VECTOR–SENSOR BEAMFORMING FOR AUTONOMOUS GLIDER NETWORKS

A Thesis Proposal
Presented to
The Academic Faculty

by

Brendan Nichols

In Partial Fulfillment
of the Requirements for the Degree
Master of Science in Mechanical Engineering

Georgia Institute of Technology
December 2015

VECTOR–SENSOR BEAMFORMING FOR AUTONOMOUS GLIDER NETWORKS

Approved by:
Dr. Karim G. Sabra, Advisor
Mechanical Engineering
Georgia Institute of Technology

David Trivett
Mechanical Engineering
Georgia Institute of Technology

François Guillot
Mechanical Engineering
Georgia Institute of Technology

Date Approved: 26 June 2015

ACKNOWLEDGEMENTS

This research was sponsored by the Ocean Acoustics Program of the Office of Naval Research under Grant No. N00014-14-1-0159. I would foremost like to thank my advisor Dr. Karim Sabra for his guidance, input and dedication to this work.

I would also like to acknowledge the contributions of Dave Trivett, James Martin, Dr. Kevin B. Smith and the crews of the R/V John Martin and R/V Fulmar for the hardware design and at–sea data collection.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

SUMMARY

Detection and localization of sound sources in an ocean environment can be achieved with a distributed array of passive acoustic sensors. Utilizing an array of autonomous littoral gliders, which offer long–term and quiet operation, and vector sensors, which measure both acoustic pressure and particle velocity, the array's localization performance can be improved. However, vector sensors are susceptible to errors induced by acoustic noise, and autonomous gliders as a sensor platform introduce positional errors. Through both simulations and at–sea data, the localization performance of four processing methods are evaluated under various noisy conditions. In both simulated and at–sea data results, a new cross–coherent method outperforms traditional methods by mitigating the effects of acoustic noise, provided sufficient positional accuracy of the array elements.

CHAPTER I

INTRODUCTION

Environmental and defense agencies are often interested in locating and tracking underwater objects. In the case of marine life, one is often interested in tracking objects such as whales or dolphins. In the case of a defense agency, one is more concerned with detecting foreign threats in an area. Either problem usually simplifies to detecting an acoustic source and then tracking, or localizing, the source.

Array signal processing, also known as beamforming, is a method of combining data from an array of sensors to detect and locate acoustic sources. Traditional arrays are comprised of hydrophones, which measure only acoustic pressure. Vector sensors, which measure both acoustic pressure and particle velocity, can improve array performance as they can infer the direction of acoustic intensity, which is typically parallel to the direction of propagation. Compared to a hydrophone array, vector sensor arrays offer improved signal detection, higher directionality for localization, and less ambiguity. As vector sensors have become smaller and cheaper to produce, they are finding their way into more array processing applications.

One application of vector sensor arrays is to mount the sensors on a mobile platform such as a littoral glider. Littoral gliders are autonomous underwater vehicles capable of long–term deployment in shallow coastal regions. The gliders move by changing buoyancy and gliding forward while either sinking or rising. This method of movement is much more silent than actively driven propulsion. Having a quiet, mobile platform capable of passively monitoring an area offers many benefits over a stationary or fixed

array; however, much of the current literature concerning vector sensor beamforming assumes a fixed array.

The objective of this work is to investigate the feasibility of extending the current vector sensor array processing formulations to a mobile array. It also aims to compare the performance of the array using different beamforming techniques. The proposed approach will be demonstrated using both simulated data and two at–sea datasets.

CHAPTER II

BACKGROUND

2.1 Vector Sensor Beamforming

Array signal processing is a field that concerns the processing of a spatially distributed array of sensors. Beamforming refers to the process of spatially filtering array data to obtain propagation direction information. The first uses of beamforming in acoustics focused primarily on source direction in narrowband hydrophone (scalar) arrays (Van Trees, 2004, Ch. 1-4). The field has since seen advances in many areas, such as broadband processing, optimal detection and adaptive methods (Van Trees, 2004, Ch. 7).

Vector sensor beamforming followed a similar progression, beginning with simple extensions of the corresponding hyrdrophone array methods. The first developments introduced an intensity–based approach, which identified the angle of arrival of a plane wave by estimating the product of pressure and velocity, resulting in the direction of energy propagation (Nehorai and Paldi, 1994). A second method, which followed more closely the concepts of hydrophone array processing, utilizes an array steering vector, or weight vector, to spatially filter the array data (Hawkes and Nehorai, 1998). As in the scalar array case, the weight vector takes on a form that mimics the propagation characteristics experienced by the array when sound arrives at a particular angle.

A notable difference in the array steering vector for vector sensors, however, is that they use both the relative delay information induced by propagation delay between

sensors as well as the directional response of the sensor to an impinging wave. Having both the propagation delay and directional information allows the weight vector to distinguish the direction of arrival (DOA) of a source more accurately than an identical hydrophone array (Hawkes and Nehorai, 1998). An additional benefit of vector sensors is the increased number of measurements due to the added velocity components (DSpain et al., 2006). This effectively increases the signal to noise ratio (SNR) by measuring the "signal" more times. Lastly, a vector sensor array also decreases ambiguities experienced by a corresponding hydrophone array (Cox and Lai, 2009). For example, a linear hydrophone array would exhibit a "cone of ambiguity", where any source located on the surface of a certain cone could have generated the array data in question. An otherwise identical array with vector sensors, however, would leverage the directionality of the sensors to pinpoint the exact direction of the source. Hydrophone arrays can also exhibit grating lobes: false peaks that result from aliasing in the wavenumber domain by not spacing array elements more closely than half–wavelengths (Cox and Lai, 2009). Vector sensors on the other hand, can diminish the effect of a grating lobe simply because sound coming from that direction will not have the correct velocity components seen by the true source direction (Chen and Zhao, 2004). Therefore, a vector sensor array offers many theoretical improvements over a hydrophone array.

In more realistic scenarios, however, vector sensors experience a wide variety of noise–induced problems. Many sensors do not measure particle velocity directly, but rather acceleration, which can be noisy to integrate to obtain velocity (Nehorai and Paldi, 1994). Vector sensors are also very sensitive to flow noise, caused by any movement of the sensor through water or currents flowing around the sensor (Lauchle et al., 2002). D'Spain et al. have developed an optimal weight vector which accounts for the increased amount of noise experienced by the velocity elements. Other adaptive methods exist, which attempt to leverage the knowledge of the noise field to improve

array performance (Hawkes and Nehorai, 2003; DSpain et al., 2006). One area of vector sensor processing which has not been fully developed yet pertains to adaptive weighting of the covariance matrix itself.

Another area of vector sensor processing which has not been fully explored is the addition of positional inaccuracies of the sensors within the array. Much of the literature on vector sensor arrays assumes perfect knowledge of the array position, orientation, and individual sensor location. This is most common, since many arrays are designed and constructed to specification. Some literature exists on calibrating the array to obtain precise sensor position information, but relies on a calibration procedure before its use (Rockah and Schultheiss, 1987a,b). The most relevant research on positional inaccuracies so far only relates to the performance of scalar arrays with random positional errors added to sensor position (Schultheiss, 1980). If the vector sensors in an array are to be mounted to an autonomous mobile platform, the exact position and orientation of each sensor will not be determinate, and the effects on array performance should be understood.

2.2 Littoral Glider Array Platform

An attractive platform for an autonomous array of vector sensors is the littoral glider (see Figure 1). The main benefits of a glider network include array mobility, long–term autonomous deployment, and quiet operation. A littoral glider is designed to operate in shallower coastal waters, using buoyancy as its means of propulsion, making it both silent and energy efficient. When surfaced, the glider is positively buoyant and is capable of obtaining a GPS fix. To move forward, the glider pulls in a piston whose outer face is in contact with the surrounding water to become negatively buoyant and thus start sinking. While sinking, wings on the glider help propel it forward. The glider changes its heading by rolling and pitching, which is achieved by moving its

center of mass relative to the center of buoyancy. When reaching its maximum depth, the glider pushes out its piston to become positively buoyant, gliding forward again on the ascent (Imlach and Mahr, 2012).

Figure 1: An autonomous littoral glider by Alaska Native Technologies (ANT–LLC), fitted with a vector sensor mounted inside a cage approximately one meter off the nose.

Despite having control over its orientation and buoyancy, the glider's exact position underwater is not always known. It can only be inferred by knowing the descent rate, pitch angle and heading history, which is integrated to estimate the current position. This method of positioning, called "dead reckoning", suffers from compounded error during integration, and lack of current/drift information, and thus is not very accurate or reliable. On the surface, however, the glider is able to obtain its position from GPS satellites quite accurately. Acoustically, the surface is not an ideal place for a vector sensor array since the sea–surface is quite noisy and a floating glider will tend to oscillate through the waves. Wave–generated surface noise and the flow of water past a vector sensor induce large amounts of noise, which can diminish a vector sensor array's performance, even below that of a single hydrophone if not accounted for in

processing (DSpain et al., 2006). The last concern of gliders as an array platform is the noise generated when the glider moves its internal mass (e.g. batteries). Though still quieter than an active propulsion system, the noise generated can drown out nearly all signals of interest. With these drawbacks carefully managed, a glider has the potential to perform well as a platform for an autonomous vector sensor array.

2.3 Problem Formulation

Given an array of gliders, each with positional inaccuracies and subjected to noisy environments, the main goal is to identify and track the location of a source of interest (see Figure 2). The purpose of this work is, firstly, to determine whether the proposed problem can be feasibly solved in a real ocean environment. Secondly, the localization performance of different beamforming methods will be compared, using both simulated and at–sea data.

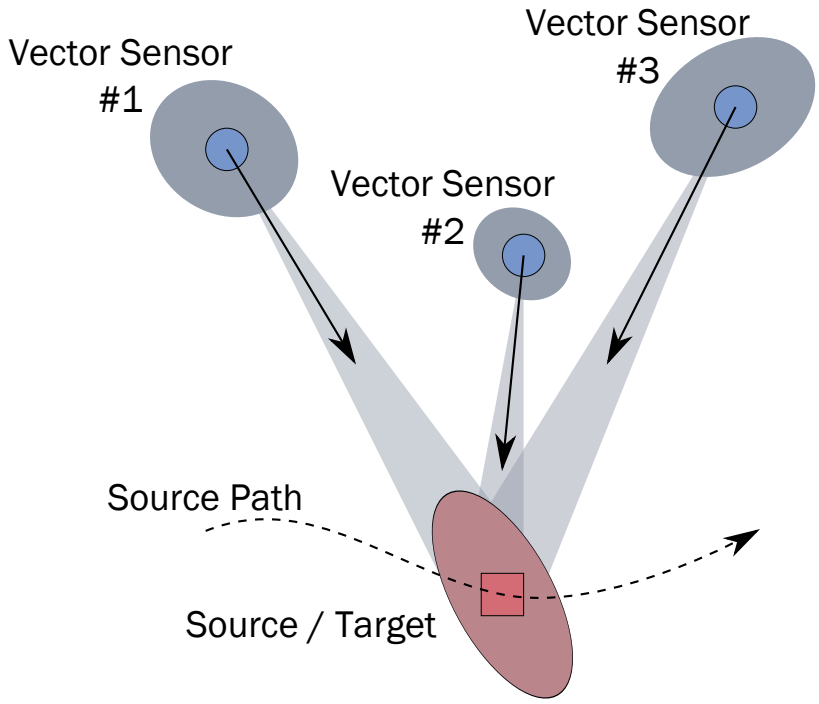

Figure 2: Schematic of the localization problem. Gliders (blue circles) are situated around a source (red square), which is following an unknown path (dashed arrow). The gliders, having positional errors (grey ellipses), attempt to locate the source's position with as little error as possible (red ellipse).

CHAPTER III

THEORY

3.1 Coordinate Systems and Conventions

The coordinate system used in the subsequent sections is a geodetic East-North-Up (ENU) system[1], where the local tangent plane to the Earth's surface represents the XY–plane. The positive X–axis points East, the positive Y–axis points North, and the positive Z–axis points up (see Figure 3). The Earth's magnetic field locally points nearly northward, but with declination (variation from true North), and inclination (variation from parallel to the local tangent plane). The gravity vector is assumed to always point downward, perpendicular to the local tangent plane.

[1]The coordinate system more common to avionics and underwater vehicles is NED, and will be used in future publications.

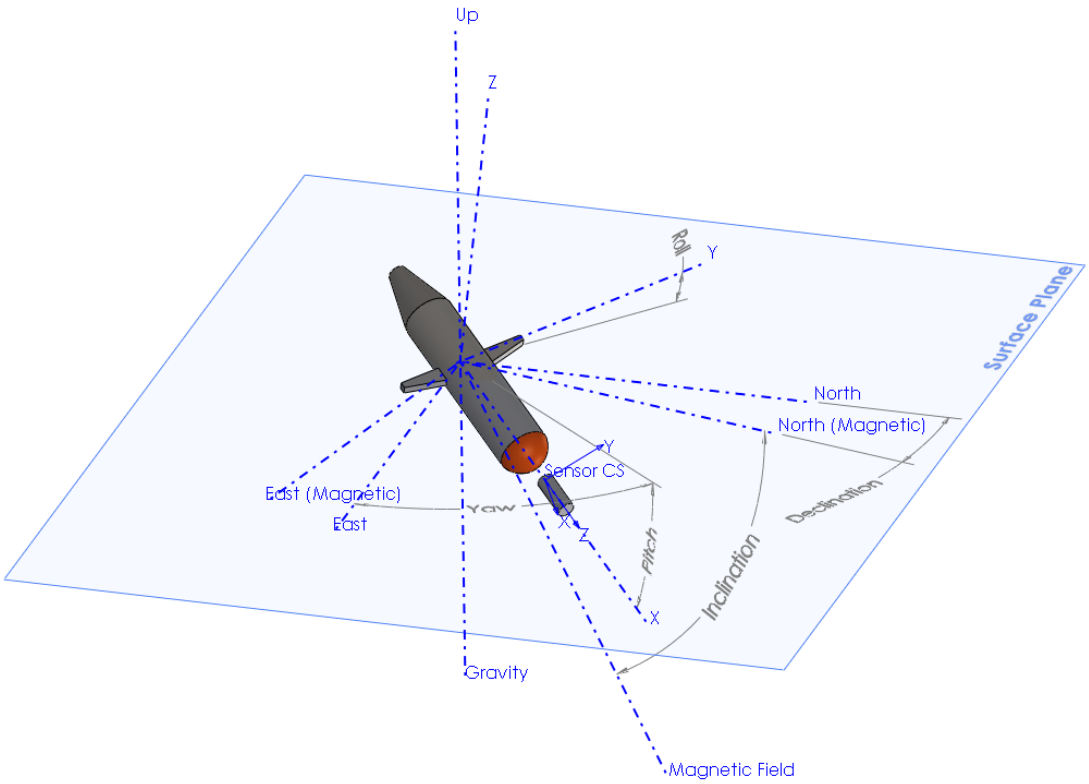

Figure 3: Geodetic coordinate system with angles used to determine glider orientation. Angles defining the local magnetic field direction (declination and inclination) are also shown.

Any glider within this local geodetic coordinate system can use the gravity and magnetic field vectors to determine its orientation. The orientation is specified by yaw (heading), pitch and roll angles, which are defined by positive right–hand intrinsic rotations about the z–, y– and x–axes (see Figure 3). The glider coordinate system is then such that the positive x–axis points towards the nose of the glider, collinear with its cylindrical axis. The positive y–axis lies in the plane of the gliding wings, pointing toward the port (i.e. left) side of the glider. The positive z–axis is perpendicular to both x– and y–axes, forming a right–hand coordinate system.

The vector sensor coordinate system is specific to the VS–301 used in experimental data, where the positive z_s–axis is collinear with the glider x–axis. The exact orientation of the x_s– and y_s–axes is variable, but forms an angle β between the sensor x_s–axis and the glider -z–axis (see Figure 4).

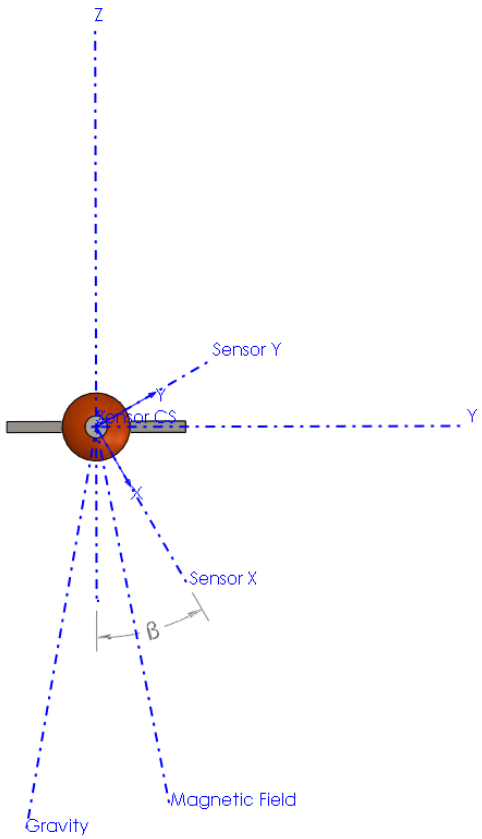

Figure 4: Vector sensor coordinate system with respect to the glider coordinate system. The viewpoint is looking from the front of the glider down its axis.

Lastly, the angle convention used for vectors within the geodetic coordinate system assumes an azimuth angle ($\theta \in [0, 2\pi)$) which is defined positive counter–clockwise from the positive x–axis; the elevation angle ($\phi \in [-\pi/2, \pi/2]$) is defined as positive towards the positive z–axis (see Figure 5).

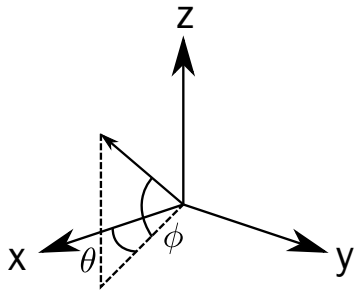

Figure 5: Angle convention used for vectors within a coordinate system. The azimuth is given by θ and the elevation by ϕ.

3.2 Vector Sensor Processing using Covariance Matrix Weighting

Considering N distributed vector sensors, the time–domain data vector obtained from the array is defined as

$$\mathbf{d}(t) = \Big[p_1(t), v_{1x}(t), v_{1y}(t), v_{1z}(t), \ldots, p_N(t), v_{Nx}(t), v_{Ny}(t), v_{Nz}(t)\Big]^T \tag{1}$$

which includes the pressure, $p_i(t)$, and three components of velocity, $v_{ix}(t)$, $v_{iy}(t)$, and $v_{iz}(t)$, from the i-th sensor ($i = 1, ..., N$). Here, we assume all desired preprocessing has been performed (see Appendix B). In the frequency–domain, the data vector is defined as

$$\tilde{\mathbf{d}}(\omega) = \Big[p_1(\omega), v_{1x}(\omega), v_{1y}(\omega), v_{1z}(\omega), \ldots, p_N(\omega), v_{Nx}(\omega), v_{Ny}(\omega), v_{Nz}(\omega)\Big]^T. \tag{2}$$

Pairwise correlation of the array element data yields an estimate of the data covariance matrix (DSpain et al., 2006)

$$\mathbf{R} = \tilde{\mathbf{d}}\tilde{\mathbf{d}}^H = \begin{bmatrix} \mathbf{C}_{11} & \mathbf{C}_{12} & \ldots \\ \mathbf{C}_{21} & \mathbf{C}_{22} & \ldots \\ \vdots & \vdots & \ddots \end{bmatrix} \tag{3}$$

where the superscript H denotes conjugate transposition and the explicit dependence on frequency is omitted for clarity. In Equation (3), the 4x4 sub–matrices, $\mathbf{C}_{ij}$ ($i, j =$

$1, ..., N$), are the cross–covariance matrices of data from vector sensors i and j and are defined by

$$\mathbf{C}_{ij} = \begin{bmatrix} p_i p_j^* & p_i v_{jx}^* & p_i v_{jy}^* & p_i v_{jz}^* \\ v_{ix} p_j^* & v_{ix} v_{jx}^* & v_{ix} v_{jy}^* & v_{ix} v_{jz}^* \\ v_{iy} p_j^* & v_{iy} v_{jx}^* & v_{iy} v_{jy}^* & v_{iy} v_{jz}^* \\ v_{iz} p_j^* & v_{iz} v_{jx}^* & v_{iz} v_{jy}^* & v_{iz} v_{jz}^* \end{bmatrix} \tag{4}$$

where the superscript $*$ denotes complex conjugation. Diagonal matrices $\mathbf{C}_{ii}$ are termed the incoherent covariance matrices, as they correspond to correlations amongst pressure and velocity channels of the same vector sensor. Off–diagonal matrices, $\mathbf{C}_{ij}$ where $i \neq j$, are termed coherent covariance matrices as they correspond to correlations across pressure and velocity channels of different vector sensors. The different weighting methods for the covariance matrix can be described by which elements of the covariance matrix are preserved in processing. This selection is achieved by element–wise weighting of the covariance matrix in Equation 3 and is given by

$$\mathbf{R}_w = \mathbf{W} \odot \mathbf{R} \tag{5}$$

where $\mathbf{R}_w$ is the weighted covariance matrix, $\mathbf{W}$ is the weight matrix, and the element–wise (Hadamard) product is denoted by the $\odot$ operator.

