Automated Optimal Processing of Phase Differencing Side-scan Sonar Data Using the Most-Probable Angle Algorithm

Val E. Schmidt and Tom C. Weber
Center for Coastal and Ocean Mapping
University of New Hampshire
Durham, NH

Arthur C. Trembanis
Coastal Sediments, Hydrodynamics and Engineering Lab
University of Delaware
Newark, DE

Abstract*— Phase-differencing side-scan sonar systems produce co-located bathymetry in addition to each side-scan amplitude measurement. Bathymetric soundings are calculated from the range to each measurement (derived from the two-way travel time) and the receive angle of the incoming signal. Because phase-differencing systems produce a seafloor sounding with each individual measurement, they are often characterized as noisy when compared to multi-beam sonar systems, whose seafloor estimates, whether by amplitude-weighted mean or sub-aperture phase difference detection, are the product of averaging several measurements. In addition, every effort is made to increase the resolution of side-scan data by increasing the bandwidth and sampling rate of the transmitted signal, often producing more than 10,000 data points per ping. This volume of outlier-prone, relatively noisy data is difficult for operators to interpret and software to process. A series of methods has been developed for the automated processing of phase-differencing side-scan sonar data producing seafloor estimates and related uncertainties optimized for the survey application. The "Most-Probable Angle Algorithm" (MPAA) has been developed for the filtering of outliers in range-angle measurements. With outliers removed, the uncertainty of the filtered measurements are estimated. Angle estimates are then calculated as an uncertainty-weighted mean where the number of measurements contributing to each estimate is determined from that required to achieve a desired depth uncertainty. The resulting swath of depth measurements contains irregularly spaced soundings, typically obtaining full spatial resolution of the side-scan data from 20-50 degrees from nadir, and combining several measurements to reduce the uncertainty elsewhere. In this way, given a survey requirement, an optimal amount of information can be extracted from the sonar data in varying conditions. *(Abstract)

***Index Terms*—sonar, side-scan, phase differencing, direction of arrival, estimation.** *(key words)*

I. Introduction

Phase-differencing side-scan sonar systems produce co-located bathymetry in addition to each side-scan amplitude measurement. These bathymetric sonar systems have been popular of late due to high across-track resolution, a promise of increased swath width and the lack of a large across-track array which gives them a convenient form-factor for tow bodies and autonomous underwater vehicles (AUVs).

Bathymetric soundings are calculated from the two-way travel time to each measurement and the receive angle of the incoming signal. Bathymetric side-scan sonars measure the receive angle with either an array of receive staves whose spacing meets the half-wavelength Nyquist sampling requirement to prevent spatial aliasing at the operating frequency of the sonar (often processed with variants of the CAATI method [1]), or alternatively, with pairs of staves whose spacing does not meet the Nyquist criterion but having spacings of non-integer wavelength multiples allowing unique determination of the receive angle from their combined measurements (the so-called Vernier Method. [2]).

Whether Vernier or CAATI type systems, phase-differencing bathymetric side-scan sonars can produce large amounts of seemingly noisy, outlier-prone data when compared with multibeam sonar systems. To understand why this is so, and why this comparison is misplaced, consider the methods by which the two systems derive seafloor measurements. Phase-differencing systems produce a seafloor sounding with each receive angle measurement that results from a return of the acoustic signal from the seafloor. These measurements may be noisy and, in general, to maintain the highest across-track resolution possible the measurements are not combined except in gridding algorithms in later stages of processing. In contrast, multibeam sonars produce seafloor estimates by amplitude-weighted mean for soundings near nadir or sub-aperture phase difference detection elsewhere. [3] In either case, methods used by multibeam sonar systems for bottom detection greatly reduce the noise in their reported soundings by the averaging of many individual measurements. Although the fundamental measurements made by phase-differencing side-scan systems and multibeam sonars are the same for (sub-aperture phase detections), due to the averaging processes the soundings reported by multi-beam systems are actually seafloor *estimates* that result from many individual measurements and therefore contain noticeably less noise.

Further complicating processing of phase-differencing side-scan systems is the fact that geometries that favor high-quality side-scan do not favor good bathymetry. Low operating altitudes above the seafloor produce long acoustic shadows of objects proud to the seafloor. It is these shadows that make objects recognizable in side-scan images. However, acoustic shadows produce no bathymetric data, measuring the water column or in shallow water, sometimes the water's surface instead. These measurements produce copious outliers with respect to the seafloor that can be difficult and time consuming to remove.

