Obstruction Detection and Data Decimation for Airborne Laser Hydrography

Gary C. Guenther, Thomas J. Eisler, Jack L. Riley, and Steven W. Perez

National Oceanic and Atmospheric Administration
National Ocean Service, Office of Coast Survey
1315 East-West Highway
Silver Spring, MD 20910

ABSTRACT

Airborne lidar technology is attractive for hydrographic surveying because of its high rate of area coverage coupled with high spatial data density, rapid-response and reconnaissance capabilities, and cost effectiveness. The National Ocean Service (NOS) has been involved in the design, implementation, and testing of such systems in conjunction with NASA, the U.S. Navy, and the U.S. Army Corps of Engineers (USACE) since 1975. Two topics of current interest are 1) the probability of detection of small objects on the sea bottom as a function of their size, and 2) the efficient reduction of the very large data sets to hydrographically representative but much smaller subsets. Results for both of these are reported here.

INTRODUCTION

Depth determination algorithms written for the USACE Scanning Hydrographic Operational Airborne Lidar Survey (SHOALS) system have been primarily focused on accurate timing of the entire bottom return. The question has arisen to what extent small obstructions or "targets" such as rocks or coral heads are capable of being discriminated and/or detected in the presence of the frequently much stronger bottom return. Target detection probabilities for typical SHOALS parameters, based on a geometric model which includes Monte Carlo simulation results, have been calculated for a variety of target heights, target surface areas, water depths, and water clarities, for air incidence nadir angles of zero and twenty degrees. Results are presented for objects with heights off the bottom between 0.5 m and 2.0 m and surface areas from 1 m^2 to 25 m^2. The effects of sounding grid density on detection probabilities are discussed, and potential depth measurement biases for the detected submerged objects are mentioned.

The processing of airborne lidar surveys involves very large data sets which contain features of hydrographic interest that are spatially oversampled. Because geographic data processing packages can be choked by extremely large data volumes, it is beneficial to pre-process the data to select a much smaller subset which contains all the necessary information. Common solutions to the problem of data reduction entail large-area models which are, themselves, computationally intensive and require significant computer resources. A new approach, which moves linearly through each swath and is not affected by the overall volume of data, has been developed to reduce the size of the survey data set by a factor of roughly ten. The procedure runs efficiently on a low-end PC and yields results which retain hydrographically important features and contain supporting soundings in the overlap regions between swaths.

BACKGROUND

Airborne lidar hydrography (ALH) is a technique for measuring the depth of relatively shallow, coastal waters from the air using a scanning laser beam. Operational systems are presently being flown in the United States, Australia, Sweden, and Canada; these are, respectively, SHOALS (Lillycrop and Banic 1993; Lillycrop et al. 1993; Lillycrop et al. 1994), LADS (Setter and Willis 1994; Nairn 1994), Hawk Eye (Steinvall et al. 1994; Koppari et al. 1994), and Larsen (Vosburgh 1994; Hare 1994). SHOALS and Hawkeye are helicopter borne in aircraft of opportunity, while LADS is in a dedicated Fokker F-27 fixed-wing aircraft, and Larsen is flown in several fixed-wing aircraft. Similar systems have been tested in Russia and China. A number of other nations are expressing interest in purchasing systems or surveys. For example, SHOALS has recently completed a contract survey offshore at Cancun, Mexico for the U.S. Naval Oceanographic Office in conjunction with the Mexican Directorate of Naval Oceanography, and one Hawk Eye is currently being flown in Indonesia.

ALH is an excellent choice for the determination of shoaling in navigation channels and of dredging needs to maintain the channels (Irish et al. 1994), for coastal engineering studies such as investigations of sediment transport and other beach processes (Irish and Lillycrop 1996), for rapid reconnaissance surveys after storms, and for coastal hydrography in both large and small project areas. As an example, in 1995 SHOALS performed a survey following Hurricane Opal to evaluate shoaling of the federal navigation channel at East Pass near Destin, Florida for the USACE Mobile District (Irish et al. 1996). It has been shown by intercomparison with sonar surveys (Riley 1995) that SHOALS can meet International Hydrographic Office (IHO) depth measurement accuracy standards under these conditions.

The general technique of ALH (Guenther and Thomas 1983; Guenther 1985; Guenther 1989; Estep 1993) involves the use of a scanning, pulsed, infrared and blue-green laser transmitter and a receiver composed of a telescope, light detectors, amplifiers, analog-to-digital converters, and tape storage, all under computer control. Approximate depths can be calculated in the air, but precise depths are determined in post-flight processing of the stored waveforms. Typical aircraft altitudes are in the 200-500 meter range, and with maximum scanner nadir angles in the 15-20 degree range, swaths 150-250 meters wide are covered under the aircraft track. Scan patterns vary from system to system. Although laser beams are commonly envisioned as highly collimated with a small cross section (as they are in space or over short distances in air), this is not the case in water where scattering causes the beam to expand into a cone whose interior angle and cross section increase with depth. The resulting increase in irradiated bottom area is beneficial to the detection probability for small objects but, as with broad sonar beams, detrimental to depth accuracy when high-relief features are present.

Conceptually, individual water depths are determined for each pulse from the time-of-flight differences between returns received from the water surface and the sea bottom. In reality, separate receiver channels are used for the surface and bottom returns (Guenther et al. 1994), and one must remove the effects of surface waves (Thomas and Guenther 1990) and a number of biases which accompany the measured surface times (Guenther 1986) and bottom times (Guenther and Thomas 1984a). Maximum surveyable depths depend strongly on water clarity and range from greater than 50 meters in very clean coastal waters to less than 10 meters in cloudy waters. Although a number of technical difficulties have been handled in different ways for the various systems, many of the system limitations involve the physics of light propagation in water and are hence common to all ALH systems. To a large extent, this is true of the ability to detect small obstructions (given appropriate software).

Obstruction detection

In past decades, during the early years of development of ALH systems, the standard hydrographic sensor was vertical-beam echo sounder, and the typical survey encompassed a broad coastal area. In those surveys, bottom topography was inferred from relatively widely-spaced sounding lines. Features of limited extent were often missed between lines. In comparison, ALH offered more nearly full bottom coverage with soundings spaced more evenly throughout the swath. Today, however, most NOS surveys are dedicated to harbors and approaches where small objects on the sea bottom are the primary concern. Vertical-beam and multibeam sonars are now used in conjunction with 200% side-scan coverage to detect features which need further resolution or "development" by additional passes or by divers.

The uses for which ALH is envisioned do not generally include the scenario in which the detection of literally every rock with dimensions on the order of a meter or two must be absolutely guaranteed. Such targets can frequently be detected, but often with less than unit probability unless an unusually dense sounding grid is utilized. Due to the recent emphasis on "100% bottom coverage", it is of interest to know more precisely to what extent small objects on the bottom can be detected from the air for typically-used survey densities. In this paper, results from a propagation model are presented for the functional forms and magnitudes of these hit probabilities for a wide range of parameters.

Data decimation

Today's techniques for shallow-water hydrographic surveying provide data sets which are more spatially dense and more uniformly distributed over the survey area than those previously acquired from classic vertical-beam sonars. One method of dealing with very large data sets containing spatially oversampled features is to institute procedures which "decimate" or select a subset from the large data base that can adequately represent the hydrographic survey

and serve as input to subsequent steps in the charting process, such as contouring. Such strategies typically organize the data from all of the survey swaths into an array of cells and systematically examine the contents of the individual cells. Approaches of this type have been successfully implemented, but they are slow, and the data volumes seem to increase faster than the ability of computer hardware to efficiently handle the required cell arrays.