The traditional incoherent processing method keeps only the diagonal correlation matrices C_{ii} by using a block–diagonal matrix given by

$$\mathbf{W} = \begin{bmatrix} \mathbb{1}_4 & \mathbb{0}_4 & \dots \\ \mathbb{0}_4 & \mathbb{1}_4 & \dots \\ \vdots & \vdots & \ddots \end{bmatrix} \tag{6}$$

where $\mathbb{1}_n$ denotes a $n \times n$ matrix of ones and $\mathbb{0}_n$ indicates a $n \times n$ matrix of zeros. Incoherent processing effectively treats each sensor individually, maintaining only correlations among pressure and velocity channels of the same sensor. As a result,

local uncorrelated noise strongly affects the incoherent method's weighted covariance matrix.

Second, the coherent processing method uses all information contained within the covariance matrix, and thus the weight matrix is given by the unity matrix

$$\mathbf{W} = \mathbb{1}_{4n}. \tag{7}$$

The coherent method is still affected by local correlated noise at each sensor through the incoherent covariance sub–matrices $\mathbf{C}_{ii}$, but also includes the coherent covariance sub–matrices $\mathbf{C}_{ij}$ $(i \neq j)$ corresponding to cross–sensor correlations.

Third, a cross–coherent processing method is introduced to attempt to reject the effects of uncorrelated noise between array elements by keeping only the off–diagonal correlation sub–matrices $\mathbf{C}_{ij}$ $(i \neq j)$ using a weight matrix given by

$$\mathbf{W} = \begin{bmatrix} \mathbb{0}_4 & \mathbb{1}_4 & \dots \\ \mathbb{1}_4 & \mathbb{0}_4 & \dots \\ \vdots & \vdots & \ddots \end{bmatrix}. \tag{8}$$

Note that the phase of the off–diagonal matrices in $\mathbf{C}_{ij}$ $(i \neq j)$ relates to the relative propagation delay between sensors. Hence, provided sufficient accuracy in the sensor positions, localization is possible using the relative propagation delays from a source. For widely distributed sensors, it is expected that the local sea–surface generated noise is nearly uncorrelated amongst sensors, and thus the off–diagonal covariance matrices $\mathbf{C}_{ij}$ $(i \neq j)$ are less affected than the diagonal matrices $\mathbf{C}_{ii}$ (Hawkes and Nehorai, 2001).

Lastly, a zero–diagonal coherent processing method is introduced to attempt to reject only noise from same sensor elements by zeroing only correlations between like elements. The weight matrix for this method is given by

$$\mathbf{W} = \mathbb{1}_{4n} - \mathbf{I}_{4n} \tag{9}$$

where $\mathbf{I}_n$ denotes a $n \times n$ identity matrix. This method is almost identical to the coherent method, however all autocorrelations (which reside on the diagonal of the covariance matrix) are neglected. In the presence of electronic noise, only the autocorrelations will be largely affected, assuming the noise is nearly uncorrelated amongst sensor elements.

3.3 Vector Sensor Beamforming

Here, a simple free space, time–delay beamformer is implemented to locate a single source of interest. After element–wise weighting with the method of choice, the weighted covariance matrix (Equation 5) is used to form the classical Bartlett beamformer output given by

$$\tilde{B}(\omega, \hat{\mathbf{r}}_\mathbf{s}) = \mathbf{w}^H(\omega, \hat{\mathbf{r}}_\mathbf{s})\mathbf{R}_w(\omega)\mathbf{w}(\omega, \hat{\mathbf{r}}_\mathbf{s}) \tag{10}$$

where $\tilde{B}(\omega, \hat{\mathbf{r}}_\mathbf{s})$ is the beamformer output power for a given frequency ω and estimated source location $\hat{\mathbf{r}}_\mathbf{s}$ (DSpain et al., 2006; Hawkes and Nehorai, 1998). The weight vector, $\mathbf{w}(\omega, \hat{\mathbf{r}}_\mathbf{s})$, is also known as the array steering vector and is formed using knowledge of the phase delay and direction of arriving plane waves at each sensor (Hawkes and Nehorai, 1998). The relative time delay τ_i of a plane wave arriving at sensor i is given by the distance between the source and sensor, divided by the constant sound speed c_0. The direction of propagation is along the vector from the source to the i–th sensor, which has an azimuthal bearing θ_i and elevational bearing ϕ_i (see Figure 6). The weight vector is then given by

$$\begin{aligned}\mathbf{w}(\omega, \hat{\mathbf{r}}_\mathbf{s}) = [&\mathrm{e}^{\jmath\omega\tau_1}, \quad \mathrm{e}^{\jmath\omega\tau_1}\cos\theta_1\cos\phi_1, \quad \mathrm{e}^{\jmath\omega\tau_1}\sin\theta_1\cos\phi_1, \quad \mathrm{e}^{\jmath\omega\tau_1}\sin\phi_1, \\ &\ldots, \quad \mathrm{e}^{\jmath\omega\tau_n}, \quad \mathrm{e}^{\jmath\omega\tau_n}\cos\theta_n\cos\phi_n, \quad \mathrm{e}^{\jmath\omega\tau_n}\sin\theta_n\cos\phi_n, \quad \mathrm{e}^{\jmath\omega\tau_n}\sin\phi_n]^T.\end{aligned} \tag{11}$$

Note that the weight vector used here only slightly differs from that in D'Spain et al., where all sensors are assumed to receive the plane wave from the same direction. In

the mobile array case, this is not a valid assumption, as the gliders may be positioned arbitrarily around the source.

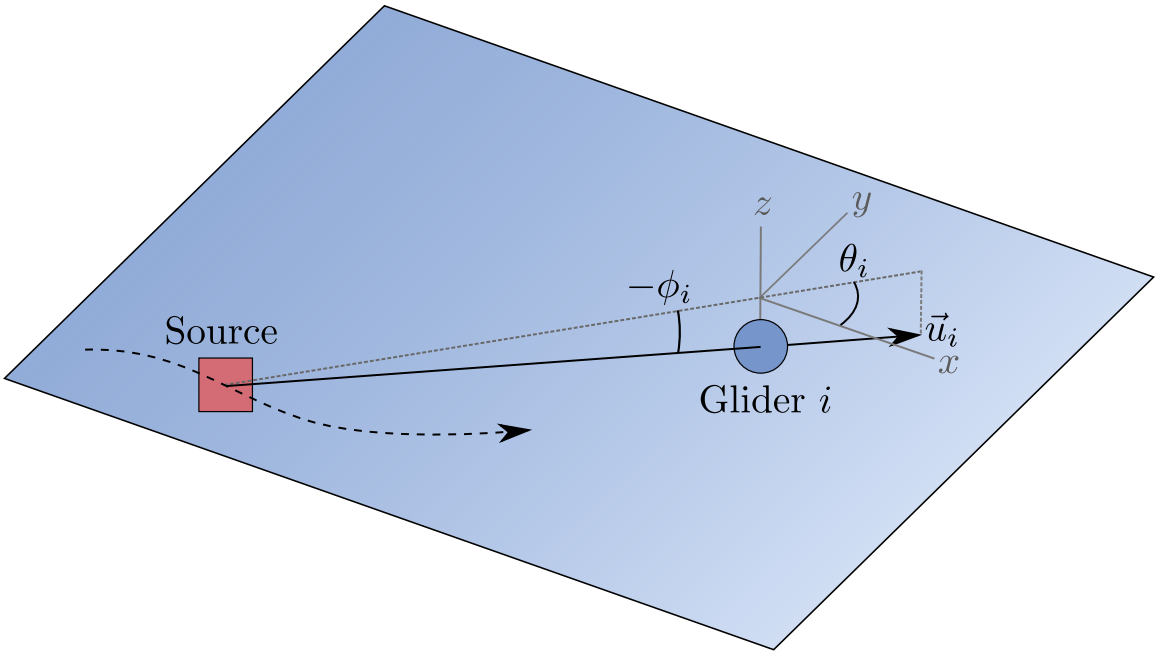

Figure 6: Schematic of propagation of a plane wave from the source to the i–th glider. The direction of propagation is given by $\mathbf{u_i}$, which forms an azimuth angle θ_i and elevation angle ϕ_i.

To estimate the likelihood a source resides at a given location $\hat{\mathbf{r}}_\mathbf{s}$, the time–domain beamformer output, obtained as the inverse Fourier Transform of Equation (10), is windowed and root square–summed to yield an ambiguity surface

$$\Pi(\hat{\mathbf{r}}_\mathbf{s}) = \sqrt{\sum_t \left[f(t)B(t,\hat{\mathbf{r}}_\mathbf{s})\right]^2}, \tag{12}$$

where $f(t)$ is a time–window with width inversely proportional to the signal bandwidth centered around $t = 0$ to select only the main peak of the time–domain beamformer output. A sample ambiguity surface shown in Figure 7 shows values of the ambiguity surface normalized between 0 and 1 as colors where darker colors denote higher likelihood of a source being present.

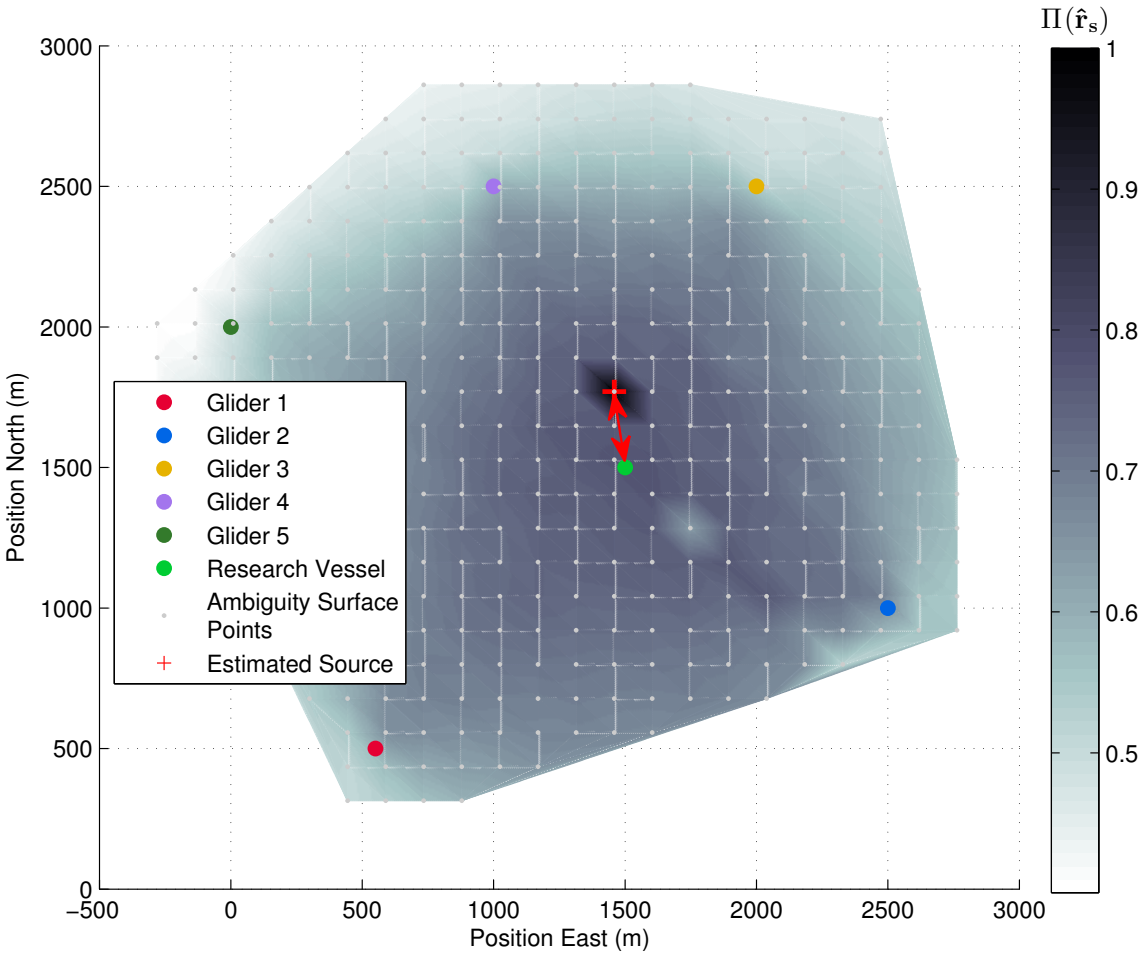

Figure 7: A sample ambiguity surface, where the evaluated points are shown as grey dots. Colors between are interpolated from the vertices, and scaled between 0 and 1, where the maximum value becomes 1. The estimated source position is shown as a red cross, and the associated error shown by a red arrow.

The estimated source location, $\hat{\mathbf{r}}_{\mathbf{pk}}$, is assumed to reside at the peak of the ambiguity surface, which represents the most likely location for the source. The accuracy of the location estimate

$$\epsilon_a = |\hat{\mathbf{r}}_{\mathbf{pk}} - \mathbf{r}_\mathbf{s}| \tag{13}$$

is defined as the distance between the estimated source location, $\hat{\mathbf{r}}_{\mathbf{pk}}$, and the true

source location, $\mathbf{r_s}$. The precision of the ambiguity surface is quantified by the logarithm of the ratio of the peak value to the average non–peak value (i.e. sidelobe). It gives a measure of the confidence in the estimated location and is defined by

$$\epsilon_p = 10 \log_{10} \left(\frac{\Pi(\hat{\mathbf{r}}_{\mathbf{pk}})}{\frac{1}{P-1} \sum_{\hat{\mathbf{r}}_{\mathbf{s}} \neq \hat{\mathbf{r}}_{\mathbf{pk}}} \Pi(\hat{\mathbf{r}}_{\mathbf{s}})} \right) \tag{14}$$

where P is the total number of discretized search points in the ambiguity surface. Together, the accuracy and precision metrics define the localization performance of the vector sensor array using conventional time–delay beamforming. Accuracy quantifies the ability of the array to locate a source reliably. Precision quantifies the ability of the array to detect whether a source is present in the first place, and also how well the array will distinguish two nearby sources. Both are important if one is concerned with detecting an object and then having reliable knowledge of its position.

CHAPTER IV

SIMULATION MODEL

Simulation of arbitrary source and glider configurations were carried out to compare different beamforming methods. For a littoral ocean environment, the depth of gliders and sources relative to the distance between them is assumed to be negligible, and thus all simulations were carried out on the sea–surface (see Figure 8).

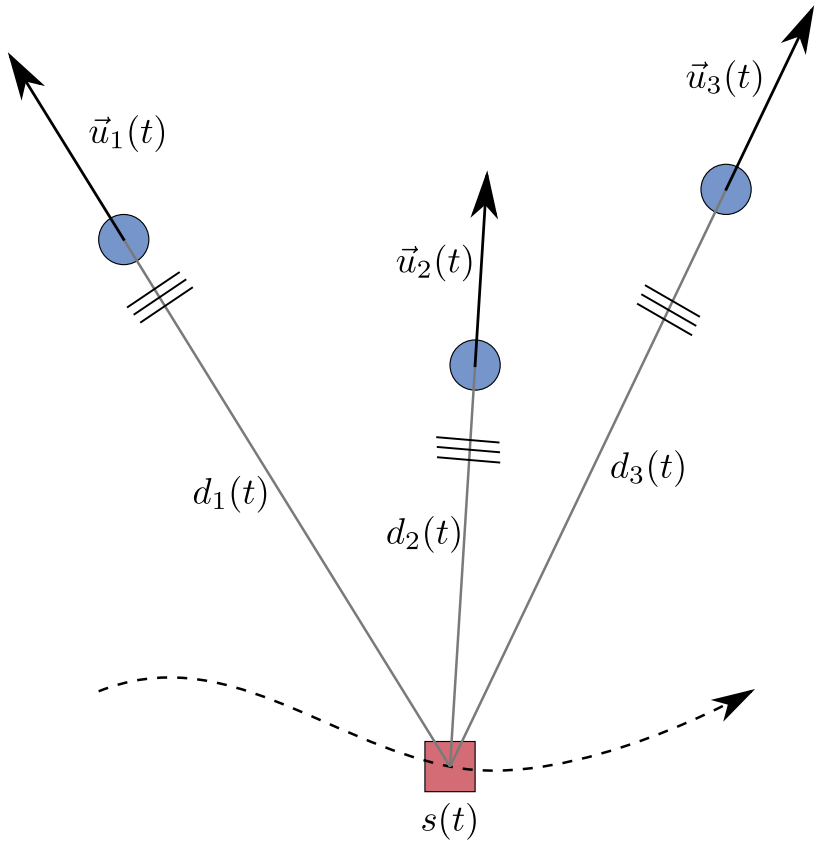

Figure 8: Schematic of simulation parameters. The source (red square), emits the signal $s(t)$, which is assumed to propagate across distance $d_i(t)$ to each glider (blue circles), arriving as plane waves with direction $\mathbf{u_i}(t)$

The surface vessel is treated as a random radiator, broadcasting a signal $s(t)$ defined as Gaussian white noise with amplitude variance arbitrarily set to unity. Assuming homogeneous free–space and unattenuated propagation in the presence of isotropic uncorrelated white noise, the received pressure at the i-th glider is then given by

$$p_i(t) = s\left(t - \frac{d_i(t)}{c_0}\right) + (SNR_p)^{-\frac{1}{2}} n_{ip}(t) \tag{15}$$

where $d_i(t)$ is the distance from the glider to the source at time t, c_0 is the speed of sound, and $n_{ip}(t)$ is additive white Gaussian noise with unit variance which is assumed to be uncorrelated to the random signal $s(t)$ or any other receiver noise $n_{jp}(t)$ $(j \neq i)$. The amplitude of the noise term $n_{ip}(t)$ is scaled such that the ratio of pressure signal power to noise power is given by the pressure signal–to–noise ratio SNR_p. Similarly, assuming a plane wave arrival at the i–th glider (see Figure 8), the received velocity vector is given by

$$\mathbf{v}_i(t) = \frac{\mathbf{u_i}(t)}{\rho_0 c_0} s\left(t - \frac{d_i(t)}{c_0}\right) + (3SNR_v)^{-\frac{1}{2}} \mathbf{n}_{iv}(t) \tag{16}$$

where $\mathbf{u_i}(t)$ is the unit vector pointing from glider i to the source at time t, and ρ_0 is the ambient density of the medium. Each component of the additive noise $\mathbf{n}_{iv}(t)$ is white Gaussian with zero mean and unit variance and is uncorrelated to any other component; its amplitude is scaled such that the ratio of velocity magnitude signal power to noise power is given by the velocity signal–to–noise ratio SNR_v (DSpain et al., 2006).

To implement Equations 15 – 16 numerically, all trajectories are linearly interpolated to the timebase of the trajectory with the most samples. The distances $d_i(t)$ and vectors $\mathbf{u_i}(t)$ are evaluated at each discrete point in time. Also note that the vector $\mathbf{u_i}(t)$ is not evaluated at $t - \tau_i$ where τ_i is the propagation time from the source to the i–th glider. It is assumed that $\mathbf{u_i}(t - \tau_i) \approx \mathbf{u_i}(t)$ (i.e. the change in source and glider positions is negligible during the propagation time).

