

THE GEOLOGICAL INTERPRETATION OF SIDE-SCAN SONAR

H. Paul Johnson
School of Oceanography
University of Washington, Seattle

Maryann Helferty
Geophysics Program
University of Washington, Seattle

Abstract. Recent developments in side-scan sonar technology have increased the potential for fundamental changes in our understanding of ocean basins. Developed in the late 1960s, "side looking" sonars have been widely used for the last two decades to obtain qualitative estimates of the acoustic properties of the materials of the seafloor. Modern developments in the ability to obtain spatially correct digital data from side-scan sonar systems have resulted in images that can be subsequently processed, enhanced, and quantified. With appropriate processing, these acoustic images can be made to resemble easily recognizable optical photographs. Any geological interpretation of these images requires an understanding of the inherent limitations of the data acquisition system. When imagery is collected, these limitations are largely centered on the concept of resolution. In side-scan sonar images, there are several different types of resolution, including along- and across-track resolution, display resolution, and absolute instrumental resolution. All of

these parameters play a critical role in our ability to calibrate and ultimately to interpret the new pictures of the ocean floor. Acoustic image processing is a new application of an old and well-established technique. Digital optical images have benefited from several decades of development in processing techniques. These relatively sophisticated techniques have been applied to photographic images from satellites and spacecraft, images which are "noisy" and difficult to obtain but extremely valuable. Side-scan sonar systems, on the other hand, have only recently been able to produce spatially correct, digital images of the seafloor. The application of digital signal-processing techniques to side-scan sonar data will now allow us to quantify what had been previously very subjective and qualitative interpretations of images of the seafloor. The goal of all this processing of acoustic images remains clear: the development of an interpretable map of the geology of the seafloor.

INTRODUCTION

Our early perception of the deep ocean floor as a featureless, static environment has undergone dramatic modification in the last 50 years. Early depth sounding with mechanical devices, and even early wide-beam acoustic echo sounders, gave us an extremely low resolution picture of the ocean basins. This early image of the ocean bottom consisted largely of a flat lying seafloor with few hills and ridges of any consequence, completely covered with a thick layer of sediment, and only an occasional, inexplicable outcrop of hard rock. The few recovered rock samples were of little value in understanding the scientific processes of the deep sea because they could not be placed within any sort of geological and morphological context.

Acceptance of seafloor spreading and the Vine/Matthews hypothesis in the 1960s altered forever our perception that the floor of the ocean basins was unchanging, at least on a geological time scale. The initial drilling efforts of the Deep Sea Drilling Project also modified our

view about the uniformity of the seafloor, even on a scale of a few hundreds (and perhaps tens) of kilometers. Higher-resolution bathymetry maps, using multiple narrow-beam echo sounders, strongly reinforced this new perception of a nonuniform and scientifically interesting seafloor [Tyce, 1986]. The initial interest grew to excitement as our perspective was extended by visual observation—in a very few places—down to the scale of a few meters by early submersible expeditions to mid-ocean ridge spreading centers [Ballard and van Andel, 1977].

Clearly, a new tool, beyond the wide-beam echo sounder, was needed to map features of the seafloor and to understand the processes at work there. To be effective, this would have to be a tool which had both a sufficiently wide "view" for tectonic context and adequate resolution for interpretation in terms of geological processes. The efficiency of sound transmission in seawater, and the development of both electronic and digital techniques capable of rapidly processing the high data rates necessary to generate images, dictated that this tool would be some form of acoustic swath mapping. The scientific need for

the appropriate instrumentation has resulted in the development of side-scan sonar into a new "swath mapping" tool (Figure 1), one that can map the physical properties of the surface of the seafloor, and do it quickly, cheaply, and over a wide area [Tucker, 1966; Edgerton, 1966; Fleming, 1976; Mudie *et al.*, 1970; Klein and Edgerton, 1968; Andrews and Humphrey, 1980; Laughton, 1981; Reut *et al.*, 1985].



Figure 1. GLORIA image from the Washington margin, showing two anticlinal hills at 2000 m depth. Bright pixels are hard acoustic returns, and dark pixels are soft returns or acoustic shadows (the traditional GLORIA polarity). The bright pixels on the left flank of the hill on the right-hand portion of the image are returns due to proposed carbonate deposits associated with the dewatering of sediments during subduction. The anticlinal hill in the lower left of the image is of similar topographic relief but does not show comparable bright reflections. (Data from U.S. Geological Survey.)

A further major advantage of side-scan sonar is that it can produce spatially correct data that can be viewed as an image. The resulting "image" is a familiar photographlike representation (Figure 2) that can be enhanced, modified, and ultimately interpreted, much as a land geologist interprets the structure of a roadside rock outcrop or an aerial photograph. The value of this ability to visually interpret large amounts of data as recognizable images is not lost on anyone who has, for example, studied multi-channel seismic records or has been disoriented in a labyrinth of potential field contour lines. While side-scan sonar as a technique has not yet achieved the full potential of the initial promise, it is not an exaggeration to say that it has radically changed our perception of the seafloor and our understanding of the geological processes at work there.

The development of side-scan sonar has evolved to the point where we can now view these acoustic data as spatially correct images. These digital image data are "correct" in the sense that all of the acoustic targets are in the same undistorted spatial relationship to each other as they are on the seafloor. A series of mathematical operations can subsequently be applied to the image to increase our ability to interpret the data in terms of geological processes; the operations which accomplish this are grouped into the general category called "image processing." Basic digital processing techniques for optical images have been in existence since the early 1960s [Andrews, 1968; Campbell, 1974; Rosenfeld, 1976; Rosenfeld and Kak, 1976]; however, application of these standard image processing techniques to acoustic images of the seafloor is a relatively new phenomenon. With the recent application of digital processing to acoustic side-scan data, we have moved to what was state-of-the-art for optical images about 25 years ago.

The direct transfer of digital image-processing techniques from optical to acoustic images suffers from several inherent problems. First, unlike the customary multifrequency optical images, where color conveys a tremendous




Figure 2. SeaMARC 150 side-scan image of a PB4Y aircraft at the bottom of Lake Washington (Seattle, Washington). Swath width for this image was 100 m, and the grid lines are 10 m apart; water depth is 27 m. The starboard wing is estimated to be 5 m above the lake bottom, causing the apparent image distortion, and the port wing is partially buried. Polarity of this image is the same as in Figure 1. (Data from J. Kosalos.)

amount of information, side-scan images are monochromatic and, at best, consist of a small number of shades of grey. Second, the acoustic images are orders of magnitude lower in resolution than those we are accustomed to observing visually. Finally, and perhaps more important, the side-scan image is not a representation of how the seafloor would look if the water were somehow removed from the ocean. Instead, it is a graphical presentation of how the seafloor interacts with acoustic energy. This conversion, from how the seafloor "sounds" to how our geological models tell us the seafloor should "look," can be a major pitfall for the interpreter of the images.

FUNDAMENTALS OF SIDE-SCAN SONAR

Side-scan sonar is a logical extension of the same basic acoustic principles used in the wide-beam echo sounder (Figure 3a). The echo sounder has served marine geology well since its development in the early 1920s [Vogt and Tucholke, 1986; Urlick, 1983]. The basic echo sounder consists of (1) a transmitter, which emits sound downward into the water column, (2) a receiver, which detects the reflected acoustic energy, and (3) a clock, which measures the elapsed time between transmitted and received pulses. Although there are many refinements to this basic procedure, these three components are the heart of the echo sounder and of any side-scan sonar system. The wide-beam echo sounder emits (and subsequently listens, usually with the same transducer) in a roughly 30° wide cone. This wide cone intersects the bottom with a "footprint" that is almost (in 3500 m of water depth) a

nautical mile in diameter (1 nautical mile = 1.852 km). The entire acoustic return is integrated to a single data point, the "depth," clearly a low-resolution image of the seafloor.

Side-scan sonar, on the other hand, uses multiple, interconnected transducers rather than a single, dual-purpose transducer used on the wide-beam echo sounder [Belderson *et al.*, 1972]. With side-scan, a linear array of transducers (usually, the same set is used both to transmit and to receive) is mounted on each lateral face of the towing body, and these transducers listen, or "scan" outward toward either "side" of the ship track. This long, narrow array of transducers produces an acoustic beam that is wide in the across-track direction and narrow in the along-track direction (Figure 3b). Figures 3a and 3b show the fundamental differences in the beam pattern between a




Figure 3a. Schematic diagram showing the large area of the seafloor ensounded by the low-resolution, 30° cone of the wide-beam echo sounder. As discussed in the text, the diameter of the cone at 3500 m water depth can be as large as 1 nautical mile (1.85 km).

Figure 3b. Schematic diagram (not to scale) of the acoustic "footprint" of a side-scan sonar system. As in Figure 3a, the grey hatched area represents the intersection of the beam pattern with the seafloor. Typical beam width for side-scan systems is

2° , so the along-track angle is greatly exaggerated. An actual beam pattern for a "real" side-scan system has many minor lobes in all three spatial directions, and the pattern shown in this figure is to illustrate the basic principles.




Figure 4. A three-phase diagram showing the process of generating image pixels from a single side-scan "ping." In the top diagram the outgoing acoustic pulse from an individual ping is reflected back from the seafloor directly under the fish, and the internal side-scan clock ($T = 0$) is started. In the middle diagram the hatched region at the left represents the outgoing pulse, and the low-amplitude returns are the time when the pulse is the

two-way travel time in the water column. After the return of the first bottom bounce, subsequent acoustic returns, due to backscatter and specular reflection, appear as peaks and valleys in the transducer voltage (shown as the rapidly varying line). These peaks and valleys in voltage are then integrated and translated into pixel values, which are presented as dark and light "grey-scale" regions in the displayed side-scan image.

wide-beam echo sounder and a narrow-beam side-scan sonar for a single transmitted ping. When a single transducer is used at a specific frequency, the directivity of the array (the beam pattern, in our case) will be exactly the same for both transmitting and receiving. Thus the shaded area of the seafloor shown in Figure 3b represents both the region that receives the maximum acoustic energy from the ship during the outgoing ping and the area to which the transducers are subsequently "listening." While the actual situation is somewhat more complicated, this description serves as an adequate first-order model of the processes involved.

