

Motion Compensation of AUV-Based Synthetic Aperture Sonar

Daniel A. Cook, James T. Christoff, and Jose E. Fernandez
Naval Surface Warfare Center
Coastal Systems Station, Dahlgren Division
6703 W Hwy 98
Panama City, FL 32407

Abstract—To date, ocean-going synthetic aperture sonar (SAS) systems have been deployed exclusively in a configuration where the sonar instrument is housed in a towed body that receives power from and exchanges information with the vessel to which it is attached. Meanwhile, recent years have witnessed the beginnings of maturity with respect to both SAS and autonomous underwater vehicle (AUV) technologies. In order to move away from the towed sonar paradigm, the Coastal Systems Station has recently taken delivery of and begun using the first AUV-based SAS. The AUV was manufactured by Bluefin Robotics and the sonar used on this vehicle is the existing CSS LF/HF SAS. This transition is not without its challenges, however, as the operation and dynamic behavior of an AUV is different from that of a towed body. In general, the AUV configuration makes the problem of unwanted platform motion more severe and more difficult to solve. This paper discusses motion compensation in the context of initial evaluations of the performance of the CSS AUV-based SAS system.

I. INTRODUCTION

The transition of ocean-going SAS from towed to AUV-borne systems is inevitable. However, only recently have both SAS and AUV technologies advanced sufficiently to allow this transition to begin. In the Spring of 2003, the US Navy Coastal Systems Station (CSS) took delivery of the first AUV-based SAS system, the SAS21/RELIANT. The hardware details of this system are discussed in a companion paper [1]. This paper will briefly discuss the imagery results from the initial testing of the system and the motion compensation problem associated with producing them.

II. AUV VS. TOWED SAS SYSTEMS

Before getting to the details of the CSS system, it would be beneficial to discuss the ways in which AUV platforms behave differently than towed platforms. Since 1996, the CSS LF/HF SAS system has been operated from a 21" diameter towed body with active control surfaces [2], [3]. The overall length of this body varied depending upon the mission configuration, but was on average approximately 25' long. The roll of this towed vehicle was controlled by active tail fins, while its depth was controlled by a wing situated over the center of buoyancy. By comparison, the Bluefin AUV is also 21" in diameter, but it is only 7' long. Moreover, it lacks fins altogether, and control of the AUV is accomplished using an articulated thruster.

The differences just described have significant effects on the dynamic behavior of the respective vehicles. For example,

pitch has become a major concern in the operation of the Bluefin where it was a nonexistent problem with the towed vehicle. Since it has no wing, the Bluefin's thruster must point the vehicle up or down in order to maintain a constant depth in the presence of negative or positive buoyancy, respectively. Although the optimal trim for the AUV would be neutral buoyancy, this is currently impossible to achieve in practice because of the tendency for salinity to vary in near-shore waters. Hence, the average pitch of the vehicle depends on local water conditions. In general, the smaller size and lack of control surfaces cause the Bluefin to be fundamentally less stable compared to the towed vehicle previously used by CSS. However, as experience is gained through working with the system, it is expected that the stability of the AUV will improve.

III. REDUNDANT PHASE CENTER MOTION ESTIMATION

The use of redundant phase centers (or displaced phase centers) is a reliable way to obtain estimates of ping-to-ping changes in range for an array used to construct a synthetic aperture. The range displacements Δ_n are those obtained by cross-correlating a redundant phase center (RPC) return with the appropriate channel from the previous ping which occupied the same place in azimuth as the RPC. The quantity of interest is typically the ping-to-ping range displacement referenced to the center of the array— not the redundant phase center. Thus Δ_n is not, in general, the quantity needed to correct the data. Figure 1 shows that the presence of a yaw angle ϕ_n can cause the distance between the overlapping phase centers, Δ_n , to be different than the distance between array centers for adjacent pings, $\Delta_{C,n}$. Assuming 2d geometry only and that the yaw is known at each ping, the range displacement relative to the array center can be computed using the following equation:

$$\Delta_{C,n} = \Delta_n - d_{\text{fore}} \sin \phi_{n-1} + d_{\text{aft}} \sin \phi_n \quad (1)$$

where d_{fore} and d_{aft} are the distances from the array center to the respective phase centers. The convention is that phase centers forward of the array center have a negative distance while those behind the array center have a positive distance.