After simulating acoustic data, the source and glider trajectories are further discretized into windows of length T_b, each starting at the k–th snapshot time t_k (see Figure 9). This window acts as a snapshot of the current source and glider locations and allows localization to be performed over a smaller time–window. Combining the localization results across snapshots provides an estimated source trajectory.

To discretize the source and glider trajectories into windows of length T_b, the assumed location of an object along its trajectory is taken as the mean location over the duration of the window. This "frozen realization" assumption ensures the weight vector (see Equation 11) remains constant over the time window, allowing for simple application in the frequency domain. For most of the following analyses, $T_b = 10$ seconds, and the maximum speed of a source is within 5 m/s. Thus, during one snapshot, the source may move approximately 50 meters, which even at a short distance of 500 meters would be a maximum heading change of 6 degrees.

The final simulation step involves adding positional error to the gliders to simulate measurement errors and uncertainties. While surfaced, the gliders obtain a GPS fix and should have known measurement uncertainty (Grimes, 2008). While submerged, it is unclear what the exact location uncertainty will be, as it depends on many variables involved in the dead reckoning. For simplicity, it is assumed the glider's depth is determinate, although there is likely some error in the pressure sensor reading and assumption of water density involved. Furthermore, the glider's uncertainty in the horizontal plane is assumed to be independent of the glider's direction of travel with equal variance in any direction. Although a simplification of the otherwise complicated positional uncertainty, it is possible to model a worst–case scenario by choosing the largest possible variance experienced by the gliders when submerged.

As both the surfaced and submerged models of positional uncertainty are assumed to be identical, having no variance in the geodetic Z–axis and identical variance

in the geodetic X– and Y–axes, a 2D white Gaussian random variable $\mathbf{e_i}$ which is independent, identically distributed (IID) is added to each glider's location in the horizontal plane. The error vector for the k–th snapshot, $\mathbf{e_i}(t_k)$, is independent from all other error vectors $\mathbf{e_i}(t_j)$ $(j \neq k)$. The standard deviation of the positional noise is denoted E (see Figure 9).

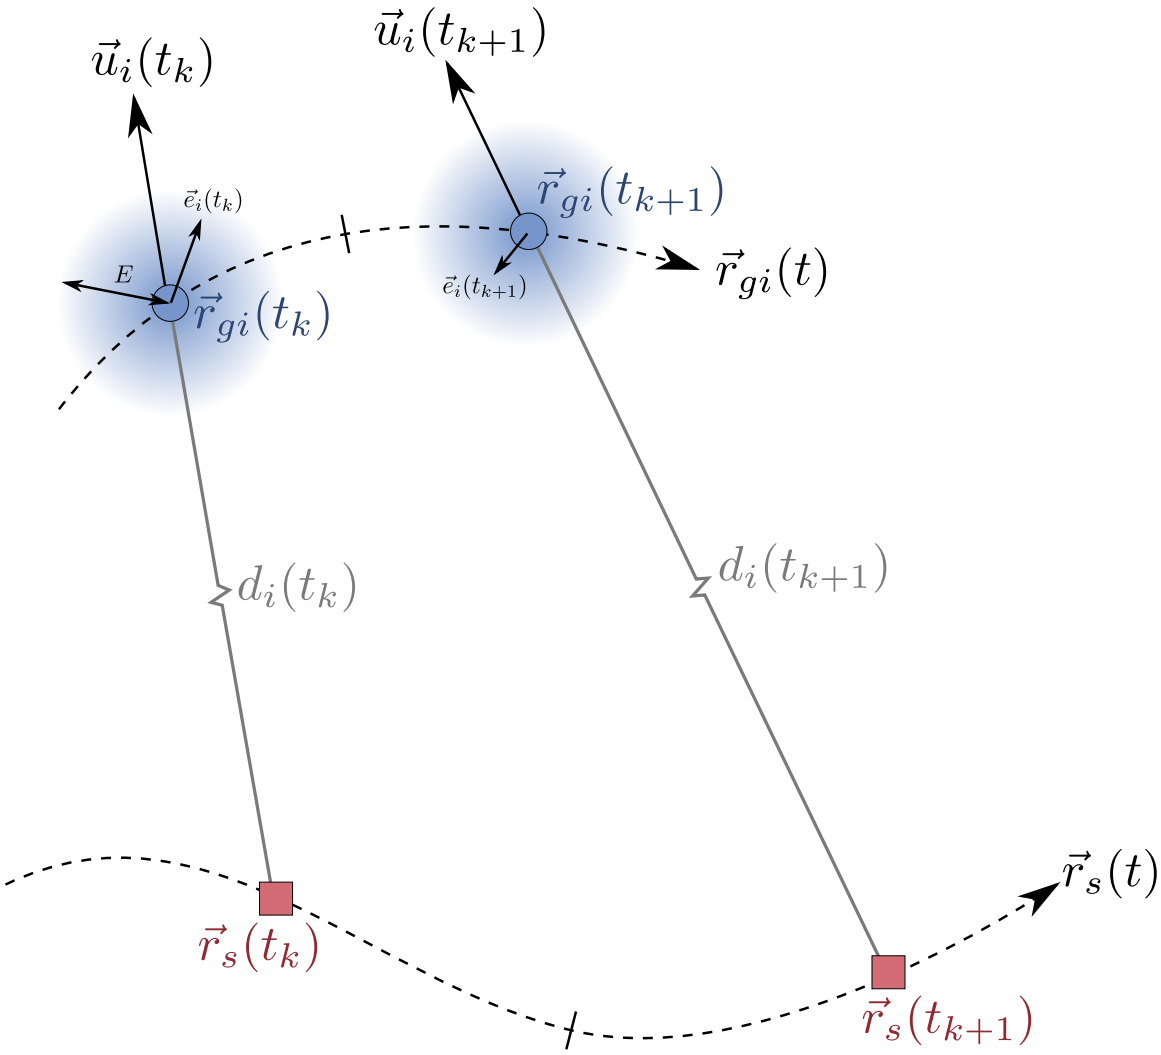

Figure 9: Snapshot simulation schematic. The source position is discretized from its continuous path $\mathbf{r_s}(t)$, with the assumed position at each discrete time t_k being the mean position during the interval. Each glider is discretized from its continuous path $\mathbf{r_{gi}}(t)$, and at each discrete point in time, a random vector $\mathbf{e_i}(t_k)$ with standard deviation E is added. All simulation parameters (distance, propagation direction) are determined from the discretized locations of the source and glider.

CHAPTER V

METHODS I: SIMULATIONS

Simulated data were generated for 10 second long snapshots with a sampling frequency of 5512.5 Hz and filtered in the band 100 – 800 Hz (to match experimental parameters used in Chapter 6). The assumed sound speed and density were 1494 m/s and 992 kg/m^3, respectively. The ambiguity surface was discretized on a 21×21 grid centered on the source location. Lastly, data was preprocessed identically to the at–sea data (see Chapter 6) for consistency, where it was filtered/whitened, clipped and then filtered (see Appendix B).

The remaining parameters, SNR_p, SNR_v and E were varied between trials to compare the performance of the methods described in Equations 6–9. The trials performed are outlined in Table 1.

Table 1: Simulation configurations and the corresponding glider positions, SNR's and positional error (E) used for each trial.

Trial	**Configuration**	SNR_p	SNR_v	E
1	2G-1	20 dB	20 dB	varied
2	3G-1	20 dB	20 dB	varied
3	5G-1	20 dB	20 dB	varied
4	5G-1	-5 dB	-10 dB	varied
5	5G-1	varied	varied	10 m

The configurations used were arbitrarily selected, with the only goal being to separate

the gliders from each other and the source. Having closely spaced sensors is not ideal, as it becomes increasingly difficult to obtain an accurate range estimate (Chen et al., 2002). The configuration names are formatted as NG-i, where N is the number of gliders and i is the instance of that configuration. Configurations 2G-1 and 3G-1 have gliders separated by approximately 2 km from the source and all located to one side (see Figures 10–11). A third glider was added to the 2G-1 configuration to obtain the 3G-1 configuration. The 5G-1 configuration is completely different, with the source centrally located and the five gliders spread evenly around it between 1 and 1.5 km away (see Figure 12).

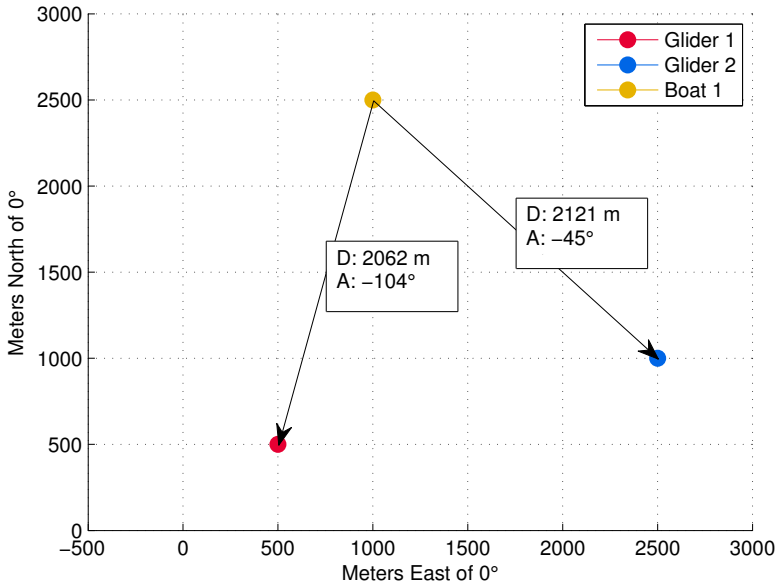

Figure 10: Configuration 2G-1, with two gliders (red and blue dots) towards the south side of the source (yellow dot). The distance (D) and angle (A) to each glider from the source is shown by the arrows with text.

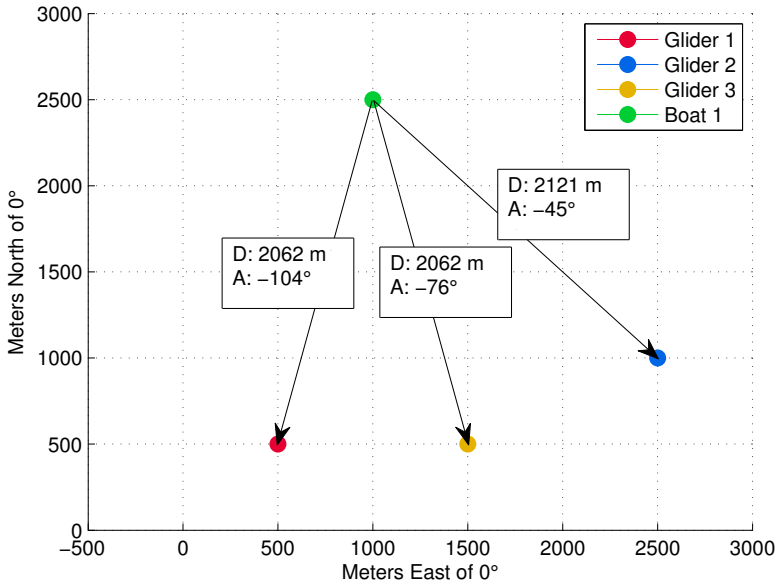

Figure 11: Configuration 3G-1, with three gliders (red, blue, and yellow dots) towards the south side of the source (green dot). The distance (D) and angle (A) to each glider from the source is shown by the arrows with text.

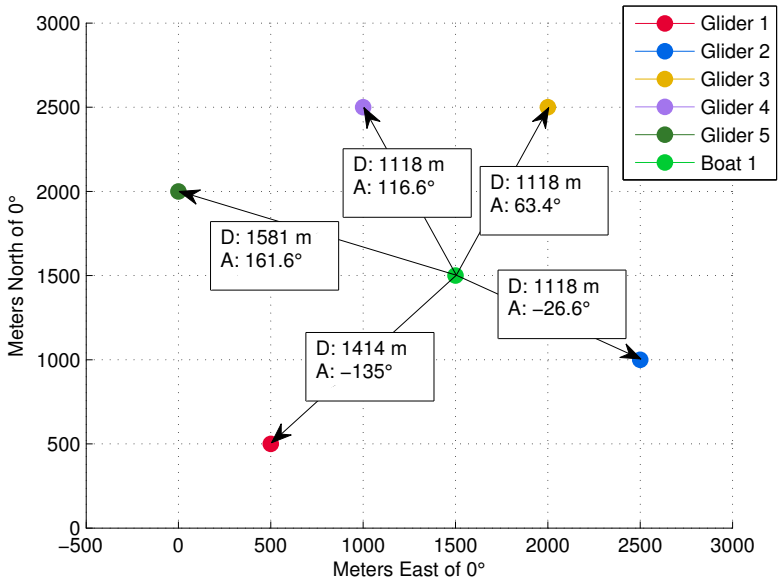

Figure 12: Configuration 5G-1, with five gliders (red, blue, yellow, purple, and dark green dots) surrounding the source (green dot). The distance (D) and angle (A) to each glider from the source is shown by the arrows with text.

For each trial, simulations were repeated 50 times to obtain averages of the localization performance. The resulting accuracy (Equation 13) and precision (Equation 14) were averaged across all 50 realizations and plotted against the standard deviation of positional error E (see Figures 13–20).

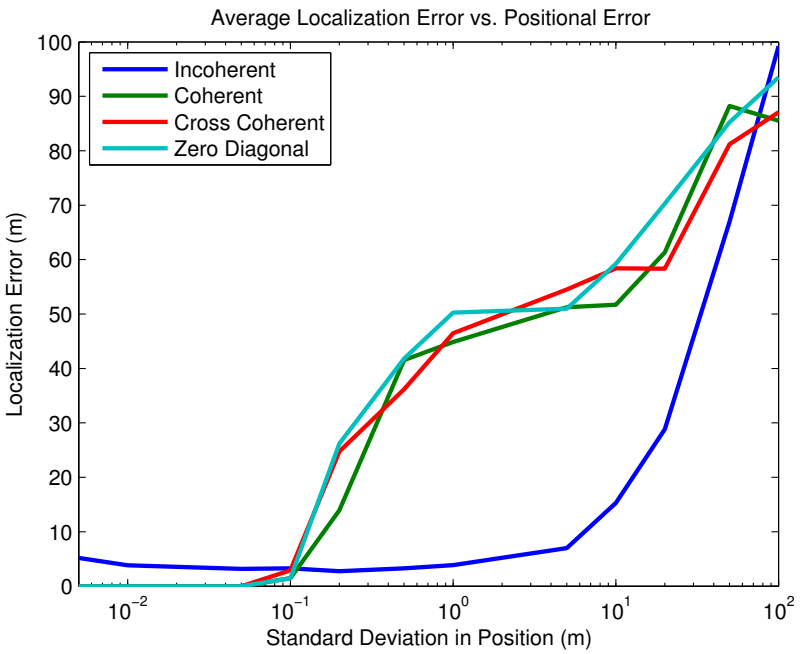

Figure 13: Average localization error plotted versus standard deviation of positional error for the incoherent (blue), coherent (green), cross–coherent (red) and zero–diagonal (cyan) weighting methods. Simulation configuration used was 2G-1, with SNR_p and $SNR_v = 20$ dB.

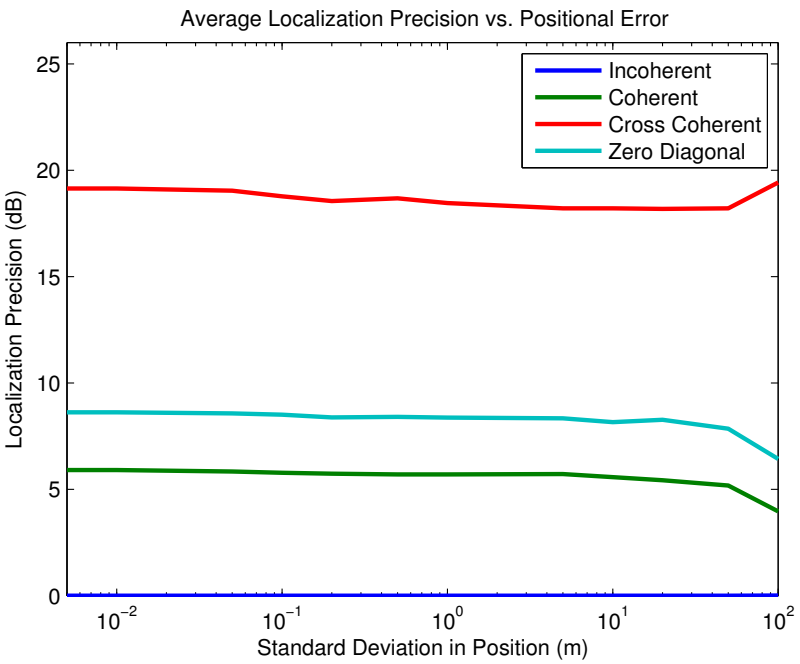

Figure 14: Average localization precision plotted versus standard deviation of positional error for the incoherent (blue), coherent (green), cross–coherent (red) and zero–diagonal (cyan) weighting methods. Simulation configuration used was 2G-1, with SNR_p and $SNR_v = 20$ dB.

The results from Trial 1 show that under low acoustic noise, the incoherent method (Equation 6) actually outperforms all the coherent methods in terms of accuracy. The coherent and its variants (Equations 7–9) all performed similarly, and appear to only have better accuracy for small positional errors. This is in fact due to an error in the preprocessing steps which caused the velocity channels to be very slightly skewed, and the later trials correctly show the incoherent error tending to zero as positional error decreases to zero.

The precision of the cross–coherent method (Equation 8) was drastically higher than all other methods at 19 dB, followed by the zero diagonal at 8 dB, the coherent at 6 dB, and the incoherent at nearly 0 dB. The incoherent method utilizes the cardioid

beampattern of each sensor to locate the source (DSpain et al., 2006), and as such the sidelobes are much larger. The larger sidelobes generate large regions around the source with large ambiguity surface values, leading to poor precision. Under low noise, however, the incoherent method proves to be most accurate, as it is not perturbed by positional errors as strongly as coherent methods. The coherent methods rely more strongly on inter–element time delays which are directly affected by errors in positioning.

Trial 2, which added a third glider to the 2G-1 configuration, was tested using a smaller ambiguity surface having only 121 points (11×11). This test aimed to confirm that the results seen with two gliders would be nearly the same with three gliders, and that the glider configuration would require a more drastic change to see different results with more gliders.

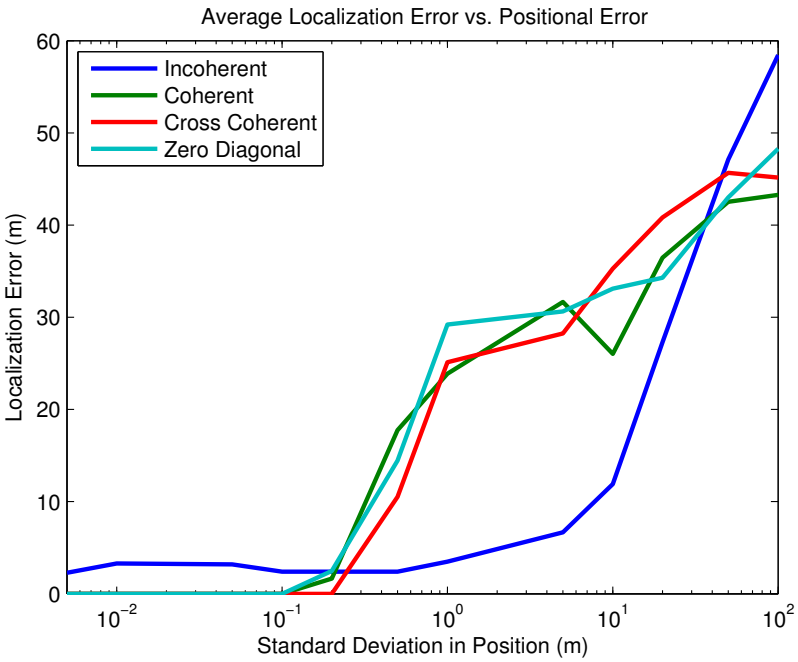

Figure 15: Average localization error plotted versus standard deviation of positional error for the incoherent (blue), coherent (green), cross–coherent (red) and zero–diagonal (cyan) weighting methods. Simulation configuration used was 3G-1, with SNR_p and $SNR_v = 20$ dB.