A procedure has been developed in the Office of Coast Survey to decimate both multibeam sonar and scanning lidar data sets. It is not limited by the size of the original data set because the algorithm moves linearly through each swath with little regard for neighboring swaths that may also be contained in the survey. The data in each swath are tested to determine if there are soundings that are sufficiently similar to be grouped into a local neighborhood which is suitably flat to be represented by a single depth at a single position. The horizontal and vertical requirements for membership in the neighborhoods are set by the operator. Other options include several types of smoothing and several types of shoal biasing. Care has been taken to ensure that local peaks and deeps are not missed. Because of this approach in which each swath is processed separately, as a linear data stream, the price is that the reduction ratio is somewhat limited, and there is some duplication of outputs in the overlap areas between swaths. The latter, however, should be considered to be a benefit, as it provides a level of supporting data for obstruction resolution and general cross checking. For SHOALS data, the reduction ratio is about ten to one for a vertical threshold of 30 cm. With the use of this decimation routine, reductions in dense data sets can be accomplished with high confidence for all spatial scales that are over-sampled during data acquisition.

OBSTRUCTION DETECTION

Introduction

The probability of detecting an object or "target" on the sea bottom using an ALH system depends on the spatial sounding density, the horizontal surface area of the object, its height above the bottom, the depth of the water, water clarity, the nadir angle of the scanning beam, and the data processing algorithms utilized. In general, this detectability depends on three considerations: 1) the geometric "search" probability of hitting at least part of the object with the scanning pulsed laser beam; 2) the probability that the target return is resolvable and strong enough to be detected given the illuminated fraction of the target, the specific location of the target within the illuminated area on the bottom, and the water clarity; and 3) the ability to discriminate the target return and/or accurately measure its location in the waveform. The Office of Coast Survey has provided depth determination algorithms, to USACE for their SHOALS system, which have the capability to recognize and report a small target return in the presence of the much stronger adjacent bottom return.

Procedure

A model has been developed to predict case-specific lidar return waveforms based on geometry and on propagation predictions which are from Monte Carlo simulation studies (Guenther and Thomas 1984b). The geometric model includes a circular disc target of area A(m^2) at a height H(m) above the bottom in water of depth D(m) with diffuse attenuation coefficient, K(m^{-1}). For most cases, the laser beam nadir angle in the water is 15 degrees (consistent with a 20-degree scanner angle in the air), but results are also presented for the nadir case which applies to a portion of the LADS "pushbroom" scan. The waveform model is run in MathCAD which provides interactive operation and rich graphics. Detection probabilities are derived for a square unit cell having the beam footprint centered in it at the bottom. Results are reported for the 4m×4m spatial sampling grid typical of SHOALS operations, but the effects of changing to 3m×3m, 5m×5m, or 10m×10m are discussed.

The nominal SHOALS incident beam diameter at the surface is roughly two meters. Based on the Monte Carlo beam propagation results, the beam diameter in the water expands such that the effective diameter for optically deep waters (large number of scattering lengths) is one-half the depth. (In optically shallow waters, this is somewhat an overestimate.) The modeled beam diameter is thus constant at two meters to a depth of four meters, and half the depth of the water thereafter. For simplicity, a uniform energy density distribution across the downwelling beam was assumed. The target and bottom are assumed to have the same reflectivity, and the relative backscattering strength for each is consequently proportional to its respective irradiated area.

The total backscattered signal to which the detection algorithm is applied consists of the sum of reflected pulses from the target and bottom. The shapes and widths used for these pulses are based on SHOALS parameters and on the Monte Carlo results; they are characteristic of real returns in typical coastal water clarities. The pulse shapes of the

laser source pulse and the impulse response functions of the water column for the target and bottom returns are of the form $(t/t_p)^5 \exp[5(1-t/t_p)]$ (Guenther 1985). Note that this normalized function is zero at t=0, peaks at one at $t=t_p$, and has a fast rise, slow decay character. The value of t_p for the source pulse, expanded to include receiver bandwidth effects, is 9 ns. The values of t_p for the impulse response functions are of the form $\rho D/c_w$, where D is the depth, c_w is the speed of light in water, and ρ is a unitless constant derived from the Monte Carlo results. A value of $\rho=0.085$ is used for the bottom return and half that (representing single scattering downward and multiple scattering upward) for the target, which is usually much more compact spatially.

The total response functions for the target and bottom returns are the convolutions of the source pulse with the respective impulse response functions. These total response functions are normalized, time shifted, weighted, and summed; this weighted sum forms the simulated received signal. The pulse time references are the respective half-amplitude times. The waveforms are first aligned at this point and then time shifted according to the required slant depth difference between the centers of the illuminated portion of the target and the remaining illuminated region of the bottom. The weights are the product of the respective illuminated target areas with exponential decay factors which depend on the slant depths and the selected water attenuation coefficients. Limited receiver fields of view are not applied to decrease the detected bottom area in order to increase the target contrast. This is not done with the real system because it decreases the maximum penetration depth and would frequently limit the amount of surveyable area.

Waveform detection criteria are consistent with SHOALS post-flight data processing algorithms and an absolute extinction criterion. The pulse detection time is determined by applying a leading-edge threshold at the half-peak amplitude level of the identified pulse. Because the sum of target and bottom returns may or may not produce a waveform with two local maxima, for the purposes of this study two types of detection are defined. When two local maxima are present, separated by a dip with at least two one-nanosecond samples lower than the first peak, a target detection clearly occurs. We call this a "type-1" detection. The two nanosecond criterion is based on the design of the SHOALS waveform processing algorithms. Type-1 detections are valid in SHOALS, even if the target return is significantly weaker than the bottom return, because the SHOALS post-flight data processing software retains both "first choice" and "second choice" bottom returns and can also be operated in a "first-pulse" mode. In systems without these facilities, not all type-1 detections may be recognized as targets, and automated detection probability could suffer. No depth accuracy criterion is applied in this case.

For small target heights (on the order of 1 m), the target and bottom returns may be merged into a single "inflected" pulse without two distinct peaks. The detection point on the waveform may or may not be representative of the target, depending on the ratio of the irradiated target area to bottom area and on target height. The object of processing is to get (nearly) the correct target depth, even if it is not possible to recognize the existence of a discrete target. In this case, we choose to consider the target detection successful when the depth, determined by applying a half-peak height threshold to the combined waveform, is not deeper than the true target depth by more than 10 cm. This is called a "type-2" detection, which will be the case for targets that are relatively large in area (compared to the downwelling beam) but not very high off the bottom. Depths shoaler than true are accepted because they err on the "safe" side. Ten centimeters was arbitrarily selected as an acceptable error magnitude; a sensitivity study indicates that increasing the threshold to 15 cm does not significantly increase the detection probability or maximum depth limit.

Target detections are also limited by an absolute extinction criterion based on water clarity, illuminated target area, and target depth. In ALH, the deepest surveyable bottom depth is generally expressed in terms of a maximum KD product (extinction coefficient); for this study, we have selected a conservative value of $KD_{max}=3$ for the bottom return. The extinction coefficient for the target is scaled exponentially, according to the illuminated target area A_t, using the expression $Kd_{max}=3+0.5*\ln[16A_t/\pi d_{max}^2]$ for $d \geq 4$ m, where "d" is the depth at the top of the target. (For $d<4$ m, the argument of the natural log is just $[A_t/\pi]$.) Target extinction coefficients are hence less than three, and smaller for smaller irradiated target surface areas.

