Article

PRESENTING AN AUTOMATED CALIBRATION PROCEDURE FOR AN AIRBORNE LIDAR SYTEM AND ITS POTENTIAL APPLICATION TO ACOUSTIC HYDROGRAPHY.[1,2]

By Lt Michael O. GONSALVES (USA)
NOAA National Geodetic Survey, Remote Sensing Division – JALBTCX
Department of Marine Science, University of Southern Mississippi

Abstract

Whether using an airborne lidar or a ship-based acoustic system, all hydrographers must contend with geometric system calibrations. A poorly aligned system leads to erroneously reported depths, diminished system resolution and internally inconsistent datasets. Most of today's calibration procedures are cumbersome and subjective enterprises that possess little statistical merit. This paper presents a least squares adjustment algorithm designed to calibrate a (presently under-development) lidar. This method is automated, objective, repeatable, and reports a confidence on the calibration values. Using simulated lidar datasets, the algorithm is explained and demonstrated. A brief modification is also proposed to expand the use to multibeam echosounders.

Résumé

Que ce soit à l'aide d'un lidar aéroporté ou d'un système acoustique embarqué, tous les hydrographes doivent faire face à des étalonnages de systèmes géométriques. Un système mal aligné conduit à des erreurs dans les profondeurs indiquées, à une diminution de la résolution du système et à des ensembles de données inconsistants en interne. La plupart des procédures d'étalonnage actuelles sont compliquées et sujettes à des tâches qui n'ont qu'un faible mérite statistique. Cet article présente un algorithme d'ajustement à l'aide de la méthode des moindres carrés conçu pour étalonner un système lidar (actuellement en développement). Cette méthode est automatique, objective, répétable et rend compte d'une confiance dans les valeurs d'étalonnage. A l'aide d'ensembles de données lidar simulées, l'algorithme est expliqué et démontré. Une brève modification est également proposée afin d'étendre leur utilisation aux échosondeurs multifaisceaux.

Resumen

Independientemente de si se usa un lidar aerotransportado o un sistema acústico embarcado, todos los hidrógrafos deben enfrentarse a las calibraciones de sistemas geométricos. Un sistema escasamente alineado conduce a errores en las profundidades indicadas, a disminución de la resolución del sistema y a colecciones de datos internamente inconsistentes. La mayoría de los procedimientos de calibración actuales son complicados y sujetos a tareas que poseen poco mérito estadístico. Este artículo presenta un algoritmo de ajuste mediante un método de mínimos cuadrados designado para calibrar un lidar (en vías de desarrollo actualmente). Este método es automatizado, objetivo, repetible, e indica una confianza en los valores de calibración. El algoritmo se explica y se demuestra utilizando colecciones de datos del lidar simulado. Se propone también una breve modificación para ampliar el uso a los sondadores acústicos multihaz.

---

1. Portions of this work were completed as part of the graduate program requirements of the University of Southern Mississippi.
2. This work is a revised version of a paper presented at the 2010 Canadian Hydrographic Conference.
3. Electronic mail: michael.gonsalves@noaa.gov

I. INTRODUCTION

Practitioners of acoustic multibeam hydrography are well-versed in the process of field calibration, referred to as the "patch test". The standard patch test seeks to resolve, through a series of coupled survey lines, the angular misalignments (pitch, roll, heading) between the Inertial Navigation System (INS) and the sonar. These lines, acquired over a particular grade of seafloor, are designed to isolate and identify a single parameter at a time. Final determination of these misalignments can be a subjective affair and is dependent upon the sound velocity (SV) and tidal characteristics being well known. In contrast, the geometric calibration of an airborne bathymetric laser, or lidar, can be performed on land, thus eliminating SV and tidal concerns. Further, rather than search the seafloor for suitable acoustic calibration targets, a lidar calibration can use cultural features like roads and gabled roofs.

The Coastal Zone Mapping and Imaging Lidar (CZMIL), a system presently under development by Optech International for the U.S. Army Corps of Engineers, will employ a prototype circular scanner using a refracting prism. This new design has the potential for geometric misalignments not previously confronted in a contemporary system and has forced its developers to rethink their calibration strategy. To this end, an automated least-squares adjustment (LSA) routine has been developed that allows all flight lines to be conducted over a single flat featureless surface (e.g. an airport runway or the sea surface).

In this paper, a brief background is presented on the current practices of multibeam echosounders (MBES) and lidar calibration, emphasizing some of the unique advantages airborne lidar has to offer. The LSA method of calibration is then discussed using synthesized datasets that simulate the CZMIL scan pattern. Preliminary results will show that the technique is so robust the calibration routine can be expanded to simultaneously adjust up to 13 calibration parameters (some unique to CZMIL, and some that would be of interest to those who work with acoustic sounders). Finally, a discussion of the feasibility of modifying the algorithms for the development of an automated multibeam calibration utility is presented.

II. COMPARING TRADITIONAL CALIBRATION TECHNIQUES

A. Multibeam Echosounders

The National Oceanic and Atmospheric Administration (NOAA) offers several good descriptions for multibeam calibration (NOAA 2010a, NOAA 2010b). The goal of this calibration, or patch test, is to determine the angular alignment between the INS reference frame and the sonar reference frame (the pitch, roll and yaw bias), in addition to any time latency between the systems. For the purposes of this paper, it is assumed the systems being discussed will use some form of precise timing protocol (PTP) like those discussed in Calder and McLeod (2007). As such, time latency will not be further considered as a calibration parameter.

Calibration lines are typically acquired in pairs in such a way that the bias of a single parameter is isolated from the others. Depending on the parameter being investigated, these survey lines can be focused on a prominent feature on the seafloor (e.g. a rock) or on a featureless bottom. *FIG. 1* shows a simple line plan to be used by NOAA on a featureless bottom. Generally speaking, the pitch lines are run in opposing directions up-and-down a sloping bottom; the roll lines are run in opposite directions over any bottom profile; and the yaw lines are run in opposite directions such that their outer beams overlap (where a sloping bottom is required for yaw determination).

