted a tenth of a radian apart, these are intended to show the effects of phase resolution and noise on the decoded value of θ. Note that as the transducer separation increases the decoded values for θ, the vertical line pairs, converge.

Fig. 5 also shows that if the results of all the receiver spacings are combined we can use the unique solution from the ½λ plot to choose the correct higher resolution solution from the 4λ spacing plot.

The result of the phase decoding process is θ (see figure 4), an angle relative to the transducer array. Range is similarly relative to the array. To translate this array relative sample pair to an absolute, geographic, position we must apply the fish's roll, pitch, yaw, pressure depth, heave etc. These are measured tens of times a second and are applied, in real-time, to give the user accurate images of the sea-bed currently being surveyed.

Each profile consists of several thousand sets of range, angle and amplitude data, a sample set approximately every 7.5 cm out along the profile. Any, or all, of this data can be made available to the user for research purposes, for example bottom texturing and object location, as well as for more mundane uses, such as sidescan output corrected for horizontal range. If the system is to be used for charting purposes, however, ten thousand samples of depth, position and amplitude every second are rather unwieldy. A ten kilometre square survey would result in around a thousand million sample sets. A useful amount of data to a researcher, possibly, but prohibitive in processing time for a commercial charting system. This density of data is still useful at real-time, though, as it allows a statistical estimate to be made of the effect of noise on the system at successive points along the profile. Thus an estimation of the quality of the data can be obtained for each sample. This noise estimate can be expressed in terms of the mean deviation of the angular error in θ, the decoded phase angle. This allows the user to set a physically meaningful limit on the noise of the system. The number of sets of range, angle and amplitude sent to tape and to the display is adjustable at run-time; the default is currently around 200 per ping, which gives a sample every 75 cm, approximately, along a 150 m profile.