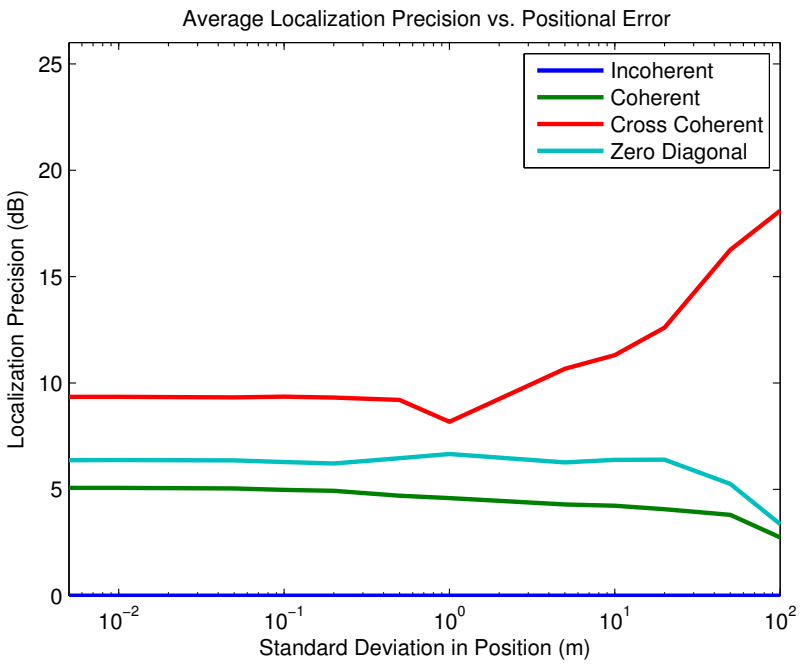

Figure 16: Average localization precision plotted versus standard deviation of positional error for the incoherent (blue), coherent (green), cross–coherent (red) and zero–diagonal (cyan) weighting methods. Simulation configuration used was 3G-1, with SNR_p and $SNR_v = 20$ dB.

The results from Trial 2 are nearly identical to those seen in Trial 1, however the accuracy values quickly hit their limit 60 meters, since the ambiguity surface region only extended 50 meters away from the source. Likewise, some of the precision values increased for large positional errors due to the estimated source location falling on or near the ambiguity surface boundary. Otherwise, the results indicate that the addition of a third glider in roughly the same configuration will not affect the performance of the array under low noise conditions. The third trial, using five gliders, should confirm the performance of the array is roughly unchanged under the low noise case.

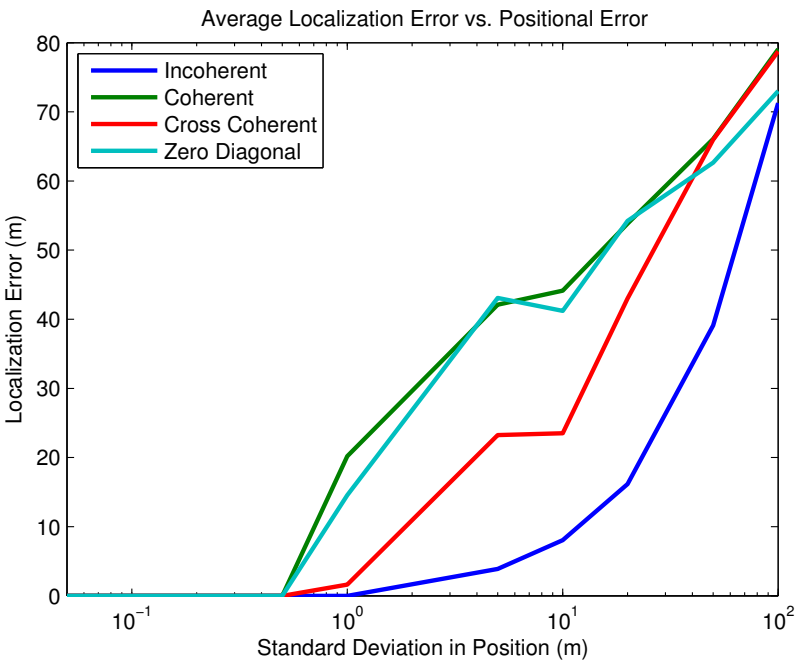

Figure 17: Average localization error plotted versus standard deviation of positional error for the incoherent (blue), coherent (green), cross–coherent (red) and zero–diagonal (cyan) weighting methods. Simulation configuration used was 5G-1, with SNR_p and $SNR_v = 20$ dB.

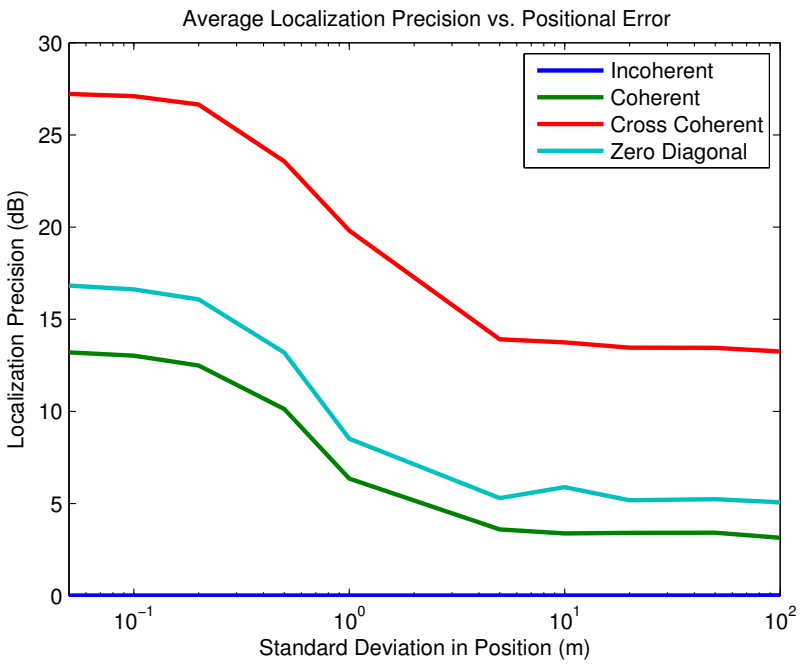

Figure 18: Average localization precision plotted versus standard deviation of positional error for the incoherent (blue), coherent (green), cross–coherent (red) and zero–diagonal (cyan) weighting methods. Simulation configuration used was 5G-1, with SNR_p and $SNR_v = 20$ dB.

Trial 3 again shows similar results to Trials 1 and 2, where the only noticeable difference is slightly better accuracy obtained from the cross–coherent method. This is due to the addition of many more gliders each with a widely differing vantage point of the source. The iso–delay hyperbolas corresponding to the differential time delay between any two gliders intersect in a more reliable fashion than with two or three gliders. Since the cross–coherent method is effectively using this method to locate the source, a notable improvement in accuracy results.

The precision with five gliders shows slightly different behavior than with two or three gliders, where the coherent methods have a clear transition from higher precision to lower precision. This is, again, likely due to the intersections of iso–delay hyperbolas

becoming more spread out around the source as each glider is perturbed from its nominal location.

The fourth trial was conducted under high acoustic noise, with SNR_p = -5 dB and SNR_v = -10 dB. The velocity noise was increased relative to the pressure noise since it is common for the velocity channels to experience higher amounts of noise (see Chapter 2.1). All other parameters were identical to Trial 3 in order to draw comparisons between the low and high noise cases.

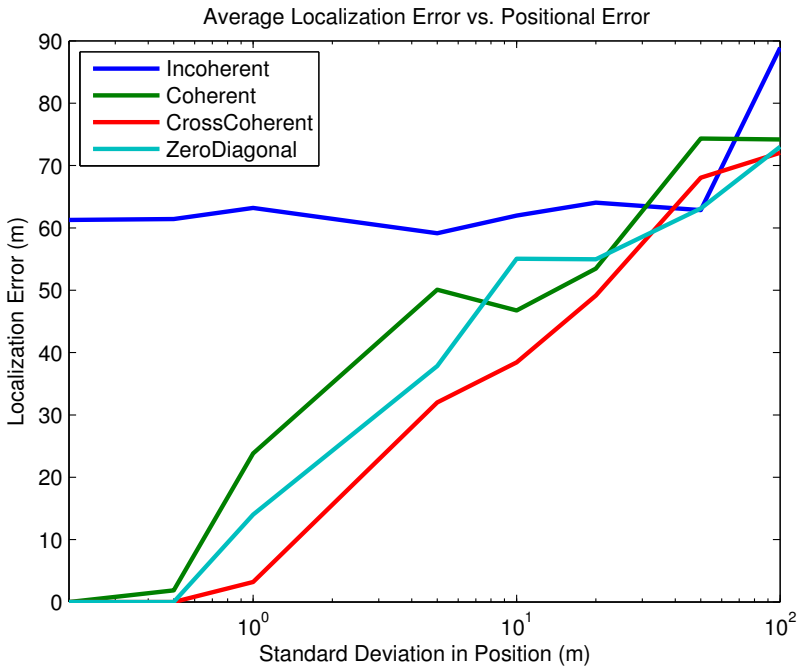

Figure 19: Average localization error plotted versus standard deviation of positional error for the incoherent (blue), coherent (green), cross–coherent (red) and zero–diagonal (cyan) weighting methods. Simulation configuration used was 5G-1, with SNR_p = -5 dB and SNR_v = -10 dB.

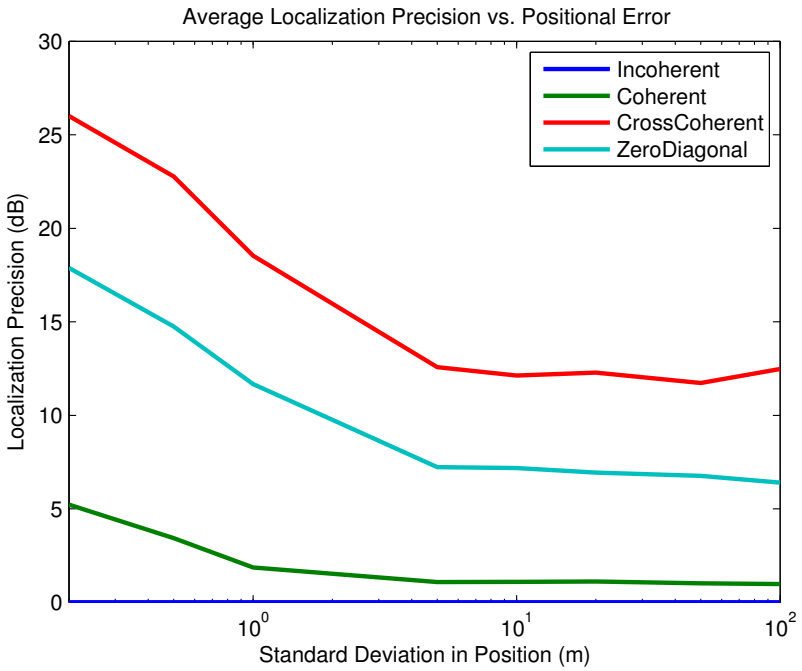

Figure 20: Average localization precision plotted versus standard deviation of positional error for the incoherent (blue), coherent (green), cross–coherent (red) and zero–diagonal (cyan) weighting methods. Simulation configuration used was 5G-1, with SNR_p = -5 dB and SNR_v = -10 dB.

The accuracy results show that the performance of the incoherent method is dramatically affected by acoustic noise. Regardless of positional error (up to a point), the incoherent method accuracy is nearly a constant 60 meters. The acoustic noise is likely the only contributing factor in the error until positional error becomes incredibly large. The coherent methods show very little decrease in accuracy from the low noise case in Trial 3. The cross–coherent and zero diagonal methods both weight the covariance matrix such that the effects of high acoustic noise are mitigated, and this is clearly seen in the accuracy results of this trial.

As for the precision, similar results are seen as those in Trial 3. All coherent methods saw an overall decrease in precision, as the effective "noise floor" of the ambiguity

surface was increased with added acoustic noise. The area of transition from higher precision to lower precision also appears to have moved left, indicating the iso–delay hyperbola intersection breaks down at lower positional errors with higher acoustic noise present.

Trials 3 and 4 show the endpoints of the acoustic noise spectrum, where in low noise cases the incoherent method offers high accuracy but low precision, and in the high noise cases the cross–coherent method offers the best accuracy and precision. To further investigate the transition, Trial 5 aims to demonstrate how the same localization performance parameters are affected by changing SNR. The positional error here was fixed to 10 meters, which is higher than the stated 3.9 meter GPS accuracy (Grimes, 2008), but takes into account the possibility the gliders are more inaccurate when submerged.

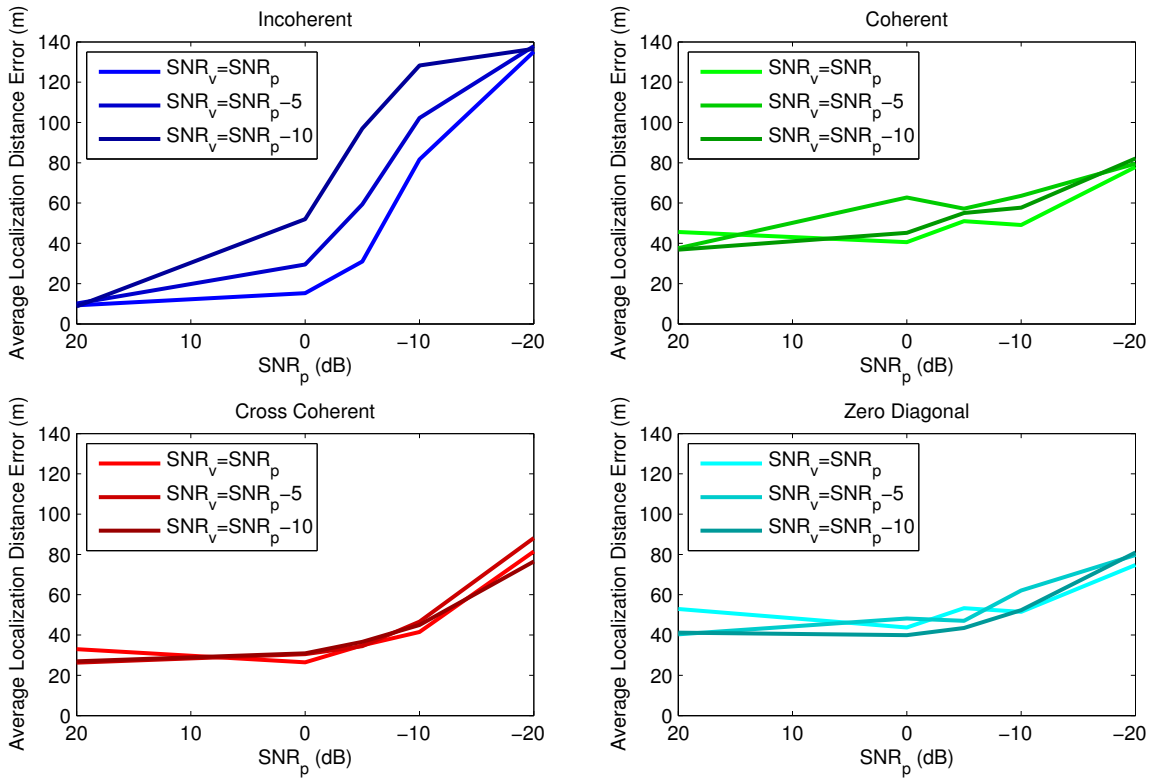

Figure 21: Average localization error plotted versus SNR_p for the incoherent (blue), coherent (green), cross–coherent (red) and zero–diagonal (cyan) weighting methods. Differing amounts of SNR_v are shown with darkened colors, and the standard deviation of positional error was fixed at 10 meters.

The results from Trials 3 and 4 are seen again here, where the incoherent method outperforms the coherent methods, but only when the SNR is greater than 0 dB and the velocity channel noise is not much higher than the pressure channel noise. The coherent methods have nearly constant error for low noise, since the driving factor is the positional noise. As the amount of acoustic noise increases, the accuracy only slightly decreases.

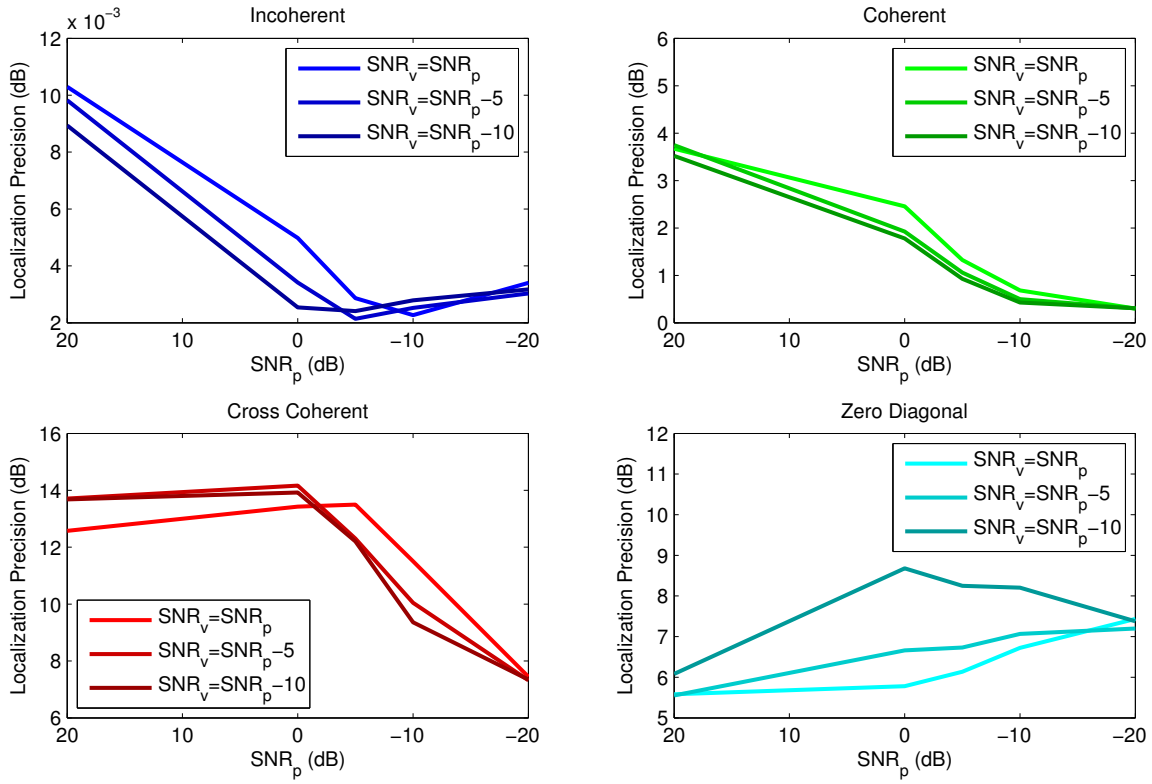

Figure 22: Average localization precision plotted versus SNR_p for the incoherent (blue), coherent (green), cross–coherent (red) and zero–diagonal (cyan) weighting methods. Differing amounts of SNR_v are shown with darkened colors, and the standard deviation of positional error was fixed at 10 meters.

The precision of each method shows a similar trend to that seen in Trials 3 and 4, where the added acoustic noise decreases the precision of all methods. In general, increasing the velocity channel noise marginally decreases the overall precision, except for the zero diagonal case, where it appears the added velocity channel noise increased precision. The exact cause of this behavior is still unknown, but as the zero diagonal method has yet to show any benefit over any other method, it is seemingly inconsequential.

The results from Trials 1–5 indicate that in high SNR cases, the incoherent method

is ideal, even if the resulting source location estimate is not very precise. However, in a more realistic passive vector sensor array where acoustic noise is expected to be much higher, the cross–coherent method appears to offer the best accuracy and precision. The only stipulation is that the gliders are separated enough from each other to obtain reliable range estimates.