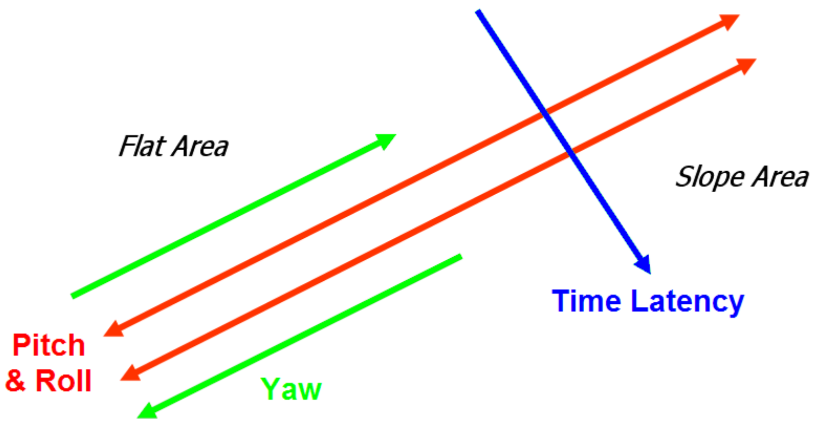

FIG. 1. *A typical line plan for the calibration of a MBES (note the second set of pitch/roll lines is included for redundancy). Modified from NOAA (2010a).*

Once the survey lines are acquired, carefully chosen subsets of the soundings are examined to systematically determine each calibration value. For example, ***FIG. 2*** shows a subset of two swaths as viewed in the across-track direction, which are used to determine the system's roll bias. The roll values for each line are manually incremented until the two adjacent swaths appear to overlap. While this can be a subjective affair, some software packages do offer the ability to semi-automate this process.

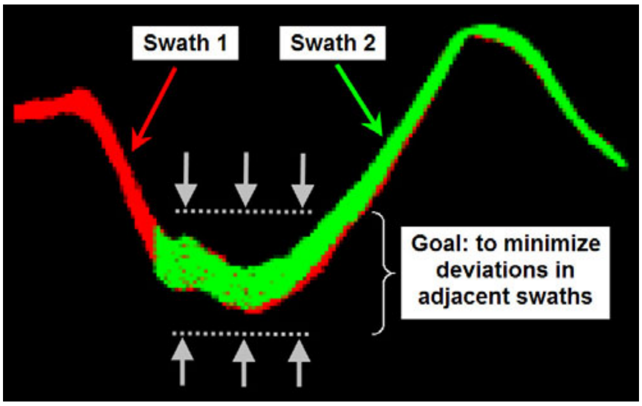

FIG. 2. *Two swaths of overlapping data are used to determine the MBES system roll bias. By steadily incrementing the misalignment, the two swaths are "rotated" until the swaths agree.*

If available, a prominent feature on the seafloor can be used as a calibration target. As depicted in ***FIG. 3***, a vessel travels over a rock going in opposite directions. The presence of a positive (forward-looking) pitch bias will result in the detection of the feature in advance of the vessel passing over the object, thus misrepresenting the objects location. By surveying in both directions, the two misrepresentations of the object can be brought into unison by adjusting the pitch bias. The scene depicted in ***FIG. 3*** can be used to adjust for a yaw bias if the rock is depicted in the swath's outer beams.

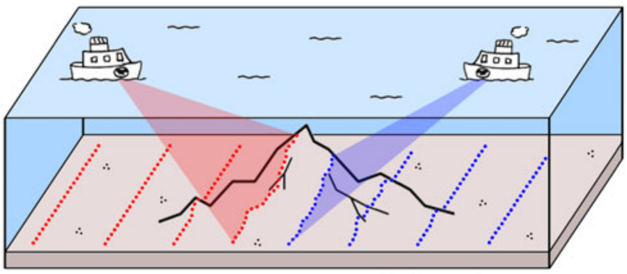

FIG. 3. *Prominent features on the seafloor are used to determine both pitch and yaw biases, as well as navigation time latency. As shown above, the two swaths are compared to determine the pitch bias. Note: due to the limited sonar ping rates, the swaths being compared are not exactly coincident; leading to an uncertainty in the calibration values.*

The drawback to using a target in the calibration procedure is that only a few "pings" of data are actually used during the adjustment. Additionally, these thin swaths may not necessarily even be at the peak (least-depth) of the feature, which is what the operator is using as a reference point in the adjustment. As a result, the confidence in the determined calibration values is diminished.

B. Lidar

The biggest difference between the calibration of a bathymetric lidar versus a multibeam is that the lidar isn't restricted to performing its calibration routine over the water. By performing the alignment on land, all uncertainties associated with sea swell, beam attenuation, and tidal effects are removed. While hydrographers must invest time in searching for appropriate study areas (flat bottom or prominent features), terrestrial targets are abundant in the form of roadways and buildings. Absolute positioning is an important aspect of survey accuracy control. It is a complicated enterprise to establish the absolute position of a feature on the seafloor; whereas, ground truthing on land can be accomplished by occupying desired calibration targets with static GPS base stations.

Most lidar calibrations are performed by acquiring data over cultural targets, like buildings (***FIG. 4***). The method of adjustment is similar to that of a multibeam calibration target in that cross-sections of the lidar swaths are examined, and the calibration values are steadily adjusted until data between overlapping strips match. Because it is difficult to establish conjugate points from one lidar swath to another, some adjustment procedures instead extract linear and planar features from the individual swaths (Habib et al 2008, Vosselman and Djikman 2001, Schenk 2001). Rather than adjust the points from one swath to another, these planar features are used instead (***FIG. 5***).

