Design considerations for achieving high accuracy with the SHOALS bathymetric lidar system

Gary C. Guenther[1], Robert W.L. Thomas[2], and Paul E. LaRocque[3]

[1]NOAA Charting Development Laboratory, National Ocean Service, Silver Spring, Maryland 20910
[2]EG&G Washington Analytic Services Center, 16156 Dahlgren Rd., Dahlgren, Virginia 22448
[3]Optech, Inc., 100 Wildcat Rd., North York, Ontario, Canada M3J 2Z9

ABSTRACT

The ultimate accuracy of depths from an airborne laser hydrography system depends both on careful hardware design aimed at producing the best possible accuracy and precision of recorded data, along with insensitivity to environmental effects, and on post-flight data processing software which corrects for a number of unavoidable biases and provides for flexible operator interaction to handle special cases. The generic procedure for obtaining a depth from an airborne lidar pulse involves measurement of the time between the surface return and the bottom return. In practice, because both of these return times are biased due to a number of environmental and hardware effects, it is necessary to apply various correctors in order to obtain depth estimates which are sufficiently accurate to meet International Hydrographic Office standards. Potential false targets, also of both environmental and hardware origin, must be discriminated, and wave heights must be removed. It is important to have a depth confidence value matched to accuracy and to have warnings about or automatic deletion of pulses with questionable characteristics. Techniques, procedures, and algorithms developed for the SHOALS system are detailed here.

Keywords: airborne laser hydrography, lidar, bathymetry, accuracy, software algorithms, SHOALS

1. BACKGROUND

The Scanning Hydrographic Operational Airborne Lidar Survey (SHOALS) system[1] was designed and constructed by Optech, Inc. for the Waterways Experiment Station of the U.S. Army Corps of Engineers (USACE) and the Canadian Department of Industry, Science, and Technology. It is presently operated by John E. Chance & Associates in a Bell 212 helicopter flown by the National Oceanic and Atmospheric Administration (NOAA) Aircraft Operations Center. Since successful field tests[2], operational missions have been flown for general-purpose hydrography and monitoring of shoaling in navigation channels[3], coastal engineering studies of sediment transport[4], and rapid-response storm damage assessment after Hurricane Opal[5]. Excellent results were obtained in an intercomparison with a National Ocean Service sonar survey[6]. In a recent survey performed for the U.S. Naval Hydrographic Office in conjunction with the Mexican Directorate of Naval Oceanography, 320 square nautical miles were charted off the coast of Cancun, Mexico.

Airborne laser (or lidar) hydrography (ALH) is a technique for measuring the depth of relatively shallow, coastal waters from the air using a scanning laser beam. The concept of ALH grew out of efforts in the mid 1960's to use the newly invented laser to find submarines. The seminal paper confirming the ability to perform near-shore bathymetry was written by Hickman and Hogg[7] based on work done at the Syracuse University Research Center. In the early 1970's a number of first-generation airborne lidar systems were successfully tested by the U.S. Navy[8-11], by the National Aeronautics and Space Administration (NASA)[12], in Canada[13,14], and in Australia[15]. Much of the early work in the Soviet Union[16] and in Canada[17] was ship-borne. Several symposia, co-sponsored by NOAA and NASA, were convened to establish design goals and user requirements for the use of the second-generation NASA Airborne Oceanographic Lidar (AOL) for hydrography[18,19]. Successful field testing of the AOL was conducted in 1977[20-22]. As a result, the existence of environmentally-induced biases in both surface and bottom returns was discovered[23]. Other second-generation systems were built and tested in Canada[24], in Australia[25,26], and in the Soviet Union[27]. The Canadian system, augmented with a scanner, was also tested in Sweden[28].

Design and testing of a number of systems, such as the U.S. Navy HALS[29,30] and the Swedish FLASH[31,32], continued through the 1980's. In Canada, the Larsen-500 was developed[33,34] and, as a result of surveys performed in the Northwest Territories, became the world's first operational ALH system[35-39]. Testing of the Australian WRELADS II was completed[40] and construction begun on the operational version called LADS[41,42]. Three multi-purpose research systems were actively tested in the Soviet Union[43-47], and work was also conducted in China[48]. The U.S. Army Corps of Engineers began the SHOALS program[49,50].

The technology matured in the late 1980's, and in the 1990's, three systems have become operational and are presently being flown in the United States, Australia, and Sweden. These are, respectively, SHOALS[51], LADS[52,53], and Hawk Eye[54,55]. The Canadian Larsen-500 system continues to be flown[56]. SHOALS and Hawkeye are helicopter borne in aircraft of opportunity, while LADS is in a dedicated Fokker F-27 fixed-wing aircraft, and Larsen-500 is flown in several fixed-wing aircraft. A number of other nations are expressing interest in purchasing systems or surveys. For example, in addition to the use of SHOALS in Mexico, one of the two current Hawk Eye systems is currently being flown in Indonesia.

2. INTRODUCTION

The general technique of ALH[57-60] 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 are calculated in the air, but precise depths, involving more detailed calculations, are determined in post-flight processing of stored waveforms. Eventually, with faster computers and more experience with processing data from diverse situations, it may become possible to do final depth determination in the air in real time. 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; both semi-circular and rectangular are in use. Maximum surveyable depths depend strongly on water clarity; they range from greater than 50 meters in very clean coastal waters to less than 10 meters in cloudy waters. 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[61] but, as with broad sonar beams, detrimental to depth accuracy, particularly when high-relief features are present.

One of the biggest problems which must be solved in the design of a bathymetric lidar receiver is the handling of the more than six orders of magnitude of amplitude dynamic range between strong surface returns and weak bottom returns. That difference, which occurs in only a matter of tens or hundreds of nanoseconds, must be compressed into the useful input range of the digitizer, which is typically only two orders of magnitude. There is no single solution; a combination of several techniques, each of which handles a portion of the range, is required. Potential approaches include avoiding near-nadir scanner angles and the use of cross polarization, partial optical blocks, logarithmic amplifiers, and variable gain photomultiplier tubes. Great care must be taken in the application of these techniques, however, because these methods also have undesirable side effects such as variable time delays and pulse shape distortion which can cause depth measurement errors.

A second major problem involves reliable and accurate determination of the location of the air/water interface[62]. Green surface returns arise from a linear combination of interface and volume backscatter returns whose relative magnitudes vary wildly from pulse to pulse[63,64]. The arrival time difference between these components is too large to permit the use of this combined return for surface timing. Infrared (IR) reflections have the advantage that the interface component dominates; the IR volume backscatter is not only generally much weaker, but because of much higher attenuation in the water[65], it arises from a volume near enough to the interface so as not to cause an unacceptable timing error. Under calm wind circumstances, however, at larger nadir angles these components can become undetectably weak in extremely clear water. The IR returns can also arise from false targets above the surface such as spray, birds, and sea smoke. For this reason, it is a great advantage to have a surface channel tuned to the green-excited Raman backscatter wavelength in the red portion of the spectrum. This inelastic process, which arises from a vibrational mode of the O-H bond in liquid water molecules[66,67], yields a weak but useable return, at least for lower aircraft altitudes. Because this return arises solely from volume backscatter under the interface, its arrival time must be corrected to the predicted location of the interface. The benefits are that it will be present regardless of sea surface wave slopes (is independent of wind speed), does not come from spray, birds, or sea smoke, and that it has a relatively small amplitude dynamic range. In order to be able to handle all circumstances and provide reliable, accurate, false-alarm free surface location, a receiver must have channels at all three wavelengths: green, IR, and red.