The target detection probability is based on the determination of a "detection area" within the square sampling cell representing the mean laser shot spacing. Since all points within the square are equally likely positions for the center of the target disc, the target detection probability is the detection area divided by the area of the cell. Within a representative cell, a 4m x 4m square, the target detection area is defined as that region within the square where a successful type-1 or type-2 detection occurs, including a test on the target extinction coefficient, when the center of

the target disc is within that region. The target detection areas for the different cases are determined manually by moving the target center to various points within the sampling cell and examining the resulting waveforms for type-1 and type-2 detections not exceeding the extinction criterion. The detection area is radially symmetric for the nadir case, but not for the more typical off-nadir case. Because target detection areas overlap into adjacent cells, the total detection area in a cell is the union of detection areas from that cell and surrounding cells. Detection probabilities are adjusted for average values of stagger between adjoining scans when overlaps occur.

Calculations have been performed for nadir angles in air, θ_a, of zero and twenty degrees within the following parameter ranges: target height above bottom ($0.5 \le H \le 2$ m); target area ($1 \le A \le 25$ m^2); depth ($0 \le D \le 30$ m); and diffuse attenuation coefficient ($0.08 \le K \le 0.40$ m^{-1}).

Results

The targets studied in most detail are flat-topped circular cylinders of 1-m and 2-m heights with surface areas of 1 m^2 and 4 m^2 (diameters of 1.13 m and 2.26 m). Detection probabilities for these four objects, for a 4m x 4m sounding grid and a twenty-degree scanner nadir angle in air, are presented as a function of bottom depth in **Figs. 1** and **2.** Four values of water clarity ranging from very clean (0.08 m^{-1}) to very dirty (0.40 m^{-1}) are represented. Not surprisingly, performance is better for larger target heights, larger target areas, and for cleaner waters. For depths at or under 5 m, the probabilities are seen to be mostly independent of water clarity, strongly dependent on target height, and somewhat less so on target area. In this depth range, probabilities for the case of H=1 m and A=4 m^2 are strongly affected by type-2 results. Probabilities rise at middle depths as the beam expands, and water clarity becomes a significant factor. Probabilities for the 2-m height are at or near unity in the 10 - 20 meter depth range for the two cleaner water cases. Results for the 1-m height are rarely above 0.6 for the A=4 m^2 case and do not exceed 0.25 for A=1 m^2. The maximum detection depth depends strongly on water clarity, as might be expected. Extinctions can occur precipitously with increasing depth as the full target area becomes too small to yield a strong enough return according to the extinction equation. For typical coastal water clarities (the two middle cases), the target extinction depths range from 8-13 meters for the 1-m^2 target area and from 10-17 meters for the 4-m^2 area. The largest detection depth is just over 21 m for K=0.08 m^{-1}.

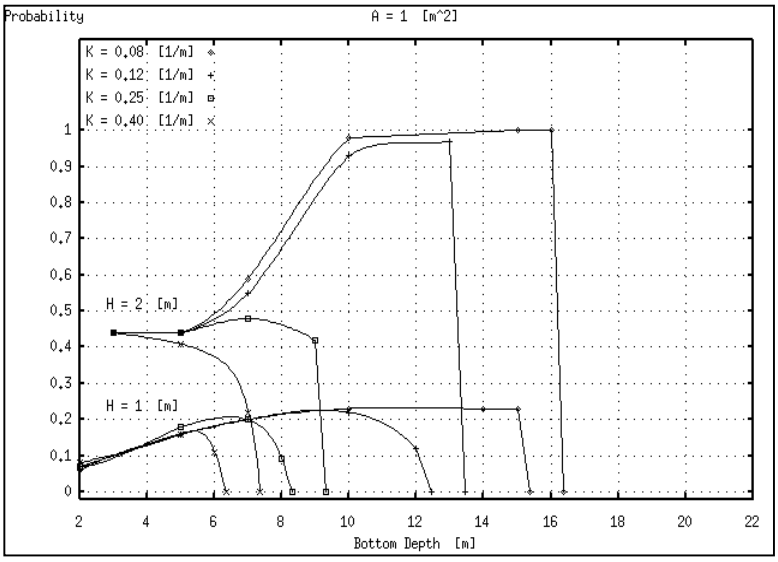

Figure 1 Detection probabilities for 1-m^2 circular cylinders in various water clarities, for 1-m and 2-m target heights, using a 20-degree scanner nadir angle on a 4m x 4m sounding grid.

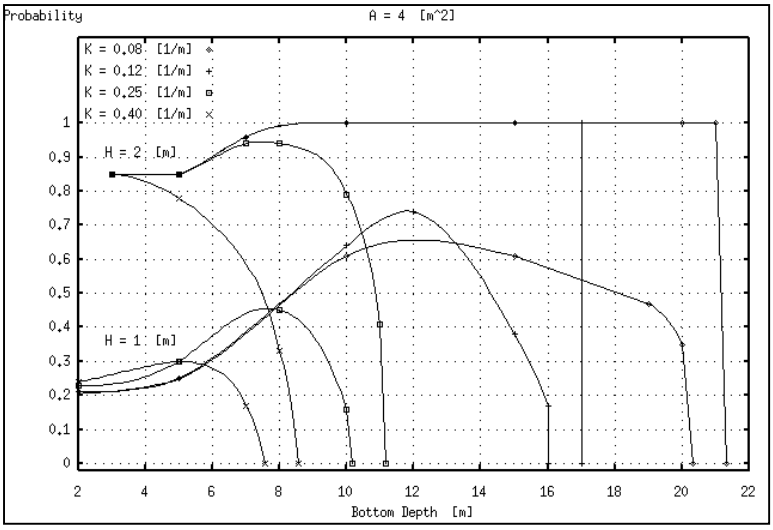

Figure 2 Detection probabilities for 4-m^2 circular cylinders in various water clarities, for 1-m and 2-m target heights, using a 20-degree scanner nadir angle on a 4m x 4m sounding grid.

Insight into the origin of the effect of target height on these probabilities for 1-m and 2-m target heights can be gained by examining a pair of detection areas, seen in **Fig. 3**, for the case of a 20-degree nadir angle with D=5 m, K=0.08 m^{-1}, and A=1 m^2. The detection area for the 2-m height is generally fairly symmetric about the beam axis. For the 1-m height it is smaller because it is shifted into the "undercutting" region on the side closer to the aircraft where the slant paths are shorter, and the resultant target times are better separated from the bottom return.