CHAPTER VI

METHODS II: EXPERIMENTAL DATA

6.1 August 2012 Dataset

The first set of experimental data was collected in Monterey Bay, California using two Alaska Native Technologies (ANT-LLC) littoral gliders. Each glider was equipped with one Wilcoxon VS–301 vector sensor mounted on a pole approximately one meter off the nose to reduce acoustic interference and scattering from the glider body (see inset of Figure 23). The research vessel John H. Martin (R/V John H. Martin Specifications), operated by Moss Landing Marine Laboratories, acted as a noise source of opportunity. GPS data for both the vessel and gliders were obtained throughout the mission, the former sampled at 0.1 Hz and the latter sampled at approximately 1 Hz. The gliders were positioned on the surface with wings level, pitched nose down about 50 degrees and spaced about 80 meters apart. The gliders drifted in this configuration while the vessel followed a planned set of maneuvers at 10 knots. The glider and source configuration is shown in Figure 23.

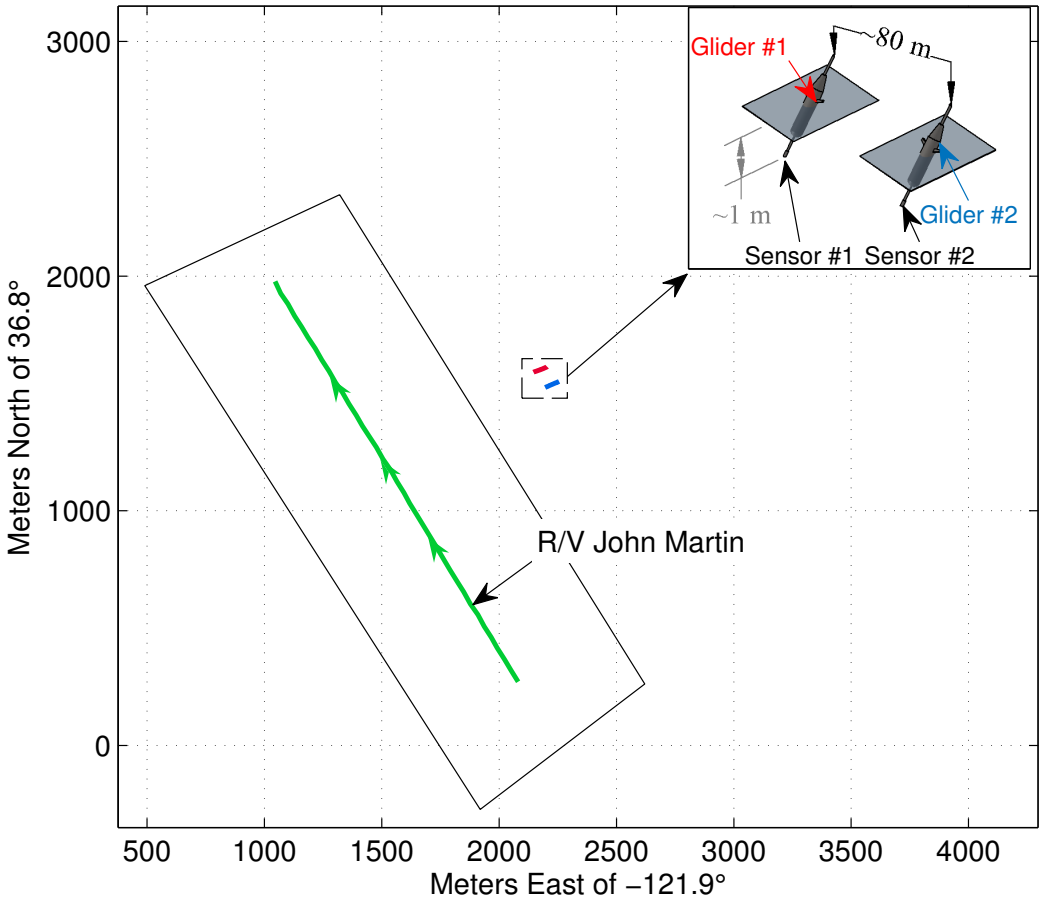

Figure 23: Schematic of the August 2012 dataset. The gliders (red and blue paths) were drifting on the surface approximately 80 meters apart (shown in inset), while the R/V John Martin (green path) was moving northwest at 10 knots. The evaluated ambiguity surface boundary is shown by the solid black area.

The selected subset of data was chosen because the source followed a straight line which spanned many different bearings to the gliders. The subset when the gliders were surfaced was also chosen for multiple reasons. Firstly, the positional accuracy obtained at the surface is drastically improved with a GPS fix. Secondly, the gliders were completely silent as they did not need to be maneuvering. Lastly, the synchronization of data between gliders was best known during this time period.

The synchronization step required for accurate coherent beamforming was done using the research vessel when it was positioned on the line connecting the gliders. Knowing the gliders' locations from GPS data and the sound speed, the propagation time delay between gliders is known. Correlating the pressure channels on each gliders' vector sensor should yield a peak at the propagation delay, however with an unknown synchronization offset. This synchronization offset was then subtracted from the NPS glider's time. Other methods of synchronization were attempted, such as using a reference source in close proximity to the gliders, but were not successful or reliable enough to use.

After synchronizing the acoustic records, the data was discretized into 10 second snapshots (see Chapter 4) and then preprocessed (see Appendix B). All acoustic data was downsampled by a factor of 8 from 44.1 kHz to 5512.5 Hz. Following downsampling, the data was filtered and whitened within the 100–800 Hz band, as this is where the majority of the research vessel noise was evident (see Appendix A). Filtering was accomplished in the frequency domain using a Tukey window with r=0.1, where r represents the fraction of the window width that is tapered (i.e. 10% of the window width is comprised of cosine–tapered ends). The whitening step normalizes the spectra so each frequency in the band has unit amplitude.

After filtering and whitening, the data were clipped to reduce the effects of sharp transient events. Any values further than 3σ from zero were replaced with $\pm 3\sigma$, where σ is the standard deviation of the data during the 10 second snapshot. Lastly, since the clipping process in nonlinear, a final filtering step is employed which simply uses the same filter from the filter/whiten step but without whitening.

With the data properly aligned and preprocessed, the source position was estimated for each 10 second snapshot using each of the weighting methods for the data covariance matrix given by Equations 6–9 and the beamforming approach defined by

Equations 10–12. All other data processing parameters were identical to the ones used in the previous chapter (e.g. sample rate, snapshot length, bandwidth, windowing function, etc.). For each snapshot, the estimated source position can be plotted with its size proportional to its accuracy, and color proportional to its precision (see Figures 24–27).

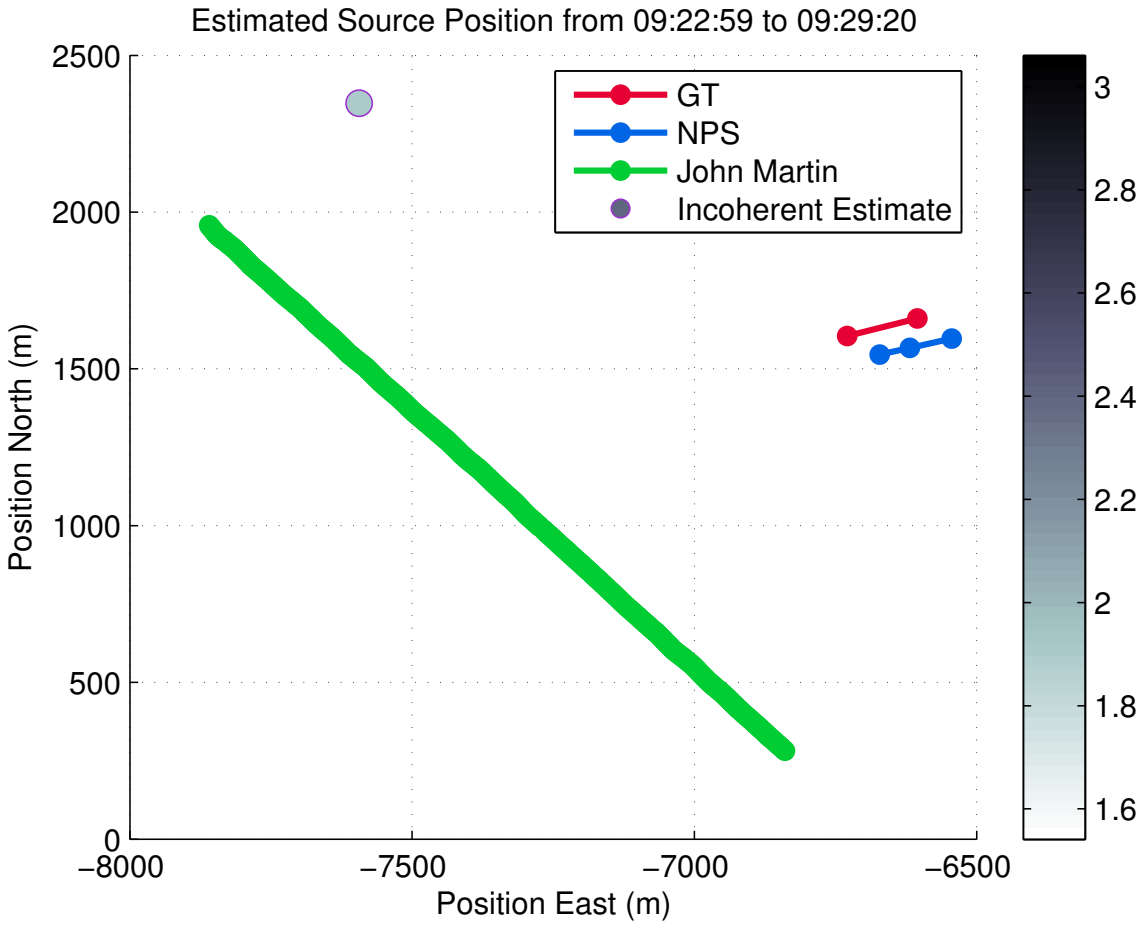

Figure 24: Localization results for the incoherent method. The position of the estimated source for each snapshot is plotted as a circle, whose size indicates the accuracy of the estimate (larger being more accurate). The color of the circle corresponds to the estimate precision (see colorbar for scale).

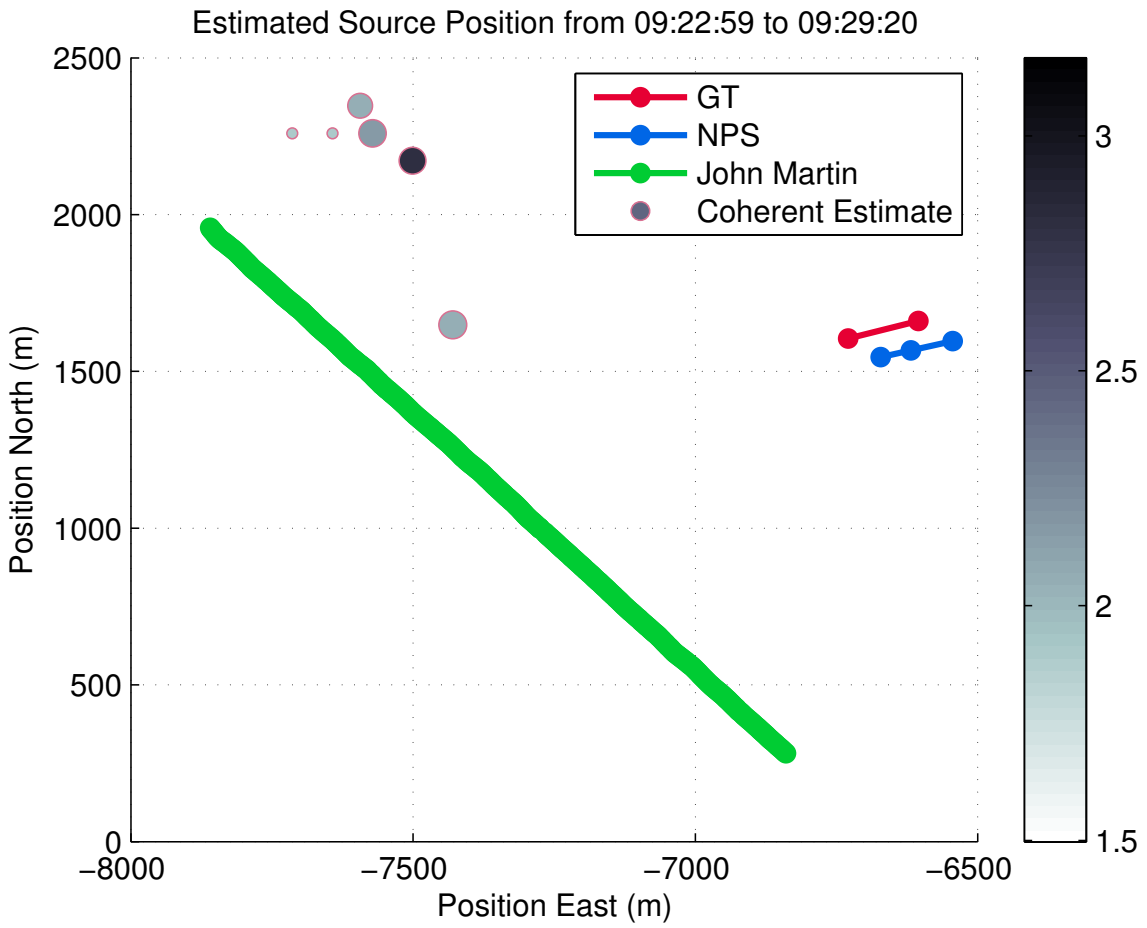

Figure 25: Localization results for the coherent method. The position of the estimated source for each snapshot is plotted as a circle, whose size indicates the accuracy of the estimate (larger being more accurate). The color of the circle corresponds to the estimate precision (see colorbar for scale).

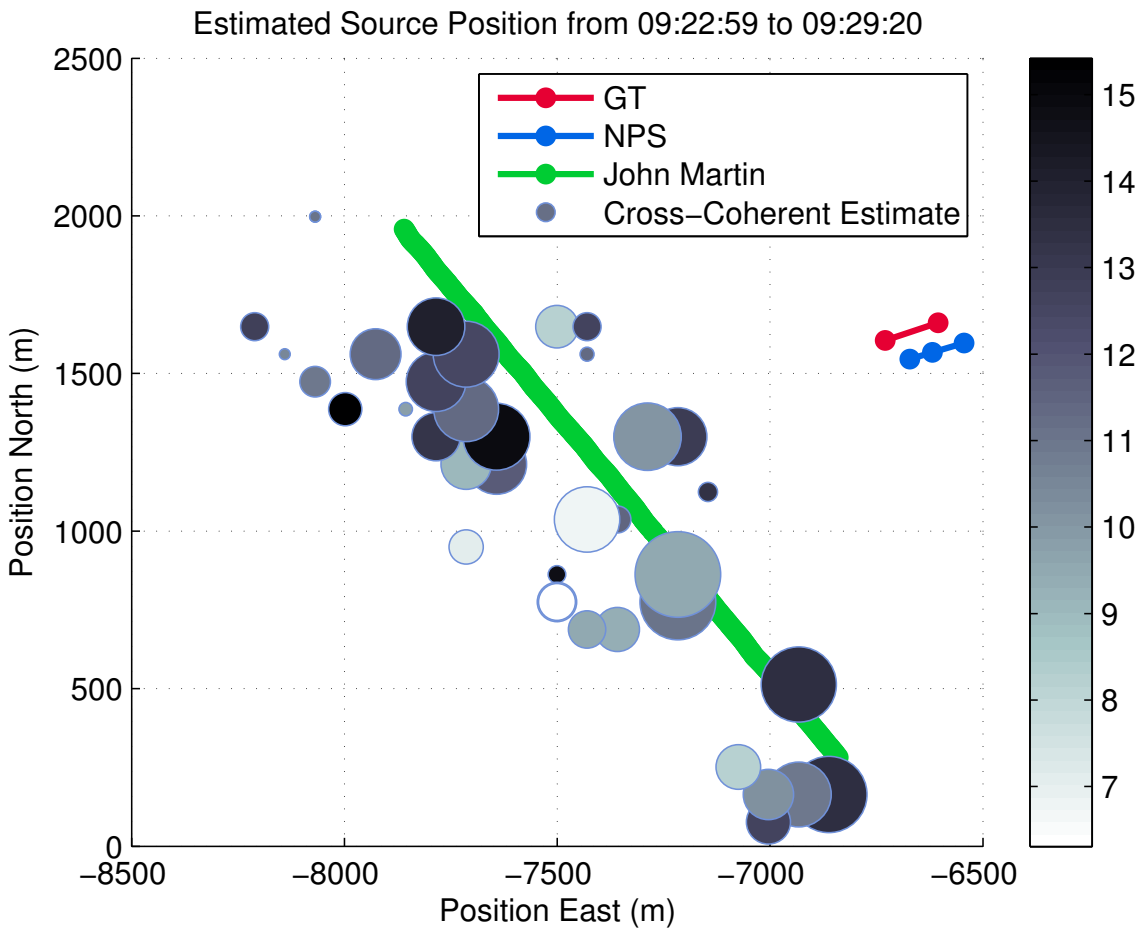

Figure 26: Localization results for the cross–coherent method. The position of the estimated source for each snapshot is plotted as a circle, whose size indicates the accuracy of the estimate (larger being more accurate). The color of the circle corresponds to the estimate precision (see colorbar for scale).

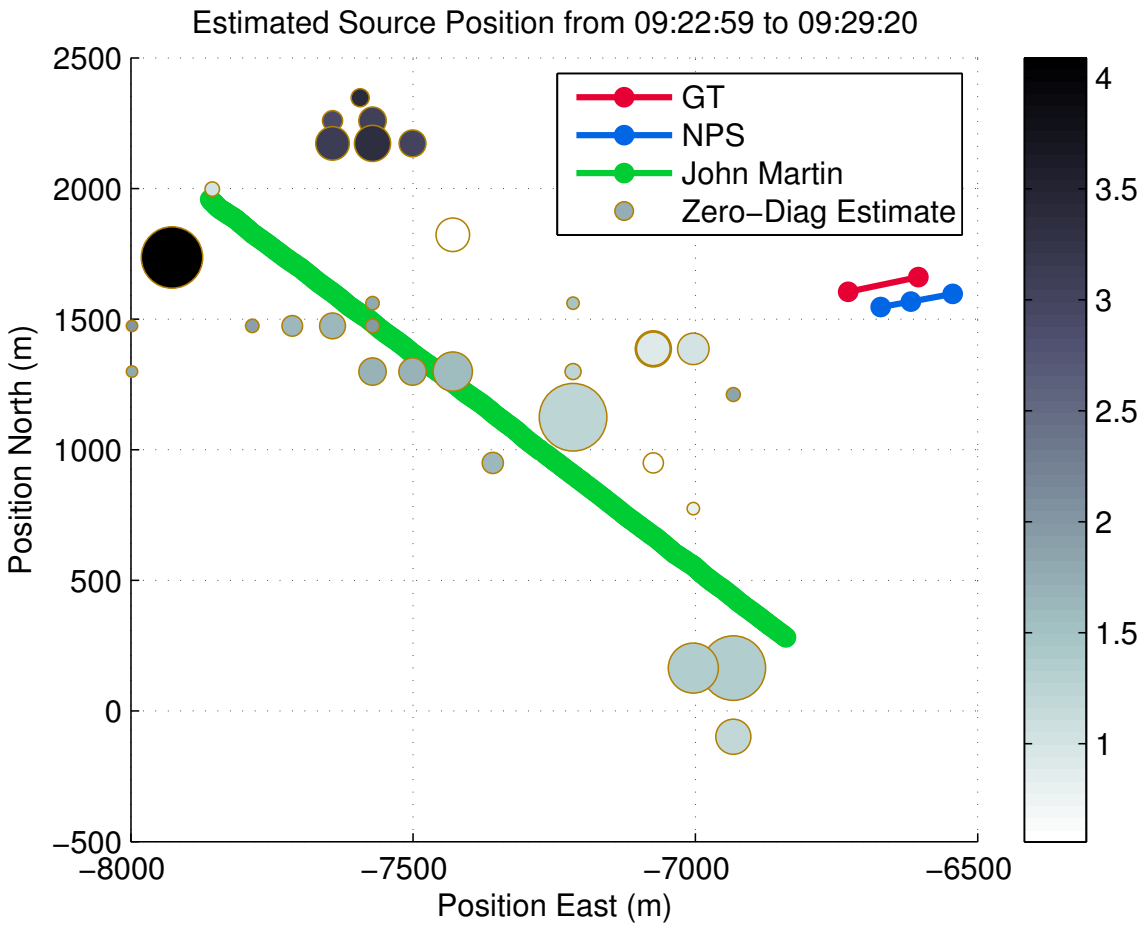

Figure 27: Localization results for the zero–diagonal method. The position of the estimated source for each snapshot is plotted as a circle, whose size indicates the accuracy of the estimate (larger being more accurate). The color of the circle corresponds to the estimate precision (see colorbar for scale).