The operational production of reliable mean water depths accurate to International Hydrographic Office standards involves detailed understanding of the characteristics of the data collection electronics and of a number of physical interactions between the laser beam and the water about which many volumes have been written. The process must begin with proper system design so that flight data have desirable characteristics, and all necessary system outputs are available and unambiguous. Optical and electronic time delays through the system must be carefully determined, particularly those that are amplitude dependent. The nadir angle of the beam must be known to very high accuracy. Timing and angle calibration

procedures must be developed. The waveform processor must discriminate the desired surface and bottom returns from noise and false targets, save at least two possible bottom returns to permit valid depths to be calculated in the presence of fish and other biota, ascertain the most precise estimate of their arrival times, correct for hardware and environmental time delays, recognize and eliminate possible erroneous data, calculate the mean water level and remove wave heights from the measured water-column depths, determine the depth in a manner consistent with the quality of available surface data, report warnings for questionable circumstances, and provide an overall level of confidence in the result which is keyed to depth accuracy. All of this must be conducted in a flexible framework which provides interactive data displays and adaptability through efficient operator involvement.

Conceptually, local water-column 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, as noted above, separate receiver channels are required for detecting the surface and bottom returns[62], and one must remove the effects of surface waves[68,69] and a number of biases which accompany the measured surface times[63] and bottom times[70-72]. Studies have shown that the most accurate and precise pulse-location algorithm is a half-peak-height amplitude threshold applied to the linear waveform[73]. It was demonstrated in the processing of data from the U.S. Navy HALS system[74] that algorithms containing heuristic rules[75] can provide excellent recognition of the desired returns along with rejection of noise, system artifacts, and some false targets in the environment[64]. Techniques and procedures employed for the SHOALS system are detailed here.

3. *SHOALS* SYSTEM HARDWARE

The laser is a Big Sky (Bozeman, Montana) Nd:YAG with an IR output of 15 mJ at 1064 nm and simultaneous, collinear, frequency-doubled green output of 5 mJ at 532 nm. The pulse repetition rate is 200 Hz and the pulse width 6 ns in the green. The beam is expanded to be eye safe at operational altitudes and to provide a reliable surface returns. The SAAB (Jönköping, Sweden) programmable scanner[31] is a flat mirror servoed under computer control in two axes to produce the desired scan pattern and to compensate for aircraft roll and pitch, as measured with a Litton LTN-90 inertial navigation system. In this way, the edges of the swath remain straight regardless of changes in aircraft attitude. This minimizes the need for overlap between swaths and reduces the likelihood of "holidays" in the coverage between swaths. The scan pattern selected for operational surveying, a segment of a circular arc aimed ahead of the aircraft, is created by maintaining a nearly constant nadir angle of twenty degrees.

A constant nadir angle is beneficial both because it minimizes the dynamic range of surface return amplitudes which can be exceptionally strong near nadir[76] and because it eliminates one major variable in the functionality of propagation-induced depth-measurement bias. It is important to avoid the strong surface returns associated with near-nadir angles because photomultiplier tube (PMT) detectors must be protected from strong signals which can cause deleterious saturation effects even when the remainder of the tube is gated off. The value of twenty degrees was selected because it provides a cost-effective swath width while minimizing the magnitude of propagation-induced depth measurement biases, and their variability as a function of unknown water optical properties, without creating unacceptable values of other geometric errors[58]. A small variation of about 1.5 degrees in the nadir angle at the ends of the scan lines is used to provide loops whose purpose is to prevent pulse pile-up and keep the sounding spacing near the desired average. Setup parameters include pulse repetition rate, planned aircraft altitude, desired swath width, and sounding spacings in along-track and cross-track directions. The scanner is then automatically programmed with the scan rate and maximum azimuthal angle needed to provide the desired characteristics. If desired, the nadir angle and pattern can be varied to meet special needs. Because highly precise knowledge of the values of scanner angles and the resulting nadir angle is crucial to system depth measurement and position accuracies, and scanner movement is rapid, the scanner angles are measured at the instant of laser firing. Latencies in attitude data have been measured and are deskewed in software.

Returning signals, collected by a 20-cm catadioptric Cassegrain telescope, are split into two green channels, two IR channels, and one red channel at the water Raman wavelength of 645 nm. All receiver options are computer controlled through a graphical user interface, and all parameters are recorded on the data tape to become part of the mission record. The outputs of the two green channels, the main IR channel, and the Raman channel are all digitized to ten bits in 1-ns samples with an Analytek 1-GHz digitizer. Digitized waveforms and all other system outputs are recorded, for post-flight data processing, on 5-GB 8-mm tape cartridges using Exabyte 8500 drives. The system includes a built-in optical simulator, using light-emitting diodes, for exercising system functionalities. A real-time depth algorithm, which is a primitive subset of the post-flight waveform processing algorithm, supplies approximate depths to the airborne operator displays. This provides a measure of

real-time quality control and permits the operator to make decisions regarding the need to alter parameters, to refly a line, or to move to a different area.

In SHOALS, the solution to the dynamic range problem includes the use of two separate green channels. The more sensitive, high-gain green channel, optimized for detection of weak, deep returns, uses a time-gated PMT detector. The PMT is generally gated “on” to its full, fixed gain only for times after the pulse has passed through the air/water interface. This avoids the associated strong surface return, which, at full gain, could cause undesirable effects such as saturation, ringing and raised baselines. Depth detection begins at about 6 meters. The channel gain is preset by selecting one of six constant values of the dynode high voltage. Because PMT’s exhibit signal transit-time delays that vary with both gain and signal amplitude, amplitude-dependent delays are corrected using a predetermined calibration vector for each fixed gain. Optical input to the PMT is further controlled by preselected, servo-driven, variable inner and outer fields of view (FOV) and optical attenuators; the maximum FOV is 50mr. The channel has been configured with a maximum depth of about 37 meters; this could be doubled for cases of extremely clear water, with some loss or precision, by reprogramming the digitizer to 2-ns bin widths.

The less sensitive, so-called “shallow” green channel uses an avalanche photodiode (APD) detector, and the digital record includes the surface return and depths to about 13 meters. This channel, with a smaller, fixed FOV, is used for detection of shallow depths, for determination of topographic heights on land adjacent to the water, and as a backup in the hardware surface return logic for controlling the digitizer. The minimum depth is approximately one meter, although studies indicate that it may be possible to ease that restriction in the future. Both the shallow and deep green channels have logarithmic amplifiers capable of compressing input signal variations of 4.5 orders of magnitude into the digitizer range. The combination of these two bottom channels provides the ability to cover a very large dynamic range while maximizing performance in both regimes, excellent flexibility, and redundancy in the 6-13 meter depth range.

The main IR channel and the Raman channel use fixed FOV’s and APD detectors, and the waveforms are digitized for near-surface times only. Signals from an adjunct, cross-polarized IR channel are used in conjunction with those from the main channel to detect the presence of land returns in real time.

Because of the above-mentioned problems associated with surface detection, SHOALS was designed with three independent surface channels: Raman, IR, and green. These three channels are designed to supply the information necessary to provide accurate, reliable air/water interface times for every laser pulse in real time. These are used for the calculation of approximate depths in real time and may be used for the determination of more accurate depths in post-flight processing. Since none of the three channels, by itself, can be guaranteed to provide the needed surface location times over the entire range of operating altitudes and environmental conditions, a prioritized “cascade” logic is used. In the most commonly used choice, “RIG”, the Raman channel is the first choice due to its insensitivity to surface conditions and to its immunity from false targets. If the Raman signal does not exceed a preselected threshold, the logic looks for an IR return. If neither Raman nor IR returns are detected, the logic defaults to the green channel. Although this return cannot yield an accurate surface time, due to its uncertain origin, it can, nevertheless, be used as a trigger to capture all three surface channels at an approximate location in the digitized records. It is also possible to use "IRG" logic or to lock the hardware surface detection to a single, selected channel. It is usually possible to extract accurate surface times from the digitized Raman and IR waveforms in post-flight software, even if they were too weak for hardware detection. The capability of reprocessing the surface waveforms allows a second chance at detecting returns which could have been lost if only hardware had been used and no waveforms preserved.