THE NATURE OF SONAR TARGETS

The fundamental purpose of a side-scan survey is to provide images of acoustic targets on the seafloor. Unlike radar images, the side-scan receiver detects sound that is backscattered from the seafloor, not reflected from large-scale planar surfaces like radar images [Chavez, 1980]. In most cases, except for the direct bottom bounce described earlier, little acoustic energy arrives at the receiver by direct reflection. Figures 4 and 5 illustrate the fate of an outgoing acoustic pulse as it interacts with the bottom. As these diagrams indicate, much of the trans-

mitted energy is reflected away from the transducers and the side-scan transducers and does not reappear during the receive portion of the cycle. Only those areas of the seafloor that have both a bottom roughness of the appropriate scale and an acoustic impedance (defined as the product of density and sound velocity) significantly different from seawater will produce substantial backscattered energy at the receiver. With a side-scan sonar pulse (Figure 5), the effective acoustic return at the receiver is a variable combination of backscattered (diffracted) and specularly reflected (as from "tiny" mirrors) sound from the seafloor.

The received side-scan signal depends mainly on the backscattered energy from the seafloor, and the strength of the acoustic return depends, in part, on the acoustic impedance contrast between the target and seawater. The amplitude of the returned signal also has a further dependency on the angle of ensonification: the angle of incidence between the sound wave front and the surface of the seafloor. This angle in turn depends on the slope of the bottom and the position of the tow body, or "fish." There are several different scales of bottom topography to consider, the most obvious two are the regional slope, which is generally much larger than the wavelength of the incident sound, and the microtopography, that scale of surface roughness that is smaller than the wavelength of the incident sound (Figure 5). Many excellent acoustics texts discuss the physics of the interaction of sound with the seafloor (see, for example, *Urlick* [1983]), and we will

limit our discussion to those features directly applicable to side-scan images of the deep seafloor.

Reflection of sound from the seafloor is straightforward to understand, but it is not the dominant process in side-scan returns. If the reflecting surface of the seafloor shown in Figure 5 were entirely flat (on all scales), then little energy would actually be returned to the transducers. Fortunately, the seafloor is rarely uniform or flat on the smallest scale, and several mechanisms ensure that sound is radiated back in nonreflected directions. The small-scale microtopography of the bottom material will, through diffraction, reradiate some small fraction of the incident sound wave back in the direction of the transducers. This diffraction of sound, from features whose horizontal scale is comparable to the acoustic wavelength, will give rise to a measurable backscatter signal. Where there is little penetration of the acoustic energy into the seafloor (i.e., a basalt flow), this surface reverberation is the major source of returned energy detected at the side-scan transducers.

The efficiency of this backscatter process is not high, with the bulk of the acoustic energy being reflected away from the side-scan transducers. It can be seen intuitively that the amount of energy backscattered by this mechanism depends on the roughness of the material on the seafloor. Materials which have a rougher surface will backscatter energy more efficiently (with a higher amplitude return at the side-scan receiver) than smooth materials with the same acoustic impedance contrast.




Figure 5. Diagram showing the fate of an outgoing acoustic pulse as it interacts with the seafloor. Within the effective range of the side-scan the returned acoustic energy is a combination of backscattered energy and specular reflection, with only a small amount of direct planar reflection energy from the seafloor. As

shown, the backscatter limit, beyond which there are no useful returns, can depend on the regional topography. The amplitude of the backscattered return, as discussed in the text, depends on the acoustic impedance contrast (density times sound velocity of each material) between the seafloor and the overlying water.

For the interpreter of the side-scan image, the slight difference in textures presented by the microreflectivity of the seafloor surface (Figure 5) is all the information that is available. The appropriate choice of side-scan frequency should be to try to match, as closely as possible, the wavelength of the sonar with the appropriate scale of the roughness of the seafloor, assuming it is known a priori. This frequency choice must be consistent with the other goals of the survey, because high-frequency sound usually means slow, near-bottom towing, coupled with the smaller spatial coverage associated with the resulting narrow swath widths.

In regions where there is substantial sediment on the seafloor, surface microreflectivity does not contribute as much backscatter energy as volume reverberation [Tyce, 1976; Fox and Hayes, 1985; Jackson *et al.*, 1986]. Significant deep-sea sediment penetration of sound occurs at frequencies of 12 kHz or lower, and in this case, a phenomenon called volume reverberation takes place. During this process, sound penetrates below the surface of the water-sediment interface, interacts with a volume of the sediments, and then is effectively reradiated in all directions, including back in the direction of the side-scan transducers [Stanton, 1984; Jackson *et al.*, 1986]. The depth of acoustic penetration, and therefore the amount of subsurface sediment that is involved in the reradiation of the sound, depends on the frequency of the sound and the physical properties of the sediments. Accordingly, low-frequency acoustic side-scan images contain more, or at least different, information about the bulk properties of the sediments that make up the seafloor than those obtained with high-frequency instruments.

ANATOMY OF A SINGLE SIDE-SCAN SONAR PING

The combined directivity of the multiple transducers in the side-scan array results in a narrow wedge-shaped footprint on the seafloor (Figure 3b), one that is narrow along track and very extended in the direction perpendicular to the ship track. This narrow strip can be thought of as two parallel time lines, one on each side of the ship, with the earliest acoustic return being from the seafloor that is directly under the ship ($T = 0$) and the latest time occurring when the sound arrives from the distal flanks on either side of the beam pattern. These time lines can be electronically subdivided, and each time slice treated as an individual "beam," thus forming multiple acoustic beams from what is, in reality, the seafloor response to a single "ping." Although Figure 4 represents the entire beam pattern as a single entity, in actual practice, each side of the track line is ensonified with its own set of transducers, and at distinct frequencies, so that there is no interaction between the two sides. As an example, the GLORIA II side-scan instrument uses a frequency of 6.2 kHz on the port side and 6.8 kHz on the starboard side [Somers *et al.*, 1978].

Figure 4 shows a schematic diagram of the entire side-scan process that begins with the initial ensonification of the seafloor and ends with the generation of an acoustic image. At the initiation of a side-scan cycle, the transducer array generates an outgoing acoustic pulse, usually a short, continuous tone of a single frequency. The GLORIA II system uses a frequency-modulated (FM) pulse, but that is the exception rather than the rule for side-scan systems [Tyce, 1986]. The length of the pulse of the outgoing sound energy is an important factor in determining the ultimate resolution, with shorter pulse lengths giving higher resolution, for a given set of system parameters. As a trade-off with resolution, longer pulse lengths contain fewer frequencies than shorter pulses, are therefore easier to filter for noise, and also contain more acoustic energy per pulse. The easier detection of these longer and more energetic pulses can increase the working swath width of the side-scan system. Typical pulse lengths range from 2–4 s for the long-range GLORIA II signal to less than 0.1 ms for the high-frequency (>100 kHz) systems.

During the transmit pulse, the receiving circuitry of the side-scan is switched off, to prevent damage or saturation of the high-sensitivity amplifiers. After the completion of the transmit pulse, the transducers are switched over to the receiving circuitry, and the continuous recording of the incoming acoustic signal begins. As the outgoing sound pulse travels through the water column, the acoustic energy encounters only midwater scattering sources (i.e., fish, temperature/velocity inversions, and particulate matter), which normally register little energy at the receiving transducers. When the bottom reflection arrives at the fish, this signal starts the high-resolution timing function that controls the generation of the side-scan data (Figure 4).

After the bottom return arrives at the transducers, it is followed by acoustic returns from the seafloor at increasing distances from the ship track. Being a direct reflection, the "bottom bounce" is invariably a very strong return. Because of the near-vertical incidence of the sound wave front as it impinges on those areas immediately below the ship (Figure 4), the sampling rate of the pixel generator (the device that divides the time scale into individual time slices) would need to be too high (almost infinite for near-vertical incidence) to be achieved practically. The net result is that the region of the seafloor directly under the side-scan tow body cannot be used in the acoustic image. Most systems automatically eliminate these near-nadir acoustic returns from the data set. These usually consist of the innermost 40 pixels out of 2048, or about 2% of the total image. For the GLORIA II system, this amounts to a linear distance of about 1200 m, out of a total swath width of 60 km [Reed, 1987; Tyce, 1986].

Once the initial "spike" of the high-amplitude bottom bounce is suppressed, the side-scan processor begins to divide the transducer voltage time series, which is produced by the subsequent bottom return signals, into unequally spaced "time slices." Because of the geometric effect illustrated in Figure 4, these time slices are ex-

remely narrow for the early returns and much wider for the later returns from a more distant slant range. Within each time slice, the varying voltage of the transducer represents the acoustic energy backscattered from a fairly large area of the seafloor, an area much larger than that represented by the pixel size on the final image. The voltage within each individual time slice is averaged (Figure 4) and then converted to a single digital number that is assigned to a specific pixel location. In practice, the conversion from uncorrected transducer voltage to spatially correct pixel values is more complicated than this description. The process varies significantly from system to system and, as described below, requires a variety of additional corrections to become an intelligible image. Modern side-scan systems now digitize the output of the receiver and spatially correct the data further so that they more closely represent a recognizable image [Blackington *et al.*, 1983; Hussong *et al.*, 1985].