Both the incoherent and coherent methods show an obvious failure in locating the source, as the estimated source location was always near the North corner of the ambiguity surface. The cross–coherent and zero diagonal methods did manage to locate the source, however, the estimated source location perpendicular to the source trajectory was never very accurate. This is because the glider locations were far from ideal, being so closely spaced. As such, obtaining an accurate range estimate

to a source is very difficult without having different bearings from each sensor to the source; however obtaining an accurate bearing estimate is still possible given the gliders' configuration (Chen et al., 2002). For this reason, the data error metric was redefined as the difference between the estimated and true bearing, rather than that given in Equation 13 in terms of absolute position. Then, the corresponding localization accuracy (with the new definition based on bearing) and precision can be plotted versus the snapshot time for each of the weighting methods (see Figure 28).

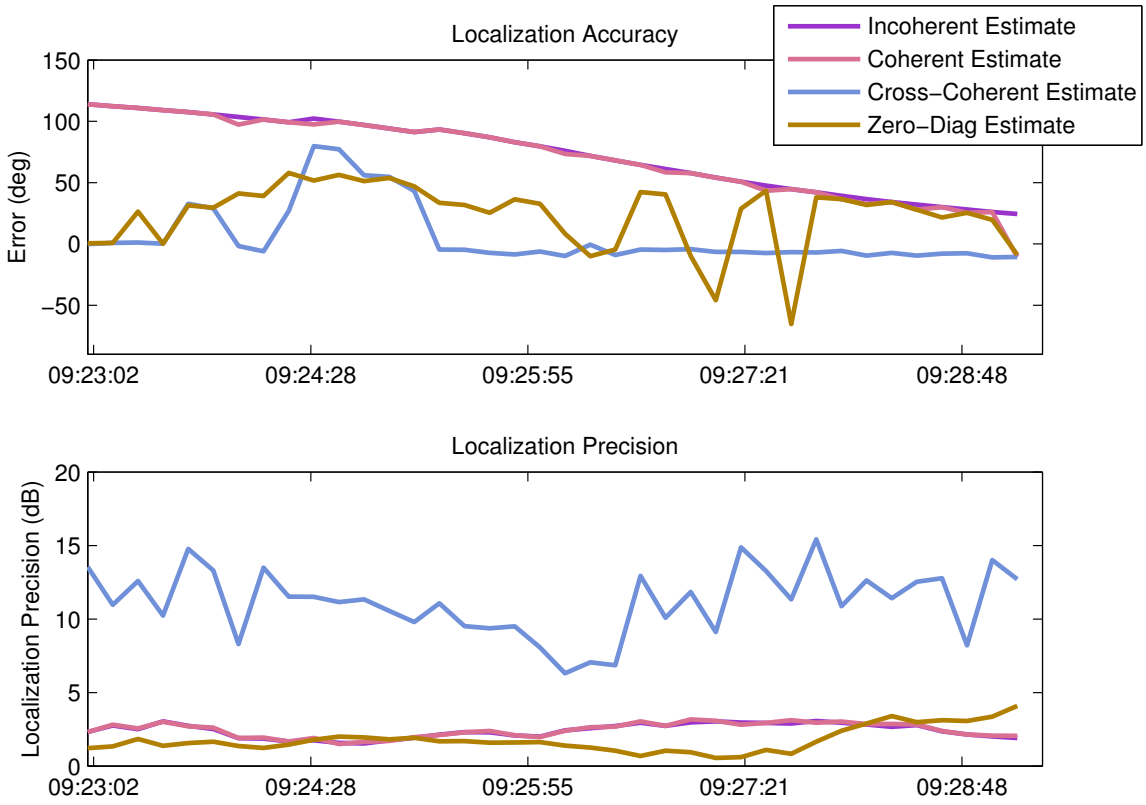

Figure 28: Localization results for the incoherent (purple), coherent (pink), cross–coherent (blue), and zero–diagonal (gold) weighting methods versus snapshot time.

Here, it becomes more clear that the cross–coherent method outperformed all other methods, and that the incoherent and coherent methods failed to locate the source at

all. The cross–coherent method shows the highest precision, as was seen in simulations (see Chapter 5). Despite performing well at the beginning and end of the dataset, the cross–coherent method exhibits a large degree of error around 09:24. This is caused by sudden jumps which occur in the GPS data; when one glider suddenly "jumps" ahead of the other, the estimated source bearing is skewed.

The glider's GPS data, although sampled at 0.1 Hz, only appears to update occasionally, likely when a new GPS fix is obtained. During the interval between readings, the glider appears to report its location based on the last known GPS location. The result is a "stair step" graph where each sudden drop corresponds to a new GPS reading (see Figure 29). To remedy this issue, the redundant positional information is removed, allowing the position between updates to be linearly interpolated from adjacent GPS readings.

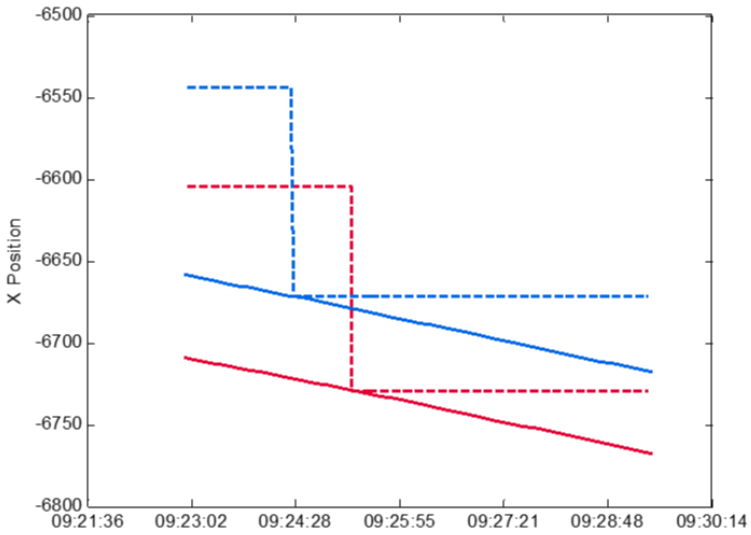

Figure 29: Raw GPS data from the gliders (dashed lines) is smoothed within the window of interest by removing repeated data points to produce a linearly–interpolated position (solid lines).

With smoothed GPS data, the beamforming methods were re–evaluated using the bearing accuracy and precision versus snapshot time (see Figure 30).

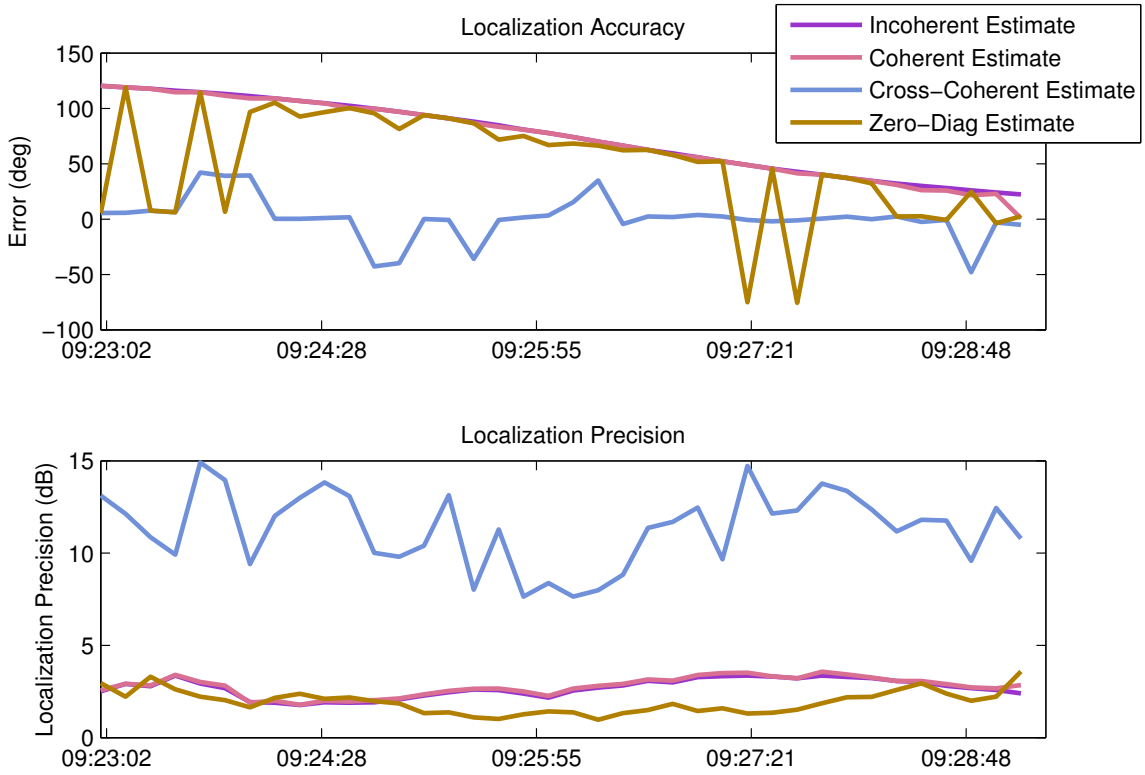

Figure 30: Localization results with GPS smoothing for the incoherent (purple), coherent (pink), cross–coherent (blue), and zero–diagonal (gold) weighting methods versus snapshot time.

The results for the incoherent and coherent method were not affected, and the zero–diagonal method appears to favor the same corner of the ambiguity surface now. The cross–coherent method does improve though, and no longer deviates as far from the true source. There are still a few sections of increased error, where the estimated bearing is off by as much as 50°. Inspecting the correlation plot between pressure channels on the gliders and overlaying the relative propagation delay which would be experienced by the estimated source shows that the culprit is likely multi–path

propagation (see Figure 31).

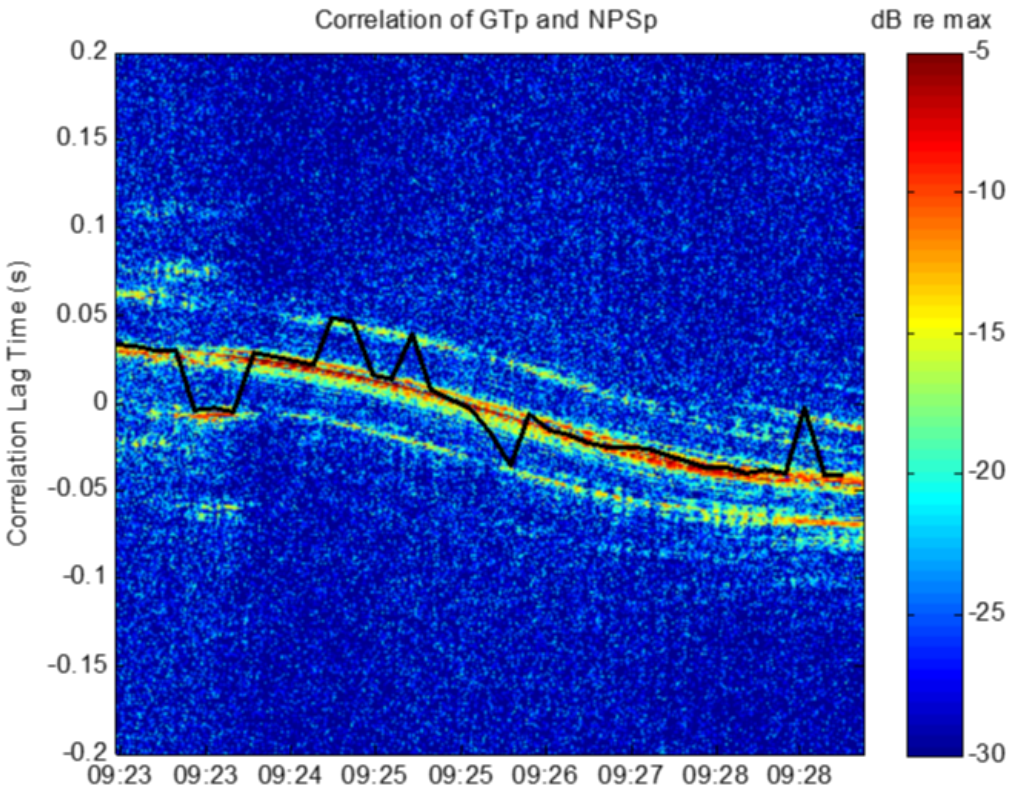

Figure 31: Correlation of 1 second of pressure channel data versus snapshot time, shown in 10 $\log_{10}$ dB scale. The propagation time–delay which would be experienced if the source was following the cross–coherent estimated source trajectory is shown in black.

The side peaks of the correlation plot appear to be strong enough at times to yield the highest beamformer output. To remedy the side peaks having such a strong effect, the correlations formed during the beamforming process (Equation 3) are Gaussian windowed, where the mean is the correlation lag time of the last known source position, and the standard deviation is 0.01 seconds. No windowing is performed on the first snapshot, as there is no known previous source position. Assuming the source is

correctly located initially, the next snapshot will window out the multi–path correlation peaks, while still allowing for the main peak to move slightly. Having introduced this form of iterative smoothing, all weighting methods were re–evaluated, still using smoothed GPS positions (see Figure 32).

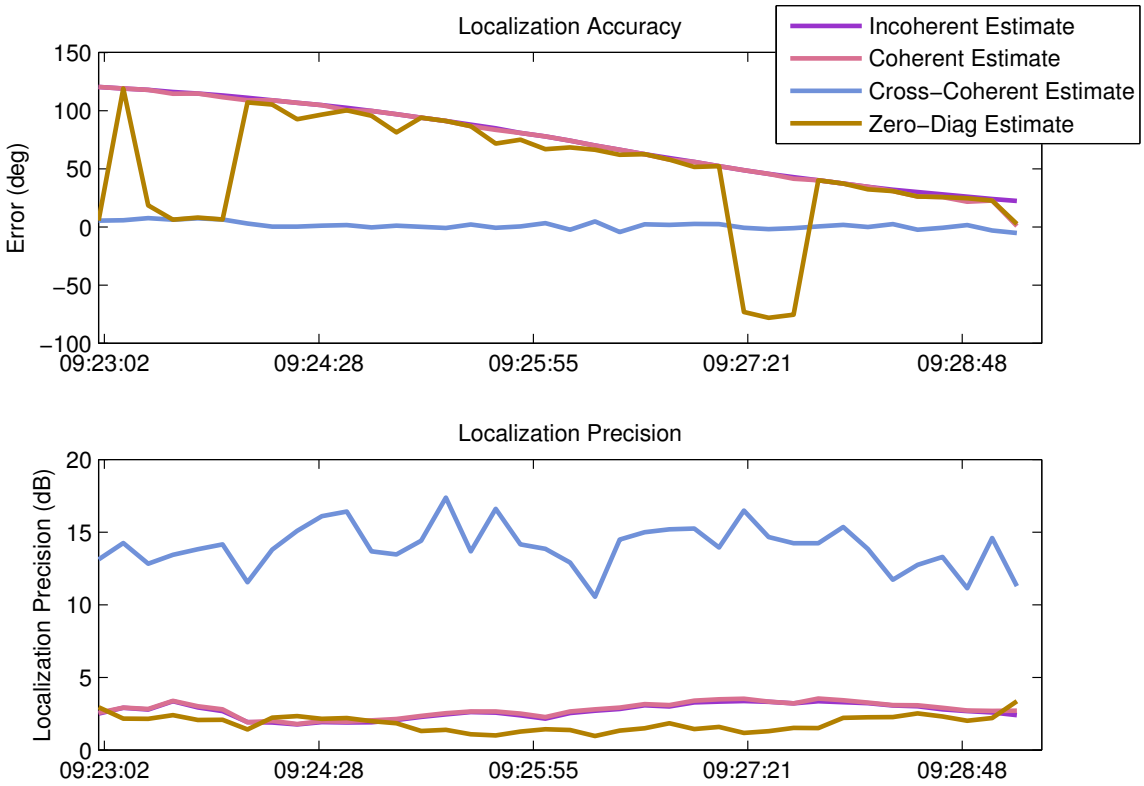

Figure 32: Localization results with GPS smoothing and iterative correlation smoothing for the incoherent (purple), coherent (pink), cross–coherent (blue), and zero–diagonal (gold) weighting methods versus snapshot time.

With both GPS smoothing and iterative correlation smoothing, the cross–coherent method finally tracks the source very well. The other methods still fail for the most part, other than the zero diagonal which did track the source for a few snapshots early on. Looking more closely at the cross–coherent performance (see Figure 33), it displays accuracy within 5° for most of the trajectory.

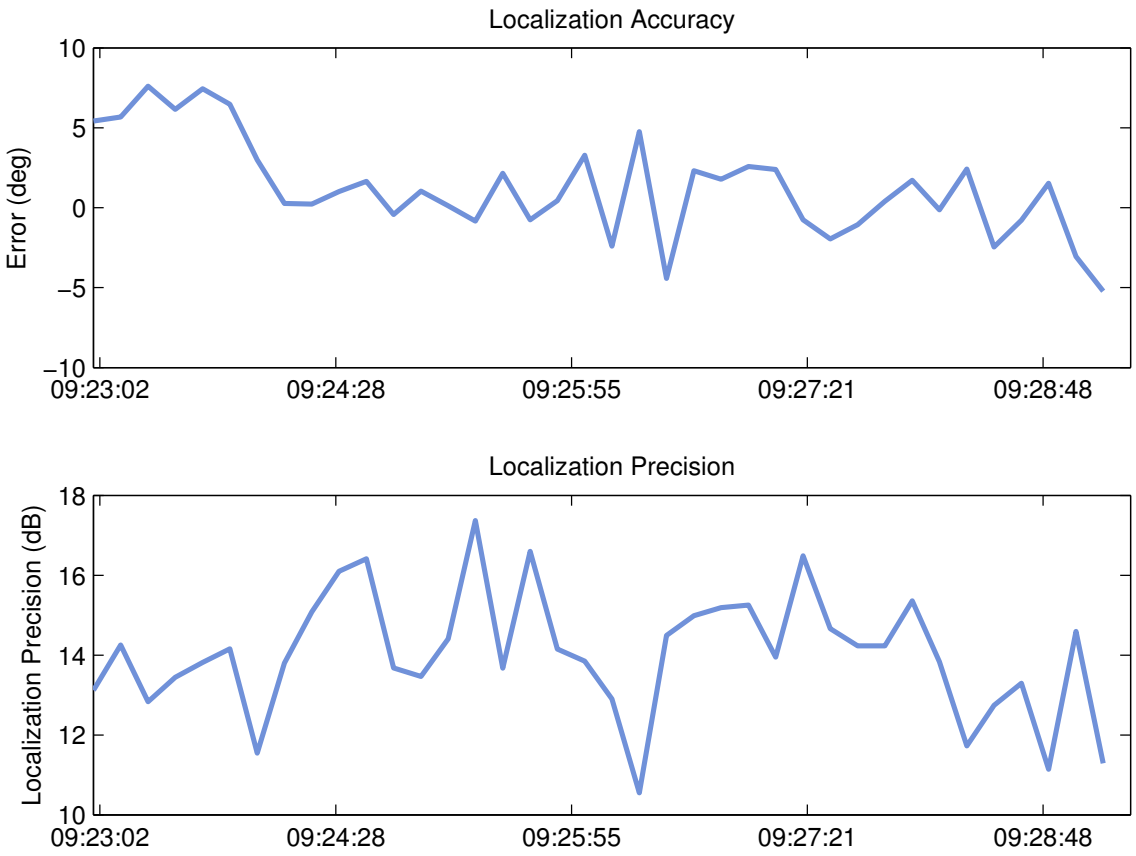

Figure 33: Localization results with GPS smoothing and iterative correlation smoothing for only the cross–coherent weighting method versus snapshot time.