The real-time surface pulse detections coming from the cascade logic are used to stop a time-interval meter which was started by the laser firing. The resulting surface slant-range times, corrected for predicted surface biases, are used to align the waveforms in the stored digital record, putting the air/water interface at a consistent, preselected bin location in each channel, and to calculate the average slant range which is used to control the digitizer if all hardware surface returns are missing. The surface slant-range times are also used in the real-time depth algorithm and stored for possible use by the post-flight algorithm. The beginnings of the stored digital records are measured relative to the laser firing; this allows the slant-range times to the surface and the bottom to be determined independently, relative to the aircraft rather than relative to one another. This permits accurate depths to be calculated when the most accurate surface location is determined, even if the real-time surface is unacceptable for a particular pulse. The surface slant-range times measured in real time are, to a certain extent, pulse shape dependent, and surface pulse shapes vary somewhat as a function of environmental conditions. Consequently, it is believed that the surface times determined in software, by processing the recorded waveforms, are a bit more accurate, and this has become the standard default.

4. *SHOALS* SYSTEM SOFTWARE

4.1 Introduction

Reliable, automated post-flight depth extraction procedures are provided which maximize throughput by recognizing and handling most problems routinely, thus minimizing the amount of human interaction with the raw data. Operator involvement necessary for handling special cases is streamlined by providing ready access to waveforms and intermediate results and the capability to quickly reprocess selected areas. Prior to each flight, a Mission Planning program is run on the HP or PC to set the coordinates of each flightline and the desired mission parameters for hardware and software. A set of standard default parameters is used unless changes are needed. All system hardware is controlled through the operator interface, and all hardware and software parameters are permanently recorded in headers associated with each tape or flightline.

Post-flight processing routines, all written in the 'C' language, consist of a tape Stripper, an Automated Processor which includes the waveform processor, sounding position determination algorithms, various calibration and utility programs, and a Manual Processor operator interface routine which provides for geographic display of the calculated depths, access to individual waveforms, and editing capability. The Stripper reads the tapes, interpolates data from low-rate sensors to the 200-Hz laser pulse repetition rate, performs time deskewing and quality checks, and parses the information into an Informix relational database. The Automated Processor, which will be described in detail in this paper, obtains inputs from the database; calculates depths, positions, and other products; corrects for tides; and writes outputs to the database. It runs at a 1:1 time ratio with data collection. Output tables stored in the database are used by the Manual Processing operator interface program, and ASCII files are created for use in subsequent geographic manipulations which are conducted with a commercial package called TerraModel.

The Manual Processing facility is used for quality control of the results from Automated Processing. It provides color-coded geographic displays of outputs such as real-time and post-processed depth, second depth, depth confidence, and land elevation for selectable portions of the survey area. At small scales, high-density colored dots merge to provide contours, while at large "zoom" scales, individual numerical values are plotted. Clicking with the mouse on any value brings up a screen containing the raw waveforms with detection times annotated by tic marks, and a number of associated output parameters for that particular pulse. An example is seen in Fig. 1. From that screen, one may choose to look at flightline information and warnings, swap first and second depths, kill a sounding, or change to a neighboring shot of choice. It is also possible to click on the waveforms to obtain time and amplitude coordinates. If desired, selected areas can be reprocessed with different parameters, to adapt to local conditions. The database can also be queried using a Wingz spreadsheet. In this way, time series or cross-plots of any recorded parameters can be plotted. Also available is digitally annotated video from a down-looking camera. The visual scene at the time of any pulse can be viewed to help in the interpretation of anomalies.

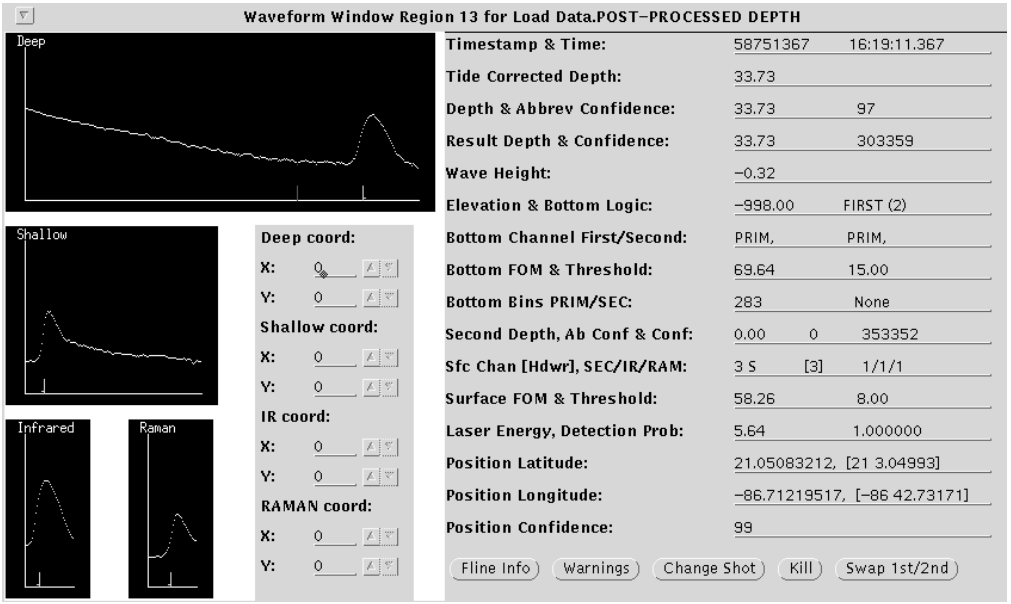

Figure 1 SHOALS Manual Processing operator interface screen.

Data processing takes place in a mobile support facility which travels from site to site with the helicopter. It is housed in an air-conditioned 40-foot tractor trailer along with space for mission planning, meetings, storage of helicopter parts, tools, tide measurement equipment, survey gear, and spare parts. The post-flight waveform processing software runs on a Sun SPARCstation System 10 with 160 MB of memory, 10 GB of disc space, and an Exabyte 8500 tape drive. Subsequent geographic data manipulations such as spatial editing, contouring, and mapping take place on a Hewlett Packard HP735 with 148 MB memory, 2 GB of disk space, a 600-MB optical drive, and 4-mm digital tape drive. The facility also includes a PC which is used for Mission Planning, general office work, and to write CD's containing the final data product.

4.2 Parameter tables

Hardware and software system parameters, collected data, and calculation products are parsed into a number of tables in the database. These various tables are called *Project Environmental Parameters (PEPS)*, *Hydrographer Accessible Parameters (HAPS)*, *hardware*, *inputs*, *outputs*, *warnings*, *results*, and *diagnostics*. Values for the first three tables are set prior to flight during Mission Planning, although several of the *PEPS* can be altered in flight. The small *PEPS* table, which may require frequent updates, contains information on environmental conditions such as water clarity, wave height, index of refraction (per salinity), and maximum depth, as well as several critical software thresholds. The *HAPS* table includes a number of software switches for selecting various possible modes of operation and parameters which may occasionally need to be updated to handle difficult cases. Access to the HAPS table is "hidden" to emphasize that changes are meant to be occasional and restricted. The *hardware* table contains a large number of system constants and mission settings such as transmitter beam divergence, receiver field of view, detector and amplifier gains and noise values, delay tables, PMT high voltage, calibration constants, sounding spacing, planned altitude and flight speed, scanner pattern, and many more. All of the data collected by the airborne system, including digitized waveforms, slant range times to the surface, digitizer values, pulse time, position, etc., is stripped from the Exabyte tapes and stored in the *inputs* table.