Fig. 1. In general, the ping-to-ping range displacement Δ_n measured by the redundant phase center technique (RPC) does not equal the actual range displacement associated with the array center, $\Delta_{C,n}$. Array yaw can lead to misleading estimates of sway. In the figure, the array travels from left to right and the black dots represent the redundant phase centers.




Fig. 2. Once the ping-to-ping range displacements have been removed, a residual yaw ϕ_n about the physical array center remains (a). For small angles, removing the yaw can be done with a linear delay across the array resulting in the proper alignment (b).

In addition, the array is moving from left to right (with increasing n), and positive yaw is taken to be counterclockwise, so the angles shown in Figure 1 are both positive.

With the $\Delta_{C,n}$ known, all channels for a given ping can be shifted in time by

$$\tau_{\text{sway},n} = -\frac{2}{c}\Delta_{C,n} \quad (2)$$

where c is the speed of sound. The result of this action is depicted in Figure 2a. At this point, the array centers are all properly aligned, but there remains a residual yaw for each ping. The final step in correcting the data is to apply another time delay across the n^{th} physical array:

$$\tau_{\text{yaw},m,n} = -\frac{1}{c}d_m \sin \phi_n \quad (3)$$

where d_m is the distance from the array center to the m^{th} phase center (approximately one-half the distance from the array center to the actual sensor element). The negative signs

in the two preceding equations are present to indicate that we want to ‘undo’ the effect of the motion. Also, the factor of 2 in Eq. (2) represents the additional two-way travel time due to the sway $\Delta_{C,n}$. This factor is absent in Eq. (3) because the transmitter is assumed to be located at the center of the receiver array. Hence, the outgoing signal is transmitted from the ‘ideal’ location and the additional path length only applies to the return signal. The redundant phase center technique is discussed further in [4]–[6] with the most recent treatment appearing in [6].

IV. MOTION COMPENSATION INCORPORATING THE DOPPLER VELOCITY LOG

A. Overview

It is generally accepted that the redundant phase center technique is unparalleled in its ability to measure ping-to-ping displacements in range. Moreover, the RPC technique is capable of accounting for phase errors from sources other than motion, such as fluctuations in the properties of the medium. However, the previous section showed that these advantages are tempered by the fact that the motion estimates are affected by the angular orientation of the array.

From the standpoint of synthetic aperture imaging, the types of motion experienced by an AUV are generally more severe than those experienced by a large towed body. In the past, the RPC sway measurements were applied without giving consideration to the yaw of the vehicle. That is, Δ_n was the correction applied at each ping. Most of the time, the results were quite successful because ϕ_n was almost always very small. However, this “blind RPC” is not an appropriate procedure for general motion compensation of AUV-based SAS. The issue then becomes one of obtaining good measurements of the yaw angle so that the true ping-to-ping displacement can be computed. It should be noted that it is possible to estimate relative changes in yaw between pings by using the time-of-arrival differences between adjacent redundant phase centers. However, experience has shown that such estimates are too noisy to be useful with the SAS21 (which is presently configured for two overlapping channels).

The approach chosen was to derive the angle estimate from the Doppler velocity log (DVL) which is an RD Instruments 300kHz Workhorse Navigator. For the SAS21, there was no common time stamp for the various types of recorded data: sonar data, Doppler log, and the inertial navigation system log (INS). The INS (Litton LN250) data had an unknown latency, and was not used for motion compensation for the Spring 2003 testing. On the other hand, the Doppler log was assumed to have no latency and was treated as if it coincided in time with the sonar data. Thus, the DVL was relied upon to establish the vehicle trajectory. This section describes the technique used to make the DVL-based yaw angle correction to the RPC motion estimate.

B. Procedure

Equation (4) expresses the yaw relative to the velocity vector in terms of the DVL measurements, and Figure 3




Fig. 3. Illustration of relevant angles and velocity components in a 2d plane.

illustrates the relationship among the quantities discussed.

$$\phi_{\text{Doppler}} = \arctan \frac{v_{\text{lateral}}}{v_{\text{longitudinal}}} \quad (4)$$

$$|\mathbf{v}| = \sqrt{v_{\text{lateral}}^2 + v_{\text{longitudinal}}^2} \quad (5)$$

The idea of motion compensating synthetic aperture data implies ‘shifting’ the data somehow in order to bring it back into alignment with the ideal straight-line intended collection path (such as the horizontal line shown in Figures 1 & 2). For a rail-mounted system, the intended trajectory is not difficult to visualize: it’s just the rail itself. However, for a freely-navigating AUV, the notion of an ideal intended trajectory is somewhat nebulous and more difficult to establish. The vehicle is programmed to follow a straight track, but the influence of the currents may cause the actual track to differ from the intended thus making the programmed track an undesirable baseline for motion compensation. Therefore, the task is to choose the best line to use as a reference for the subsequent motion compensation.