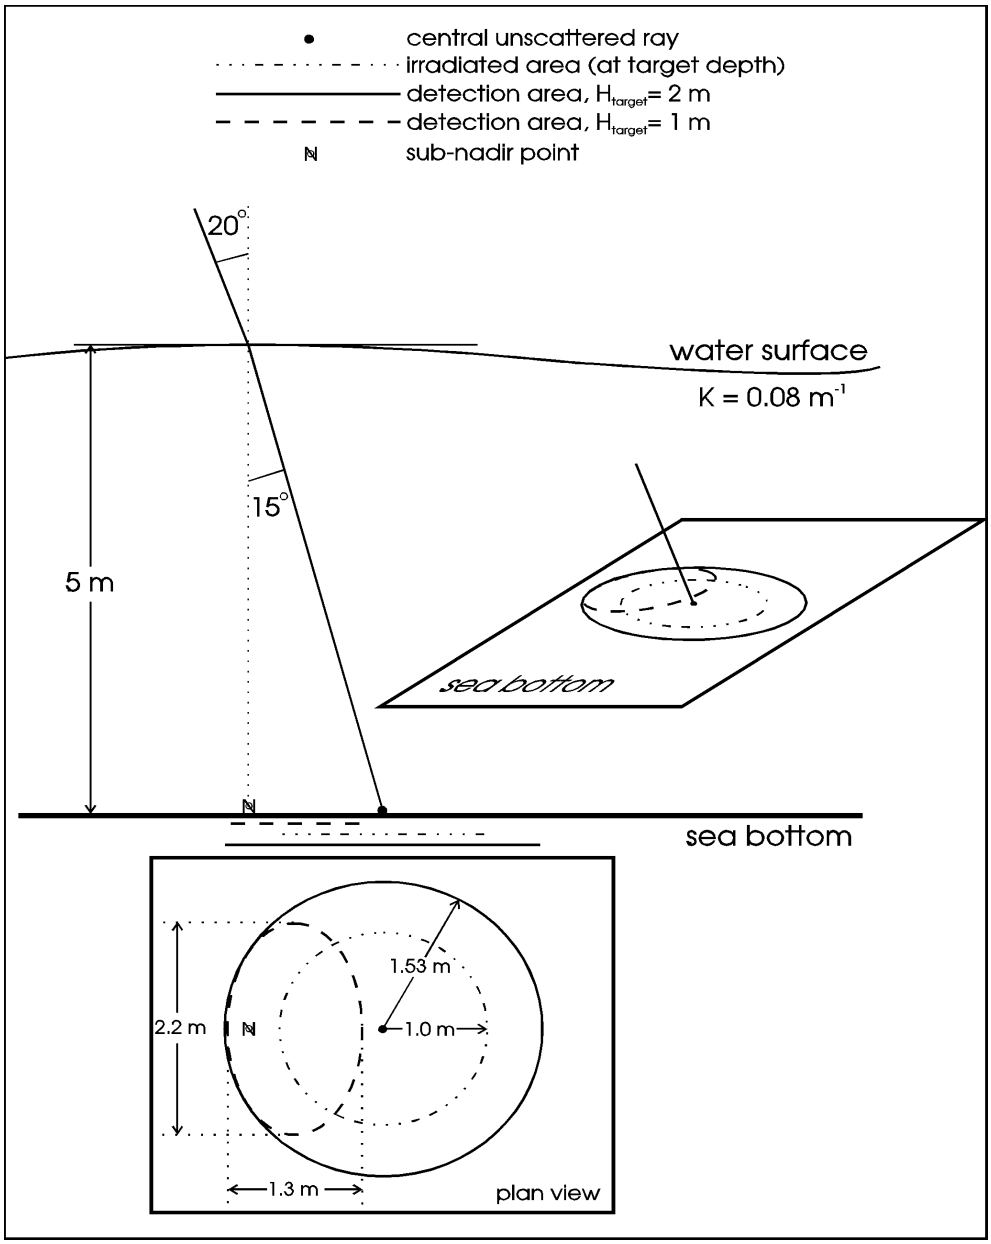

Figure 3 1-m and 2-m target height detection areas for the case of a 20-degree nadir angle with D=5m, K=0.08 m^{-1}, and A=1 m^2.

Details on intermediate target heights are presented for θ_a=20^o and K=0.08 m^{-1} in **Figs. 4a** and **4b**, for A=1 m^2 and A=4 m^2, respectively. Note that the effect of target height is quite nonlinear, particularly for the A=4 m^2 case, where the 1.3-m target height produces results nearly as good as for the 2-m height. For A=1 m^2, target height has a larger effect, and the H=1.3 m probabilities are significantly lower than for H=1.5 meters.

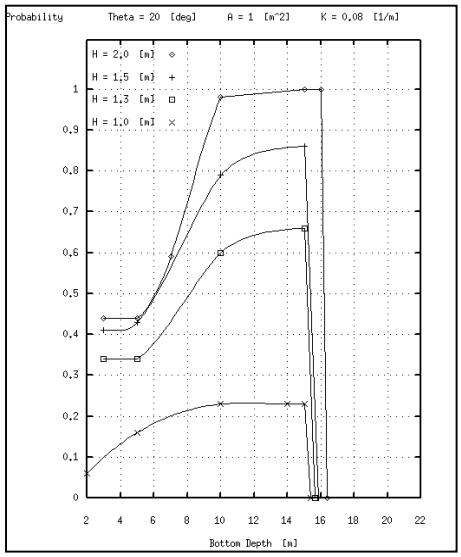

Figure 4a Detection probabilities for 1-m^2 circular cylinders of various heights, using a 20-degree scanner nadir angle on a 4m x 4m sounding grid; K=0.08 m^{-1}.

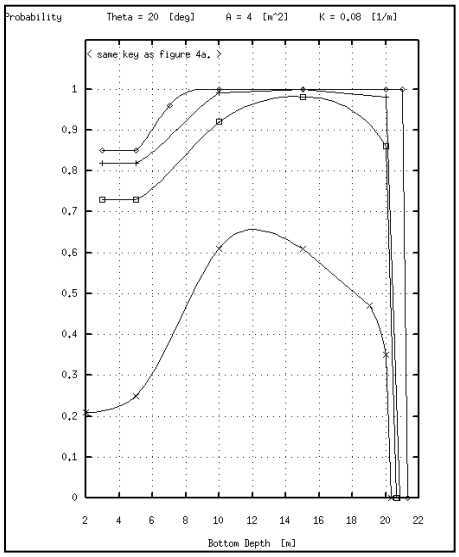

Figure 4b Detection probabilities for 4-m^2 circular cylinders of various heights, using a 20-degree scanner nadir angle on a 4m x 4m sounding grid; K=0.08 m^{-1}.

Results are seen for the nadir case (θ_a=0) in **Fig. 5** for A=1 m^2, K=0.08 m^{-1}, and the same four target heights as in **Fig. 4**. The nadir case provides virtually no detection at all for a 1-m target height at A=1 m^2 and under 30 percent even at A=4 m^2 (not shown). For a 1.3-m target height, the detection probabilities for both target areas are much smaller at depths beyond 8 m than for the θ=20 degree off-nadir case. For H=1.5 m, the θ=0^o results climb more rapidly than for θ=20^o at intermediate depths but then drop sharply beyond 10-m depths. For H=2 m, probabilities for both nadir angles reach unity, although the nadir results climb a bit more quickly, particularly for A=4 m^2.

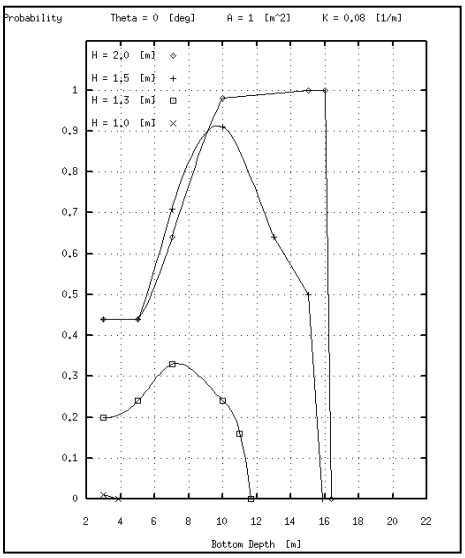

Figure 5 Detection probabilities for 1-m^2 circular cylinders of various heights, using a 0-degree scanner nadir angle on a 4m x 4m sounding grid; K=0.08 m^{-1}.

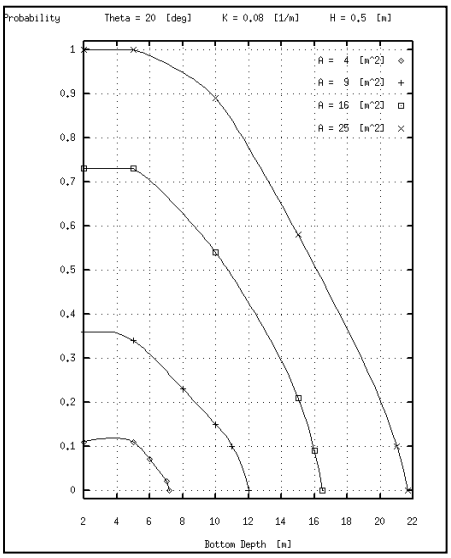

Figure 6 Detection probabilities for 0.5-m high targets of various areas, using a 20-degree scanner nadir angle on a 4m x 4m sounding grid; K=0.08 m^{-1}.