Preprocessing of the Data

In order for the side-scan returns to become a recognizable image, the pixels need to be corrected for a variety of effects. These include slant range correction (compensating for the unequal time slice intervals), absorption of sound by seawater, the geometric effect of spreading (time-varying gain amplification), and variable ship speed. The final product is an image that has a 1:1 aspect ratio (i.e., square pixels) and one that has the sonar targets in roughly the same location on the chart recorder as they are on the seafloor. The necessary corrections consist of two basic operations: putting the pixels in the "right" place in the image (the water column and slant range corrections) and giving the pixels the "correct" amplitude values (the time-varying gain correction for spreading and absorption losses) [Chavez, 1986; Reed and Hussong, 1989].

Water column corrections are straightforward operations that attempt to take into account the fact that most side-scan receivers begin acquiring data immediately following the blanking pulse associated with the transmit part of the cycle. Correcting for the time that the outgoing pulse is in the water column consists simply of subtracting a constant offset in time, such that the display processor does not start the construction of the actual image until the sound from the bottom has arrived at the receiver. Suppression of the near-nadir pixels also occurs during this part of the process.

The slant range correction can be thought of as the trigonometric calculation necessary to convert the actual measured straight line (the slant range) distance of a given piece of seafloor through the water column (Figure 4) to a horizontal distance along the seafloor, from the nadir of the fish to the target. Slant range corrections also compensate for the fact that equal time slice intervals do not correspond to uniform intervals of distance from the ship track and therefore do not represent true horizontal distance intervals from the inner edge of the image.

The spreading correction takes into account the fact that the outgoing sound pulse becomes reduced in intensity as

it moves away from the transmitter (and also as the backscattered energy transits from the seafloor back to the receiver). For a spherical wave, this reduction in intensity varies as the inverse square of the distance to the target. For the two-way travel of the narrow side-scan acoustic beam, this spreading loss is a complex function of the beam width and is usually determined empirically. Seawater, like any other medium that transmits waves, also absorbs energy from the sound, decreasing the amplitude of the wave. In salt water, this sound absorption is due largely to the presence of dissolved magnesium sulfate and, to a lesser extent, boric acid. Additional energy is lost to the wave packet owing to scattering within the water column by small suspended particles, bubbles, and occasionally fish and other organisms.

All of these losses can be corrected by application of a time-varying gain (TVG) to the returned signal. This TVG is an amplifier which has a gain that increases nonlinearly with time after the start of a side-scan cycle. Figure 6 shows a hypothetical signal amplitude that decreases with increasing time owing to the geometric, absorption, and water column losses. Application of an appropriate TVG corrects for these losses. Since the signal level decreases strongly with time, a covarying increase in amplification can produce the constant average signal level needed to produce a recognizable image. The TVG settings can vary spatially, usually owing to water temperature conditions or as a function of time, because of transducer "aging" or changes in the beam shape. The appropriate TVG settings are usually determined empirically by surveying a region of the seafloor (usually heavily sedimented) that is assumed to be absolutely featureless and adjusting the TVG until the displayed image appears uniformly bright. A TVG setting that precisely deconvolves the signal transmission losses is rarely achieved, and this inability to attain the "ideal" can be a source of serious frustration for the side-scan user.

Image Construction

Each area of the seafloor that is within the swath of the acoustic beam (Figure 7) can be assigned a location in side-scan "space." This location consists of a record number (one for each ping) and a pixel number (the number of pixels, or distance, from the ship track in the image). For most side-scan systems, there are approximately 1024 pixels per side, or 2048 total pixels in the full swath. Associated with this location pair (the record and pixel numbers) is an individual pixel value, a single number that is the average of the voltage that occurred during the relevant time slice. Depending on the system, each pixel value is usually an 8-bit integer (ranging from 0 to 255, or 256 possible shades of grey), which represent the value of the received acoustic echo after detection and after the electronic low-pass filtering associated with the time slice averaging.

A complete description of the side-scan record of an individual region of the seafloor in the side-scan swath at




Figure 6. Application of a time-varying gain (TVG) amplification to a nonconstant signal. (Top) The natural decrease with time of a returned acoustic signal that decays because of geometric spreading and absorption. Since the actual signal is a rapidly varying function, this curve is just the upper envelope of the amplitude curve. If this side-scan signal were displayed without correction, the image would be too strong in the near-fish returns and too weak in the far-field returns. (Middle) The systematic increase in amplification of the TVG. The increase in amplification with time is designed to compensate for (and deconvolve) the decrease in amplitude in the signal due to geometric spreading and absorption. (Bottom) The ideal, correct TVG applied to the signal. In this case the TVG exactly compensates for the signal losses, and the signal level presented in the image is range-independent. Backscatter amplitudes in this ideal case now depend on the properties of the seafloor material, not distance from the side-scan fish.

this point consists of (1) the location of the ship, and therefore the location of the side-scan fish, in latitude and longitude space, (2) the record number of the ping and the pixel number on the image, (i.e., record 4032, pixel 0823, starboard side), and (3) the amplitude of the pixel value, generally a positive integer from 0 to 255, representing the relative intensity of acoustic backscatter. For shallow-toned systems, like MARC II and GLORIA, the ship and fish locations in item 1 are generally assumed to be the

same, within the range of errors of the ship positioning system. For deep-towed systems used in substantial water depths (i.e., for the SeaMARC I, AMS, and Klein systems), the location of the fish with respect to the ship presents a major uncertainty. In these cases, because the fish is close to the bottom and there is substantial distance between them, the ship and fish locations can differ by several kilometers (Figure 7). This offset in location can change dramatically in a single track line, depending on the changes in the ship track.

For item 2 a number of corrections still need to be made before the sonar targets are accurately located in space, and such corrections represent a major postcruise processing effort. Figure 7 also shows some of the more obvious difficulties that can occur in processing and interpreting the resulting images when the ship track is not straight, an ideal that is difficult to achieve in the real ocean. Deviations in the cruise track from a straight line not only add uncertainty to the actual fish track, but the ensonification pattern becomes "incomplete" on the outside of a curved




Figure 7. Diagram showing the registration of data in a horizontal plane of "side-scan space." Each data point, or pixel value, is uniquely identified by a record number (one for each ping) and a pixel number (generally 1024 pixels on each side). This diagram also shows the difficulty in the spatial registration of data on a real, nonstraight track line and the incomplete and redundant coverage that results on the outside and inside of turns.

track line, while at the same time it oversamples the area on the inside of the turn.

Finally, the relative amplitude of the individual pixels contains all of the geological information in the side-scan record. These pixel numbers are only relative values and have meaning only with respect to each other. Within a given side-scan system, the amplitude signals that become pixel values are dealt with consistently, which makes spatial comparison within and between individual swaths possible. The uncalibrated nature of the transducer output, however, seriously limits the ability to make quantitative comparisons between side-scan systems, even over the same regions of the seafloor.

To account for the variations in ship speed that occur in a "real" survey, each image that is displayed needs some correction factor to provide a proper aspect ratio. The correct aspect ratio is one where the distance represented in the image in the X direction is the same as that in the Y direction, a ratio of 1:1. In side-scan images, this is done by controlling the number of repeated lines that are added to the image after each data line. Each side-scan cycle consists of an initial "ping" followed by a stream of data that represents a time-sequence of returns at increasing distance from the ship track. If these time series data were simply presented "as is," without any correction for ship speed, the resulting image would appear extremely compressed and distorted in the along-track direction.

In order to provide the necessary speed correction and the ability to view pixels that are "square" in aspect ratio, image display programs insert duplicate lines after the initial data line, with the number of repeated lines proportional to the ship speed. Largely because of the historical use of shipboard graphic recorders, most imaging systems add a single duplicate line for each knot of ship speed (1 knot = 1.85 km/h). As an example, side-scan data taken at 7 knots are normally viewed as one data line and six repeated duplicates of that line. Because only integer multiples of the data lines are possible (either six or seven repeat lines can be added, not 6.5), it is not possible to correct for speed variations smaller than 1 knot. This is a limitation to our ability to spatially correct the data that is significant only for slow speed, near-bottom systems. This replication of the data lines does not change resolution or the image processing functions, but is only an artifact of the display process.

POSTPROCESSING CORRECTIONS TO THE DATA

In addition to real-time shipboard data manipulations, postprocessing corrections are usually applied to side-scan sonar data. These postcruise modifications fall into the same general categories as the shipboard corrections, i.e., spatial or geometric corrections, which change the location of the pixel values within the image but do not change their value, and radiometric corrections, which change the value of specific pixels. Several phenomena that can

adversely affect side-scan images have been recognized [Belderson *et al.*, 1972; Tyce, 1986; Chavez, 1986; Reed, 1987; Reed and Hussong, 1989]. Reed and Hussong [1989], for example, describe an elegant "background subtraction" technique that corrects for many of the system artifacts that are generated in the image. These artifacts include bottom and surface "echos" (acoustic energy that arrives at the transducers after multiple reflections from these surfaces), incorrect bottom detections, and non-uniformities in the beam pattern. The technique of Reed and Hussong [1989] calculates the mean and standard deviation of along-track swath data over a large portion of the image (outside the region where there are artifacts) and then uses these "expected" values to modify the image where there are artifacts in the data. While successful in removing known artifacts from the data, this background subtraction technique is a fundamental modification to the image data and may make between-image comparisons difficult.

Bottom Slope Corrections

Regional, or large-scale, slope of the seafloor in the area of a side-scan survey can play an important role in the appearance of an image, an effect that can require both geometric and radiometric corrections. Bottom slope or topography can modulate both the amplitude of the return, at the pixel level, and the apparent texture of the seafloor acoustic targets, at the image level. Radiometric corrections that alter the pixel values of targets located on inward and outward facing slopes must be applied to allow a spatially correct interpretation. Large-scale regional changes in the slope of the seafloor can induce substantial geometric errors in the actual location of acoustic targets within the image. Reed and Hussong [1989] discuss this "layover error" at some length, and Reed [1987] provides a fortran program which can correct for the effects of this error. While the effects of (and remedies for) the layover correction are adequately described in the literature, it is probably useful to review the causes of the phenomenon, to be able to estimate the magnitude of the effect, and to be able to apply the correction to the data "by hand" if necessary.