Strictly speaking, one could use almost any straight line as a baseline for motion compensation. In practice, the choice is not truly arbitrary, as the data are typically corrected by simple shifts back and forth in range (effected by the time-shift property of the Fourier transform). Thus, the best line will be the one for which this type of correction introduces the least amount of error.

Equations (6) show how the vehicle position is computed using the DVL velocity measurements:

$$\begin{aligned} x_n &= |\mathbf{v}_n| \cos \phi_{\text{Doppler},n} dt + x_{n-1} \\ y_n &= |\mathbf{v}_n| \sin \phi_{\text{Doppler},n} dt + y_{n-1}, \end{aligned} \quad (6)$$

where $|\mathbf{v}_n|$ is the magnitude of the velocity at the n^{th} ping and dt is the ping period (the reciprocal of the ping rate in Hz). For the results presented, Equations (6) were computed using




Fig. 4. Trajectory computed using the Doppler velocity log (DVL).

initial values of $(x_0, y_0) = (0, 0)$, and Figure 4 shows a sample trajectory obtained using this method. Note that the horizontal and vertical axes are not to scale. The plot also shows the least-squares fit to the trajectory. This line will henceforth be referred to as the *nominal trajectory*, or NT.

Once the nominal trajectory has been computed, its angle relative to the x -axis, ϕ_{NT} , is subtracted from the measured yaw angle yielding the yaw of the array relative to the NT at each ping:

$$\phi_n = \phi_{\text{Doppler},n} - \phi_{\text{NT}}. \quad (7)$$

At this point the vehicle sway (i.e., its range displacement) perpendicular to the NT can be computed according to Eq. (1). The data are corrected by shifting the sonar returns by a time $\tau_{m,n}$ given by:

$$\tau_{m,n} = \tau_{\text{sway},n} + \tau_{\text{yaw},m,n}. \quad (8)$$

Recall that n is the ping index while m indexes the elements of the actual physical array.

Close examination of this technique will reveal its somewhat *ad hoc* nature. For example, the velocity components measured by the DVL are always relative to the vehicle centerline; thus the ‘natural’ coordinate system for the DVL velocities actually changes from ping-to-ping. The required coordinate transformation is not accounted for in this paper. As a result, the method described here is useful primarily as an augmentation to RPC motion compensation and cannot be expected to provide a comprehensive solution to the motion compensation problem. As the SAS21/RELIANT is developed further, more types of motion measurement will be available and a more robust and complete motion compensation scheme will be developed.




Fig. 5. Image of cylinder and sphere with no motion compensation applied to the synthetic aperture data.




Fig. 7. Image of cylinder and sphere with yaw-corrected RPC motion compensation.




Fig. 6. Image of cylinder and sphere with blind RPC motion compensation.




Fig. 8. Image of cylinder and sphere with yaw-corrected RPC motion compensation. This image is identical to Figure 7 except that a tapered (Hanning) window has been applied in the frequency domain.

V. RESULTS

The technique described in the previous section was applied to the SAS data collected by the SAS21 AUV-based system in the Spring of 2003. Figures 5–7 show the ‘progression’ of image improvement that comes from: doing no motion compensation, applying the blind RPC correction, and from augmenting the RPC technique with yaw angle measurements as described in the previous section. Figure 8 is the same as Figure 7 except that a Hanning window was applied in the frequency domain as opposed to the rectangular weighting used in the previous three images. All images were created using the $\omega - k$ (or wavenumber) algorithm for image reconstruction [7], [8].

This set of images depicts a cylindrical mine-like object situated near a fluid-filled sphere designed to behave as a point scatterer. Image quality is indicated by the sharpness of the sphere and the highlights associated with corners and protrusions of the cylinder. Image quality can also be gauged by the definition and contrast of the shadow behind (downrange of) the cylinder. By these measures, the technique described improved upon the results obtained by the blind RPC.