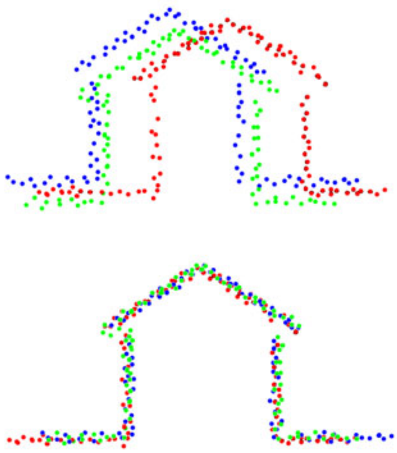

FIG. 4. (*top*) *A lidar point cloud (colored by acquisition) line as acquired over a gabled roof. The disagreement among lines reveals a poor alignment of the sensor. (bottom) By adjusting the calibration values, the building is brought into focus.*

The method of least-squares is increasingly used in lidar calibrations; the sections that follow give a brief overview of a new adjustment model. Rather than trying to best-fit neighboring strips or adjusting extracted roof tops to each other, the proposed model will fit the entire dataset to a single planar surface. This surface could be a cultural feature (like an airport tarmac) or the surface of a lake or ocean

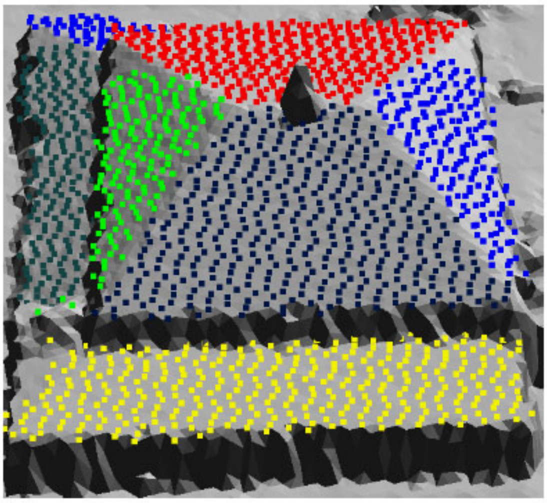

FIG. 5. A lidar point cloud acquired over a building. A planar extraction algorithm was performed to identify each of the surfaces of the building's roof. Reproduced from Freiss (2006).

III. A LEAST SQUARES APPROACH

A. A historical perspective

The application of an iterative least squares adjustment procedure is not unprecedented in the oceanographic world. In a pre-GPS constellation world, establishing a long baseline (LBL) acoustic positioning network was considered the most accurate technique for deep ocean positioning of a vessel and the only method for positioning equipment at or near the sea floor (McKeown 1975). The technique involves determining the position of a rover station (ship, submersible, etc.) through a series of acoustically-determined range observations from three or more deployed transponders of known (relative or absolute) position (**FIG. 6**).

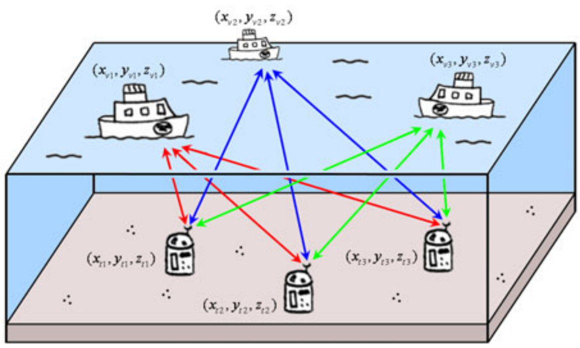

FIG. 6. A long baseline acoustic network: a collection of transponders of known positions (x_t,y_t,z_t) used to position a vessel in the water column or at the surface. (x_y,y_y,z_y)

When the transponders are first deployed, their positions relative to each other must be determined. To that end, several calibration lines are performed by a surface vessel: cloverleaves over each transponder to determine their least depth, and transect lines to determine baseline lengths for each transponder pair. These datasets establish a relative network and, with surface vessel position data collected in conjunction with the transponder ranges, establish "absolute" fixes of the network nodes. The principle drawbacks to this method are that it requires excessive ship time and tends to produce biased depth measurements.

A more efficient method proposed by Lowenstein (1965), invokes a least-squares adjustment. Under this method, the calibration lines are omitted and the vessel immediately begins its intended operations. During operations, the vessel logs the ranges to all the transponders. Once a nominal number of measurements are taken, the LSA adjusts the positions of the transponders until a best fit of all the measured slant ranges is determined.

Hydrographers may be familiar with a similar adjustment model in performing a static vessel survey. Here redundant angle and range measurements are taken among system components (GPS antennas, sonar head, vessel reference point, etc.) using a total station. These measurements are entered into an adjustment model which then estimate the relative positions of the components.

A least-squares adjustment procedure offers several advantages besides automation. Not only will systematic errors be identified, but analysis of the covariance matrix will provide estimates of the random uncertainty of each input parameter. Examination of the residuals can be used to detect blunders in the measurements. Lastly, and of critical importance in estimating the sounding confidences, an LSA provides uncertainties for the calibration values which can be used to compute estimated errors of the final depth measurements.

B. The adjustment model

To describe the least-squares adjustment model, first consider a generic function:

$$f(\vec{\ell}, \vec{x}) = 0 \tag{1}$$

which has a first-order approximation:

$$f(\vec{\ell}, \vec{x}) \approx \underbrace{f(\vec{\ell}_0, \vec{x}_0)}_{\vec{g}} + \underbrace{(\vec{\ell} - \vec{\ell}_0)}_{\vec{r}} \underbrace{\left.\frac{\partial f}{\partial \ell}\right|_{\vec{\ell}=\vec{\ell}_0, \vec{x}=\vec{x}_0}}_{\mathbf{D}} + \underbrace{(\vec{x} - \vec{x}_0)}_{\vec{\delta}} \underbrace{\left.\frac{\partial f}{\partial x}\right|_{\vec{\ell}=\vec{\ell}_0, \vec{x}=\vec{x}_0}}_{\mathbf{A}} \approx 0 \tag{2}$$

where:

f = the observation equation

$\vec{\ell}$ = observables' true values (laser range, vehicle position, etc.)

$\vec{\ell}_0$ = observables' measured values

$\vec{x}$ = adjusting parameters' true values (pitch bias, roll bias, etc.)