Layover Correction

One of the basic assumptions that is made in the processing of side-scan data is that the seafloor is both flat and horizontal (Figure 8). To the extent that this assumption is not true, and the real bottom is uneven and slopes, an error in across-track position is introduced in the placement of the acoustic targets. To the geologist, uneven seafloor is interesting geology, and small-scale corrections or the topography are neither necessary nor desirable. Large-scale regional slope, like a continental margin or seamount flank, however, can cause substantial errors in the placement of acoustic targets. The qualitative effect of this error can be described as follows: where the seafloor slopes up from the nadir of the side-scan fish, the




Figure 8. Diagram showing how regional slope can cause the target on the seafloor to appear to be displaced closer to the track line than it actually is. The first acoustic return, which the side-scan system assumes is directly below the ship, is from the sloping bottom (point A), which is closer to the ship than the point directly beneath it (point B). This starts the side-scan clock too early and causes an effective horizontal displacement of the baseline from point A to point B. When the return from the seafloor arrives at the fish, it is displaced in the image by the timing signals, from point C to point D; the layover correction moves the target image to the correct location at D. As shown, when targets are uphill from the fish, they are displaced closer to the image. Conversely, when targets are downslope from the fish position, they are displaced farther from the actual track line.

target will appear to be closer to the ship than it really is; where the bottom slopes away (down) from the fish track, the targets appear farther away.

Figure 8 shows a diagram of a target on a sloping seafloor, a situation that violates the flat, horizontal bottom assumption. As mentioned previously, when the first "hard" return arrives at the fish, the side-scan processing clock begins the precision timing of the subsequent events. If the bottom has substantial slope, as shown in Figure 8, then the first "bottom bounce" will be from that portion of the seafloor that is closest to the ship. With a sloping bottom, this will not be from directly below the fish but from the point of nearest approach at one side of the track line. The processing algorithms for the side-scan, which assume that the bottom is flat, incorrectly take this position, point A, as the nadir of the fish and move the baseline of the image to the bottom at point B in Figure 8. When the acoustic energy from the target arrives at the fish, it is assumed to be at the same depth as the original bottom return, and the correction process places it in the image at point D, rather than at its "real" location in the image shown at point C. Note that even the "real" location of the target in the image would be on a projected horizontal plane, not on the sloping bottom. In this case of a bottom sloping up from the transducer, the uncorrected effect of slope is to place the target closer to the track line than it should be. Conversely, if the side-scan is looking "downhill," the target will appear farther from the track line than it would if the bottom were flat.

The correction for this effect is conceptually simple but actually quite difficult in practice; the main problem is in determining the true water depth of the target, away from the ship track line. Application of the layover correction to the actual distance of the acoustic target (Figure 8) requires the determination of the slope of the seafloor and the height of the fish off bottom. The successful application of

a computer processed "layover correction" as described by *Reed and Hussong* [1989] depends critically on the coregistered swath bathymetry generated simultaneously by SeaMARC II. As an illustration of the magnitude of this effect, *Reed and Hussong* [1989] cite examples of SeaMARC II data where specific acoustic targets are displaced by the layover effect as much as 1000 m, and fault scarps on slopes which show an apparent angular rotation by 20° . This example of a basic data correction confirms that although side-scan data have become much more "recognizable" images of the seafloor, there are clearly significant effects inherent in the technique that must be either corrected, or at least considered, in the interpretation of these records.

OPERATING PARAMETERS AND SYSTEM CHARACTERISTICS

Different side-scan systems vary in ensonification frequency, height of the fish off the bottom, ping repetition rate, pulse length, and swath width. The latter four dependent variables are largely controlled by the operating frequency and, in turn, control the pixel size, the resolving power, the size of the area ensonified, and ultimately the quality of the image that will be interpreted for geology. In order to evaluate the impact of these different parameters, it is necessary first to consider the role that the acoustic frequency plays in the quality of the final side-scan image.

Side-Scan Frequency

Ensonification frequency is the primary independent variable in choosing an appropriate side-scan sonar system. Existing side-scan sonar systems are single-frequency instruments, and the choice of the operating frequency

determines many of the other system parameters. Because of the physical properties of sound waves, particularly attenuation, these dependent parameters include towing altitude (using either a deep-towed or surface-towed configuration), the below-surface depth of interaction of the backscattered sound (the penetration), and image resolution. While the velocity of sound in water is largely independent of frequency, the absorption of acoustic energy in the water column, the ability to penetrate sediments, and the practical limitations on pulse length are all strongly dependent on frequency.

Strong absorption of high-frequency sound by seawater limits surface-towed deep-water systems to frequencies of less than 30 kHz. Side-scan frequencies currently in use in the open ocean range from the 6-kHz (GLORIA II) and 12-kHz (SeaMARC II) instruments used by surface-towed systems, to the 30-kHz (SeaMARC I) and 110- to 150-kHz (Scripps Deep-Tow, AMS-120, SeaMARC 150, Klein) instruments used in deep-towed systems. Existing side-scan sonar systems are deployed in a "fish," towed behind the hull of the surface ship. In addition to system portability, this configuration is used to decouple the transducers from the motion of the ship, to reduce the amount of ship-generated noise at the receivers, and, most important, to get below the strong velocity gradients associated with the thermocline in the surface waters.

Towing Altitude

Figure 9 shows the basic towing configuration of a side-scan system; this figure is very schematic but generally represents existing deep- and surface-towed systems. In this representation, the fish is attached to the ship through an armored coaxial cable that provides both the strength member for towing and the necessary electrical connection to the surface ship. Power and control signals are sent down the cable; side-scan and telemetry (altitude, pitch, yaw, depth) signals are sent up. The neutrally buoyant fish is attached to the towing cable through a depressor weight, a heavy mass that effectively

translates the up-and-down vertical motion of the ship to an oscillatory horizontal motion that causes less distortion to the image. As examples of towing configurations, the SeaMARC II system is towed 50-100 m below the sea surface, regardless of the water depth. In contrast, the SeaMARC 1A system is towed either at a constant depth over the bottom, generally 100 m above the highest upward projection of the seafloor in the survey area, or in the "draped mode," at a constant altitude of 100-200 m above the varying bottom.

Surface-towed systems, with their lower frequency, can be towed faster (7-8 knots (13-15 km/h) compared with the 1-3 knots (1.85-5.55 km/h) in the deep-tow mode) and have a wider swath width (40-60 km for GLORIA, compared with 1 km or less for the deep-towed systems). As a practical rule-of-thumb, swath widths and tow-fish altitudes are roughly related by a factor of 10; a 1-km-wide swath SeaMARC 1A survey would normally be flown at an altitude of 100-200 m. Technical comparisons of the properties of different side-scan systems are abundant in the literature, with good reviews by Tyce [1986], Davis *et al.* [1986], Mazel [1985], and Chavez [1986].

RESOLUTION

Resolution in an image is the ability to distinguish closely spaced objects as individual features. In side-scan sonar images, there are several different types of "resolution" that must be considered in the interpretation of the image. As the sonar beam pattern differs fundamentally in the along- and across-track directions, side-scan images are intrinsically anisotropic. It is therefore necessary to understand the concept of resolution in both the along- and across-track directions before attempting to interpret the geological features in an image. As an additional complication, this "directionally sensitive" instrumental resolution is further modified in the final display by the finite pixel size of the displayed image. This range-




Figure 9. General configuration of the towing systems used for side-scan sonar systems. As described in the text, the heavy depressor weight (1000-2000 kg) decouples the side-scan fish from the motion of the ship and converts the vertical wave mo-

tion into horizontal variations in the towing speed. Values for X and Y depend on the system used and can vary from several hundreds of meters for surface-towed systems to several kilometers for deep-towed systems.

independent pixel size is calculated for a given side-scan system from the ratio of the swath width to the number of pixels in the image, a value suitably termed the "display resolution." An appreciation of how the instrumental resolution and display resolution interact will allow users both to design appropriate survey strategies and to better evaluate acoustic target identification in the final images.

Across-Track Resolution

Three factors control the ability to resolve sonar targets in the across-track direction: pulse length, the width of the acoustic beam in the along-track direction and, indirectly, the range or distance from the track line. Independent of their range, two targets are theoretically resolvable if their separation is one-half the length of the incident sound pulse multiplied by the sound velocity. For the SeaMARC II system, for example, this theoretical limit is approximately 7 m. In practice, the effective resolving power of the system will always be worse than this. The actual projection of the sound wave on the seafloor increases in width at greater distances from the tow fish (Figure 10). This dependence of the footprint size on the distance from the side-scan fish, in the across-track direction, will be important in our discussion of ultimate system resolution.

tion, but once the optimum ping rate is reached, a further increased rate only provides more redundant data. As a practical rule-of-thumb, along-track resolution is usually much worse than the across-track resolution [Tyce, 1986].

In the along-track direction, a fundamental control on resolution is the ping rate. Two physically distinct objects cannot be distinguished in the along-track direction if they are detected only by a single ping. Individual detection of the two objects by two adjacent pings is a minimum criterion. If we make the assumption that only one ping should be in the water column during a single side-scan cycle, then we can estimate the maximum ping rate that can be used for a particular type of side-scan. Using SeaMARC II as an example, with a swath width of 5 km on a side, the time required for the sound to travel the full 10-km roundtrip distance is approximately 7 s. This is a basic upper limit to the ping rate for this system and one that is dependent solely on the desired swath width. In actual practice, SeaMARC II uses a ping rate with a 10-s interval.