VI. CONCLUSIONS

A conspicuous omission from this discussion of AUV-based SAS motion compensation is a treatment of the full three-dimensional problem. Although the SAS21/RELIANT vehicle was equipped with an inertial navigation system during the Spring 2003 testing, the recorded INS log lacked a usable time stamp with which to align the INS data with the sonar data. It was known that a latency existed in the INS log, but it could not be quantified. Hence, all attempts at motion compensation were restricted to using the sonar data and the DVL which appeared to have no problems with latency. Future tests are expected to incorporate a common time stamp on all data. The suitability of the INS as an aid to motion compensation will be studied at that time.

The question naturally arises of whether or not the full 3d case needs to be considered at all. Figure 9 shows typical INS angular measurements over the course of a 50-meter SAS data collection run. This plot clearly shows that the physical array used to construct the synthetic aperture undergoes fairly large variations in yaw and pitch as one might expect. The average pitch, for example, is roughly -8° indicating that the vehicle maintained a ‘nose-down’ orientation in order to counteract its positive buoyancy. Likewise, a yaw with nonzero mean would imply that the vehicle faced into a cross-current in order to follow its programmed track.

If uncorrected, these angular errors can cause grating lobes and defocusing. Additionally, the roll component is apparently small, but is of such a high frequency that it could conceivably cause a general blurring in the final SAS image. The effect of these angular errors is difficult to quantify as it is a function of range (*i.e.*, synthetic aperture length) and sonar parameters (beamwidth and wavelength). CSS is currently investigating




Fig. 9. Sample angular data recorded by the SAS21/RELIANT inertial navigation system.

these motion effects using actual INS motion data as input to an ideal point-target simulation.

Another area of investigation suggested by the results presented is that of optimizing the AUV control algorithms such that a vehicle carrying a SAS would be more likely to follow a ‘benign’ track if the local conditions didn’t allow it to closely follow the intended track. That is to say, the AUV could be programmed to deviate from the intended track in order to pursue a different trajectory that would be closer to an ideal straight line path, albeit in a different direction.

The results shown in Figures 5–8 are representative of the image quality improvement achieved for the limited amount of data analyzed for the Spring 2003 SAS21/RELIANT testing. The deficiencies found in the system during these tests are currently being evaluated and corrected. Subsequent tests are envisioned to include more complex vehicle trajectories thus inducing motion problems beyond the capability of the technique presented here. For example, additional information will be required in order to put the DVL measurements into a common coordinate system. Furthermore, the overall accuracy of the DVL unit will be considered and weighed against other methods of motion measurement.

Overall, the imagery from these initial tests does not match the quality and consistency of the imagery obtained from SAS21 in its towed-body configuration. However, the results achieved are adequate and encouraging given the complexity and novelty of combining SAS and AUV technologies. This marriage marks a milestone in underwater acoustic imaging technology.

ACKNOWLEDGMENT

This work was supported by the Office of Naval Research.

REFERENCES

- [1] J. E. Fernandez, J. T. Christoff, D. A. Cook, "Synthetic aperture sonar on AUV," MTS/IEEE Oceans 2003 Proceedings, San Diego, CA, 2003.
- [2] G. Sammelmann, J. E. Fernandez, and J. T. Christoff, "High frequency/low frequency synthetic aperture sonar," SPIE, 1997.
- [3] R. W. Sheriff, T. C. Montgomery, J. E. Fernandez, and J. T. Christoff, "High resolution synthetic aperture sonar," UUV Showcase 98 Conference, 1998.
- [4] R. S. Raven, "Electronic stabilization for displaced phase center systems," U. S. Patent 4244036, January 1981.
- [5] R. W. Sheriff, "Synthetic aperture beamforming with automatic phase compensation for high-frequency sonars," IEEE Proceedings of the 1992 Symposium on Autonomous Underwater Vehicle Technology, pp. 236 – 245, June 1992.
- [6] A. Bellettini and M. A. Pinto, "Theoretical accuracy of synthetic aperture sonar miconavigation using a displaced phase-center antenna," *IEEE Journal of Oceanic Engineering*, Vol 27, No. 4, October 2002.
- [7] R. Bamler, "A comparison of range-Doppler and wavenumber domain SAS focusing algorithms," IEEE Transactions on Geoscience and Remote Sensing, vol. 30, pp. 706–713, July 1992.
- [8] D. W. Hawkins, "Synthetic aperture imaging algorithms: with application to wide-bandwidth sonar," Ph.D. Thesis, University of Canterbury, Christchurch, New Zealand, 1996.