For a target height of 0.5 m, the target returns are merged with the bottom returns, and the boundary of the detection area is determined by type-2 detection. As seen in **Fig. 6**, large target areas are required for high probability of detection in shallow water, and the probabilities decrease more or less linearly with increasing depth. These curves tend to tail off more slowly with increasing depth than for larger target heights; this is because these are terminated by geometry rather than by extinction limits. For targets with the smaller surface areas in shallow water, the low detection probabilities reflect the fact that the laser spot would have to fall solidly on the target for its minimum depth to be successfully distinguished from that of the surrounding bottom. Beyond the bounds of these reported detection areas, the existence of such objects will cause reduced depths to be reported, but the depths will be in error by more than 10 cm.

The effects of survey density can be inferred by considering the ratios of other cell areas to the 16 m^2 used to obtain the above results. These ratios are 1.78 for 3m×3m, 0.64 for 5m×5m, and 0.16 for 10m×10m. The detection probabilities for the other sounding densities can be estimated by multiplying the 4m×4m results presented here by these ratios. This only applies, of course, when the result is less than unity. One can ask, for example, whether it would be worthwhile to decrease the spacing to 3m×3m. The largest effective benefit of the factor of 1.78 occurs for 4m×4m probabilities around 0.5, because these will be increased to near unity. This is true for both the H=1 m, A=4 m^2 case and the H=2 m, A=1 m^2 case. For the H=1 m, A=1 m^2 case, probabilities remain below 0.5. For all but the smallest targets, it would, therefore, be worthwhile to reduce sounding spacing to 3 m. This can be achieved by flying with a narrower swath. Systems with 10m×10m spacing will have detection probabilities only one-sixth as large as those in **Figs. 1-6**. In that case, the probabilities rarely exceed 0.2.

Discussion

The general character of the probability curves is first an increase with increasing depth and then, in many cases, a decrease prior to extinction. At small depths, the probabilities are limited by the fraction of the sounding grid cell irradiated by the laser beam. The probabilities for 2-m target heights are dictated primarily by the chance of hitting the target, while for the 1-m height case, the smaller probabilities reflect the increased difficulty of obtaining path lengths short enough to yield a target return distinct from the bottom return. Probabilities for small depths are greater for the A=4 m^2 case because of the existence of type-2 detections which are missing in the A=1 m^2 case (for which the weaker returns are below the half-peak amplitude used by the timing algorithm). As the depth increases, the detection probabilities increase as the downwelling laser beam in the water expands due to scattering and covers a larger area on the sea bottom. At larger depths, most probabilities decrease due to two factors: the need for a larger fraction of the target to be illuminated to avoid extinction (smaller target separations from the beam axis) and a need for shorter path lengths to maintain separation of the target and bottom pulses as the off-axis path lengths become more nearly equal to the on-axis path for a constant offset from the axis. The curves go quickly to zero at extinction where even the full target area is too small to provide a usable return from that depth.

The nadir case has a simple geometry in which the detection areas are circles due to complete symmetry. As the target moves radially away from the center of the downwelling beam in any direction, its path length increases and causes the target return to merge with the bottom return. The minimum detectable target height is strictly limited by the target height itself. The off-nadir situation is much more complex than the nadir case due to asymmetry in the directions toward and away from the aircraft. The tradeoff is that while detection area is lost in the down-beam direction (away from the aircraft) due to the longer paths causing the target return to merge with the bottom return, the opposite is true in the up-beam direction where the separation is enhanced. At larger depths, this extended up-beam detection area overlaps the missing down-beam area of the adjacent pulse, thus permitting the detection probabilities to be quite large.

The detection area in the up-beam direction is generally semi-circular because it is constrained by the extinction criterion (having enough of the target within the illuminated region). The detection area on the down-beam side typically appears semi-elliptical because it is based on a circular arc drawn not from the beam center, but rather from the aircraft sub-nadir point. This is the locus of points having a constant path length and hence arrival time with respect to the bottom return. For the case of small target heights, such as for H=1 m in **Fig. 3**, the center of the downwelling beam is not in the detection area. This is why there is detection at this off-nadir angle but not for the nadir case. In certain cases, the target returns are plainly visible in the waveform, but "no detection" is indicated because the type-2 depth error criterion is exceeded. This can occur for targets located in both the up-beam and down-beam locations, but for different reasons.

These results depend on a large number of parameters and assumptions such as the received pulse widths and shapes, the beam diameter model, equal target and bottom reflectivities, the extinction criterion, and the pulse detection algorithms. While those used here are believed to be representative, different models will provide different numerical results. Even so, we believe that the relative trends and relationships for different values of H, A, and K should generally hold. The sensitivity to parameters is greater for the borderline cases, particularly for small target heights where the detections are often weak, and at small depths where a relatively large shift in target location is needed to affect the result. It is expected that there is a large sensitivity to the pulse detection/location algorithm, particularly to the 2-ns dip criterion for type-1 and to the use of an accuracy criterion for type-2. It is believed unlikely that alternative reliable algorithms could yield significantly higher detection probabilities, and many will not do as well.

One area where the model could be improved is in the description of the spatial character of the downwelling beam. The next step would be to change from a hard-limited, flat distribution to something more continuous like a Gaussian. One might argue that the model overestimates probabilities at depths in the 5-10 meter range where the scattering model expands the beam more than may be the case in relatively clean waters. On the other hand, the probabilities viewed in nature will actually be larger, because the real beam has a significant amount of energy outside of our fixed diameter criterion.

The depth determination algorithms are carefully tuned to provide the correct result for a return from a flat bottom. When a small object is located at an arbitrary position within the beam footprint, the arrival time of its return pulse depends on its slant depth which, in turn, depends strongly on the spatial relationship of its location to the beam center and the angle of the beam axis. For the nadir case, all target slant ranges off the beam axis are longer. This can cause the target depth to be biased deep, particularly for targets far off axis in the beam fringe. This is undesirable from a navigation safety point of view. This is primarily a problem for small, relatively high targets (type-1 detections). For shorter targets, the target return tends to be lost in the bottom return. The error decreases with increasing depth. The off-nadir situation is more complex. Results for the cross-beam case are similar to the nadir case. In the down-beam direction, the deep bias is even larger, although since slant ranges become longer very quickly, the target return tends to merge with the bottom return, and detections are quickly lost (but will be picked up on the short side of the neighboring pulse). In the up-beam direction, slant ranges are shorter to fringe targets, and shoal biases can result. These are, at least, navigationally safe.

Conclusions

Small objects on the bottom can frequently be detected by an ALH system, but detection cannot generally be guaranteed unless the density of soundings is higher than is currently considered normal. The concept of detection in this paper includes both the recognition of the existence of an object, and the accurate determination of the least depth if recognition is not possible. Because target returns are frequently much weaker than the adjacent bottom return, they could easily go unrecognized unless the waveform processing software is specifically designed, first, to detect small objects on the bottom, and, second, to retain information on the detections. For many ALH applications, it is not necessary to guarantee detection of small objects.

Detection probabilities for 4m×4m sounding spacing are close to unity only for depths greater than 10 m and only for targets with a height of at least 2 m. Results are severely impacted for cases of poor water clarity. Probabilities are highly nonlinear with respect to target height. At this spacing, the detectability of a one-meter cube is below 25 percent. Results are significantly different for beam incidence at nadir and at twenty degrees off nadir. The off-nadir case is generally superior, particularly in the case of borderline target heights on the order of 1 m. Objects less than 1-m high are not frequently detectable, regardless of nadir angle, unless they are relatively large and fall mostly within the laser illuminated area. Performance can be improved by increasing the density of soundings or reducing the laser pulse width. Significant gains can be obtained in many cases by decreasing the average linear sounding spacing to 3 m. Systems with a 10-m spacing will miss all but a small fraction of the target sizes studied here. The depths determined for a small object can be significantly in error, much as with many sonars, if the object is off the axis of the downwelling beam. Operation at nadir yields only potentially dangerous deep biases while for off-nadir beam entry angles, the biases may be either deep or shoal, but are frequently on the "safe" shoal side.