Parameters in four data categories can be produced by the waveform processor. The *outputs* table contains the primary products of processing such as depth, depth confidence, altitude, wave height, surface and bottom detection figures of merit, and digitizer bin numbers for detection times in each channel. The *warnings* table contains specific flags which may be set by the processor to indicate a status or that certain tests were failed and implying the possible presence of errors or undesirable properties in the data. Some of these *warnings* can cause the depth for a sounding to be deleted or the depth confidence to be reduced, while others are simply for operator attention. There are presently 28 of these. When the processor is run, there are three possible levels of output. The first produces only *outputs* and *warnings*; the second provides *results*, in addition; and the third adds *diagnostics*, as well. Since the latter two consume more resources, they are only run when necessary. *Results* are specially calculated, value-added parameters which contain additional information about system performance, such as time differences, wave height statistics, and surface tracker statistics. *Diagnostics* are the normally lost intermediate products of existing calculations which are helpful when debugging software changes and in tracking down the causes of unusual results. Because this latter set includes normalized waveforms, it consumes a lot of memory and is only run on short data segments.

4.3 Waveform processing

The function of the SHOALS post-flight "waveform analyzer" program, which resides within Automated Processing, is to analyze digitized waveforms from PMT, green APD, IR, and Raman channels and supporting hardware timing parameters to estimate for each pulse the mean water depth and its associated confidence factor, along with a variety of warnings and ancillary statistical and diagnostic outputs aid understanding of the performance of the algorithms. Even seemingly small corrections are made to prevent their sum from causing a problem. The program, which is executed once for each laser pulse, is divided into four major components; these are 1) estimation of necessary parameters, checking for saturated waveforms, surface tracking, and waveform normalization; 2) determination of surface return times; 3) detection of bottom return times; and 4) wave correction, estimation of depth (or land height), and bias correction. These are discussed in following sections.

4.3.1 Preliminary calculations

The first step is to perform calculations that will only need to be done once per flightline. This includes selection of start and stop bins for each channel to define the regions of the waveforms which will be processed and calculation of parameters for later use in depth calculations. The "shallow" green (APD) channel has a fixed *surface_stop_time_apd* to avoid picking a

bottom return as the surface, in the unlikely event that the surface return is missed, and a *bottom_start_time_apd* to avoid picking the surface return as the bottom. The latter is designed to permit direct control of the minimum depth potential with the HAP called *safety_margin_apd*. It is normally set to accept depths greater than about 1.1 meters. The "deep" green (PMT) channel has its own *bottom_start_time_pmt*, controlled by *safety_margin_pmt*, to avoid the early peak caused by gating the tube on. Both shallow and deep green channels have *bottom_stop_times* which can limit the waveform search in waters known to be shallow, based on a HAP called *max_apriori_depth*. The IR and Raman waveforms are scanned in their entirety. Among the constants calculated, based on the *HAPS* and *hardware* values in use, are the minimum acceptable pulse risetimes for each channel; the predicted surface biases based on the apriori diffuse attenuation coefficient; the relationship between time and digitizer bins; the arrays, for each channel, of conversion factors (*pe_per_count_ra*) between digitizer counts and input photoelectron count per effective bandwidth time at the photocathodes; and the hardware noise terms for the signal-to-noise ratio calculations. The values of the *pe_per_count_ra*'s are based on the detector and amplifier gains, channel impedance and bandwidth, digitizer parameters, and electronic constants. A table of selected values is created for viewing by the operator. From this point on, all calculations are performed once for each pulse.

Because the hardware land detector has not been reliable, it has been augmented by an additional software test. The IR channel is tested for saturation. Based on field experience, if more than five bins are saturated, the land flag is set. Possible conflicts between land flags and hardware Raman detections are noted. Statistics on the detections of surface returns in the three surface channels are collected for later display.

Surface slant ranges are corrected for predicted surface biases. Because hardware surface slant ranges may be used in the surface component of depth calculations, severe quality control is applied to the input data. A surface tracking algorithm is applied to the hardware surface return data to edit out spurious values which may arise due to detection of birds, spray, or "sea smoke" above the true surface or from hardware glitches. The aircraft height for each pulse is estimated by multiplying the surface slant range by the cosine of the nadir angle. The mean aircraft height and standard deviation of heights is calculated in a cyclic window of selected size (a HAP generally set to 100 pulses). A minimum number of good heights in the window is required. New heights are tested against a threshold of a selected number of standard deviations from the mean (a HAP typically set to 3). If they are good, they are used to update the height buffer. If a height flier is detected, indicating a bad slant range, either an *early_surface_flag* or a *late_surface_flag*, as appropriate, is set, along with a *no_sfc_data_flag*, and the height is rejected. Because early surface detections shift actual surface and bottom pulses to later-than-expected times in the digitizer, it would be possible for the software bottom detector to mistakenly pick the surface return. To prevent this, the *bottom_start_time_apd* is augmented by the magnitude of the timing error which caused the *early_surface_flag*. Other consistency checks are also performed which, if failed, will set a *no_sfc_data_flag*. This flag will be used later in the wave correction algorithm to exclude that slant range from use and to switch the depth calculation to a different mode.

It is not desirable to perform detection and timing on the raw waveforms which, except for the Raman, have been compressed by logarithmic amplifiers. The raw waveforms are linearized or "normalized" by converting them to units of the number of electrons per effective bandwidth time at the detector photocathodes. This is accomplished by multiplying by the raw waveform elements by the *pe_per_count_ra* elements, calculated earlier, appropriate to each raw digitizer count. The value used for the Raman channel, which is already linear, is a scalar rather than a vector. This also permits the estimation of absolute signal-to-noise ratios and associated detection probabilities and false alarm rates. Warnings are generated if inappropriate data is recognized. Moving forward in time, the measured risetimes of all potential candidate pulses in all channels are determined by using an algorithm which measures the length of "runs" of increasing digitizer counts, subject to a set of rules which add a measure of noise immunity. All channels are tested for saturation; the shallow green channel is also tested separately for saturated bottom returns. Appropriate flags are set. The surface channel to be used for depth calculations is determined based on which hardware channel was picked by the data collection software and on two HAPS which select whether hardware or software surface data is chosen, and, if software, which channel. Error-ridden green surface return times are never used for wave height or depth calculation under any circumstances.

4.3.2 Determine surface times

Normalized waveforms in the three surface channels are scanned forward in time to find the peak bins and risetimes of the first surface returns which meet minimum risetime and signal-to-noise ratio (SNR) requirements. The use of the first return is particularly important in the Raman channel where unexpectedly strong bottom returns have been detected for the case of extremely shallow water. The SNR is calculated as the signal above the background level divided by the square root of the no-signal noise variance. Signal is calculated above a level representing the mean amplitude prior to the start of the pulse.