In addition, phase-differencing side-scan sonars typically produce far more data per ping than multibeam systems. To produce high-resolution sides-can data, every effort is made to increase the bandwidth and the sampling frequency of a system. Increasing the bandwidth of the system, either by transmitting shorter pulses or by transmitting chirp or other encoded pulses, reduces the size of the ensonified footprint on the seafloor and hence increases the system's ability to distinguish adjacent objects. With a smaller ensonified area, the sample rate may be increased while maintaining statistically independent adjacent measurements. With this increase in sample rate comes an increase in the volume of data. Systems may record more than 10,000 measurements per ping.

For multi-beam operators accustomed to relatively noise and outlier free data having less than 1000 soundings per ping, processing phase-differencing sonar data can be quite challenging. The ratio of noise and outliers to actual useful measurements may be as high as 50% overwhelming any processing routine accustomed to far less. Moreover, the volume of data alone can be more than both physical hardware and common processing suites can handle. Indeed some software tools do not support the number of soundings reported per ping by phase-differencing systems and must decimate the data before ingesting it. To be sure, many systems allow for some filtering of data as it is collected to reduce this processing burden later, but systems operated from AUVs cannot do so having no operator on board leaving a laborious task when the survey is complete.

When multibeam systems began to revolutionize the art of bathymetric map making from lead-line surveys and single-beam echosounders a generation before, a similar challenge was presented. The volume of data increased exponentially and individual soundings could no longer be scrutinized. Rather statistical methods were devised to guide surveyors to measurements that were not self-consistent and warranted further investigation. The Combined Uncertainty Bathymetric Estimator (CUBE) [4] is perhaps the most successful of these methods, combining measurements based on their predicted uncertainty and producing multiple estimates of the seafloor location when measurements are made that are not statistically consistent with the bulk of measurements at a given location. For CUBE to be successful at estimating the seafloor and reducing operator workload, a survey must have a propensity of statistically independent seafloor measurements over outliers and noise. Because every outlier becomes a hypothesis to investigate, CUBE alone is insufficient to process raw phase-differencing sidescan sonar data.

In an effort to overcome the challenges produced by the great quantities of relatively noisy, outlier prone data produced by phase-differencing side-scan sonar systems an automated method of filtering has been developed. The process involves several steps, including filtering the data for outliers, empirical estimation of the uncertainty of the remaining measurements, combining measurements in an optimal way to reduce noise and to meet a user specified uncertainty estimation and finally the use of CUBE or other suitable gridding algorithm to create a seafloor surface. Section II describes the filtering method and the estimation of uncertainty. Section III describes the method of optimally combining measurements and Section IV discusses the results.

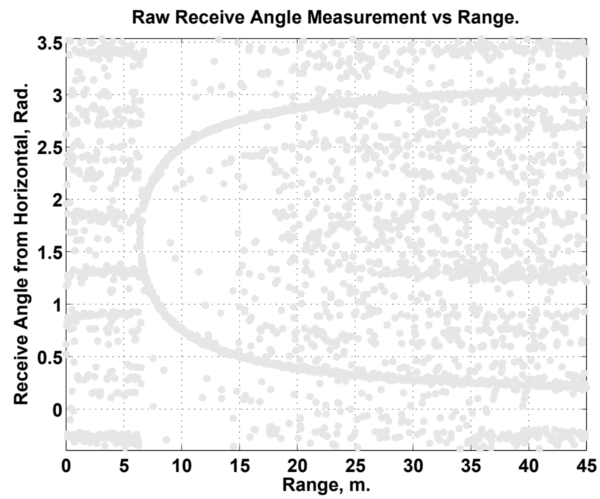

Fig. 1. Typical raw receive angle measurement vs. range from a phase differencing side-scan sonar. The sonar is 6 m above the bottom and the seafloor is apparent as a hyperbolic curve of points extending to distant ranges to either side.

II. Filtering of Data

Figure 1 illustrates a typical ping of unfiltered, receive-angle vs slant range data from a Vernier-type phase-differencing sidescan sonar system. While the seafloor is clearly evident in this plot as a hyperbolic curve of dense data points, the quantity of outliers and noise is significant, even for this relatively flat seafloor.