$\vec{x}_0$ = adjusting parameters' initial guesses

$\vec{r}$ = corrections to $\vec{\ell}_0$ (residuals)

$\vec{\delta}$ = corrections to $\vec{x}_0$ (adjustments).

Equation (2) can be rewritten as:

$$f(\vec{\ell},\vec{x}) \approx \vec{g} + \mathbf{D}\vec{r} + \mathbf{A}\vec{\delta} \approx 0 \quad (3)$$

which, when applying the least squares model and solving for $\vec{\delta}$ yields:

$$\vec{\delta} = \left(\mathbf{C}_{\vec{x}_0}^{-1} + \mathbf{A}^{\mathrm{T}}\left(\mathbf{D}\mathbf{C}_{\ell}\mathbf{D}^{\mathrm{T}}\right)^{-1}\mathbf{A}\right)^{-1}\mathbf{A}^{\mathrm{T}}\left(\mathbf{D}\mathbf{C}_{\ell}\mathbf{D}^{\mathrm{T}}\right)\vec{g} \quad (4)$$

where:

$\mathbf{C}_{\vec{x}_0}^{-1}$ = a priori estimate of uncertainties in $\vec{x}_0$

$\mathbf{C}_{\ell}$ = a priori estimate of uncertainties in $\vec{\ell}$.

To perform the geometric calibration, the lidar data will be acquired over a flat surface. The above least-squares adjustment will then be performed which will adjust the lidar calibration values to best fit the point cloud to that planar surface. A similar model was suggested by Freiss (2006), in which Equation takes the form:

$$f(\vec{\ell},\vec{x}) = \vec{n}\cdot(x_{OBS} - x_P) \quad (5)$$

where:

$\vec{n}$ = vector normal to planar surface

x_{OBS} = 3D coordinates of laser point from laser location equation

x_P = 3D coordinates of fixed point on planar surface

In studying equation (5), one should note the dot product of two vectors is zero if the vectors are orthogonal (perpendicular). Given $\vec{n}$ is already normal to the plane (*FIG. 7*), that implies the vector $(x_{OBS} - x_P)$, and thus the point x_{OBS} , must be on the plane. Accordingly, equation fits the laser points x_{OBS} to a planar surface.

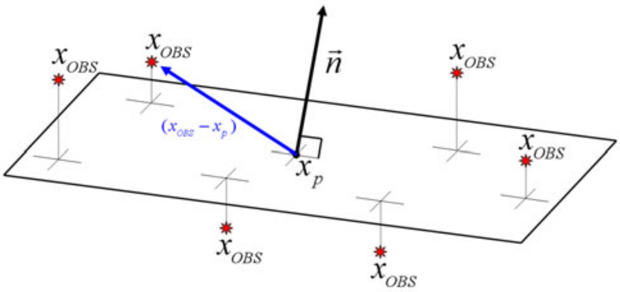

FIG. 7. In fitting the point cloud to a planar surface, the dot product of the planar normal vector, $\vec{n}$, is taken with respect to the offset vector of each laser point, $\vec{x}_{OBS}$, and a fixed point on the plane, $\vec{x}_p$.

C. A simplified calibration procedure (pitch, roll and yaw)

On initial inspection, one might think it is impossible to extract pitch, roll and yaw boresight misalignments from a featureless planar surface. What follows is an abbreviated geometric argument showing such a technique is possible. First, consider a circular-scanning lidar with no geometric misalignments that is flown over a horizontal planar surface (FIG. 8 – left). In a well-aligned system, the measured point cloud will perfectly describe the planar surface, and assuming a level flight, every laser pulse will report the same range between the laser and the laser footprint (within the measurement noise level).

Now consider if, unknown to the operator and processing algorithms, the laser is pitched 10° towards the nose of the aircraft (FIG. 8 – right). Under such a configuration, the forward-looking beams will travel a greater distance, while the aft-looking beams will travel a relatively shorter distance. The operator, still believing the lidar to be properly oriented, will interpret the longer-length forward beams, coupled with the shorter aft beams to mean the system is acquiring data down the backside of a hill. The key is that the biased point cloud will no longer describe a planar surface, but a helix with a vertical deflection that is proportional to the bias in the pitch boresight angle. This deflection from a planar surface is what will be resolved with the least-square adjustment.

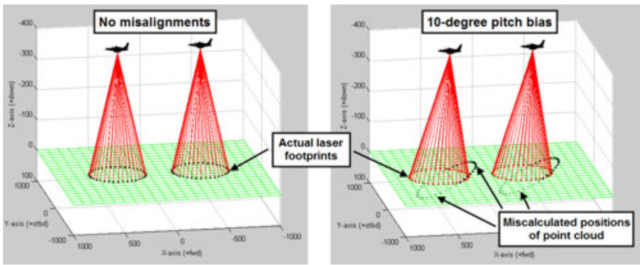

FIG. 8. (left) Two revolutions of the laser scanner with no boresight misalignments – notice both circular traces are co-planar. (right) With a 10° forward (i.e. towards the nose) pitch boresight bias, two revolutions are again depicted with the actual laser footprints shown in red and the miscalculated point cloud shown in black – notice the biased points are no longer co-planar.

Already, the circular-scanner design shows its advantages over a lateral swath design with regard to a geometric calibration routine; in the former case, a single level flight line will reveal the pitch boresight angle. Not all calibration parameters will immediately reveal themselves with such a flight itinerary. For example, introducing a roll boresight misalignment would cause the point cloud on the ground to also roll en masse (***FIG. 9***). Though the biased points from a single flight line would be *tilted*, they would still be co-planar, and thus immune to an LSA which seeks to adjust the point cloud to a planar surface. However, much like the acoustic patch test, if a second line of data is acquired with an opposing heading from the first, the resulting biased point cloud will not be co-planar with the first biased point cloud (***FIG. 9***). Thus the LSA can be applied jointly to these two datasets, adjusting the roll boresight calibration angle until the data is aligned with a single planar surface.