Noise variance is the quadrature sum of shot noise, amplifier noise, and detector noise; thermal noise is negligible. Acceptance flags are set for successful detections, and the SNR becomes the figure of merit (FOM). For each channel, the times associated with the detected returns are determined by applying a backwards-looking (from the peak) half-peak-height threshold to the normalized waveforms. [Matched filter and other energy-based pulse-location algorithms were rejected because they are too sensitive to pulse shape, which can vary considerably[77], particularly the pulse tails. Peak detectors are not used because, being based on a zero-slope point, they are very noisy. Half-peak-height is the location of maximum slope which provides the highest precision.] The detection time is interpolated between the amplitudes of the bins immediately above and below the threshold. Corrections for surface bias time, calibration constants, and amplitude-dependent time delays are applied to each channel. Pulse detection times are converted from relative to absolute by summing digitizer times with the digitizer start times relative to pulse firing.

As a check on system consistency, intercomparisons are made among software surface times for the three surface channels and between the hardware-derived value and its software counterpart. Assorted timing and FOM statistics and software surface return probabilities are calculated and reported in a table available to the operator at the end of processing. *Warnings* are set for apparent problematic behavior. The software surface detectors are more sensitive than the hardware surface detectors. In the rare event that the hardware surface detection was missed in all channels, an average slant range is used to position the waveforms in the digitizer, and the location of the waveforms is uncertain but frequently acceptable. If this occurs, the location of the IR software surface return, if detected, is used to determine the corrector for the *bottom_start_bin*. If the use of software surface times rather than hardware surface times has been selected via HAPS for use in depth calculations, the hardware slant range times and heights are replaced by their software equivalents. A height is also calculated in the shallow green channel which can be used to determine the topographic height of land returns above the mean water level.

4.3.3 Determine bottom times

Green waveforms in the shallow and deep channels are analyzed between start and stop bins for potential bottom returns. The peak bins of the best two bottom return pulse candidates (primary and "second") are determined for each channel using SNR and minimum acceptable measured pulse risetime as selection criteria. Their SNR's and acceptability flags based on SNR thresholds are stored along with measured risetimes, peak locations, and noise variances. Because the bottom returns frequently ride on a steep backscatter slope, the value of the peak "signal" is estimated above a fit to the backscatter curve under the bottom return. If the slope is positive rather than negative, the fit is replaced by a constant equal to the value of the waveform at the beginning of the pulse. This situation can occur for the case of a small target sitting on the bottom. SNR is then calculated by dividing the signal by the square root of the no-signal noise variance, as for the surface returns. Tests are made for distorted returns based on the shape and duration of their leading edges. Up to two identified potential bottom returns in each of the shallow and deep green channels are time tagged by applying half-peak-height thresholds and their parameters stored for further use. The time tags are passed to Manual Processing where they are used to apply tic marks to the waveforms. In the case of extremely shallow depths where the bottom return is merged with the surface return, a modified procedure is invoked which estimates the predicted time of threshold crossing based on the peak location and known pulse characteristics. Amplitude-dependent time delay correctors and calibration constants are applied for each channel. Pulse detection times are converted from relative to absolute by summing digitizer times with the digitizer start times relative to pulse firing. It is of note that except for the (following) wave corrector, the algorithms contain no filtering or smoothing, because these processes tend to distort waveforms and can lead to bias errors.

Depending on the setting of the HAP *bottom_logic_switch*, either "strongest-pulse", "first-pulse", or "last-pulse" logic is applied to select one of the four potential bottom returns as the one to be used for the depth calculation and to store the associated parameters such as its *total_bottom_time* and *bottom_confidence*. The strongest pulse mode selects the pulse with the largest SNR, given that it exceeds the minimum SNR criterion. This is usually the gross bottom return, but it could occasionally be a fish school or marine mammal. In this mode, small objects (such as rocks and coral heads) sitting on the bottom may be missed[61]. The first-pulse mode is used to maximize the probability of finding such targets. Here, as seen in Fig. 2, the earliest of the four potential bottom pulses which meet the required minimum SNR criterion becomes the selected bottom return. The problem with this mode is that fish detections will be selected even if the underlying bottom return is strong. The last-pulse mode uses the latest of the four potential bottom returns, which meets the minimum SNR requirement, for the depth calculation. This mode is meant to punch through fish or other biota returns to detect the underlying bottom. It is probably not a good mode to use for first processing, but it can be used to reprocess areas or pulses containing such problems. Regardless of the mode, a "second" bottom return with all its parameters is also retained in the *results* table to be presented to the operator. It is always selected as the return with the highest SNR of those not selected to be the primary return, regardless

of whether it meets the minimum SNR criterion. In Manual Processing, this second depth could eventually be switched to become the reported depth, without any need for further processing, if it meets the minimum SNR requirement. Tests ensure that the primary and secondary depths do not originate from the same physical target in the depth overlap region common to the shallow and deep channels. Other tests are made for consistency between channels and for distorted returns, various *warnings* are reported as required, and flags may be set which will reduce the depth confidence.

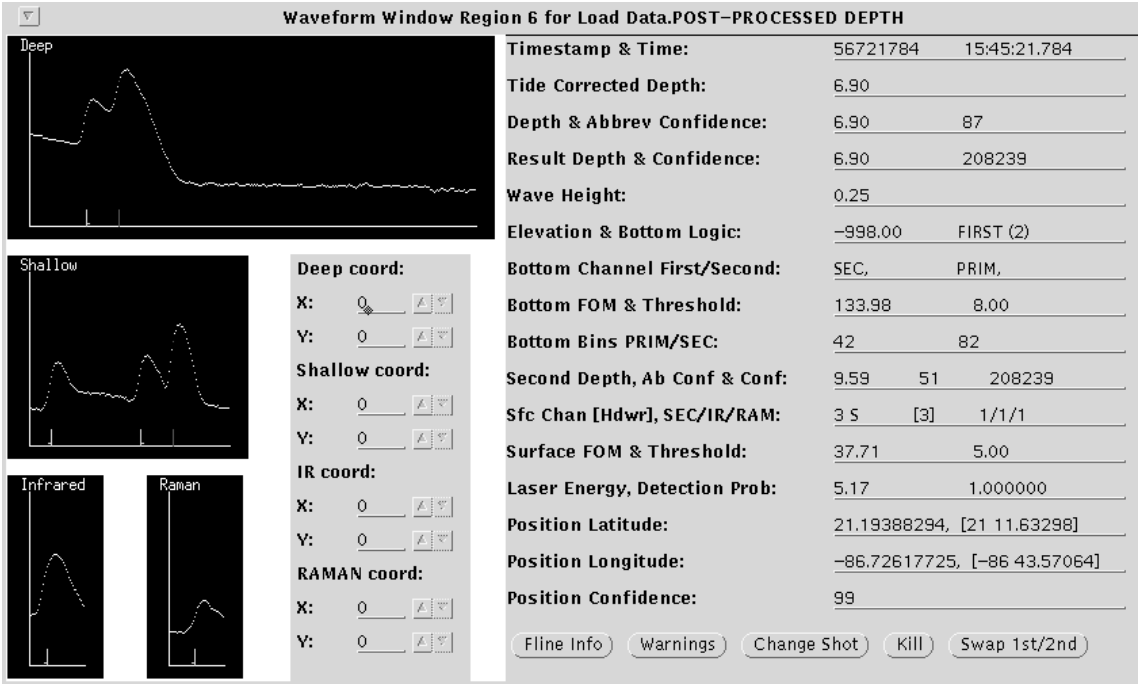

Figure 2. Demonstration of the automatic detection of a shallow target in first-pulse mode.