To gain some insight into the filtering methodology proposed here, consider the empirical equation (Equation 1) proposed by Lurton [5], which gives the predicted standard deviation in phase-difference measurements from a pair of signals in Gaussian noise as a function of signal to noise ratio, d.

$$\delta\Delta\Phi = \sqrt{2}\left[\frac{6}{\pi^2} + \left(\frac{1}{d} + \frac{1}{d^2}\right)\right]^{-1/2} \quad (1)$$

Equation 1 extends previous estimates of phase difference uncertainty for high SNR, to that for low SNR, by considering that the phase difference measurement between two sinusoids is ambiguous when the measurement is spatially aliased, that is, when the measuring devices have a physical separation greater than a half wavelength. Without additional knowledge the measured phase differences cannot be larger than +/-π, even though the actual phase difference may be several integer wavelengths larger. Therefore as signal to noise decreases, Equation 1 asymptotes to $\pi/\sqrt{3}$, which is the standard deviation of a distribution of measurements uniformly distributed over the possible measurement range of - π to π.

This addition is profound, in that it provides a fundamental limit to the knowledge that can be gained from phase-differencing measurements as SNR decreases. In general, the variance of a measurement is considered a metric indicating the amount of information contained in that measurement. However in the case of phase difference measurements at low SNR, as the SNR and hence true information of the signal decreases, the system becomes, in a sense, saturated as the ambiguous nature of the measurement does not cause the variance to increase proportionately. The variance, therefore, is no longer a measure of the information in the measurement in the traditional sense, and measurements approaching this threshold can no longer be combined in uncertainty weighted gridding algorithms like CUBE in a statistically meaningful way.

Therefore, the first step in filtering the data is to apply an SNR filter to remove measurements whose SNR is sufficiently low that the variance is no longer indicative of the uncertainty. While many systems implement an amplitude threshold in an attempt to achieve a similar result, such a filter produces poor results in systems that apply a time-varying gain (TVG). In these systems, an appropriate amplitude threshold at close range fails to filter noise at distant ranges, as the noise floor of the system is invariably amplified above the filter threshold by the TVG. An SNR filter is different in that it compensates for time-varying gains applied to low-amplitude values by measuring the amplified noise floor of the system as a function of time after transmit (range) and compensating the amplitudes of each measurement for it. To measure the noise floor of a system one may place the system in receive-only mode, for systems that provide it, measuring several hundred ping-cycles of data for the estimate. Alternatively, one may operate the system in deep water such that nothing but the water column is measured for a similar amount of time. When amplitudes are reported in linear units, the noise floor estimate as a function of range can be calculated from Equation 2.

$$Nf(R) = \frac{\sum_{i}^{N} A(R)^2}{N} \qquad (2)$$

where A is the amplitude value at range, R, and N is the number of measurements at that range. Although the noise floor of a system is, to some extent, environmentally dependent, in practice, one may usually make the measurement once and store the noise floor vs. range pairs to be used for subsequent data collected from the system (barring major modifications). In applying the SNR filter, noise floor values are interpolated to each amplitude measurement with range, *A(R)*. Each measurement's SNR is then calculated from Equation 3.

$$SNR = \frac{A(R)^2}{Nf(R)} \qquad (3)$$

Curves of phase angle variance vs. SNR [5] indicate that SNR threshold values from 10-20dB are appropriate to ensure the measurement is not variance-saturated and therefore an estimate of the variance of the measurement is indicative of its uncertainty.

After SNR filtering the next step is to segment the data to gain some statistical power from which to decide which measurements to retain and which to reject. Port and Starboard sides are considered separately. Data is segmented into equal sized horizontal bins (assuming a flat seafloor) in the across-track direction. This method gives fewer measurements per bin near nadir where SNR is high, while giving more measurements per bin at distant ranges where SNR is low and greater statistical power is desired. Moreover near nadir the measurement is less statistically stationary with time, changing quickly as the signal propagates across the seafloor at close range. Smaller bins improve on the stationarity assumption there. For shallow water systems 1 m bins are chosen due to consideration for the common hydrographic survey requirement for detection of a 1 m object.

Next, an empirical probability density distribution (a histogram) of the data in each range bin is calculated as a function of receive angle. The receive angle bin size is set to 1 degree for this calculation and the bin having the largest number of measurements is identified as the *most probable angle*.

When measurements are Gaussian distributed the most probable angle is coincident with the mean, and such a calculation would provide a far simpler method for finding the most probable angle. However phase differencing data is not commonly Gaussian distributed at this point. Correlated noise in the system, returns from the surface and other confounding factors may result in a bimodal (or more) distribution of measurements, neither of whose central tendency or median coincide with the propensity of data. Choosing the peak of the probability distribution ensures that the most probable receive angle for that bin is chosen and more accurately reflects the preponderance of data, whatever the source.