The results of the August 2012 dataset indicate that it is certainly possible to locate a source using a distributed vector sensor array mounted on gliders. Even when surfaced and in the presence of flow noise, wave noise and other surface effects, the cross–coherent weighting method is able to determine the bearing towards the source. As in simulations, it obtained the highest accuracy as well as precision. Comparing other methods to simulations was not possible though, as the local noise overwhelmed the incoherent covariance matrices, causing all methods which relied on them to fail.

6.2 March 2014 Dataset

The second set of experimental data were also collected in Monterey Bay, California using the same glider and sensor setup as in the August 2012 experiment. The research vessel used instead was the R/V Fulmar,(R/V Fulmar Specifications), operated by the Monterey Bay National Marine Sanctuary. The gliders were deployed on the surface with wings level, pitched nose down about 50 degrees and positioned about 10 meters apart. They were then instructed to dive in a northeast direction towards a waypoint while the research vessel drove quickly away from the gliders (see Figure 34).

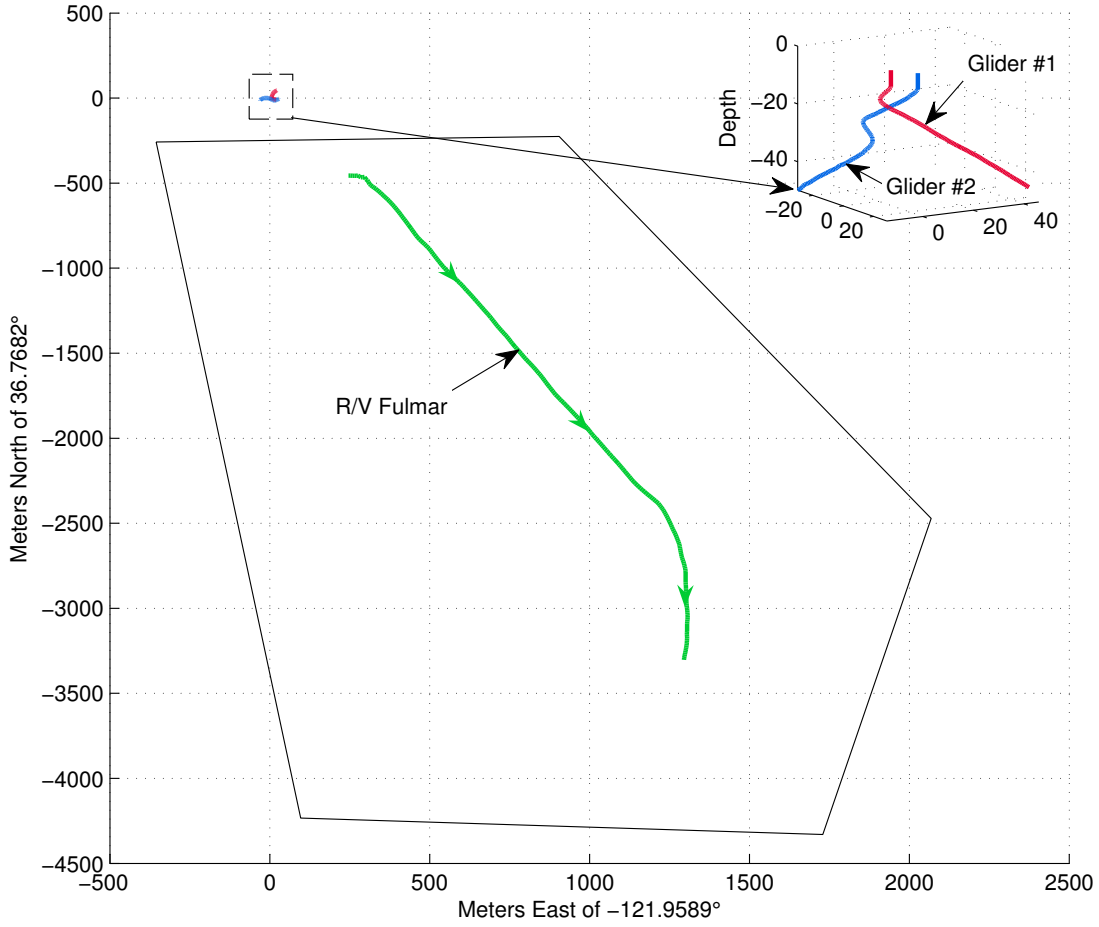

Figure 34: Schematic of the March 2014 dataset. The gliders (red and blue paths) were diving from an initial separation of 10 meters (see inset), while the R/V Fulmar (green path) was following a path southeast. The evaluated ambiguity surface boundary is shown by the solid black area.

This dataset, as was the August 2012 dataset, was one of opportunity and thus not ideal for localizing a source, since the gliders were very closely spaced and the vessel was traveling radially outward. However, one benefit to this dataset was the addition of an electronic synchronization chirp pulse applied simultaneously to the acoustic records. The first set of two pulses was applied one hour prior to deployment, and the second set of two pulses applied immediately before deployment. It was intended

to apply a third set of pulses at the end of the record, however an error in the glider operation reset the acoustic data collection during the mission. Correlating the synchronization pulses allowed for more accurate synchronization between gliders, but it is unclear whether this synchronization was held for the majority of the mission.

The acoustic data obtained in this test was sampled at 39.0625 kHz and downsampled by a factor of 8 to 4882.8 Hz. The processing band used was changed to 100–1500 Hz to capture the majority of noise seen in the acoustic spectra (see Appendix A). All other parameters (snapshot length, bandwidth, sample rate, density, speed of sound, preprocessing steps, etc.) used were identical to the August 2012 experiment and simulations. As was done in later iterations of the August 2012 dataset, iterative correlation windowing was applied. GPS smoothing was not necessary though, since the glider's reported positions were based on dead reckoning while diving, and were updated every second. All four weighting methods were then evaluated at each snapshot and the estimated source position was plotted with a circle whose size indicates the estimate accuracy, and whose color indicates the estimate precision (see Figures 35-38).

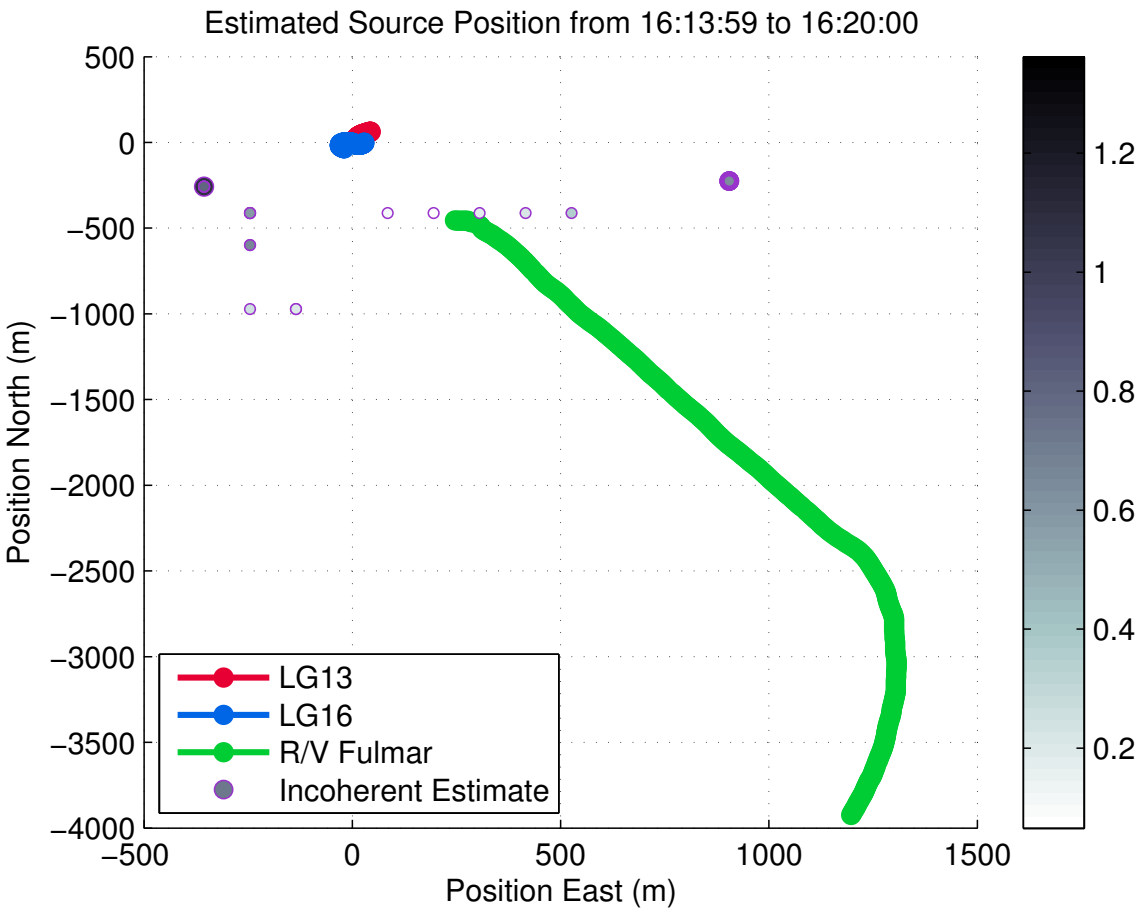

Figure 35: Localization results for the incoherent method. The position of the estimated source for each snapshot is plotted as a circle, whose size indicates the accuracy of the estimate (larger being more accurate). The color of the circle corresponds to the estimate precision (see colorbar for scale).

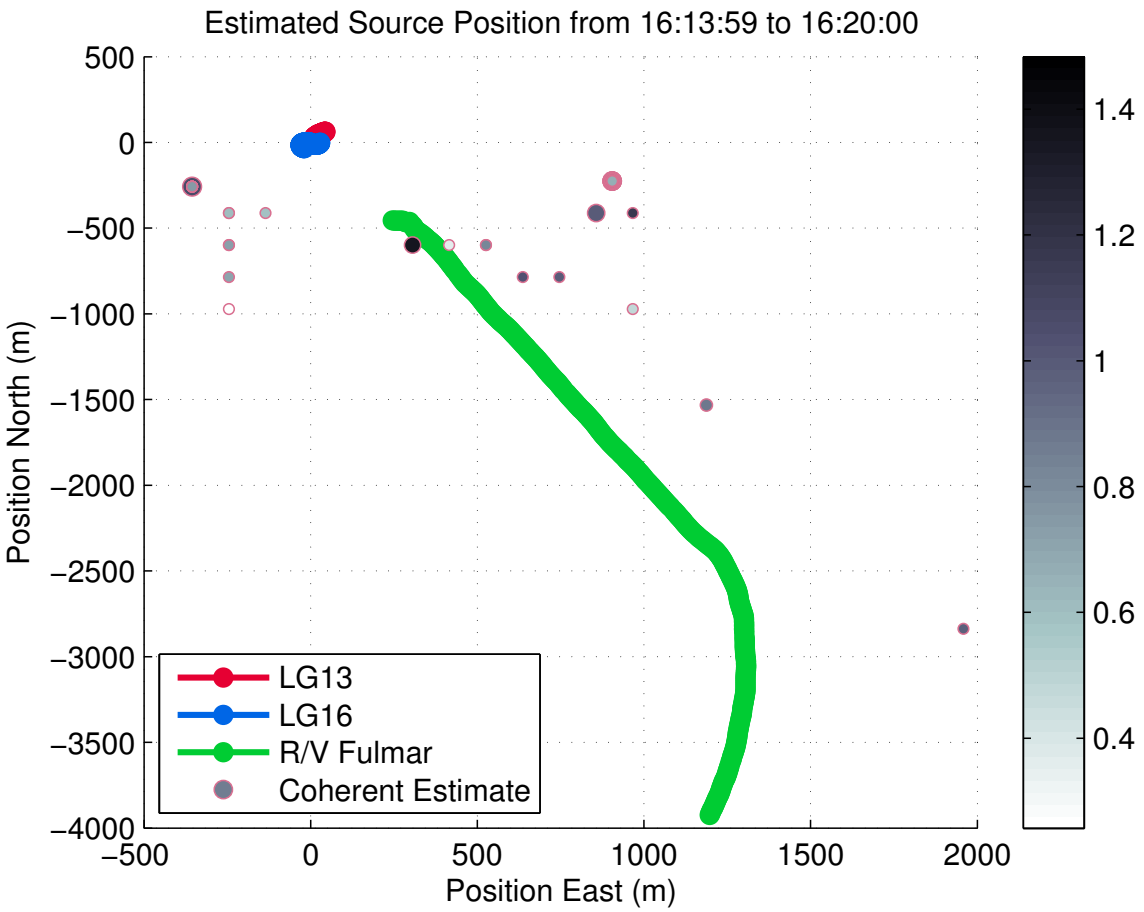

Figure 36: Localization results for the coherent method. The position of the estimated source for each snapshot is plotted as a circle, whose size indicates the accuracy of the estimate (larger being more accurate). The color of the circle corresponds to the estimate precision (see colorbar for scale).

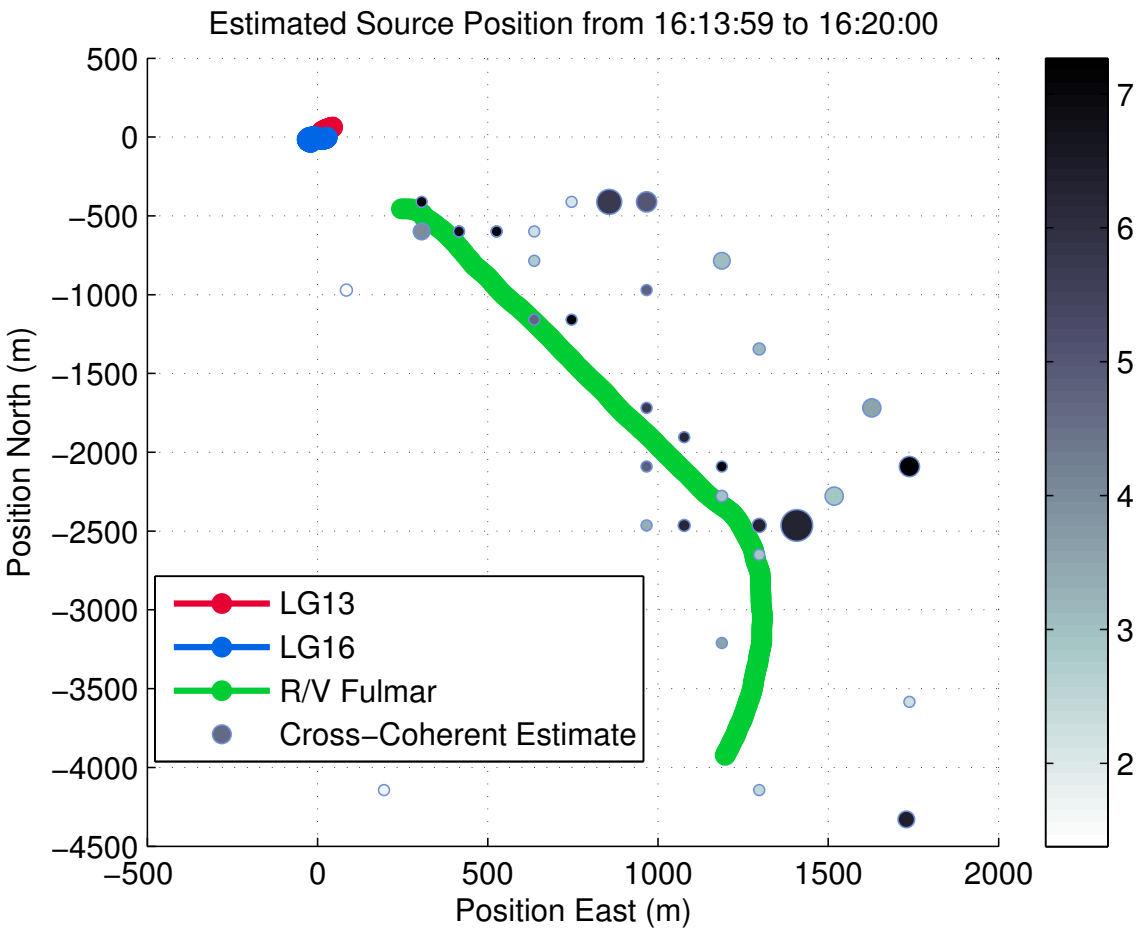

Figure 37: Localization results for the cross–coherent method. The position of the estimated source for each snapshot is plotted as a circle, whose size indicates the accuracy of the estimate (larger being more accurate). The color of the circle corresponds to the estimate precision (see colorbar for scale).

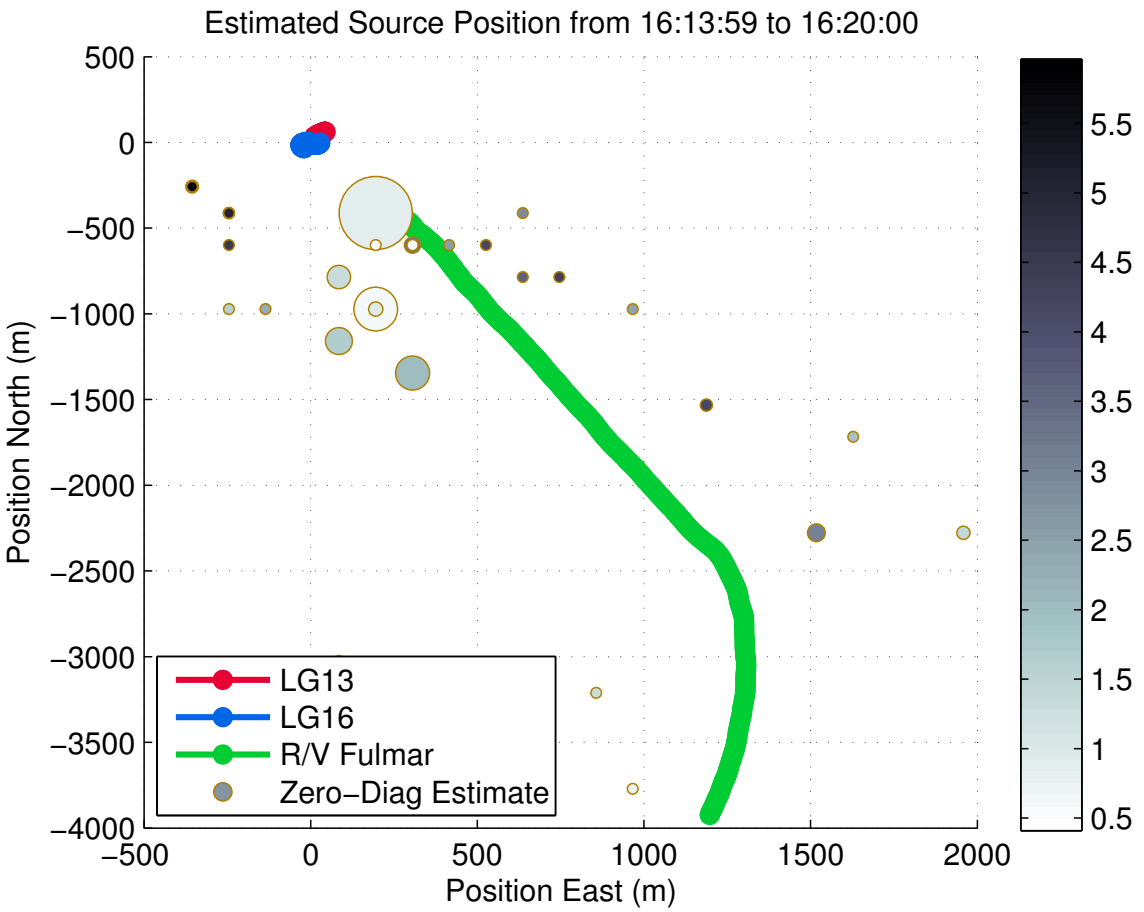

Figure 38: Localization results for the zero–diagonal method. The position of the estimated source for each snapshot is plotted as a circle, whose size indicates the accuracy of the estimate (larger being more accurate). The color of the circle corresponds to the estimate precision (see colorbar for scale).