The previous discussion of roll encapsulates the flight line strategy associated with this adjustment procedure. The vessel must be maneuvered in such a way that any misalignments manifest themselves as a deviation from the otherwise flat planar surface.

If a vessel were to fly level over a horizontal ground (as in ***FIG. 8*** – left), then any yaw boresight misalignment would cause the point cloud to experience a radial shift across the ground; but would still be co-planar. To determine the yaw angle, the vessel must have a change in attitude. As depicted in ***FIG. 10***, if a vessel is pitching nose up, with no boresight misalignments, then the greatest measured laser range will be produced from the forward-most beam. Were this same vessel to have a yaw boresight misalignment, then the range from the forward-most beam would be erroneously assigned to an azimuth rotated by the yaw angle from the forward direction. The end result is a biased cloud that is sometimes below and sometimes above the actual ground plane, which permits solving for the yaw bias by adjusting to the planar surface model.

As an example of the algorithms' effectiveness, two 20-second flight lines were simulated flying in opposing directions (heading 0° and 180°) with vessel pitches of ±10° respectively (***FIG. 11***). The simulation added roll, pitch and yaw boresight misalignments of 10°, 15° and 20°. The calibration procedure converged to the correct misalignments with uncertainties (1-sigma) of 0.0022° in roll, 0.0023° in pitch and 0.0118° in yaw. The reported uncertainties are those output by the LSA algorithm. In all cases, the algorithm's calculated calibration values agreed (to within the predicted tolerances) of the "real" calibration values used in the simulation.

Regarding the simulation, noise was added to all the observations based on the manufacturer's specifications of the hardware. The planar surface was assumed to be flat, however even a rough surface can produce satisfactory results (Gonsalves 2010a). The reader should not be distracted by the large magnitude of the misalignments used in the simulation (10°, 15° and 20°); the algorithm performs equally well on a system with misalignments of only a few tenths of a degree.

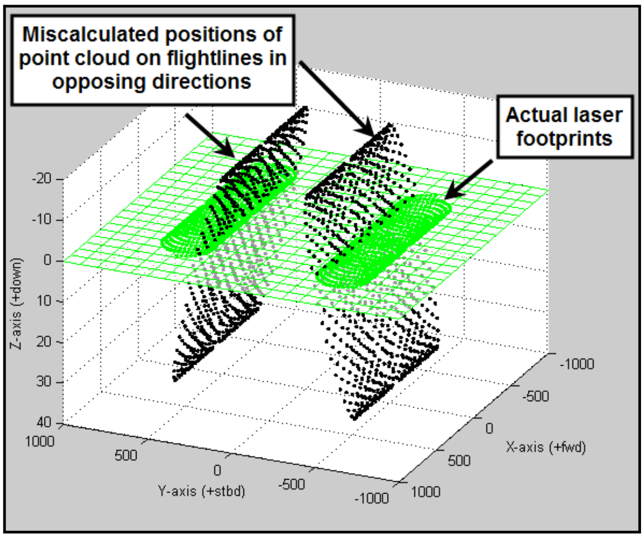

FIG. 9. *Two survey lines acquired with an unknown roll bias and opposing headings. The actual point cloud is shown in green; the miscalculated point cloud is shown in black. Notice the biased points for any given flight line are respectively co-planar, but are not co-planar with each other*

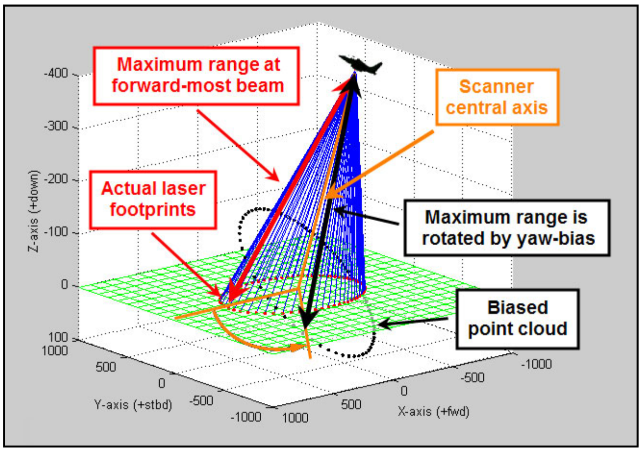

FIG. 10. *A vessel pitching nose-up will measure the longest slant range in its forward-most beam (indicated by red arrow). If this same vessel, however, has a yaw misalignment (i.e. a rotation about the scanner's central axis – in orange), then the range previously associated with the forward-most beam will be rotated by the yaw angle bias (indicated by black arrow).*

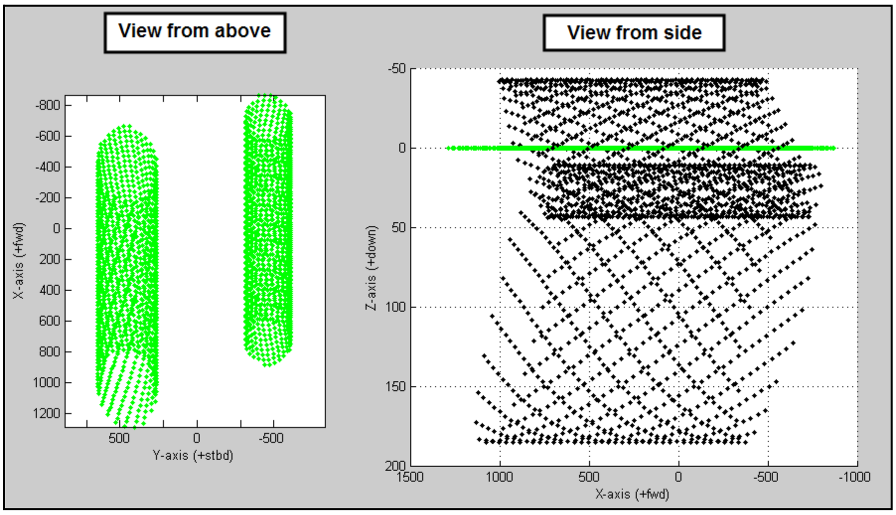

FIG. 11. *(left) Top view of two simulated flight lines used in a roll-pitch-yaw boresight calibration. (right) The incoherent point cloud pre-calibration (black) shown with the post-calibrated data (green) which has been fit to a planar surface.*