4.3.4 Estimate depth

The surface returns have all been processed and corrected to yield the slant range from the aircraft to the air/water interface at the location of each sounding. Because, these slant ranges include the effects of wave heights, a wave-correction algorithm must be applied to provide a mean water level from which the water depths will be measured. This is by far the most complex algorithm in the processor[69]. The goal is to remove the effects of surface waves from the measured depths while permitting vertical aircraft motion to be distinguished from the effects of long-period swell on the water surface.

In order for the wave corrector for a given pulse to be based on a mean water level, the wave height must be calculated from a window containing both earlier and later surface slant range data. While all the preceding code runs on a "one pulse in at time t', one result out at the same time" basis, that cannot be the case here. A great deal of data must be buffered in "cyclic" buffers so that wave heights can be calculated and utilized at some delay after the time t' at which the data is input from processing by earlier routines. A HAP time window sets the size of the buffers; ten seconds of data (2000 pulses) is most commonly used. The key parameter from which wave heights are determined for each pulse is the "corrected local height" (CLH) which is the height of the aircraft above the detected air/water interface minus the double integral over time of the vertical aircraft acceleration. For the steady state solution, the wave heights are estimated recursively at a special time within the window by employing a least-squares fit to the CLH data. The algorithm consists of a set of "states" which accomplish certain purposes; these are: 1) filling buffers; 2) performing a fit to the input CLH's and supplying wave heights and mean altitudes for times up to the steady state point; 3) steady-state processing (one in, one out); and 4) closing out to the end of the window, with a least-squares fit, when input data are exhausted or when an error condition requires reinitialization. Statistics are collected on the wave conditions and aircraft dynamics. Solutions are designed to be valid for aircraft motions up to and including constant acceleration. The effects of rapid changes in acceleration (formally termed "jerk"), which could invalidate the solution, are tested for, and a *warning* (*jerk_flag*) can be set which kills depths for the associated pulses. A test is also conducted to ensure

that rapid changes do not take place in any acceleration bias if one is present. In order to test for a possible "surfing" condition, wave height statistics are also accumulated within a number of longitudinal "stripes" into which the swath is divided. The stripe waveheight means are examined with a Student-t test to determine if they are significantly different from zero. If any of them is, a *surfing_flag* is set, which will cause the *depth_confidence* to be reduced.

The aircraft altitude to the mean water surface level and the wave height are estimated for each pulse for use in the calculation of the mean water depth for that pulse. Provision is made for predicting the mean altitude for a limited number of pulses with missing surface returns, based on the fit to the good data in the window. Given the selected bottom pulse return time, depths are calculated in one of three possible ways, depending on the quality of available surface data both for a given pulse and for that in the entire window. Four "modes" are defined; these are: 0) good fit for local altitude and good surface return for that pulse (and hence a good wave height estimate for that pulse); 1) good altitude but missing surface return; 2) bad fit for altitude but good local surface; and 3) neither good fit nor local surface. The depth calculation uses a speed of light appropriate for the salinity of the water. The mode-0 depth is the desired depth to the mean water surface, and the depth error is one-quarter of the error in the wave height estimate. In mode 1, the depth estimated is also the desired depth to the mean water surface, but the error is one quarter of the physical wave height because time to the surface was estimated as the time to the mean surface, and the fraction of the total path to the bottom in each medium is thus in error by the wave height. For mode 2, a mean does not exist, and the depth estimated is the local water column depth. The mode-2 error is the physical wave height. In mode 3, no depth can be calculated. The second depth is calculated in like manner from the second bottom return. Surface heights over land calculated in the shallow green channel are converted here into topographic heights above the mean water level.

Following the depth calculation, the *depth_confidence* is determined for the primary and second depths. Points are subtracted from an initial basis of 99, for predicted wave height estimation errors and various warning flags, at a rate of roughly one point per centimeter of expected depth error. When this value drops below 50, the depth for that pulse is killed. Pulses with SNR's which do not meet the minimum SNR requirement begin with a basis of 49. The final step is the application of three types of small depth measurement biases. These are 1) differences between the shallow and deep green channels due to differing fields of view and alignment (obtained from a one-time interchannel calibration as a function of scanner azimuth); 2) from one-time field-test calibration based on intercomparison with sonar surveys; and 3) for predicted propagation-induced pulse stretching[70,71] as a function of depth and nadir angle. Part of the second item involves a predictable surface geometry bias[58]; there are plans to handle this as a separate, modeled component in the near future.

A modification presently being undertaken is implementation of the capability to use kinematic, differential GPS carrier-phase tracking[78] with on-the-fly lane ambiguity resolution[79]. This will permit the altitude of the aircraft with respect to the ellipsoid to be determined to within several centimeters while the vertical distance from the aircraft to the sea bottom is being simultaneously measured by SHOALS. This will permit depths to be determined without wave correction. Additionally, if the tidal datum at a given location is also measured with respect to the ellipsoid using GPS, the water depth with respect to that tidal datum can be determined without the need of measuring the tidal heights at the time of the hydrographic survey. It is important to note, however, that even if GPS techniques are adopted for obtaining depths by measuring the distance to the bottom relative to the ellipsoid, the distance to the water surface must still be accurately measured, because the light pulse travels a fraction of the distance to the bottom in air and the remainder in water for which the speed of light is significantly slower. For either method, given a detectable surface return, the resulting depth error is roughly one-quarter of the error in the surface height measurement plus three-quarters of the error in the altitude estimate[69]. Without a water surface time, the depth error for a given pulse would be at least one-quarter of the local wave height relative to the mean water level. For higher wave conditions, this would not be acceptable. The benefit of using the GPS approach is that the wave correction procedure, which is complex and computationally intensive, would not be necessary. The wave corrector approach, however, with its calculation of a mean water level, does have an advantage of permitting depth estimation, with the above depth error, for a limited number of pulses whose surface returns are missing or from an above-the-surface target.

5. RESULTS

An intercomparison was conducted[6] between a SHOALS survey at the approaches to Tampa Bay (Florida) and vertical-beam sonar data collected in an operational survey by the NOAA Ship MT MITCHELL. Depth differences were determined between 30,000 sonar depths and a Triangulated Irregular Network (TIN) surface created from over 5.5-million SHOALS depths. The standard deviation of the depth differences was 0.20 m. Assuming equal error contributions from both systems,

the random depth error attributed to each system is then 0.14 m. This meets both current and proposed future IHO requirements. The mean difference was 0.07 m. This bias has been used to determine a new, permanent calibration constant for the system, and hence it will not appear in any future data.

6. CONCLUSIONS

Of the many challenges involved in the design of a successful airborne laser hydrography system, the two which are most pervasive are the accurate and reliable determination of the location of the air/water interface, and the handling of the extremely large amplitude dynamic range between strong surface returns and weak bottom returns. Because many of the solutions to the latter problem are nonlinear and entail amplitude- and gain-dependent time delays, great care is needed in both hardware and software design to mitigate resulting timing errors. The solution to the surface problem is the use of results from several independent channels at different wavelengths in a logical framework. Although an eventual goal is to perform the final calculation of depths in the air during data collection, that is not currently possible both because of the need for greater computer resources and for greater experience in the interpretation and handling of unusual waveforms. Because of the high sounding rate, it is important that the post-flight data processing software be able to automatically deal with many difficult circumstances without operator intervention. On the other hand, because not all cases can be handled automatically, it is also important to have a flexible operator interface which permits efficient interaction with the raw data.