In the next step, the most probable receive angles are de-spiked in both the along and across track directions to remove outliers that result during long acoustic shadows. To de-spike the data, a median filter is applied separately in each direction. For a median filter, objects one-half the size of the filter window or larger tend to be retained while objects smaller tend to be rejected. Therefore window sizes are chosen corresponding to approximately twice the minimum size object one desires to detect. For example, for a system with an along track ping-to-ping distance of 30 cm desiring to detect an

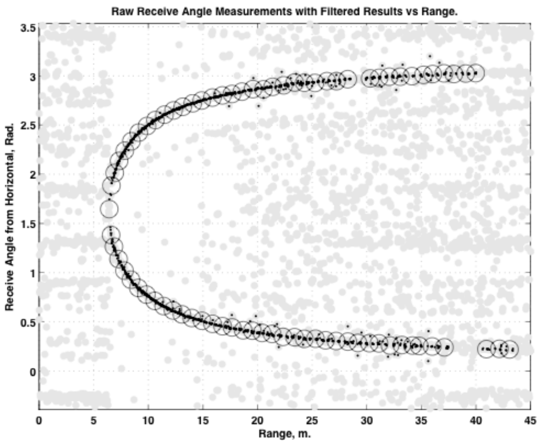

Fig. 2. Raw receive angle measurements (gray) with MPAA picks (circles) and individual retained measurements (black points).

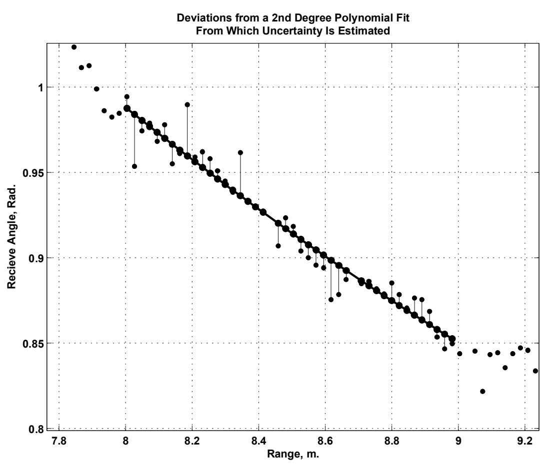

Fig. 3. Uncertainty estimation is accomplished by calculating the mean square deviation from a polynomial fit to the data over 1 m segments

object of 1 m in size, the along track filter window may be set to 6 measurements. When the across-track spacing bin size has been set to 1 m, an across track filter window of 3 measurements suffices. Most probable angle values are then retained whose predicted depth is within 1 m of the median filtered data in the along track direction and 3 m in the across track direction. The effect of this de-spiking process is to require 3 pings of data on any object whose nominal horizontal dimension is 1 m or smaller and whose vertical dimension is greater than 1 m from the surrounding seafloor.

In the final filtering step, most probable angle values are used to guide the filtering of the original individual measurements. Measurements within each bin that fall outside a band centered on the most probably angle estimate, whose width is equal to the ambiguity sector that results from the largest stave spacing, are removed along with all the data in any bin whose most probable value was filtered altogether during de-spiking. In addition, it can happen that the SNR filtering step will remove nearly all the data in a range bin. When this happens there can be insufficient data to produce the most probable angle histogram. When fewer than 25% of the original measurements remain in the range bin all the data are removed.

To explain the choice of filtering of measurements that fall outside a phase ambiguity, consider Equation 4 which gives the uncertainty in receive angle as a function of aperture spacing, L, and phase difference uncertainty, *δΔ*□, for a pair of transducers measuring a signal of wavelength, λ, arriving at an angle, θ, from baseline.

$$\delta\theta = \frac{\delta\Delta\phi}{2\pi}\frac{\lambda}{L}\frac{1}{\cos(\theta)} \qquad (4)$$

Note that as the SNR decreases, causing a corresponding increase in the phase difference uncertainty, *δΔ*□*,* the uncertainty in receive angle, δθ, increases first for those pairs of staves having the largest spacing, L. Hence, phase difference measurements from the largest stave spacing will be the first to asymptote to the +/-π phase difference variance saturation limit described previously. Therefore, a phase ambiguity corresponding to the largest stave spacing was chosen as natural threshold.

Figure 2 illustrates the result of applying these filters to a single ping. Both most probable angle estimates and retained individual measurements are shown. With the filtering complete, remaining measurements are statistically likely to result from a real object and their measurement uncertainty can now be assessed.