Because the distance traveled by the fish between pings is a fundamental limitation for along-track resolution, this parameter is of interest and is easy to calculate. Taking SeaMARC II as an example and assuming a survey speed




Figure 10. Greatly exaggerated diagram of the butterfly-shaped beam pattern of the side-scan fish as it intersects the seafloor at increasing distance from the fish. In this hypothetical illustration for SeaMARC II, the 2° wide beam intersects a linear distance of 34 m at a slant range of 1 km, and 175 m at the full 5-km half-width of the swath. The upper diagram shows the pattern of overlap between two adjacent pings which must be considered in determining the along-track resolution for any system. None of the dimensions in this figure are to scale.

Along-Track Resolution

Resolution in the direction parallel to the ship track will depend largely on the width of the acoustic beam but will also depend on fish height and a combination of the ping rate and towing speed. Like across-track resolution, distance from the ship track determines the ability to resolve distinct targets which are aligned parallel to the track line. Since the two sides bounding the acoustic beam are not parallel, the beam spreads with increasing distance, and the ability of the side-scan to separate sonar targets decreases further from the tow fish. Faster ping rates and slower towing speeds also give better along-track resolu-

tion of 7 knots (13 km/h) and a ping rate interval of 10 s, this means that the ship travels 36 m along the track line in the interval between pings. This is the ultimate limit on the along-track resolving power of the system, but, as we will see, even this level of resolution is not reached in practice.

The horizontal angle of the acoustic beam is small but still finite, and therefore the width of the beam in the along-track direction is not constant but spreads with increasing distance from the ship. This expansion of the beam with increasing distance from the track line causes a changing footprint in the along-track direction, and the changing dimension results in decreasing along-track

resolution. This can be seen intuitively with simple trigonometric arguments (Figure 10). This simple geometric expansion of a beam pattern of finite angular width results in a range-dependent along-track resolution similar to that previously discussed for the across-track resolution.

Display Resolution

The along- and across-track resolution described above provide the fundamental limitations to the resolving power of side-scan systems. Unfortunately, none of these inherent limits on the resolution of the side-scan system have any real meaning for the interpreter until the data set is presented as an image. The technique used to create and display the image which is actually viewed acts as a "final filter" to the data, a filter that can integrate, modify, artificially enhance, or (more commonly) blur the final image. The final image that is displayed can never be made sharper than the resolution limits described above and can be substantially worse.

In side-scan sonar the most common form of resolution used is that of pixel size, the figure obtained by simply dividing the distance of the full swath width by the total number of pixels. In the SeaMARC II example this is 4.9 m (10 km/2048 pixels). This figure has been labeled "image processing resolution" [Tyce, 1986] and is in easy

"figure-of-merit" value to calculate for the comparison of different side-scan systems. Although correctly described by Tyce [1986], this range-independent figure is frequently misinterpreted as the actual instrument resolving power of side-scan systems. While it has the correct dimensions (length) associated with resolution, and it is certainly a characteristic length in the image, this value is not a true measure of the resolving power of the technique. In contrast, actual instrumental resolution must take into account the area of the seafloor that is the source of the acoustic return that is ultimately integrated into a given pixel. For lack of a better term, we call this concept the "pixel ensonification area." Since both along- and across-track resolution vary with distance from the ship track, it is also necessary to consider the range dependency of that area of ensonification.

Instrumental Resolution

Figure 11 shows a representation of areas of the seafloor that provide the total integrated acoustic energy that is averaged into a single pixel, for several different distances from the track line. This diagram, modified from Karlin and Johnson [1987], is based on the calculations for a SeaMARC II system by Kosalos and Chayes [1983]. As shown in this figure, all of the backscattered sound that is received from the shaded area is averaged into the single




Figure 11. Diagram showing the size and shape of the areas of the seafloor that are integrated into a single pixel value as a function of across-track distance from the fish. This figure uses SeaMARC II as an example and is based on the calculations of Kosalos and Chayes [1983]. As with Figure 10, none of the dimensions in this figure are to scale. Since the size and shape of the area integrated into each adjacent pixel determine the ability to distinguish seafloor targets, the resolution power of the systems will vary directly as the areas shown.

pixel value. If we assume a pixel size of 5 m by 5 m for the SeaMARC II, then in the area nearest to the track line, pixel number 50 (out of 1024 total) contains backscattered energy from an area of seafloor that is 5 pixels by 9 pixels in size, or an area that is 25 m by 44 m in actual dimensions. For this example the returned acoustic energy that is integrated into a single pixel value comes from a pixel ensonification area of over 1100 m², not the 25 m² (5 m × 5 m) represented by the pixel size in the image. In the region that is closest to the track line, the active area of detected and integrated acoustic return is roughly equidimensional, a circumstance that is dramatically not the case at the distant edge of the image.

To see how this active region varies with distance from the track line, Figure 11 also shows a region at the full 5-km distance of the swath width (i.e., pixel number 1024). Here the pixel ensonification area that provides the returned acoustic backscatter—that is ultimately integrated into a single pixel number—is now only 1 pixel wide across track but almost 20 pixels long in the along-track dimension. In this distal region example the total active pixel ensonification area is 588 m² compared with the 25 m² that the pixel represents in the image. The more distant area is not equidimensional but is elongated 20:1 in the along-track direction. Two intermediate areas, representing pixels numbers 150 and 400, are also shown in Figure 11 to demonstrate how actual resolution varies with distance from the track line. While this example is drawn using the SeaMARC II system, with appropriate scaling the same conclusions can be applied to side-scan systems in general.

Clearly, the implications of the different across-track areas represented in Figure 11 are important in the interpretation of side-scan images. It can be seen that the effective resolution for features that are distributed parallel to the track line is much greater at the distal edge of the swath than close to the center line. This figure also shows why some types of geological features seem to have a fundamental anisotropy when they are viewed from ship tracks with different “look” directions. When the elongate areas of ensonification are parallel to the texture of the feature, the textures can be resolved; when the textures are oblique, they are unresolvable. Further, Figure 11 shows the fundamental difference between “image-processing resolution,” or pixel size, and actual resolution. Finally, the elongate nature of the ensonification areas should have an impact on how survey lines are planned; small-scale features parallel to track lines will be visible in the image, while features perpendicular may not, depending on their size. Many survey track lines are now laid out at 45° to the known strike of geological features, such as spreading centers [Davis *et al.*, 1986; Kong *et al.*, 1988] to compensate for this anisotropic resolution.

Pixel Overlap and Redundant Data

In discussions of resolution, consideration of the size and shape of the areas of active ensonification bring up the

subject of data redundancy, or pixel overlap. Because the active ensonification region of each pixel is much larger than the pixel in the image that represents it, and the pixels are (by definition) immediately adjacent to each other in the image, it follows that some of the area integrated for one pixel will also be included in the adjacent pixels. As with resolution, the redundant areas are different for along- and across-track adjacent pixels. Figure 12 shows both the along- and across-track overlap for the respective pixels, again using the SeaMARC II as an example. In the across-track direction the overlap is between adjacent pixels in a given data line, for a given ping. In the along-track direction the overlap is between pixels at the same distance from the track line (i.e., the same pixel number, for a straight track line) but for successive pings. In the across-track mode, shown in Figure 12, there is a great deal (89%) of redundancy in seafloor ensonification for adjacent pixels near the track line and very little overlap (10%) for pixels at the distal end of the swath width. For the inner pixels the large region of overlap between adjacent pixels reduces the effective resolution; only those acoustic targets that are represented in one pixel, but not in the adjacent pixel, can be resolved.




Figure 12. Summary diagram showing how adjacent pixels overlap in the along- and across-track directions. In the across-track direction the overlap is between adjacent pixel numbers; in the along-track direction the overlap is between successive pings. As with Figure 11, the small squares in the middle of each larger square indicate the relative size of the pixel in the final image as compared with the actual area that is the source of the acoustic backscatter energy.

Clearly, overlap continues to degrade the resolving power of the inner pixels of the swath. The variation in across-track overlap for intermediate positions in the swath, at pixel numbers 150 and 400 (out of 1024) are also shown in Figure 12.

Figure 12 also shows similar overlap of adjacent pixels (due to successive pings) in the along-track direction. In this example the range variation of the degradation of resolution is the opposite of that for the across-track overlap. The innermost pixels have an overlap of only 20%, while the adjacent pixels at the edge of the swath overlap by 64%. Unlike most other effects on resolution, this effect has the most impact on the outer pixels of the swath and the least effect on those in the inner region. This phenomenon has a direct impact on the appearance of the image, with images (and their shadows) appearing elongated in the along-track direction, particularly in the far field of the image, distant from the ship track. By using point migration techniques similar to those used in seismic data processing, *Reed* [1987] has developed a digital correction technique for the effects of this target elongation on the side-scan image.

Reflection on the actual resolving power of side-scan sonar systems, as discussed above, can be a discouraging exercise. Instead of the ability to resolve (for the SeaMARC II example) objects with dimensions of the order of 5 m—a capability that consideration of the pixel size alone would lead us to believe—we can, at best, only resolve objects that are several tens of meters in linear dimension. Further, this resolving power is both range-dependent and orientation-dependent and varies both with distance from the ship track line and with the angle between the track line and the texture-orientation of the feature. In extreme cases, for objects with an unfavorable across-track orientation and a location at the edge of the swath, we only have a resolving ability of the order of 100 m.