Many of the components of an airborne laser hydrography system are still considered state-of-the-art and are not available "off the shelf". The technology is being pushed to the envelope in areas such as laser characteristics, digitizer capabilities, detector specifications, optical filter parameters, and computer speed. Given the parameters of typical components in use today and the physical laws which govern the complex nature of the interaction of a light beam with a natural body of water, it is difficult to build an ALH system which can meet International Hydrographic Office accuracy standards, even for flat bottoms. Great attention to detail is required. A large number of potential error sources, arising from hardware, software, and the environment, must be recognized, modeled, quantified, and eliminated or corrected. From the first day of design, the SHOALS team has been dedicated to producing a system with the utmost accuracy potential. Design decisions were always resolved in favor of accuracy, not cost or simplicity. The result is a flexible system which, we believe, has achieved nearly the limiting accuracy dictated by uncontrollable and unknowable environmental parameters.

7. ACKNOWLEDGMENTS

The authors would like to express gratitude for funding by Headquarters, U.S. Army Corps of Engineers, Operations, Construction, and Readiness Division and by the Canadian Department of Industry, Science, and Technology.

8. REFERENCES

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

2. W.J. Lillycrop, L.E. Parson, L.L. Estep, P.E. LaRocque, G.C. Guenther, M.D. Reed, and C.L. Truitt, "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, 1994.

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

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

5. J.L. Irish, J.E. Thomas, L.E. Parson, and W.J. Lillycrop, "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, CA, (in press).

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

7. G.D. Hickman and J.E. Hogg, "Application of an airborne pulsed laser for near-shore bathymetric measurements," Remote Sens. of Env., 1, Elsevier, New York, 47-58, 1969.

8. L.M. Ott, H. Krumboltz, and A.K. Witt, "Detection of submerged submarine by an optical ranging and detection system and detection of pulses by a submarine," Proc. 8th U.S. Navy Symp. of Military Oceanography (Vol. II), May 18-20, 1971, Naval Postgraduate School, Monterey, CA (CONFIDENTIAL), 1971.

9. L.L. Cunningham, "Test report on Pulsed Light Airborne Depth Sounder (PLADS)," Naval Oceanographic Office Tech. Note 6620-102-72, U.S. Navy, 53 p., 1972.

10. M. Rankin, "Naval Air Development Center program," Laser hydrography user requirements workshop minutes, January 22-23, 1975, Rockville, MD, L.R. Goodman, (ed.), National Aeronautics and Space Administration, Wallops Island, VA, 49-74, 1975.

11. A.K. Witt, J.G. Shannon, M.B. Rankin, and L.A. Fuchs, "Air/underwater laser radar test results, analysis, and performance predictions," Report No. NADC-76005-20, Naval Air Development Center, Warminster, PA, 293 p. (CONFIDENTIAL), 1976.

12. H.H. Kim, P. Cervenka, and C. Lankford, "Development of an airborne laser bathymeter," NASA Tech. Note TND-8079, National Aeronautics and Space Administration, Washington, D.C., 39 p., 1975.

13. M. Bristow, "CCRS program," Laser hydrography user requirements workshop minutes, January 22-23, 1975, Rockville, MD, L.R. Goodman, (ed.), National Aeronautics and Space Administration, Wallops Island, VA, 25-34, 1975.

14. R.A. O'Neil, V. Thomson, J.N. de Villiers, and J.R. Gibson, "The aerial hydrography program at CCRS," Proc. Coastal Mapping Symp., Aug. 14-16, 1978, Rockville, MD, 125-132, 1978.

15. R.H. Abbot and M.F. Penny, "Air trials of an experimental laser bathymeter," Tech. Note WRE-TN-1509, Weapons Research Establishment, Dept. of Defence (Australia), Salisbury, South Australia, 39 p., 1975.

16. A.P. Ivanov, A.L. Skrelin, and I.D. Sherbaf, "Study of optical characteristics of water media using pulsed sounding," ZhPS, 17, 2, 232-240, 1972.

17. A.I. Carswell and S. Sizgoric, "Underwater probing with laser radar," Proc. of The Uses of Lasers for Hydrographic Studies, Sept. 12, 1973, Wallops Station, National Aeronautics and Space Administration, Wallops Island, VA, 123-140, 1974.

18. L.R. Goodman (ed.), Laser hydrography user requirements workshop minutes, January 22-23, 1975, Rockville, MD, National Aeronautics and Space Administration, Wallops Island, VA, 143 p., 1975.

19. L.R. Goodman (ed.), Laser hydrography technical review workshop minutes, August 25-26, 1976, Rockville, MD, National Oceanic and Atmospheric Administration, Rockville, MD, 127 p., 1976.

20. G.C. Guenther and L.R. Goodman, "Laser applications for near-shore nautical charting," Proc SPIE Ocean Optics V, Vol. 160, 174-183, 1978.

21. G.C. Guenther, L.R. Goodman, F. Hoge, R.N. Swift, R.W.L. Thomas, and D. Bright, "AOL project: results to date," Proc. Airborne Laser Hydrography Symp. III, October 5-6, 1977, Rockville, MD, National Oceanic and Atmospheric Administration, Rockville, MD, 62-103, 1979.

22. F.E. Hoge, R.N. Swift, and E.G. Frederick, "Water depth measurement using an airborne pulsed neon laser system," Appl. Opt., 19, 6, 871-883, 1980.

23. G.C. Guenther and R.W.L. Thomas, "Accuracy and penetration measurements from hydrographic trials of the AOL system," Proc. 4th Laser Hydro. Symp., ERL-0193-SD, Sept. 30 - Oct. 3, 1980, Australian Defence Research Centre, Salisbury, South Australia, 108-150, 1981.

24. R.A. O'Neil, "Field trials of a lidar bathymeter in the Magdalen islands," Proc. 4th Laser Hydro. Symp., ERL-0193-SD, Sept. 30 - Oct. 3, 1980, Australian Defence Research Centre, Salisbury, South Australia, 56-84, 1981.

25. M.F. Penny, "Laser hydrography in Australia," Proc. Int'l. Conf. on Lasers '81, Dec. 14-18, 1981, 1029-1042, 1981.

26. R.H. Abbot, "WRELADS II trials," Proc. 4th Laser Hydro. Symp., ERL-0193-SD, Sept. 30 - Oct. 3, 1980, Australian Defence Research Centre, Salisbury, South Australia, 188-215, 1981.

27. V.N. Balandin and R.D. Volodarskiy, "Laser instruments for measuring the depth of shallow water," Geodeziya i kartografiya,2, 58-61, 1979.

28. O. Steinvall, H. Klevebrant, J. Lexander, and A. Widen, "Laser depth sounding in the Baltic Sea," Appl. Opt. 20, 19, Oct. 1, 3284-3286, 1981.

29. M.M. Harris, G.D. Hickman, and R. Booker, "Development of the Airborne Bathymetric Survey system," Proc. Hydro USA '86, March 25-27, 1986, Norfolk, VA, National Oceanic and Atmospheric Administration, Rockville, MD, 50-55, 1986.

30. T. Curran, T. Keck, V.M. Contarino, M.M. Harris, and S.P. Haimbach, "Digital ABS Laser Sounder bathymetry," Proc. SPIE Ocean Optics IX, Vol. 925, 242-249, 1988.

31. R. Axelsson, O. Steinvall, and P. Sundberg, "Programmable scanner for laser bathymetry," Int'l. Hydro. Rev., 67, 1, 161-170, Jan. 1990.

32. O. Steinvall, K. Koppari, and U. Karlsson, "Experimental evaluation of an airborne depth sounding lidar," Proc. SPIE Lidar for Remote Sensing, Vol. 1714, 108-126, 1992.