To estimate receive angle measurement uncertainty, a 2nd degree polynomial is fit to the retained receive angle measurements within each horizontal range bin as a function of range and the mean square deviation from this curve is calculated (Fig. 3). This method is borrowed from the approach taken by Lurton in deriving a “quality factor” for phase detections of multibeam and phase differencing side-scan systems [6] and provides a step-wise ping-to-ping empirical estimation of measurement uncertainty across the swath.

III. Creating Seafloor Sounding Estimates

Seafloor measurements retained from the filtering process can be combined to produce sounding estimates, which are lower in noise than the individual measurements themselves. This combination of measurements invariably comes at a loss of resolution and several methods (e.g. constant angle bins, constant range bins and horizontal across track distance bins) have been considered.[7] An additional method is proposed, in which the estimation process is optimized using the empirically derived uncertainty estimates in an across-track uncertainty-weighted mean.

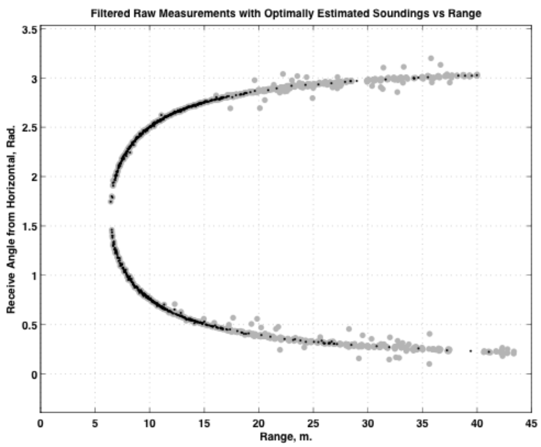

Fig. 5. Individual retained soundings after filtering (gray) along with seafloor estimates created using an uncertainty weighted mean with a 0.1 m depth uncertainty limit.

To create optimal receive angle measurements, a desired depth uncertainty limit is first specified. Adjacent receive angle measurements within a ping are then combined in an uncertainty-weighted mean (Equation 5) until the uncertainty (Equation 6) in that mean would produce a depth uncertainty less than the specified limit. Figure 4 shows the raw filtered receive angle measurements vs. range along with the seafloor estimates generated using this method for a single ping where the specified depth uncertainty was set to 0.1 m.

$$\hat{\theta} = \frac{\sum_i^N \frac{\theta_i}{\sigma_i^2}}{\sum_i^N \frac{1}{\sigma_i^2}} \quad (5)$$

$$\hat{\sigma}^2 = \frac{1}{\sum_i^N \frac{1}{\sigma_i^2}} \quad (6)$$

The effect of producing seafloor estimates limited by the uncertainty in the individual measurements is to produce a variable number of soundings per ping whose across-track spacing varies with the information gained about the seafloor. Near nadir and at the outer edges of the swath many individual measurements are combined to reduce the estimate uncertainty, while in between, where uncertainty is low, estimates may contain only a single measurement obtaining the full spatial resolution of the system. By combining soundings in this way resolution is optimized for the given desired uncertainty.

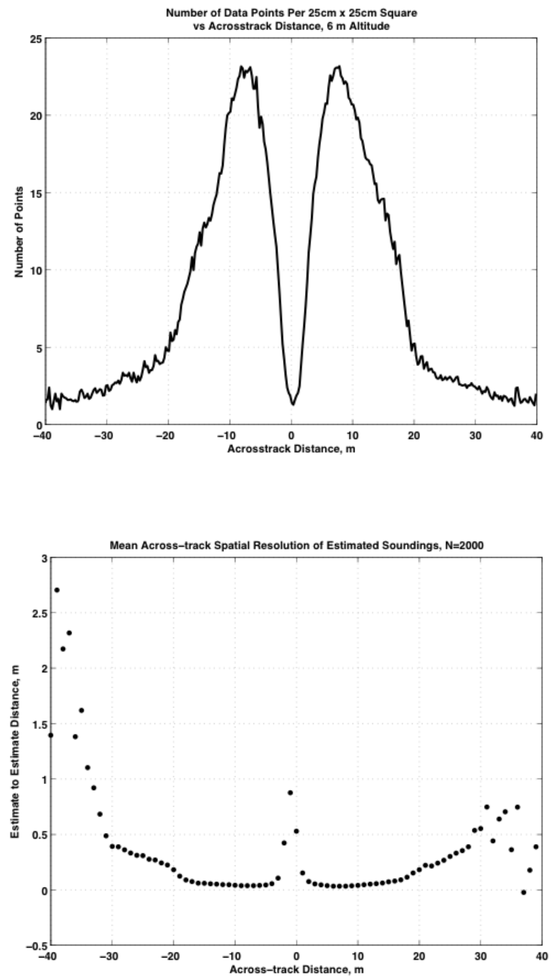

a)

b)