The reported detection probabilities depend on the validity of the model, on the input parameters, and on the detection algorithms. Values for different algorithms could vary considerably; it is believed that the algorithms used here are nearly optimal. It is felt that these results, although clearly approximations, fairly reflect trends in nature and should be representative for SHOALS. The uncertainty of the results is greater in shallow water, because it takes a fairly large change in the target location within the illuminating beam to have a significant effect on the total target response function which, in that domain, is dominated by the laser source pulse. The area where the model can be improved is in the description of the spatial beam profile. Because parameterization and pulse detection procedures vary from system to system, detailed results for similar systems will vary. It is expected, however, that the general trends will hold because they depend primarily on natural phenomena. Finally, systems not specifically designed to find these target returns generally will not.

DATA DECIMATION

Procedure

The basic function of this simple, fast decimation routine for swath data is to scan linearly through the input data and keep only those points which represent a significant change in depth or horizontal distance, as determined in comparison with vertical and horizontal thresholds. It has the ability to work with depths or elevations and to handle positive and negative elevations seamlessly; this permits it to work through the datum or onto land. The input data are sequential depths and positions from the SHOALS output file which contains the temporal history of the data acquisition process. The amount of free memory required to run the program, on even very large data sets, is minimal because the program blocks the SHOALS data into consecutive, non-overlapping groups of 200 points. This blocking provides the capability for the algorithm to dynamically change in concert with changes in the local character of the survey data.

After the program setup selections have been made and the input data set(s) specified, the program blocks the first 200 data points. Within each block, the routine first picks the deepest and the shoalest depths and computes the standard deviation of the depths, and the average horizontal separation between successive points. The first data point in the first block of 200 is taken as the first "base point" and saved to the decimated data set. Base points are used as a basis for comparison with subsequent data points or "test points" to determine if a given data point is to be discarded or kept as an additional base point. A "comparison depth" is assigned to each data point and used whenever depths are compared. It is either the value of the raw input depth for that point or a value obtained by smoothing (see below). The next test point in the data set is compared to the previous base point. In this comparison, the horizontal distance difference and the "comparison depth" difference between the base point and the test point are determined. If either of those parameters exceeds the associated threshold value specified during program setup, the test point becomes a new base point, and the raw depth (never the smoothed depth) of the previous base point is written to the output file (subject to a running check on base points, as discussed below). If neither of those parameters exceeds its threshold, the next sequential data point becomes the new test point, and the base point is unchanged. The shoalest and deepest raw depths for each 200-point data block are also written to the output file, except for the case of "strongly shoal bias" (defined later) when only the shoal depths are written. This operation is repeated until the entire user-specified data set has been processed. The base point and the test point may be in separate blocks of 200.

In addition, the routine performs two running checks on base points, as follows. One can add data points to the output file, and the other can suppress some. The first check is to avoid missing local "peaks" or "deeps" because these are important for fully retaining the envelope of the bottom topography throughout the survey. A local "deep" will occur where the sequential base points are going deeper and then change to moving shoaler. A local "peak" occurs where the sequential base points are going shoaler and then change to becoming deeper. Of the four newest base points, let the base point which was selected nearest the start of the data set be designated point number four and the most current base point be designated point number one. Points number two and three fall between points number one and four, in order. Consider the case of the four most recent base points where base point numbers four, three, and two are all moving in the same direction (either getting shoaler or deeper), but base point number one does not continue the trend. In this case, the routine will search for the data point with the extreme raw depth between base points number three and one. If a point with a raw depth more extreme than the raw depth of base point number two is found in the original data set, that extreme point becomes base point number "two-A". If the distance between base points two-A and three and the distance between base points two-A and one are both less than the maximum

allowable separation, then base point two-A replaces base point two. If either of those separations is greater than the maximum allowable separation, then base point two-A is retained in addition to base point number two.

The second check is to avoid retaining more points than necessary on very large spatial features with gentle slopes. The routine checks to see if the four newest base points are all moving in the same direction -- either getting shoaler or deeper. If so, one of them may be removed, as follows. Using the same point numbering system as above, if the horizontal separation between points number four and two does not exceed the horizontal distance criterion and all four points are on the same slope (up or down), then point number three is removed from the decimated data set. When this happens, the judgement is that the original four base points infer a spatial feature that can be adequately represented using only base point numbers four, two, and one.

There are four procedures which impact the performance for which operator input is required; these are 1) fixed or adaptive vertical threshold, 2) fixed or adaptive maximum horizontal separation threshold. 3) possible smoothing of the bottom model comparison depths, and 4) possible shoal biassing of results. These are described as follows.

1) The operator must either specify a fixed vertical threshold for testing changes in the comparison depths or select an automatic, data-based vertical threshold. The automatic vertical threshold is defined as 1.645 times the standard deviation of the depths in each block of 200 sequential depths. The one-sided tail of a Gaussian distribution beyond 1.645 sigma is 5%; this means that the algorithm will select one point in 20, on average, on the shoal side. As seen below, the effective depth threshold used on the deep side is set by the type of "shoal bias" desired, if any. With no shoal bias, one would also expect one point in 20 to be selected on the deep side. The net result is that the current, automatic, data-based vertical threshold is expected to select one out of every ten points for the decimated data set, assuming a normal distribution. If, on the other hand, a fixed vertical threshold is specified, the reduction ratio will depend strongly on that value. The threshold selected depends strongly on bottom topography and on needs; fixed values might possibly be as small as 0.25 m, but would more likely be in the 0.50 to 1.0 m range.

2) The maximum-allowed horizontal separation between the decimated points is set either by a fixed, operator-input maximum separation threshold or by an automatic separation threshold based on nominal cross-track swath width. In the fixed mode, the two-dimensional distance between each test point and the base point is compared against the given threshold, and a new base point is created when the threshold is exceeded. For the automatic mode, the maximum separation distance between depths is set at 50 percent of that nominal cross-track swath width. On average, in the absence of depth variation, this will save at least two data points to the output decimated data set for each scan. A typical SHOALS semi-circular scan (100-m swath width, 4-m spot spacing) contains about 30 data points. The minimum number of points sent to the decimated data set, based on automated separation threshold alone, would thus be one in fifteen, which is fairly large.

3) For the tests to identify the junctions between adjacent "neighborhoods", the operator must select to use "spatially smoothed", "non-spatially smoothed", or "unsmoothed" comparison depths. If either form of smoothing is selected, the comparison depths used for the neighborhood tests will be calculated from five weighted depths. The two sequential points in the data set preceding a data point, the two sequential points following that point, and the point itself are included in the smoothing which yields a comparison depth for the data point. "Unsmoothed" is the simplest approach, in which the comparison depth is set equal to the raw input depth. "Non-spatial smoothing" involves the use of "boxcar" smoothing based on equal weighting of the five points; i.e., the normalized weights are all one-fifth. The comparison depth is the simple average of the five sequential input depths. The "spatial smoothing" function uses weights based on the normalized horizontal distances of points from the center point. The center depth is weighted unity; weights for the four other depths are one half times the inverse of the normalized distance between the point and the center point expressed in "nominal separation units". The nominal separation unit is a distance normalizing factor which is equal to the average of the separations between successive soundings in the current block of 200 sequential data points. For example, a depth separated by one nominal spatial unit from the central depth is weighted 0.5, and a depth separated by two nominal separation units from the central depth will be weighted 0.25. Weights are truncated at unity for normalized distances of less than one half. Since these weights are unnormalized, the sum of the weighted depths is divided by the sum of the weights to provide normalization of the weighted depths for use as the comparison depth.