Target Detectability

It is important not to confuse the concept of resolution with the ability to detect objects in the side-scan swath. Specifically, resolution is the ability to distinguish two separate targets on the seafloor, while detectability is the ability of an object on the seafloor to make a visible record on the side-scan image. This confusion arises because side-scan sonar has the ability to detect objects on the seafloor that are much smaller than the resolving power of that system, particularly in an area of low ambient backscatter. The fundamental criterion for detectability is different from that of resolution; if an object can provide sufficient backscattered acoustic energy to the receiving transducers and if the ambient background backscatter is both uniform and low (to provide adequate contrast with the hypothetical target), then the object can be “detected” as a sonar target even if it is considerably smaller in dimension than the resolving power of the system. The clear visibility in the side-scan record of small fissures, fault scarps, and cracks on the seafloor, features with

dimensions much smaller than the footprints shown in Figure 11, are an obvious set of examples. For high-resolution side-scan systems towed near the bottom, using 100-kHz frequencies and narrow swath widths, objects as small as a single piece of line or cable, lying on a smooth, sedimented bottom, have been detected [*Mazel*, 1985]. Conversely, where the acoustic impedance of the sonar target is similar to that of water (e.g., water-logged pine wood [*Mazel*, 1985], large, high-relief objects, such as ships, can sometimes be virtually undetectable by the side-scan.

IMAGE ENHANCEMENT AND PROCESSING

Three distinct operations are loosely associated with the term “image processing” as opposed to the geometric data corrections described in the previous section. First, image enhancement consists of those operations that make the image look “good,” or at least more like what we expect. These operations make the data display more “pleasing” by modifying the overall appearance of the image. Some examples of enhancement are the variation in amplitude gain, increasing or decreasing the contrast, and application of threshold values (high- or low-pass filters) to the pixel amplitudes. The second class of image processing, image analysis, consists of operations which provide compact numeric information based on the data within the image. These processes distill the “essence” of the data in a image and present it in short-hand numeric form. Application of these image analysis operations obviously does not change the appearance of the original image, and the results are usually presented in a nonimage format, such as a table of numbers. Finally, the third category of image processing, image coding, includes results of mathematical analyses that are sufficiently complex to be presented as a secondary image of the original data. Pattern recognition techniques, including the “feature recognition” and image classification techniques of *Reed and Hussong* [1989], fall into this category.

HISTOGRAM ANALYSIS

A major tool in image processing that is used for both enhancement and analysis is the image histogram. A histogram is simply a plot of the frequency distribution of pixels contained within an image (Figures 13a and 13b). In addition to being an easy-to-understand representation of the range of values that exist within an image, histograms provide a mechanism for visualizing the statistical techniques that are applied to pixel amplitude values. Most of the major enhancement techniques that are applied to images can be considered as specific mathematical operations on pixel amplitudes, which can be visualized as histogram changes. Thus histograms provide a quantitative component to what is normally a very subjective




Figure 13a. A SeaMARC II image of a recent lava flow (dark) onto sediments (lighter) at the northern end of NSR segment, Juan de Fuca Ridge. In this image, traditional polarity convention is observed, with strong acoustic reflectors shown as dark pixels and weak returns as white. (Left) The acoustic image. (Right) An interpreted line drawing. In the figure, significant geological features are the lava flow (horizontal hatches), a low-reflectivity halo surrounding portions of the flow (diagonal hatches), several rounded and sedimented seamounts on the right side and lower center of the image, and both active (dark) and inactive (light grey lineations) faults. The direction of ensonification is from the left edge of the image.

operation. Since acoustic images are monochromatic, we are concerned here only with what is termed as the “grey level histogram.”

Side-scan images are generally arrays of 8-bit numbers, values which represent 256 (i.e., 2^8) intensity levels, or shades of grey. This means that the horizontal axis of the histogram varies from 0 to 255, while the vertical axis displays the total number of pixels in each column. Figure 13a shows a side-scan image that contains both hard acoustic reflectors (newly erupted basalt) and soft acoustic reflectors (unconsolidated sediments). The clear separation of the pixel amplitude distributions within each region illustrates the power of the histogram to define the acoustic reflectivity for homogenous sonar targets (Figure 13b). Unfortunately, the majority of side-scan images from the seafloor produce much more complicated histograms, with pixel values distributed over the entire range of possible values. While extremely valuable in defining the distribution of pixel values, the histogram does not provide any information about the spatial distribution of these pixels within the image.

Two additional concepts are important in any discussion of image enhancement: contrast and dynamic range. Both of these concepts can be easily interpreted in terms of pixel distribution within the image histogram. Image contrast

can be defined as the “separation” of pixel values within a histogram: a low-contrast image has pixels narrowly grouped about one central value of the histogram (Figure 14a), while a high-contrast image has a bimodal distribution of pixels, with the two high-density distributions located at opposite ends of the histogram (Figure 14b). Generally, an image that is pleasing to the eye has pixel values that are uniformly distributed throughout the histogram, although this may not produce the desired representation of specific geological features.

The dynamic range of pixel values within an image is a measure of the width of the occupied portion of the histogram in comparison with the total number of pixel amplitude values available. An image with pixel values




Figure 13b. Histograms of the areas marked A, B, and C in Figure 13a. Area A is the low-reflectivity region of modified sediments adjacent to the recent flow, where the backscatter has been altered by some unknown process. Area B is the lava flow (which has been both dredged and photographed with a deep-tow camera). Area C was chosen as a region that contains both sediments and flows. In these diagrams, pixel amplitude value 0 is the strongest possible acoustic return, and 255 is essentially no return.


narrowly grouped around any specific value (Figure 14a, for example) is said to have a low dynamic range. Very high contrast images (Figure 14b) can also be considered to have low dynamic range because of their narrow (even if bimodal) distribution of pixel values. High dynamic range images (Figure 14b) are generally considered more pleasing, because they are closer in appearance to natural optical images and are commonly perceived by the viewer to contain more information than low dynamic range images [Baxes, 1984]. In side-scan sonar, images with a low dynamic range are a common problem because of the nature of the acoustic reflectivity. Reallocation of the

pixel values throughout the full dynamic range of possible values, by histogram sliding and/or stretching, is a technique that is often applied early in the processing stage.


Histogram Sliding and Stretching

Histogram sliding and stretching are two processes that redistribute the pixel amplitudes over the full range (0–255) of possible values. While this can sometimes make a more pleasing and recognizable image (as can be seen in the comparison of Figures 15a, 15b, and 15c), it is important to understand that the pixels redistributed during


a Low Contrast Lavaflow
on Lavaflow



Mean 207.9
Median 212
Mode 217
Entropy 6.43789
Stand. Dev. 22.95031



b High Contrast Lavaflow
on Sediments



Mean 193.5
Median 205
Mode 239
Entropy 7.16549
Stand. Dev. 44.12181




Figure 14. Two SeaMARC II images and associated histograms showing examples of high and low contrast in side-scan images. (a) Image from the southern Juan de Fuca Ridge, showing the low contrast associated with lava flows on top of only slightly older lava flows. The histogram of the marked area shows what

appears to be only a single population of pixel values, and even with some enhancement, the flows are difficult to distinguish from the background. Lava flows contrasted with sediments. The histogram of the marked area clearly shows two distinct populations of pixel values.



Mean 202.8
Median 214
Mode 232
Entropy 6.97256
Stand. Dev. 38.79581

A




Figure 15a. SeaMARC II image of the same area shown in previous figures, but now with the histogram of the entire image shown. This histogram, which from the spatial relationships in the image has three distinct populations of pixel values, shows that the overlap in populations obscures any distinction.



Mean 131.1
Median 140
Mode 146
Entropy 6.61874
Stand. Dev. 30.77029

B




Figure 15b. The same image as shown in Figure 15a, but now subjected to a pixel amplitude "slide" of -80 (subtracting 80 from each pixel value), which makes the entire image darker and lower in contrast.

histogram manipulation do not contain any more information than the original image. Histogram sliding is a process that moves the "center-of-mass" of the pixel amplitude distribution to the middle of the range. Histogram sliding is the simple addition or subtraction of a constant value to all the pixels within the image. Since the acceptable range (for display) for 8-bit images is within 0–255, this operation can, for certain offsets, produce new values that are outside this range. The over-ranging of the new image pixels is a common problem that occurs in image-processing operations, and simple arithmetic operations are used to move the pixel array values back within acceptable display values.

Similar to histogram sliding, histogram stretching is a simple mathematical operation that redistributes the pixel values within a histogram, an operation generally done to

increase both the dynamic range and the contrast of the image. Mathematically, a histogram "stretch" is the simple multiplication of the pixel amplitudes by a constant value, usually, but not necessarily, an integer. If NPV and OPV are the new and old pixel values and K is some constant, then

$$\text{NPV}(n) = K \times \text{OPV}(n)$$

The effect of this multiplication is to increase (or decrease) the distribution of the pixels throughout the entire range of the histogram. A "linear stretch," for example, is simply taking the difference of the maximum (PVmax) and minimum (PVmin) occupied histogram values and dividing each old pixel value by the factor that expands the



Figure 15c. The same image as shown in Figures 15a and 15b, but now subjected to a pixel amplitude “stretch” of 0.5 (dividing the pixel amplitudes by 2, which is actually a histogram compression). Again, the image becomes darker and much lower in contrast. These operations do not, obviously, improve the image and are shown as examples of specific histogram manipulations. Histogram equalization, which is a combination of sliding and stretching, is the operation normally done for image enhancement in the initial stages of processing.

distribution to fill the entire range. Computationally [Pratt, 1978; Reed, 1987], the new pixel values can be generated by

$$NPV = [(OPV - PV_{min}) / (PV_{max} - PV_{min})] \times 255$$

As with histogram sliding, histogram stretching can also produce pixels that are outside the range of values (0–255) acceptable for display. While the stretching operation almost always improves the appearance of low-contrast, small dynamic range images, exactly the same amount of information is contained in the new image as was in the old.