No. of flight lines used (Acquisition scheme - Laser pulse rate)	No. of laser strikes used	No. of parameters adjusted	Confidence in both the calibration parameters and maximum point cloud propagated uncertainty (horizontal and vertical)				
			1σ-roll	1σ-pitch	1σ-yaw	Horiz. TPU	Vert. TPU
Two lines (opposing direction - 50Hz)	~1,900	3 (roll/pitch/yaw)	0.0022°	0.0023°	0.0118°	0.074m	0.017m
One line (slow change in roll and heading - 50Hz)	~1,900	3 (roll/pitch/yaw)	0.0155°	0.0061°	0.0370°	0.251m	0.051m
One line (slow change in roll and heading - 10000Hz)	~188,000	3 (roll/pitch/yaw)	0.0005°	0.0002°	0.0011°	0.008m	0.002m
Four lines (2 rolling/yawing, 2 pitching/heaving - 10000Hz)	~751,000	11 (boresight, IMU offset, etc.)	0.0014°	0.0014°	0.2415°	0.012m	0.002m

Table 1.
Performance of calibration algorithm for various sized datasets and flight configurations

D. A more robust calibration

In the previous discussion, a means of determining the roll boresight misalignment was proposed by conducting two flights in opposing directions. Such a technique is in keeping with the traditions of the acoustic patch test in which a pair of coupled survey lines are designed to isolate a single parameter at a time. Interestingly, two separate flight lines are not required to determine the roll angle. What is required is merely a flight line where the heading changes. A single flight line in which the pilot makes a slow turn to the left or right provides enough linearly independent information to extract both the pitch and roll calibration values. Similarly, by also rolling the vessel, all three boresight angles can be determined from a single flight line (FIG. 12). For the purposes of the simulation, the vessel will roll 5° to the right, then 5° to the left, then return to a level attitude (a similar 5° oscillation will be imposed on the vessel heading).

The results of the calibration from the single dynamic flight line discussed above are shown in Table 1 – 2nd entry. Notice with the same number of data points the confidence (reported by a smaller standard deviation) is greater for the flight lines conducted in opposing directions. This is because flight lines in opposite directions present a stronger geometric alignment in which the biases due to roll are most pronounced. Ultimately the field personnel will decide whether to trade the greater confidence provided in multiple flight lines with the cost and time savings of a single wiggly line. It should be noted, however, that only 1/200th of the available data were used in the previous calibrations. If the full 10,000Hz dataset is included in the adjustment, then a 20-second flight line can successfully determine the three boresight angles to within a thousandth of a degree (Table 1 – 3rd entry).

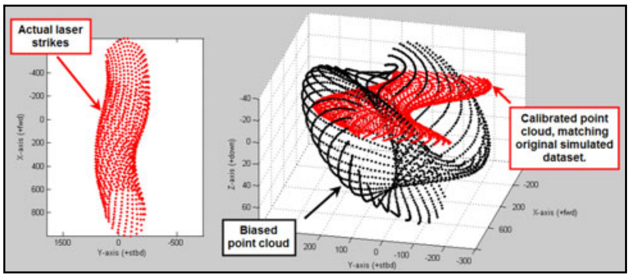

FIG. 12. *(left) Top view of a simulated flight line where the vessel exhibits a slow roll and change in heading. (right) Before calibration the point cloud is incoherent (black), but after application of the LSA, the calibration values are determined and the point cloud is restored.*

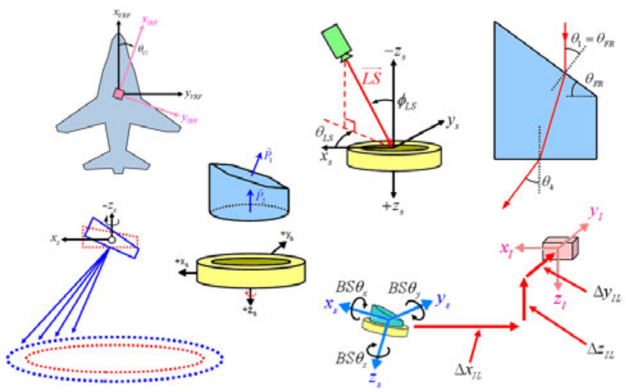

FIG. 13. *Some of the calibration parameters being adjusted for CZMIL include vessel-to-INS heading bias (upper left), scanner-to-prism alignment (lower left), prism slope (upper right) and INS-to-laser offsets (lower right).*

More important than the calibration values themselves is how the calibration uncertainties carry forward to the ultimate location of the soundings (as derived using the general law of the propagation of variances). For the 10,000Hz trial, the uncertainties in the calibration parameters will only contribute 0.008m (1σ) to the soundings horizontal uncertainty and 0.002m (1σ) to the vertical uncertainty (Table 1 – 3rd entry). When compared to the uncertainties of either ellipsoidal positioning or tides, the uncertainty of the calibration values are negligible.

As mentioned earlier, CZMIL is a prototype lidar with a novel scanner design. The impetus for pursuing a new method of calibration is in anticipation of having to calibrate a system which may exhibit geometric misalignments not previously seen. Going beyond just the roll/pitch/yaw boresight calibration, a total of 15 parameters are being investigated for calibration (***FIG. 13***). Using conventional "patch test" wisdom, in which a pair of survey lines are used to decouple each parameter, then one could conservatively anticipate having to run 16 survey lines to estimate all the parameters. The ability to solve for several calibration parameters at once (as demonstrated in ***FIG. 12***) showcase the flexibility afforded by the LSA approach.