4) The operator must select between a reduction of depths which is "unbiased", "weakly shoal biased", or "strongly shoal biased". The three cases are explained as follows. (a) Unbiased reduction is the case where the thresholds for

both positive and negative changes in the comparison depths, which define the junctions between two adjacent neighborhoods, are equal in magnitude (to the vertical threshold). The data points with the deepest and shoalest raw depths in each non-overlapping block of 200 successive depths are also included in the output file. (b) Weakly shoal biased is the case where the effective threshold for a deeper comparison depth change is 1.2 times the vertical threshold used for a shoaler depth change. The deepest and shoalest raw depth data points in each block of 200 are also included in the output file. The weak shoal bias results from the 1.2 factor between the two thresholds. (c) Strongly shoal biased is the case where the effective threshold for a positive (deeper) comparison depth change defining the junction between two adjacent neighborhoods is 1.5 times the vertical threshold. The shoalest raw depth data point in each block of 200 is also included in the output file. The strong shoal bias results both from the 1.5 factor between the two thresholds and from the deepest point in each block of 200 not being included in the output file. The selection of one of the shoal bias options provides a reduced data set which is less dense on the deep side. This might be appropriate for hydrography (i.e., nautical charting); it would not be appropriate for bathymetry (i.e., mapping the bottom contours).

The output file contains the decimated data set followed by a trailer with the parameters employed in the reduction algorithm. This provides an audit trail for the reduced data set. The routine also retains the parameters employed in its most recent utilization. When using the program on multiple data sets, this setup history is used to prompt the operator in setting the parameters on the program's application to the next data set.

Results

As seen in Table 1, the reduction ratio depends strongly on the values used for vertical and horizontal thresholds, and also on whether smoothing is used for comparison depths. It also, of course, depends on the roughness of the bottom topography and the noise level of the measurement system.

Table 1. Reduction ratios for various parameter combinations.

vertical threshold	horizontal threshold	smoothing on/off	Reduction ratio	smoothing on/off	Reduction ratio
15 cm	half scan	off	3.9	on	6.9
15 cm	100 m	off	4.1	on	8.1
15 cm	150 m	off	4.1	on	8.9
30 cm	half scan	off	8.0	on	10.9
30 cm	100 m	off	10.0	on	15.6
30 cm	150 m	off	10.1	on	17.0
1.645 sigma	half scan	off	8.4	on	13.0
1.645 sigma	100 m	off	10.5	on	21.1
1.645 sigma	150 m	off	10.5	on	21.3
50 cm	half scan	off	12.4	on	15.2
50 cm	100 m	off	20.0	on	28.9
50 cm	150 m	off	20.2	on	29.8

It can be seen that a vertical threshold of 15 cm is too small to achieve a significant reduction, except for the smoothed cases with a larger horizontal threshold, given the bottom topography and the system noise level. For the unsmoothed case, reduction ratios of around 10:1 are obtained both for a 30-cm fixed vertical threshold and for the automated vertical threshold which is expected to pick about ten percent, given a Gaussian distribution. Smoothing of the comparison depths provides a significant improvement in performance. For the case of the automated vertical threshold, the implication of the over 20:1 ratios is that the smoothed comparison depths no longer have a Gaussian distribution. If a 50-cm vertical threshold is satisfactory, ratios of 20:1 to nearly 30:1 can be achieved for the larger, manual horizontal thresholds. It is noted that the automated half-scan (50 m) horizontal threshold seriously limits the reduction ratio for all of the practical vertical thresholds. If larger minimum horizontal separation is acceptable, significant gains are achieved. For horizontal thresholds beyond 100 m, there is little difference.

A sample of points representing the sounding locations for a section of a full-density, standard SHOALS data set is seen in **Fig. 7a** with a 25-m square grid overlaid. The associated decimated set for a fixed vertical threshold of 30 cm and a horizontal threshold of 100 m is seen in **Fig. 7b**. The reduction ratio for this case is 10:1. It should be noted that because the decimated data set is sparse, the existence of complete coverage must be confirmed in the full data set prior to decimation.

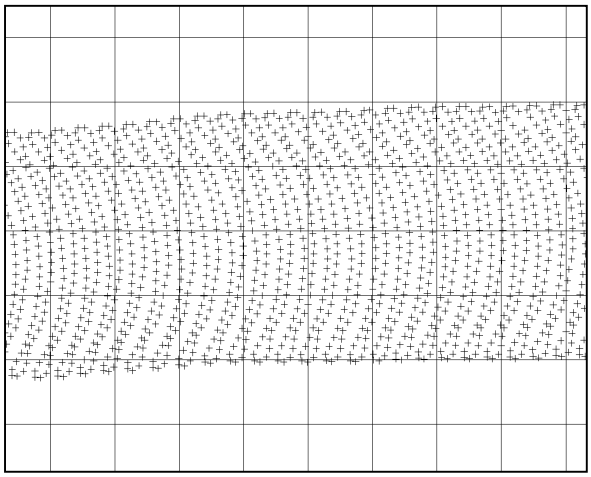

Figure 7a Portion of SHOALS 100-m swath; average sounding spacing is 4m x 4m; grid is 25m x 25m.

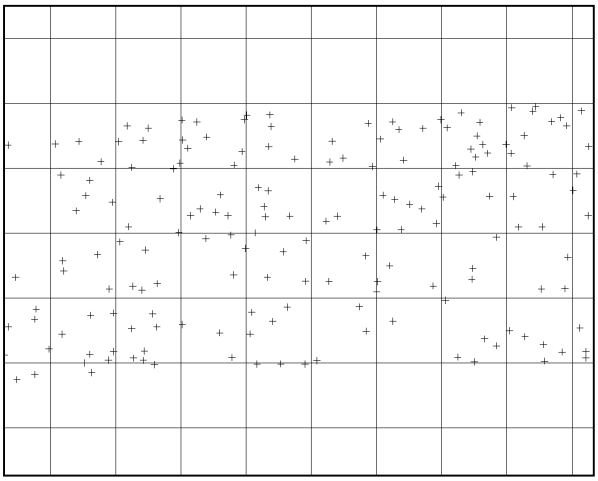

Figure 7b Portion of decimated SHOALS 100-m swath; grid is 25m x 25m (10:1 reduction ratio: 30-cm vertical threshold, 100-m horizontal threshold, unbiased, unsmoothed).

Three simulated shoal points were added in an overlap region between three SHOALS swaths, one in each swath. A portion of this data is seen in **Fig. 8a** overlaid by an 8-m square grid. The simulated feature can be noted on the far right at mid height. All three of these points are found in the decimated set in **Fig. 8b.** This added support from crosslines or overlaps is felt to be a beneficial trait. A more sophisticated, fully two-dimensional decimation routine would not have kept all three points. [A 15-cm vertical threshold, which would not normally be used, was applied simply to put a few more points on the paper; otherwise it would have been virtually blank except for the three points representing the feature.]