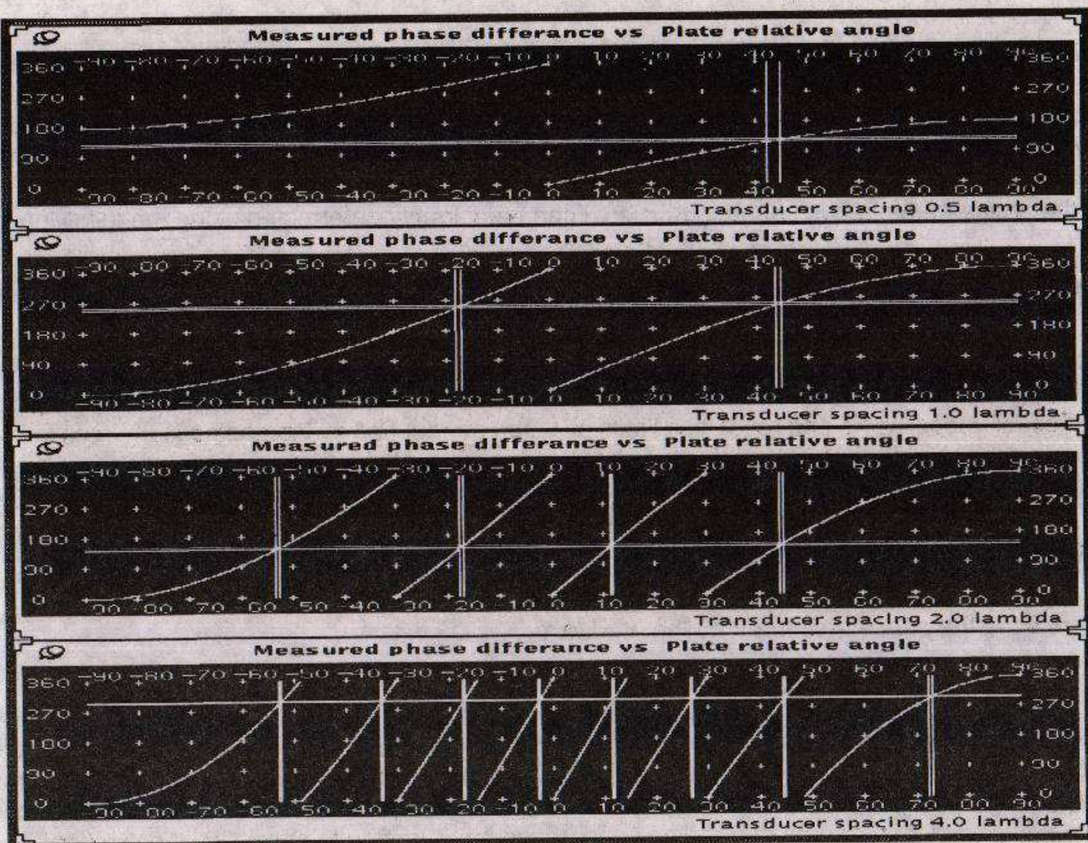

Fig. 5: Measured Phase Difference vs. Plate Relative Angle

The gain laws used to set the amplifiers for the amplitude signal digitisers are calculated at run-time by software running on one of the Transputers in the wet unit. This allows the gain-law to be optimised for the sonar conditions encountered. The system records, several thousand times a ping, the absolute amplitude of the returned signal. This allows the data to be reprocessed to extract different aspects of the data. The availability of the angle of the reception angle, θ, allows the signal to be corrected for the beam-pattern in the vertical axis. With depths available every 7.5 cm along the surface an attempt can be made to estimate, and if desired to remove, the effect of the orientation to the surface in three dimensions of the sonar beam. This allows researchers and users to get closer to the pure effects of the bottom surface on the returned data.

5. POST-PROCESSING

5.1 Purpose of a Postprocessing System.

The availability of fast, distributed computer hardware makes it possible to process the sonar data in real-time. However, post-processing is still necessary, for several reasons:

some data is not available in real-time (particularly tide);

the uncertain nature of marine work requires more flexibility and "fault-tolerance" than is possible in real-time;

the achievement of the best accuracy and quality control will always require a certain amount of involvement by the surveyor; and building a surface model requires the combination of data from many different times and places.

5.2 Principal Outputs of the Post-Processing System.

The main aim of the post-processing software is to convert the sonar, attitude, position tide and other environmental data into a numerical model of the surveyed sea-bed. This model includes depth, and whatever information is extracted from the amplitude data, such as surface reflectivity or bottom type. The user is then provided with the tools required to extract the data which he or she requires for the particular application.

The surface model is built in two stages. The first stage converts the sonar data into depth, amplitude and position data, on a ping-by-ping, swath-by-swath basis. The output of this stage is in the form of "swath files", and this stage is referred to as "swath generation". The swath files are in a format which may be used directly by the user, if required. However, missing out the surface model generation stage skips some important quality control and calibration features.

The swath files are combined together into a surface model. This model is in the form of a "grid" of square "bins", each of which contains all the information from the swath files which falls within its area. The final surface model is thus contained in a "grid file", and the process generating it is the "gridder".

As well as displays and other quality control information available to the user throughout the post-processing, the gridded surface model (called the "grid" from now on) may be displayed, plotted and printed in a number of formats. The grid files themselves are in a form accessible to the user, but software tools are available to extract the data, or sub-sets of it, in a variety of formats. In addition, software "bridges" to industry-standard hydrographic "end-user" software are under development.

The software also provides for "calibration" of the system by cross-correlating data from adjacent swaths. This calibration data is then fed back to the swath generation stage for a second pass through the data.

5. 3 Features of the Post-Processing System

The main design aims of the post-processing system are:

- accuracy and data confidence: the large amount of overlapping data ensures that each data point is cross-checked against many other samples, vastly reducing the possibility of error;
- ease of use: the software uses an industry-standard graphical user interface throughout;
- speed: the software is written for powerful UNIX workstations;
- size independence: wherever possible, the software places no constraints on data quantity; and
- flexibility: the processing system is not be affected by changes in the surveying hardware, and provides a degree of "fault-tolerance" through data correction and choices of data sources.