While most of the parameters shown in ***FIG. 13*** are beyond the scope of this paper, one has received some attention in the hydrographic literature: aligning the vessel reference frame (VRF) with the INS reference frame (IRF). Any misalignments along this vertical axis will lead to cross-talk between the INS-sensed pitch and roll (Hilster 2008). For example, consider a vessel with a laser mounted one meter forward (measured with respect to the VRF) of the INS (***FIG. 14 - #1***). If this vessel were to strictly pitch, then the laser head would pivot up and aft (***FIG. 14 - #2***). However, should the INS be misaligned, while the vessel is pitching, the INS senses that it is mostly pitching along with a slight roll (***FIG. 14 - #3***).

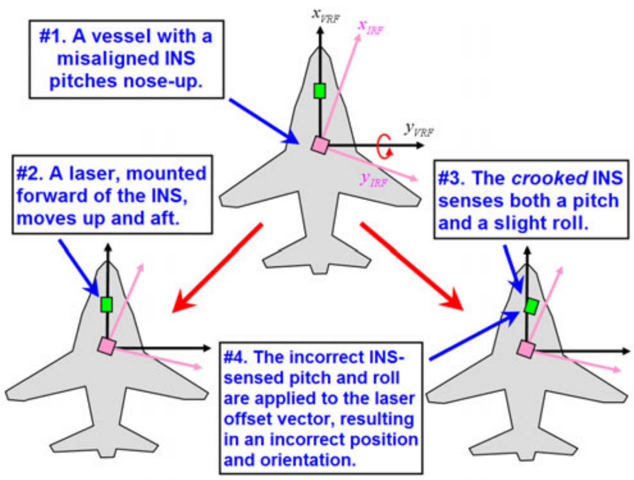

FIG. 14. *The effects of cross-talk in a poorly-mounted INS. Incorrect rotations are applied to the lever arms resulting in both translational shifts and angular biases.*

This sensed data is recorded and later applied when the vessel's trajectory is computed. By applying these incorrect rotations to the laser head it a) is computed to be in the wrong spot and b) will have an incorrectly computed orientation (FIG. 14 - #4). These induced errors are relatively minor and until recently, with system noise and poor GPS resolution, have been considered inconsequential (Hughes Clark 2003). With improved positioning techniques (real-time kinematics), which can achieve positional accuracies on the order of centimeters, these errors are rising above the noise. Finding a means of addressing this misalignment is important because, in the acoustic world, the conventional patch test methodologies do not provide any means of aligning a vessel reference frame to that of the INS (Hilster 2008, Hughes Clark 2003).

Returning to Table 1, one final calibration trial was simulated, this time attempting to calibrate eleven parameters at once (including the three boresight angles discussed previously). The greater number of parameters requires a more ambitious flight plan. In this case, four lines are flown: two crossing lines (one experiencing a slight change in roll attitude and heading and the second line experiencing a change in pitch and heave) and a second set of similar lines acquired at a different altitude. The calibration software succeeded in determining all eleven of the simulated misalignments. As an aside: among those calibration parameters was the INS-to-laser offset vector. Determining this vector through an LSA is equivalent to performing a static survey on a ship's sonar without ever hoisting the vessel from the water.

The observant reader will notice the larger reported uncertainty for the yaw boresight angle in the final calibration trial (0.2415° – 1σ). This increase in uncertainty can be attributed to the large correlation with the parameter modeling the VRF-IRF misalignment shown in *FIG. 14*. When both the variances and the covariance of these two terms are taken into account, the contribution to the point cloud uncertainty from all the calibration parameters is again negligible (0.012m horizontal and 0.002m vertical - 1σ). For a complete discussion of the calibration uncertainties and the affects of covariance on the sounding accuracy, the reader is directed towards Gonsalves (2010b).

IV. FUTURE WORK

A. A multibeam calibration proof-of-concept

Even with an idealized piece of sea floor available (completely flat and featureless), there may still be some concerns whether a calibration routine, as discussed in this paper, is possible. That is, the lidar considered has beams looking both forward and backward, as well as port and starboard. The question is whether a multibeam echosounder, with its nadir-directed fan of beams is geometrically interesting enough to be calibrated in such a manner. The short answer is: yes, the geometry is present to develop a calibration routine similar to that outlined above.

To test the feasibility of a multibeam calibration routine, first a simplified sonar simulator is created, (*FIG. 15*). This simulation permits the input of any alignment between the INS reference frame and that of the sonar. For the purposes of this proof-of-concept, the water column is assumed to be of uniform density in which the sonar pulses do not refract. It is important to note that no form of beam stabilization (e.g. roll compensation) was incorporated, as the dynamics of a rotating swath are what feed the calibration.

Respective misalignments of 5°, 10° and 15° were introduced into the roll, pitch and heading mounting angles of the sonar with respect to the INS.

The virtual vessel was then cast off for one minute in a sea state that induced a ±5° roll every 10 seconds and a ±10° pitch every 20 seconds (data collected at 1 Hertz). The results of the calibration are shown in *FIG. 16*. Not only was the least squares algorithm capable of correctly determining the system misalignments, but it did so with a confidence of greater than one decimal place for all misalignments (standard deviations of 0.02°, 0.04° and 0.05° for roll, pitch and yaw – results as reported by the LSA and though agreement between the predicted and "actual" misalignments). Similar to the lidar calibrator, preliminary simulations suggest the calibration works equally well whether the misalignments are large (as shown above) or only a few tenths of a degree.

While these results are by no means definitive, they are compelling. Plans are presently in development to acquire and calibrate actual sonar datasets. It will be interesting to see whether the calibration algorithms are robust enough to handle the additional complications of sound speed ray tracing, tides, vessel dynamic draft, and the intrinsic system noise.