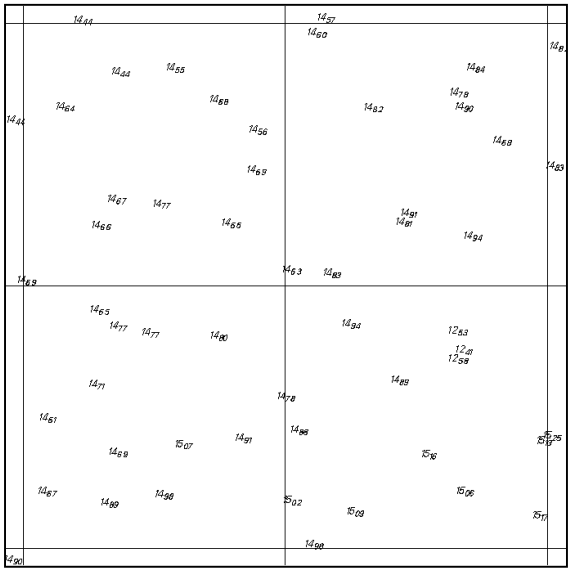

Figure 8a Overlap region between three SHOALS swaths; grid is 8m x 8m. Three simulated shoal points (12.53m, 12.41m, and 12.58m), each from a separate swath.

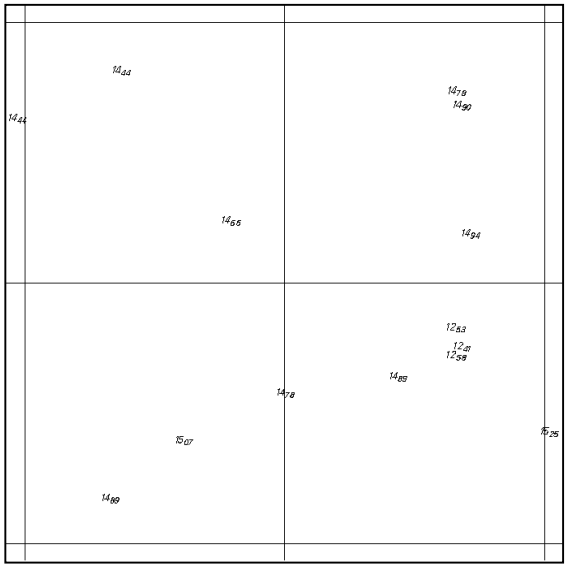

Figure 8b Overlap region between three decimated swaths; three shoal points preserved. (3.9:1 reduction ratio: 15-cm vertical threshold, 50-m horizontal threshold, unbiased, unsmoothed.)

Conclusions

A simple, efficient routine has been developed to significantly reduce the size of a SHOALS data set without loss of any important information. This procedure performs the desired data reduction task in a satisfactory manner. Peaks and deeps in SHOALS data sets are properly identified. As an added bonus, features are confirmed by supporting soundings for crosslines or swath overlaps. A fixed, 100-m horizontal threshold is recommended.

ACKNOWLEDGMENTS

The authors would like to express gratitude for partial funding of this effort by Headquarters, U.S. Army Corps of Engineers, Operations, Construction, and Readiness Division.

REFERENCES

Estep, L., 1993. "A review of airborne lidar hydrographic (ALH) systems," The Hydrographic Journal, No. 67, 25-42.

Guenther, G.C. and R.W.L. Thomas, 1983. "System design and performance factors for airborne laser hydrography," Proc. Oceans '83, Aug. 29 - Sept. 1, 1983, San Francisco, Calif., IEEE/MTS, 425-430.

Guenther, G.C. and R.W.L. Thomas, 1984a. "Prediction and correction of propagation-induced depth measurement biases plus signal attenuation and beam spreading for airborne laser hydrography," NOAA Tech. Report NOS 106 CGS 2, National Oceanic and Atmospheric Administration, Rockville, Md., 121 p.

Guenther, G.C. and R.W.L. Thomas, 1984b. "Effects of Propagation-induced Pulse Stretching in Airborne Laser Hydrography," Proc. SPIE Ocean Optics VII, Vol. 489, 287-296.

Guenther, G.C., 1985. "Airborne Laser Hydrography: System Design and Performance Factors," NOAA Professional Paper Series, National Ocean Service 1, National Oceanic and Atmospheric Administration, U.S. Department of Commerce, Washington, D.C., 385 p.

Guenther, G.C., 1986. "Wind and nadir angle effects on airborne lidar water 'surface' returns," Proc. SPIE Ocean Optics VIII, Vol. 637, 277-286.

Guenther, G.C., 1989. "Airborne laser hydrography to chart shallow coastal waters," Sea Technology, Vol. 30 #3, 55-59.

Guenther, G.C., P.E. LaRocque, and W.J. Lillycrop, 1994. "Multiple surface channels in SHOALS airborne lidar," Proc. SPIE Ocean Optics XII, Vol. 2258, 422-430.

Hare, R., 1994. "Calibrating Larsen-500 lidar bathymetry in Dolphin and Union Strait using dense acoustic ground truth," Int'l Hydro. Rev., Monaco, LXXI(1), 91-108.

Irish, J.L., W.J. Lillycrop, L.E. Parson, and M.W. Brooks, 1994. "SHOALS system capabilities for hydrographic surveying," Proc. Dredging '94, Nov. 1994, Lake Buena Vista, Florida, 314-321.

Irish, J.L. and W.J. Lillycrop, 1996. "Monitoring New Pass, Florida with high density bathymetry," J. Coastal Res. (in press).

Irish, J.L., J.E. Thomas, L.E. Parson, and W.J. Lillycrop, 1996. "Monitoring storm response with high density lidar bathymetry: the effects of Hurricane Opal on Florida's panhandle," Proc. Second Int'l Airborne Rem. Sens. Conf.,

June 24-27, 1996, San Francisco, California, (in press).
Koppari, K., U. Karlsson, and O. Steinvall, 1994. "Airborne laser depth sounding in Sweden," Proc. U.S. Hydro. Conf. 1994, The Hydrographic Society Spec. Pub. No. 32, 124-133.

Lillycrop, W.J. and J.R. Banic, 1993. "Advancements in the U.S. Army Corps of Engineers hydrographic survey capabilities: the SHOALS system," Marine Geodesy, Vol. 15, 177-185.

Lillycrop, W.J., L.E. Parson, and G.C. Guenther, 1993. "Processing lidar returns to extract water depth," Proc. Int'l Symp. Spectral Sens. Res., Nov. 1992, Maui, Hawaii.

Lillycrop, W.J., L.E. Parson, L.L. Estep, P.E. LaRocque, G.C. Guenther, M.D. Reed, and C.L. Truitt, 1994. "Field testing of the U.S. Army Corps of Engineers airborne lidar hydrographic survey system," Proc. U.S. Hydro. Conf. '94, The Hydrographic Society Spec. Pub. No. 32, 144-151.

Nairn, R., 1994. "Royal Australian Navy Laser Airborne Depth Sounder, The first year of operations," Int'l Hydro. Rev., Monaco, LXXI(1), 109-119.

Riley, J.L., 1995. "Evaluating SHOALS bathymetry using NOAA hydrographic survey data," Proc. 24th Joint Mtg. of UJNR Sea-Bottom Surveys Panel, Nov. 13-17, 1995, Tokyo, Japan.

Setter, C. and R.J. Willis, 1994. "LADS -- From development to hydrographic operations," Proc. U.S. Hydro. Conf. 1994, The Hydrographic Society Spec. Pub. No. 32, 134-139.

Steinvall, O., K. Koppari, and U. Karlsson, 1994. "Airborne laser depth sounding: system aspects and performance," Proc. SPIE Ocean Optics XII, Vol. 2258, 392-412.

Thomas, R.W.L. and G.C. Guenther, 1990. "Water surface detection strategy for an airborne laser bathymeter," Proc. SPIE Ocean Optics X, Vol. 1302, 597-611.

Vosburgh, J., 1994. "Larsen lidar -- past experience, present system, continuing refinements -- from a hydrographic perspective," Proc. U.S. Hydro. Conf. 1994, The Hydrographic Society Spec. Pub. No. 32, p.206.