Fig. 4. The 2-dimensional data density of estimated soundings from a swath, collected at 6 m altitude from an AUV operating at 3 knots, is illustrated in a) as the average number of data points per 25 cm square vs. across track distance. Plots such as this give some guidance for survey planning and system performance. Mean spacing between adjacent estimates is shown in b). Soundings achieve the full across-track spatial resolution of the system (4.8 cm) from just off nadir to approximately 12 m range. Resolution decreases near nadir and at distance ranges as multiple soundings must be combined to achieve the desired uncertainty limit.

IV. Conclusion

A method has been presented to automate the removal of outliers from phase differencing side scan sonar bathymetric measurements, estimating the uncertainty in the remaining measurements and using this uncertainty to create lower noise seafloor estimates. The method builds statistical power by binning measurements in the across-track direction and then removes outliers guided by the physics of the measurement process, i.e. the ambiguous nature of the phase difference measurement corresponding to the largest stave spacing. Seafloor sounding estimates are calculated by setting a desired sounding uncertainty and then combining receive angle measurements in a way that optimizes the resolution of the system within the uncertainty constraint.

The result of these processes produces a data product similar in quality and extent to more commonly used multibeam sonar systems. Approximately 800 estimated soundings result from a 4000-measurement swath when a 0.1 m standard deviation is specified, greatly reducing the volume of data and the subsequent processing burden. Because these estimated soundings are relatively outlier free, they may be combined in statistical gridding algorithms such as CUBE to generate a seafloor surface and, where sufficient evidence exists, multiple hypothesizes requiring further investigation. Moreover, the high across-track resolution of these systems is maintained where the uncertainty is low enough to support it and decreases automatically to maintain the required uncertainty where it is not. This combining of uncertain measurements allows extension of the swath width (albeit at lower resolution) far beyond 3-5 altitude heights providing useful information and a great increase in efficiency where lower resolutions can be tolerated. Care must be taken, however, to apply SNR filtering to prevent the variance saturation effect described above.

This method of combining measurements based on a desired uncertainty while optimizing across-track resolution may also be advantageous to multibeam systems. Multibeam sonars maintain constant across track sounding spacing in angle or distance throwing away resolution when the uncertainty is low and producing a data density inconsistent with the actual resolution when the uncertainty is high. Methods presented may similarly optimize the creation of soundings for those systems.

ACKNOWLEDGMENT

The authors wish to thank X. Lurton for numerous discussions regarding many aspects of this paper. Funding for this work was provided under NOAA grants NA05NOS4001153 and NA10NOS4000073.

REFERENCES

[1] P. H. Kraeutner and J. S. Bird, “Beyond interferometry, resolving multiple angles-of-arrival in swath bathymetric imaging,” vol. 1, pp. 37–45 vol.1, 1999.

[2] P. N. Denbigh, “Signal processing strategies for a bathymetric sidescan sonar,” *IEEE Journal of Oceanic Engineering*, vol. 19, no. 3, pp. 382 –390, Jul. 1994.

[3] X. Lurton, *An Introduction to Underwater Acoustics: Principles and Applications*, 2nd ed. Springer, 2010.

[4] B. R. Calder and L. A. Mayer, “Automatic processing of high-rate, high-density multibeam echosounder data,” *Geochemistry Geophysics Geosystems*, vol. 4, no. 6, p. 1048, Jun. 2003.

[5] X. Lurton, “Précision de mesure des sonars bathymétriques en fonction du rapport signal/bruit,” *TS. Traitement du signal*, vol. 18, no. 3, pp. 179–194, 2001.

[6] X. Lurton and J. Augustin, “A Measurement Quality Factor for Swath Bathymetry Sounders,” *IEEE Journal of Oceanic Engineering*, vol. 35, no. 4, pp. 852 –862, Oct. 2010.

[7] X. Lurton, “Theoretical modelling of acoustical measurement accuracy for swath bathymetric sonars,” *The International hydrographic review*, vol. 4, no. 2, pp. 17–30, Aug. 2003.