33. N. Anderson, P. Bellemare, M. Casey, K. Malone, R. MacDougall, D. Monahan, R. O'Neil, and S. Till, "Beginning the second hundred years -- the laser sounder," Proc. Centennial Canadian Hydro. Conf., Spec. Pub 67, Fisheries and Aquatic Services, Ottawa, Ont., April 1983.

34. J. Banic and S. Sizgoric, "Scanning lidar bathymeter for water depth measurement," Proc. SPIE Laser Radar Tech. and Appl., Vol. 663, 187-195, 1986.

35. M.J. Casey, "Deploying the lidar on hydrographic surveys," Proc. 9th Canadian Symp. on Rem. Sens., Memorial University, 165-175, 1984.

36. M.J. Casey, R.A. O'Neil, and P. Conrad, "The advent of Larsen," Proc. Canadian Hydro. Conf., Halifax, N.S., 7-12, 1985.

37. M.J. Casey and J. Vosburgh, "Chartmaking with Larsen," Canadian Surveyor, 40, 3, 1986.

38. P. Conrad, "Reaping the harvest -- the processing of Larsen data," Internal report, Canadian Hydrographic Service, Sidney, B.C., 25 p., 1986.

39. J. Vosburgh and J. Banic, "Airborne laser surveys of the Northwest Passage," Proc XIIIth Int'l. Hydro. Conf., Monaco, 1987.

40. M.F. Penny, R.H. Abbot, D.M. Phillips, B. Billard, D. Rees, D.W. Faulkner, D.G. Cartwright, B. Woodcock, G.J. Perry, P.J. Wilsen, T.R. Adams, and J. Richards, "Airborne laser hydrography in Australia," Appl. Opt. 25, 13, 1 July, 2046-2058, 1986.

41. M.F. Penny, B. Billard, and R.H. Abbot, "LADS -- the Australian Laser Airborne Depth Sounder," Int'l. J. Rem. Sens., 10, 9, 1463-1479, Sept. 1989.

42. J.S. Compton and M.A. Hudson, "New charting technology in Australia: the Laser Airborne Depth Sounder," Int'l. Hydro. Rev. LXV(2), Monaco, 145-157, 1988.

43. A.F. Bunkin, D.V. Vlasov, A.S. Galumyan, D.V. Mal'tsev, D.M. Mirkamilov, and V.P. Slobodyanin, "Versatile airborne laser system for remote probing of ocean, atomosphere, and farmland," Sov. Phys. -- Tech. Phys., 29, 11, 1284-1287, 1984.

44. A.G. Abroskin, A.F. Bunkin, D.V. Vlasov, A.L. Gorbunov, and D.M. Mirkamilov, "Full-scale experiments with laser aerial sounding at the Chayka facility," Works of the General Physics Institute: Remote sensing of the ocean (Trudy IOFAN: Distantsionnoye zondirovaniye okeana), Vol. 1, Moscow: Nauka, 29-47, 1986.

45. A.I. Abramochkin, V.V. Zanin, I.E. Penner, A.A. Tikhomirov, and V.S. Shamanaev, "Airborne polarization lidars for atmospheric and hydrospheric studies," Optika atmosfery, v.1, 2, 92-96, (in Russian), 1988.

46. E.A. Tsvetkov, "Lidar related shipboard and aircraft measurements," Technical seminar on issues in lidar and ocean., Oct. 1-2, 1991, Naval Ocean Systems Center, San Diego, CA, 1991.

47. V.J. Feigels and Yu.I. Kopilevich, "Russian airborne lidar for oceanography," Proc. Symp. on Russian airborne geophysics & remote sensing, SPIE Vol. 2111, Sept. 13-17, 1992, Golden, CO, 127-141, 1993.

48. Z.S. Liu, "Estimate of maximum penetration depth of lidar in coastal water of the China sea," Proc. SPIE Ocean Optics X, Vol. 1302, 655-661, 1990.

49. J. Pope and W.J. Lillycrop, "Development of a helicopter lidar bathymeter system," Proc. U.S. Army Corps of Engineers Surveying Conf., Fort Belvoir, VA, 213-216, 1988.

50. J. Banic, S. Sizgoric, and W.J. Lillycrop, "Second-generation airborne lidar system for hydrographic applications," Proc. Oceanology Int'l., Brighton, England, March 1990.

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

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

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

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

55. K. Koppari, U. Karlsson, and O. Steinvall, "Airborne laser depth sounding in Sweden," Int'l. Hydro. Rev. LXXI(2), September, Monaco, 69-90, 1994.

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

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

58. G.C. Guenther, "Airborne Laser Hydrography: System Design and Performance Factors," NOAA Professional Paper Series, National Ocean Service 1, National Oceanic and Atmospheric Administration, Rockville, MD, 385 p., 1985.

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

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

61. G.C. Guenther, T.J. Eisler, J.L. Riley, and S.W. Perez, "Obstruction detection and data decimation for airborne laser hydrography," Proc. 1996 Canadian Hydro. Conf., June 3-5, 1996, Halifax, N.S., (in press).

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

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

64. G.C. Guenther and H.C. Mesick, "Analysis of airborne laser hydrography waveforms," Proc. SPIE Ocean Optics IX, Vol. 925, 232-241, 1988.

65. J.E. Tyler and R.W. Preisendorfer, The Sea, M.N. Hill, Ed., Wiley-Interscience (New York), 1962.

66. G.E. Walrafen, "Raman spectral studies of the effects of temperature on water structure, J. Chem. Phys., 47, 1, 114-126, 1967.

67. D.A. Leonard, B. Caputo, and F.E. Hoge, "Remote sensing of subsurface water temperature by Raman scattering," Appl. Opt., 18, 11, 1732-1745, 1979.

68 B. Billard, "Estimation of a mean sea surface reference in the WRELADS airborne depth sounder," Appl. Opt. 25, 13, 2067-2073, July 1986.

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

70. G.C. Guenther and R.W.L. Thomas, "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., 1984.

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

72. B. Billard, R.H. Abbot, and M.F. Penny, "Modeling depth bias in an airborne laser hydrographic system," Appl. Opt. 25, 13, 2089-2098, July 1986.

73. G.C. Guenther and R.W.L. Thomas, "Error analysis of pulse location estimates for simulated bathymetric lidar returns," NOAA Tech. Rpt. OTES 01, National Oceanic and Atmospheric Administration, Washington, D.C., 51 p., 1981.

74. S.P. Haimbach, H.C. Mesick, H.J. Byrnes, and G.D. Hickman, "Optical bathymetry for the U.S. Navy: a field measurement program," Proc. SPIE Ocean Optics IX, Vol. 925, 214-221, 1988.

75. G.C. Guenther and H.C. Mesick, "Automated lidar waveform processing," Proc. U.S. Hydro. Conf. '88, Spec. Pub. 21, April 12-15, 1988, Baltimore, MD, 52-59, 1988.

76. K.J. Petri, "Laser radar reflectance of Chesapeake Bay waters as a function of wind speed," IEEE Trans. Geoscience Electronics GE-15, 2, 87-96, 1977.

77. B. Billard and P.J. Wilsen, "Sea surface and depth detection in the WRELADS airborne depth sounder," Appl. Opt. 25, 13, 2059-2066, July 1986.

78. B.W. Remondi, "Kinematic GPS results without static initialization," NOAA Technical Memorandum NOS NGS-55, National Geodetic Information Center, Rockville, MD, 1991.

79. S.L. Frodge, R.A. Barker, D. Lapucha, B.W. Remondi, and B. Shannon, "Results of real-time testing of GPS carrier phase ambiguity resolution on-the-fly," ION GPS93, Inst. Nav., Salt Lake City, Nevada, 1993.