Histogram Equalization

Histogram sliding and stretching are usually operations that are performed together, as a combined operation called “equalization,” to manipulate the pixel distribution for a more recognizable image. In practice, this operation is a histogram modification procedure that assigns equal numbers of pixel values to each of the divisions in the new image. This operation can be thought of as a mapping process between the old pixel values of the original image and the modified values of the new image (Figure 16). The process of equalization adjusts the new pixel value assignments such that there are an equal number of pixels with each new amplitude value. For example, if there were a total of 1600 pixels in the original image and 16 grey scale levels (pixel amplitudes) that could be displayed, histogram equalization would assign an “equal” number, 100 pixels in this case, to each of the individual grey scales in the new image. The effect of equalization is to expand the contrast between grey levels that occur frequently and to decrease the contrast between those pixel amplitudes that occur infrequently, using the limited amount of available image contrast where it is most effective. The mathematical operations necessary to perform image equalization are described by Pratt [1978], and the appropriate algorithms to accomplish it with acoustic images are developed by Reed [1987].

FILTERING

Spatial filtering can be generally defined as that analytical process that separates the individual frequency components of an image, where the term “spatial frequency” is used to describe the variation in relative amplitude values of adjacent pixels. A region of an image in which adjacent pixels have widely varying values is said to have high spatial frequency components, while an image with relatively little local variation is said to have low spatial frequency components. Filtering, or convolution, is the most common procedure used in image enhancement. It is an operation typically applied to the image to produce such effects as smoothing (the removal of high-frequency noise) or edge detection (the emphasis of specific linear features). Many of the operations thought of as general “image processing” are, in fact, simply applications of different types of filters.

A filter can be described by a two-dimensional array of numbers, known as the “kernel.” In general, kernels are equidimensional, and typical sizes for filters are 3×3, 5×5, and 7×7, although any dimensions may be used, and there is no real requirement for filters to be the same size in both directions. Convolution is a signal-processing operation that combines two signals to produce a third array of values. In image processing, one of these signals may be the original image array, and the other is the filter matrix. The filtering, or convolution process, involves sequentially replacing each cell (or pixel) of the image with the dot product of the kernel




Figure 16. Application of histogram equalization to a SeaMARC II image east of Axial Seamount on the central Juan de Fuca Ridge. The image is low in dynamic range and contrast, and it is difficult to distinguish off-axis lava flows even on top of a turbidite sediment background. After histogram equalization,

using the area shown in the boxes, the image becomes higher in contrast and much easier to interpret. In the histograms below the images, only the scale on the right-hand diagram is changed. This is to illustrate the equalization process that places equal numbers of pixels in each grey-scale level.

and the pixel's neighborhood values (Figure 17). To illustrate this, consider the general 3×3 filter:

$$F = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix}$$

To apply this filter to an image, we replace each pixel $I(x,y)$ of the image as follows:

$$I(x,y) = I(x-1,y-1)*a + I(x,y-1)*b + I(x+1,y-1)*c \\ + I(x-1,y)*d + I(x,y)*e + I(x+1,y)*f \\ + I(x-1,y+1)*g + I(x,y+1)*h + I(x+1,y+1)*i$$

Note that for two signals, one of size N and the other of size M , the output signal has size $N + M$.

It is also useful to note that this filtering operation is a straightforward process until the edge of the image is reached, where there are no pixels "adjacent" to those on the outer boundary. This effect causes degradation around the margin of a filtered image, an effect which can be lessened by adding artificial data outside the original image boundary.

SPECIFIC FILTERS AND ANALYTICAL TECHNIQUES

Low-Pass (Smoothing) Filter

A low-pass filter (or smoothing filter) removes high-frequency variations in an image. This technique can be used to remove the "grainy" or "snowy" texture in noisy side-scan images and is a useful first step in image enhancement. Low-pass filters come in many forms. In the simplest case, all kernel values are equal.




Figure 17. Schematic of the operations associated with filtering. In a normal 3×3 filtering operation the pixel values in an image array are convolved with the kernel values in a 3×3 mask. In the upper diagram a 3×3 filter mask with values a through i is placed over a portion of the image. The pixel value in each image cell is multiplied by the corresponding value of the filter. The sum of these products (7 in this example) is mapped into the center of the filtered image in the lower diagram.

The following filter is a simple 3×3 smoothing filter.

$$F = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

Applying this filter to an image (Figure 18) causes each cell to be replaced by the sum of the cell and its eight immediate neighbors. This is essentially computing the mean of the pixels within the kernel and placing that mean in the pixel location of the central value. This filtering process will eliminate any large local variations in pixel value and produce a "smoothed" image. The primary disadvantages of the above technique are the "blurring" of sharp features within the image and the loss of dynamic range. This loss of information can be reduced by weighting the filter such that the central cell values have more influence on the result than distant cells.




Figure 18. Application of two standard filtering techniques to the same SeaMARC II image used in previous figures. In this figure the outer margin of the image is unfiltered. The upper box, above the dark line, has been filtered using the median smoothing filter described in the text. The insert area in the lower box, below the dark line, has a Gaussian-weighted smoothing filter applied. Median filters are effective in removing spikes but maintain both the image specularity and sharp feature edges. Gaussian filters produce a smoother image but also blur the edges of the features.

Median Filter

A nonconvolution technique known as median filtering offers a smoothing technique without either the blurring or the loss of dynamic range associated with smoothing filters. The median filter replaces a cell with the median value of the neighborhood, rather than the mean value. Two points should be noted; first, the process of finding the median cannot be performed as a convolution and is somewhat more computationally expensive. Second, a basic median filter is not described by a kernel; no kernel values are involved, and only the dimensions of the search area need to be specified. A primary use of the median filter is in "noise spike" removal from the image, such as those specular reflections that occur in recent lava flows. In some cases, however, the application of this filter can degrade the image by loss of detail in the high spatial frequencies.

High-Pass Filters (Edge Detection)

High-pass filters emphasize local variations in pixel values, rather than absolute pixel values. Such filters are useful for analyzing texture and for emphasizing linear features and boundaries. One common example, the Laplacian, computes a second derivative of an image. The following matrix is a simple Laplacian filter:

$$F = \begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix}$$

Interpreting the results of a Laplacian filter is not immediately intuitive and perhaps requires some experience. As an enhancement technique, however, the image file that is the result of a Laplacian filter may be added back into the original image, effectively sharpening edge and boundary features. Interpreting this combined image is more intuitive than viewing the filtered image alone. The filter to perform this combined operation is computed as the sum of the identity filter and the Laplacian filter.

$$F = \begin{bmatrix} \text{[identity]} & + & \text{[Laplacian]} & = & \text{[combined]} \\ \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} & & \begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix} & & \begin{bmatrix} 0 & 1 & 0 \\ 1 & -3 & 1 \\ 0 & 1 & 0 \end{bmatrix} \end{bmatrix}$$

The Laplacian filter is particularly susceptible to noise, in the sense that high-frequency, high-amplitude variations such as specular reflections are emphasized. This property, which makes this type of filter valuable as an edge enhancement tool, also makes it of limited utility in very noisy images (Figure 19). In such cases it may be necessary to run a smoothing filter before applying the Laplacian, and in many cases the smoothing filter must be applied several times sequentially to be effective.

Histogram Statistics

As we have noted previously, pixel amplitudes for a given image are a finite set of values and can be displayed in a frequency distribution diagram. For well-defined, homogeneous acoustic targets, these histograms show a Gaussian, or normal, distribution of pixel amplitudes (an example is shown in Figure 13b). This Gaussian distribution of pixel amplitudes, from a single type of target, can be utilized quantitatively to both characterize the backscatter and to separate the amount of acoustic return from two different targets. In using the normal distribution properties of histograms, one caution needs to be applied: only those targets which are both homogeneous and uniform have a Gaussian distribution of pixel values. The presence within the portion of the image under consideration of non-Gaussian backscatter targets, including such features as fissures, acoustic shadows, or system artifacts, perturbs the normal distribution and invalidates the statistical analysis.

The normal distribution of pixel values allows a Gaussian curve to be fit to the data and allows the traditional statistical parameters to mean, median, mode, variance, and skewness to be calculated. Equations and




Figure 19. Application of a Laplacian (second derivative, edge enhancing) filter to a SeaMARC II image. As with Figure 18, the margins of the image are unfiltered. The upper portion shows the result of a Laplacian filter alone applied to the image. The lower portion shows the Laplacian-filtered image plus the original image, added together. The technique of adding multiples of a filtered image to the original image is common in optical image processing and can be used to accent certain types of features such as lineations and edges.

algorithms for calculating these parameters are readily available, and their graphical representations are reviewed in Figure 13*b*. This figure shows the application of Gaussian curve fitting to a SeaMARC II image (shown in Figure 13*a*) that contains both hard (basalt flow) and soft (surrounding sediments) acoustic reflectors. Application of a Gaussian curve-fitting routine to the basalt flow (image A in Figure 13), and calculation of the statistical parameters, shows that they are significantly different from those of the sediments in image B.

The more typical region of "mixed" acoustic reflectors, containing both sediments and part of the basalt flow, is shown in image C of Figure 13*b*. Standard curve-fitting routines can fit multiple Gaussian curves to the distribution, which can then be quantitatively analyzed for their relevant parameters. These Gaussian curve-fitting routines for the grey scale histograms are perhaps the most primitive of the textural analysis techniques that can be applied to images, but they illustrate the power of even simple quantitative analysis.