5. 4 Swath Generation

After the real-time data is read from tape, it is filtered to reduce "noise". The tape data is in the form of sets of elevation angles and sample times, relative to the sonar transducers. Each set is converted to a Cartesian vector, using the measured speed of sound. It is then transformed to a depth and position vector, using the attitude data. As the "sidescan" amplitude data is also recorded with sample time, the depth data and amplitude data are very closely coupled.

Before writing to the swath file, the depth data may be corrected for speed of sound variations in the water-column, and filtered to remove any data which differs strongly from neighbouring data. Each data point is tagged with a quality-factor, or "weight", which reflects the quantity and quality of tape data from which it was derived. Finally, each "ping" in the swath file is given a summary of the attitude and position data which was used in its calculation, thus allowing for a degree of reconstruction of the tape data from the swath files.

5. 5 Gridding

Having defined the parameters of the grid model, the data from selected swath files is merged into it, taking account of the data weight factors. This process takes allows for the large swath width of ISIS when compared to other systems, permitting the data to be overlapped in a controlled manner.

5. 6 Calibration

The calibration process requires two adjacent swaths to be gridded into separate grids. Using the position and attitude summary recorded with each ping in the swath files, a line of depths is extracted from each grid where the swaths overlap, for the same position on the sea-bed. These two lines are compared for angle and overall depth, and the comparison data for many such lines along the grid are tabulated and displayed. This process not only corrects for angular offset errors in the transducer mounting and attitude measurement, but also points to errors in position and tide measurements. The results of the calibration process are fed back into a second pass though the processing.

5. 7 Amplitude data

Although the angles and ranges obtained from the interferometer contain sufficient information to generate a depth model, the amplitude portion of the sonar data enables additional use to be made of the data. For example, the amplitude signal helps to identify artifacts on the sea-bed, such as wrecks or pipelines. It also gives information on the geological make-up of the sea-bed.

A number of possibilities exist for the combination of amplitude and depth data. The first is at the level of the swath files, before the gridded surface model is generated. At this stage, the amplitude data may be used "as is", but presented in the correct plan-position, accounting for depth variations along the swath. This is not possible with normal sidescan systems, as only the slant-range to the sea-bed is known. This data could be presented graphically with depth and amplitude on the same display, using, for example, colour, brightness, and contours together to represent all the data.

The more challenging task is to combine amplitude data into the surface model. In this case, data from several overlapping pings must be merged into the same grid bin, with very different "views" in both elevation and azimuth. Each bin could be given a reflectivity value, a "texture" value, a set of factors representing the back-scatter law at that point, or a combination of these. Whatever algorithm is used, all the data possible is available, as the gain value applied to the amplifier is available along with the signal itself, thus allowing the transducer signal to be calculated. As the beam-pattern of the transducer is known, and the angle of incidence on the transducer is measured, the signal level at the transducer face may be reconstituted directly from the tape data. Also, the location and attitude of the transducer is known, as is the location of the patch of sea-bed which returned the sonar signal.

The actual processes used will finally depend on what the end-user requires from the data. There is still much research to be done in this area.

6. REFERENCES

[1] P. N. DENBIGH, "Stereoscopic visualisation and contour mapping of the sea bed using a bathymetric sidescan sonar (BASS)", The Radio and Electronic Engineer, Vol. 53, No. 7/8, pp.301-307 (1983)

[2] CLOET, EDWARDS, DUNCAN, "The University of Bath Swath Sounding System", Proc. Inst. Acoustics, Vol. 7, Pt. 3, (1985).

[3] C. MALLIE, "Wide Swathe Sounding in Netherlands North Sea", Hydrographic Symposium, IHO, 11.3.1 (1992).