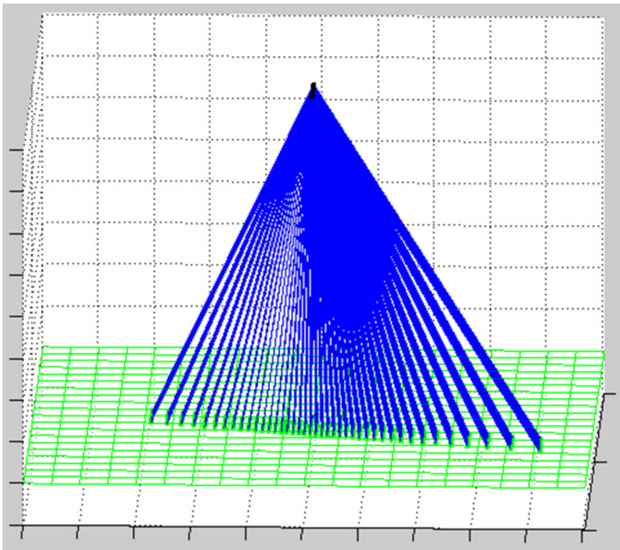

FIG. 15. A simulated sonar scan pattern from a dynamic vessel.

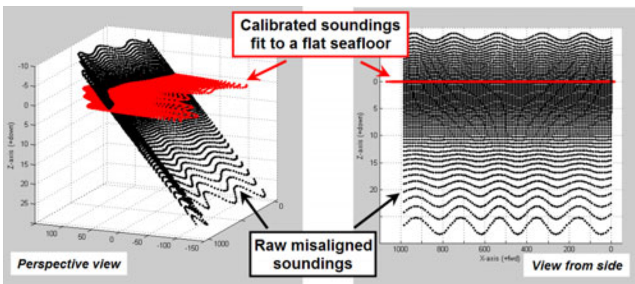

FIG. 16. A simulated sonar dataset shown both before (black) and after (red) application of the least squares calibration routine.

B. Closing remarks

It should be emphasized that thus far the methodology presented in this paper has only been performed in a simulated environment. Once CZMIL is delivered, the algorithms can be tested in an operational setting. With the brunt of the work already done, the lidar simulator and calibrator can be easily adapted for other lidar (or sonar) scanner designs.

Because the proposed LSA technique only requires a sea surface return (which is always present in the case of a bathymetric lidar) and a dynamic vessel attitude (which is provided by the atmosphere and the natural motion of the aircraft), production lines may contain all the information necessary for a calibration. This would imply an end to dedicated calibration lines, resulting in more time "on-project". Further, a calibration routine could always be running in the background during survey, testing the calibration solution and warning the operator should a misalignment be detected – this changes the philosophy of calibration from simply being a pre-survey check to being a real-time quality assurance tool. Trajectory files of past survey flights will be processed to determine if they are dynamic enough to be used for calibration.

Also, while the calibration surfaces proposed in this paper were airport runways or the ocean surface, the author believes (as demonstrated in the above proof-of-concept) minor modifications could be performed to adapt the technique to use the sea floor instead. Mud flats, the broad continental shelf or areas with small-to-moderate sand waves all provide a "flat-enough" reference surface. If the sea floor can provide an initialization for the LSA, then a method of automated calibration for a multibeam echosounder is possible.

ACKNOWLEDGEMENTS

The author wishes to acknowledge Dr. David Wells for his relentless feedback in this research and NOAA's Remote Sensing Division for encouraging his doctoral pursuits.

REFERENCES

- Calder, B. R. and McLeod, A. (2007). Ultraprecise absolute time synchronization for distributed acquisition systems. *IEEE Journal of Oceanic Engineering*, 32(4), 772-785.
- Friess, P. (2006). Toward a rigorous methodology for airborne laser mapping. *International Calibration and Orientation Workshop EuroCOW*. Castelldefels, Spain.
- Gonsalves, M.O. (2010a). Using a dynamic ocean surface to perform a geometric calibration of a bathymetric lidar. *Proceedings of the MTS/IEEE Oceans 2010*. New Jersey: IEEE.
- Gonsalves, M.O. (2010b). *A Comprehensive Uncertainty Analysis and Method of Geometric Calibration for a Circular-Scanning Airborne Lidar*, Mississippi: University of Southern Mississippi.
- Habib, A.F., Kersting, A.P., Ruifang, Z., Al-Durgham, M., Kim, C. and Lee, D.C., 2008, Lidar strip adjustment using conjugate linear features in overlapping strips. *Proceedings of the XXI ISPRS Congress. Part B1.* Publisher: ISPRS, pp. 385-390.
- de Hilster, N. (2008). The MRU and SRF Aligned. *Hydro International*, November 2008, pp. 24-28.
- Hughes Clark, J.E. (2003). A reassessment of vessel coordinate systems: what is it that we are really aligning?, US Hydrographic Conference 2003.
- Lowenstein, C. D. (1965). *Computations for transponder navigation*. Scripps Institute of Oceanography, Report No. MPL-U-2/65.
- McKeown, D. L. (1975). Survey techniques for acoustic positioning arrays. *Journal of The Institute of Navigation*, 22(1), pp. 59-67.
- National Oceanic and Atmospheric Administration. (2010a). *Field Procedures Manual*, Maryland: U.S. Department of Commerce.
- National Oceanic and Atmospheric Administration. (2010b). *NOS Hydrographic Surveys Specifications and Deliverables*, Maryland: U.S. Department of Commerce.
- Schenk, T., 2001, *Modeling and analyzing systematic errors in airborne laser scanners* (Technical Notes in Photogrammetry No. 19), Publisher: Ohio State University.
- Vosselman, G. and Dijkman, S., 2001, 3D building model reconstruction from point clouds and ground plans. *Proceedings of the ISPRS Workshop: Vol. XXXIV-3/W4. Land surface mapping and characterization using laser altimetry.* Editor M.A. Hofton, Publisher: ISPRS – 2001, pp. 37-43.

BIOGRAPHY OF THE AUTHOR

LT Michael O. Gonsalves is a NOAA Corps officer currently assigned to NOAA's Remote Sensing Division (National Geodetic Survey), stationed at the Joint Airborne Lidar Bathymetry Technical Center of Expertise (JALBTCX) in Southern Mississippi.

He is presently pursuing his doctorate in Marine Science, with an emphasis in Hydrography, at the University of Southern Mississippi.