CONCLUSIONS

Side-scan sonar has become increasingly important in our understanding of the geological processes at work on the seafloor. Clearly, we need to further develop the ability to process these acoustic signals, to enhance the resulting images, and to interpret them unambiguously in a geological sense. In the process of this development, we must be aware that there are several caveats and limitations associated with the technique, but not let these warnings blind us to the potential. Side-scan images, particularly after extensive processing, can look very much like optical images, and, of course, a photograph is a type of data presentation that we are accustomed to interpreting, without hesitation, in our everyday life. It should be kept in mind that regardless of what the images look like, they are representations of the acoustic backscatter from the seafloor: they are not "poor quality" underwater photographs.

Second, the power of image processing has now progressed to the point where it is possible to make the images "look" like almost anything we want them to, including our preconceived ideas of how the seafloor "should" appear. As with other types of geophysical analysis, it is increasingly important that geologists using the data be made fully aware of all the processes that have been previously applied to an image that is to be interpreted. As both the necessary vocabulary and the general knowledge of the acoustic image-processing techniques becomes more widespread, this can be done with increasing efficiency. The use of side-scan sonar has now reached the level of scientific significance where we need to adhere closely to this high level of responsibility.

Less philosophical but equally important pitfalls lie within the day-to-day interpretation of the images. The

viewer of side-scan images needs to be continually reminded of the range dependency of the technique; i.e., the area that is the source of backscatter information, which is compressed into a single pixel, is not the same in the along- and across-track directions. This area varies dramatically in size and shape with distance from the ship track. For reasons that are basic to the way that side-scan systems function, acoustic images of the seafloor are fundamentally anisotropic, and a single target will not look identical if viewed from different look angles or track line distance or if ensonified with a different acoustic system at a different frequency. This conclusion has strong implications for geological interpretations based on the apparent spatial variation of seafloor targets. The directional and range sensitivity of resolution, for example, would argue against the construction (either digital or analog) of any side-scan mosaic which omitted the ship track. Without this guidepost, users of the data could not judge the effects of changing resolution within the images.

Finally, "ground truth" of side-scan images, the verification of the interpretation of an image by independent data, has never been successfully accomplished, although it has been approached on at least one occasion [Hugget and Somers, 1988]. The reasons this valuable procedure is so difficult to accomplish are mainly related to the problems in getting independent ground truth data, of the appropriate scale, from the same area of the side-scan image. For GLORIA II data, for example, with a display resolution (pixel size) of 30 m and an effective ensonification dimension for that pixel of over 150 m, a camera image that is 1 m by 1 m can be useful but is certainly not definitive ground truth data.

Side-scan sonar has come a very long way in the last decade; from the early uncorrected analog images—that required experience, intuition, and a high frustration threshold to interpret—to the high-resolution, digital images which, today, are corrected, processed, and admired like glossy photographs of a favorite relative. As scientists we are fortunate to be present during the first halting applications of an innovative instrument, one which will help illuminate the unknown. We are exploring the ocean basins with a new tool; with some effort, patience, and skill, it will tell us something new about the seafloor.

ACKNOWLEDGMENTS. The research reported in this review was supported by ONR contracts N00014-87-K-0160 and N00014-89-J-1022. G. Carpenter wrote the computer programs used in this study and provided much of the insight into classical image processing. T. Reed kindly provided a copy of his thesis. Useful reviews and assistance were provided by M. Holmes, D. Jackson, J. Pariso, M. Neugebauer, J. Gardner, A. Stevenson, and B. Halbert. D. Scholl served as editor for this article, and deserves both credit and gratitude.

REFERENCES

- Andrews, H. C., *Computer Techniques in Image Processing*, Academic, San Diego, Calif., 1968.
- Andrews, J., and P. Humphrey, SWATHMAP: Long-range sidescan sonar mapping of the deep seafloor, *Mar. Geod.*, 4(2), 141–159, 1980.
- Ballard, R. D., and T. H. van Andel, Morphology and tectonics of the inner rift valley at lat. 36°50' on the Mid Atlantic Ridge: *Geol. Soc. Am. Bull.*, 88, 507–530, 1977.
- Baxes, G. R., *Digital Image Processing*, Prentice-Hall, Englewood Cliffs, N.J., 1984.
- Belderson, R. H., N. H. Kenyon, A. H. Stride, and A. R. Stubbs, *Sonographs of the Seafloor*, Elsevier, New York, 1972.
- Blackington, J. G., D. M. Hussong, and J. Kosalos, First results from a combination sidescan sonar and seafloor mapping system (SeaMARC II), in *Offshore Technology Conference, OTC 4478*, pp. 307–311, Society of Petroleum Engineers, Dallas, Tex., 1983.
- Campbell, F. W., Contrast and spatial frequency, *Sci. Am.*, 231(5), 106–114, 1974.
- Chavez, P. S., Automatic shading correction and speckle noise removal mapping techniques for radar image data, in *Radar Geology, An Assessment, JPL Publ. 80-61*, pp. 251–264, Jet Propulsion Laboratory, Pasadena, Calif., 1980.
- Chavez, P. S., Processing techniques for digital sonar images from GLORIA, *Photogramm. Eng. Remote Sens.*, 52(8), 1133–1145, 1986.
- Davis, E. E., R. Curie, B. Sawyer, and J. Kosalos, The use of swath bathymetric and acoustic image mapping tools in marine geoscience, *Mar. Technol. Soc. J.*, 20(4), 17–27, 1986.
- Edgerton, H. E., Exploring the sea with sonar, *Discovery*, Sept. 1966.
- Fleming, B. W., Sidescan sonar: A practical guide, *Int. Hydrogr. Rev.*, LIII(1), 65–88, 1976.
- Fox, C. G., and D. E. Hayes, Quantitative method for analyzing the roughness of the seafloor, *Rev. Geophys.*, 23, 1–48, 1985.
- Hugget, Q. J., and M. L. Somers, Possibilities of using the GLORIA system for manganese nodule assessment, *Mar. Geophys. Res.*, 9, 255–264, 1988.
- Hussong, D. M., et al., SeaMARC II survey of the Peru trench and continental margin (abstract), *Eos Trans. AGU*, 66, 377, 1985.
- Jackson, D. R., D. P. Winebrenner, and A. Ishimaru, Application of the composite-roughness model to high-frequency bottom backscattering, *J. Acoust. Soc. Am.*, 79, 1410–1422, 1986.
- Karlin, R., and H. P. Johnson, Textural and structural analysis of SeaMARC acoustic imagery of the Juan de Fuca Ridge (abstract), *Eos Trans. AGU*, 68, 84, 1987.
- Klein, M., and H. E. Edgerton, Sonar—A modern technique for ocean exploration, *IEEE Spectrum*, June 1968.
- Kong, L., R. Detrick, P. Fox, L. Mayer, and W. Ryan, The morphology and tectonics of the Mark area from Sea Beam and Sea MARC I observations (Mid-Atlantic Ridge 23°N), *Mar. Geophys. Res.*, 10, 59–90, 1988.
- Kosalos, J. G., and D. N. Chayes, A portable system for ocean bottom imaging and charting, in *Proceedings of Oceans 83*, 649–656, Marine Technology Society, Washington, D. C., 1983.
- Laughton, A. S., The first decade of GLORIA, *J. Geophys. Res.*, 86, 11,511–11,534, 1981.
- Mazel, C., Side scan sonar training manual, Klein Associates, Inc., Salem, N. H., 1985.
- Mudie, J., W. Normark, and E. Cray, Direct mapping of the seafloor using side scanning sonar and transponder navigation, *Geol. Soc. Am. Bull.*, 81, 1547–1554, 1970.
- Pratt, W. K., *Digital Image Processing*, 614 pp., John Wiley, New York, 1978.
- Reed, T. B., Digital image-processing and analysis techniques for SeaMARC II side-scan sonar imagery, Ph.D. dissertation, 256 pp., Univ. of Hawaii, Honolulu, 1987.
- Reed, T.B., and D. Hussong, Digital image processing techniques for enhancement and classification of SeaMARC II side scan sonar imagery, *J. Geophys. Res.*, 94, 7469–7490, 1989.
- Reut, Z., N. G. Pace, and M. J. P. Heaton, Computer classification of seabeds by sonar, *Nature*, 314(4), 426–428, 1985.
- Rosenfeld, A., *Digital Picture Analysis*, Springer-Verlag, New York, 1976.
- Rosenfeld, A., and A. C. Kak, *Digital Image Processing*, 2nd ed., vols. I and II, Academic, San Diego, Calif., 1976.
- Somers, M. L., R. Carson, J. Revie, R. Edge, B. Barrow, and A. Andrews, Gloria II—An improved long range sidescan sonar, in *Proceedings of IEEE/IERE Subconference on Offshore Instrumentation, Oceanology International '78*, technical session J, pp. 16–24, BPS Publications, London, 1978.
- Stanton, T. K., Sonar estimates of seafloor microroughness, *J. Acoust. Soc. Am.*, 75, 809–818, 1984.
- Tucker, M. J., Sideways looking sonar for marine geology, *GMT*, Oct. 1966.
- Tyce, R. C., Near-bottom reflectivity of 4 kHz acoustic reflectivity and attenuation, *Geophysics*, 41, 673–699, 1976.
- Tyce, R. C., Deep seafloor mapping systems—A review, *Mar. Technol. Soc. J.*, 20(4), 4–16, 1986.
- Urick, R. J., *Principles of Underwater Sound for Engineers*, 3rd ed., McGraw-Hill, New York, 1983.
- Vogt, P. R., and B. E. Tucholke, Imaging the ocean floor: History and state of the art, in *The Geology of North America*, vol. M, *The Western North Atlantic*, edited by P. R. Vogt and B. E. Tucholke, pp. 19–44, Geological Society of America, Boulder, Colo., 1986.

M. Helferty, Geophysics Program, University of Washington, Seattle, WA 98195.

H. P. Johnson, School of Oceanography, WB-10, University of Washington, Seattle, WA 98195.