As was seen in the August 2012 dataset, the coherent and incoherent methods appear to favor the edges of the ambiguity surface. When the estimated source bearing was roughly accurate, the estimated range was always much too small. This problem of glider separation and obtaining an accurate range estimate again lends itself to defining a bearing error metric for accuracy, where the difference in the source bearing and estimated source bearing is the localization accuracy. With the updated definition of accuracy, the localization accuracy and precision is plotted versus the snapshot time

(see Figure 39).

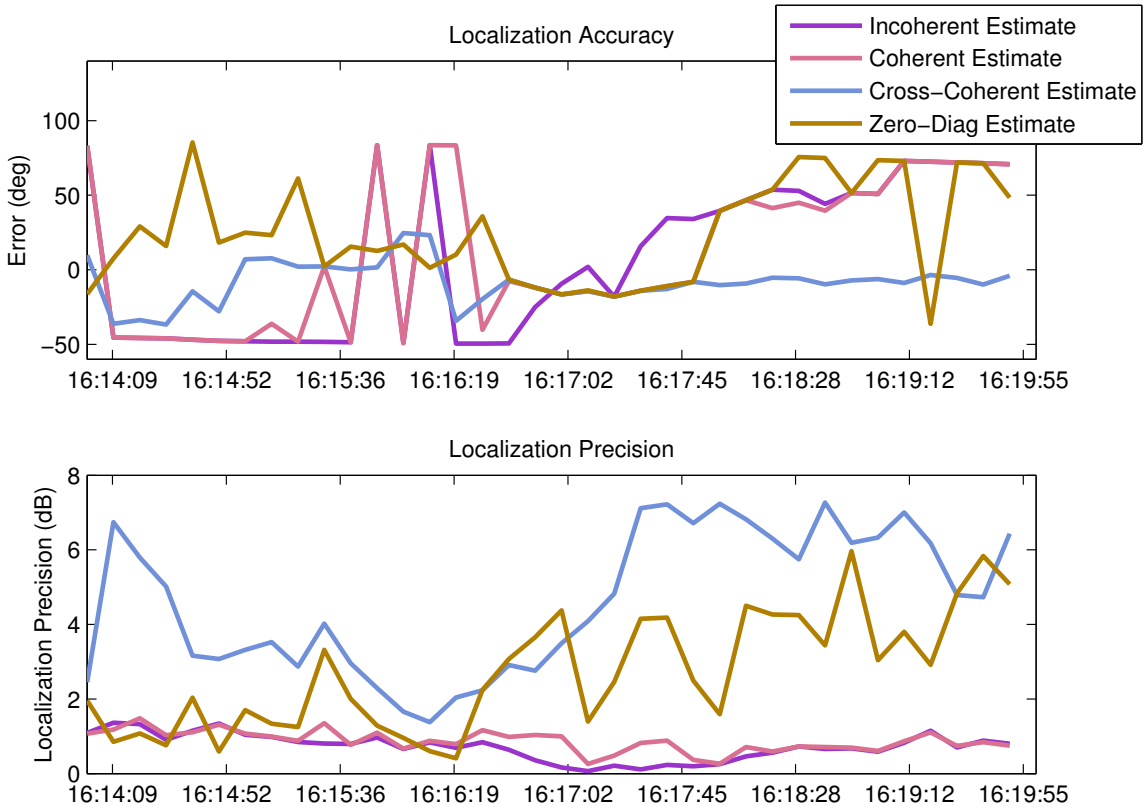

Figure 39: Localization results with GPS smoothing and iterative correlation smoothing for the incoherent (purple), coherent (pink), cross–coherent (blue), and zero–diagonal (gold) weighting methods versus snapshot time.

A noticeable difference in the incoherent and coherent methods' estimates is now visible, whereas in the August 2012 dataset, they mostly followed each other throughout. There still exists sections where the incoherent and coherent estimates were on the edge of the ambiguity surface and pointing in the wrong direction, however the occasional point around 16:17 seemed to be pointing in the right direction at least. The cross–coherent and zero diagonal methods performed slightly better, having fewer snapshots with drastically high error. The zero diagonal method appears to switch to following the same pattern as the incoherent and coherent methods after 16:18,

which might indicate an event which overwhelmed the incoherent covariance matrices with noise. The cross–coherent method is again the highest performing method, yet its precision is dramatically lower than seen in simulations or past experimental results. Having only two gliders means the cross–coherent method produces large outputs along the iso–delay hyperbola defined by the propagation delay between gliders. Since there is only one such hyperbola and its asymptote is pointed along the source path, a large number of points in the ambiguity surface have large values, thus decreasing the precision. Another potential issue with the cross–coherent method is the large errors seen at the beginning and middle of the dataset. This is likely explained by the behavior of the research vessel, where it had yet to start moving at the beginning, and it faded out for two periods of time around 16:16. A correlation plot across pressure channels shows the source peak does in fact take some time to appear, then disappears again for two periods around 16:16, finally reappearing and fading out as the source is too far away and enters port (see Figure 40).

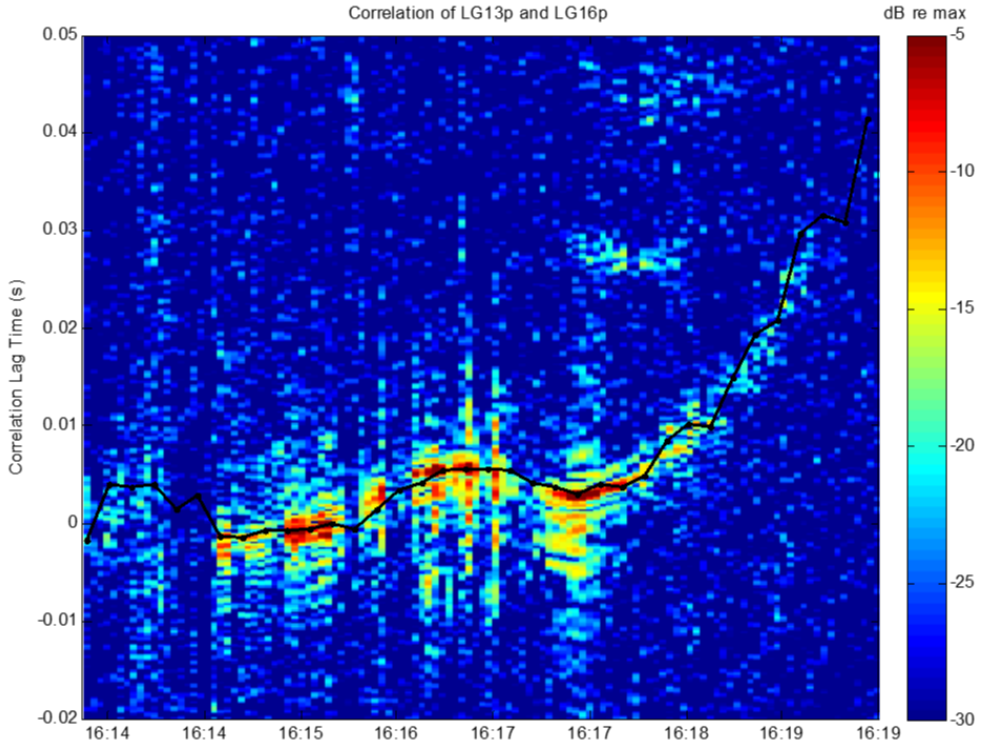

Figure 40: Correlation of 5 seconds of pressure channel data versus snapshot time, shown in 10 $\log_{10}$ dB scale. The propagation time–delay which would be experienced if the source was following the cross–coherent estimated source trajectory is shown in black

The cross–coherent estimated source delay shown in black on Figure 40 indicates it was able to track the source peaks, when present. Thus, the only explanation for bearing errors seen by the cross–coherent method is an error in positioning. If the relative position and orientation of the gliders is inaccurate, even if the propagation delay is known, the corresponding iso–delay hyperbola will not intersect with the true source. Figure 37 indicates the gliders favored an eastward bearing, which was likely introduced by positional uncertainties.

6.3 Comparison to Simulations

To attempt to verify the validity of the simulations, the August 2012 cross–coherent results were compared to simulations done with the same source and glider locations. Simulations used the same parameters used in experiments (snapshot length, bandwidth, sample frequency, etc.), although the values for positional error and acoustic noise were not measured in experimental data. The standard deviation of positional error was assumed to be 3.9 meters, which is common for commercial GPS devices (Grimes, 2008).

To determine the acoustic noise power, many methods were attempted, but the most reliable method used a five minute section of quiet data to establish a noise reference. Assuming the noise power was unchanged, the SNR was computed across each 10 second snapshot and plotted versus time. This process was repeated for two different selections of quiet data (see Figures 41-42)

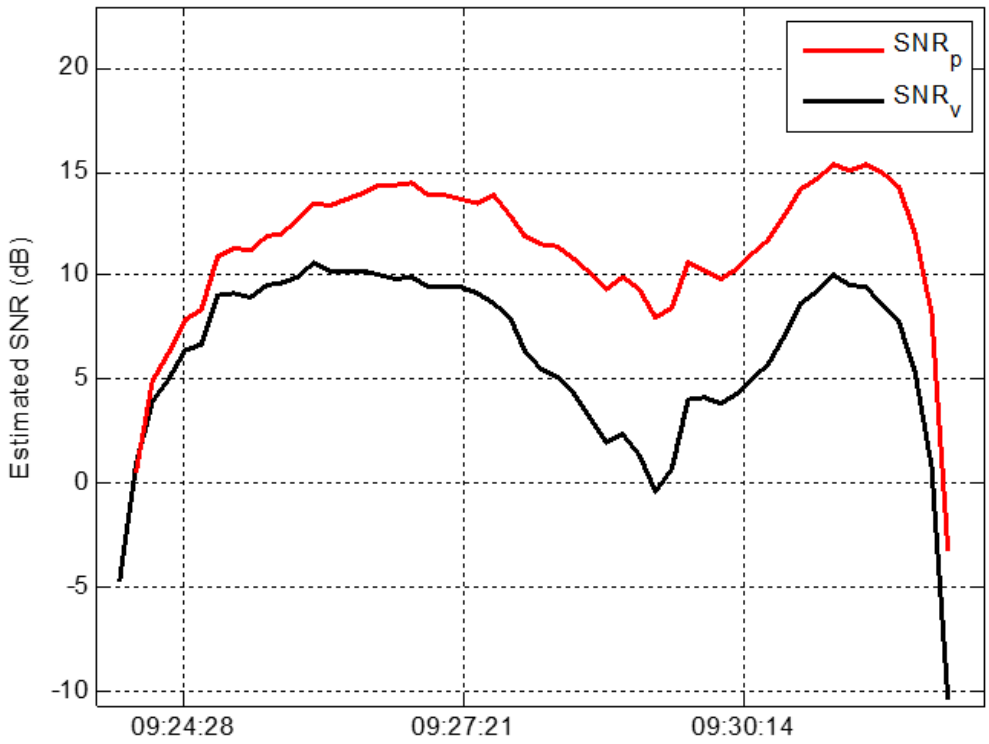

Figure 41: Estimated experimental SNR_p (red) and SNR_v (black) for 10 second snapshots versus snapshot time.

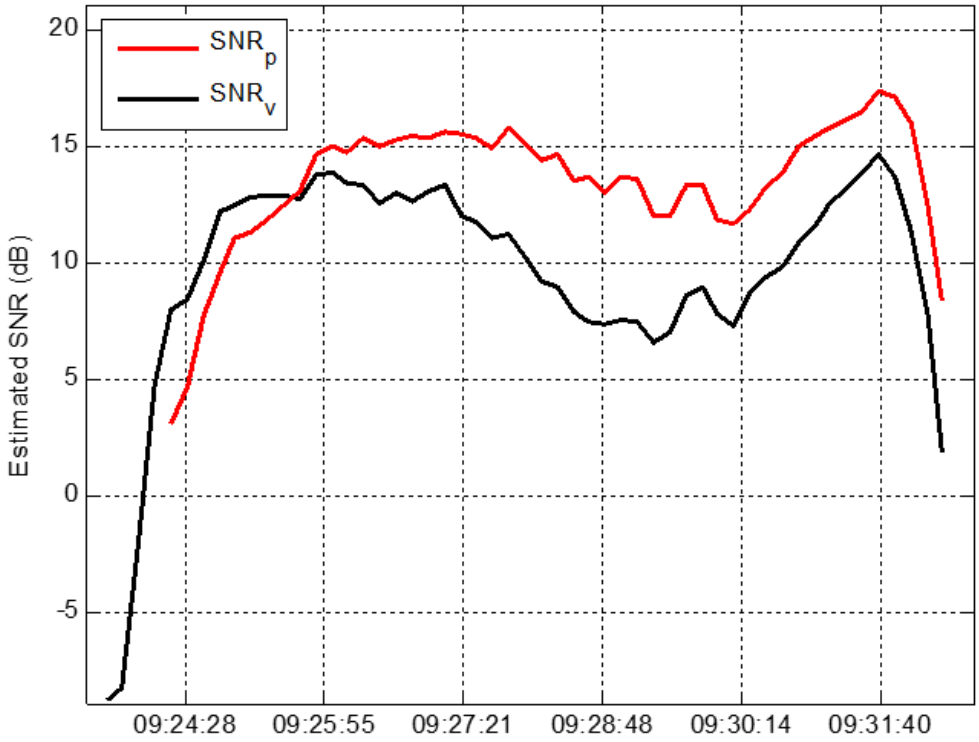

Figure 42: Estimated experimental SNR_p (red) and SNR_v (black) for 10 second snapshots versus snapshot time.

From the two noise references, it was assumed the acoustic SNR was about 10 dB. The relative difficulty of estimating the actual SNR accurately did not warrant using the variable SNR in simulations, but rather just using a constant value throughout. Performing localization with the cross–coherent weighting method and using iterative correlation windowing on 50 different realizations of positional noise and acoustic noise yields a set of localization accuracy and precision. As the accuracy tended to average to zero degrees, the standard deviation of all 50 errors was computed, σ_b. To compare to experimental data, 68%, 95%, and 99.7% confidence intervals are defined by $\pm n\sigma_b$ where $n = \{1, 2, 3\}$, respectively. These confidence intervals are plotted against time and indicate the variability which would be expected in the experimental data (see

Figure 43).

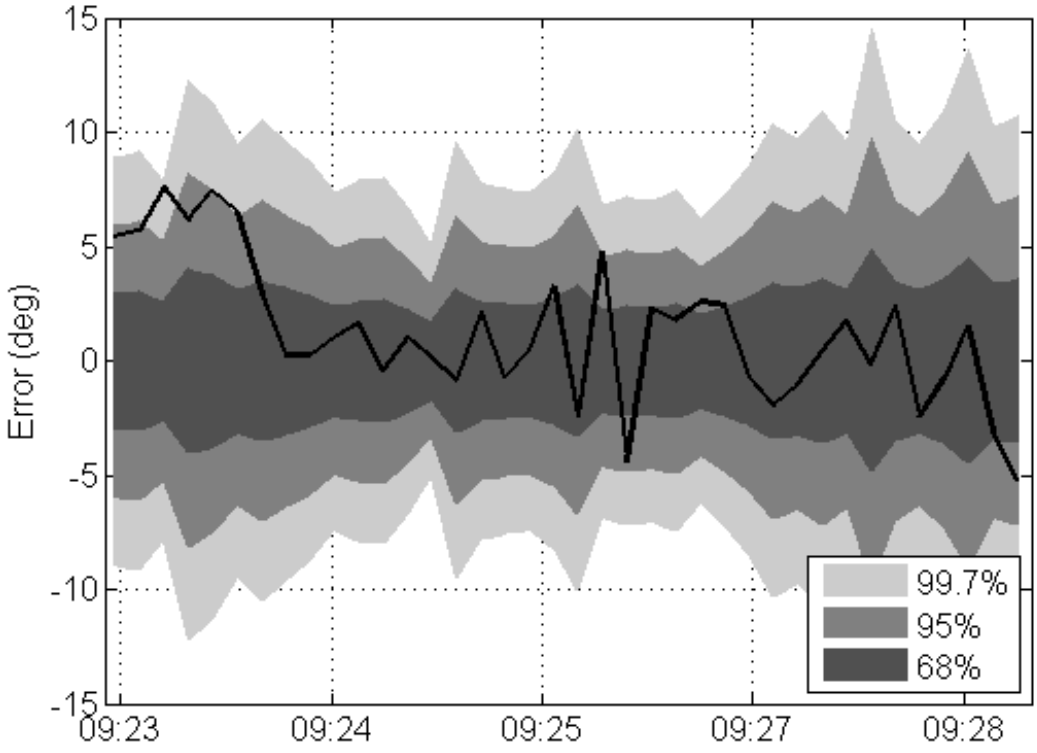

Figure 43: Experimental bearing error from the August 2012 cross–coherent estimate (black line) compared to the confidence intervals defined by the standard deviation of the simulated bearing error using the same method and glider configuration.

Overall, the estimated experimental bearing error across the boat trajectory falls within the 99.7% confidence interval. It also most often lies within the 68% confidence interval, which indicates that the simulation methodology presented in Chapter 5 is at least roughly applicable to real data, although the specific noise mechanism which caused the incoherent and coherent methods to fail was not explicitly included in the simulations. Possible improvements could be made to account for environmental variability and multipath effects in simulation as well as in the replica vector of Equation 10. Furthermore, localization performance could be improved through the

use of an adaptive matched–field processing approach. Lastly, experimental data could yield better localization results with more gliders, and also more ideal placement of the gliders.

CHAPTER VII

CONCLUSION

The results of this work demonstrated that a spatially distributed array of gliders equipped with vector sensors can be used to track a source of opportunity such as a surface vessel. Simulations indicate that in the presence of acoustic noise, the cross–coherent weighting method is more accurate and precise than the traditional incoherent or coherent methods. In experimental data, the cross–coherent method also outperformed all other methods, matching well with simulations when possible to compare. The localization performance of each method under different scenarios, both real and simulated, can help guide future development of autonomous vector sensor arrays and possibly improve the traditional methods of vector sensor beamforming used for source localization.

APPENDIX A

DATA SPECTRA

This appendix provides spectrograms and spectra of signals acquired in at–sea experiments.

A.1 August 2012 Dataset

The representative spectra of boat noise for both pressure and velocity channels are shown in Figures 44 and 45.

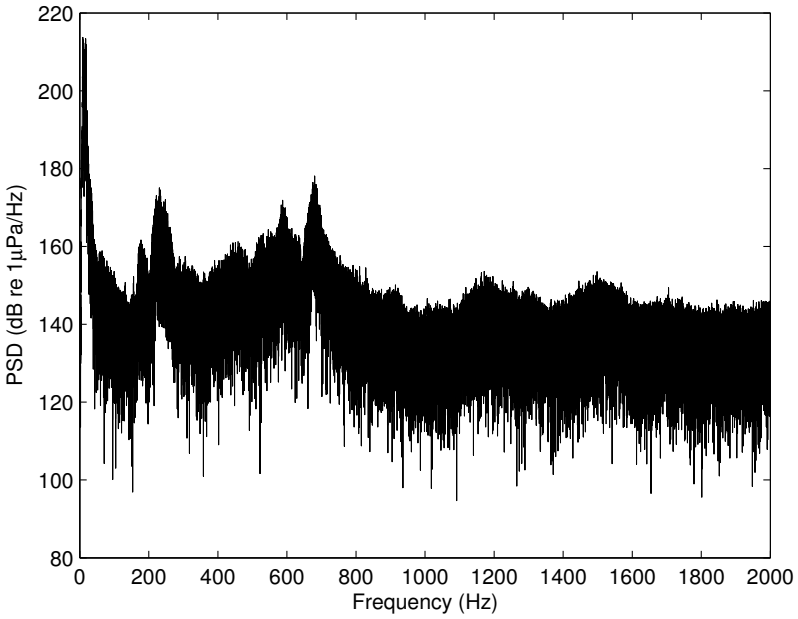

Figure 44: Power spectral density of one minute of pressure channel data taken from the August 2012 dataset.

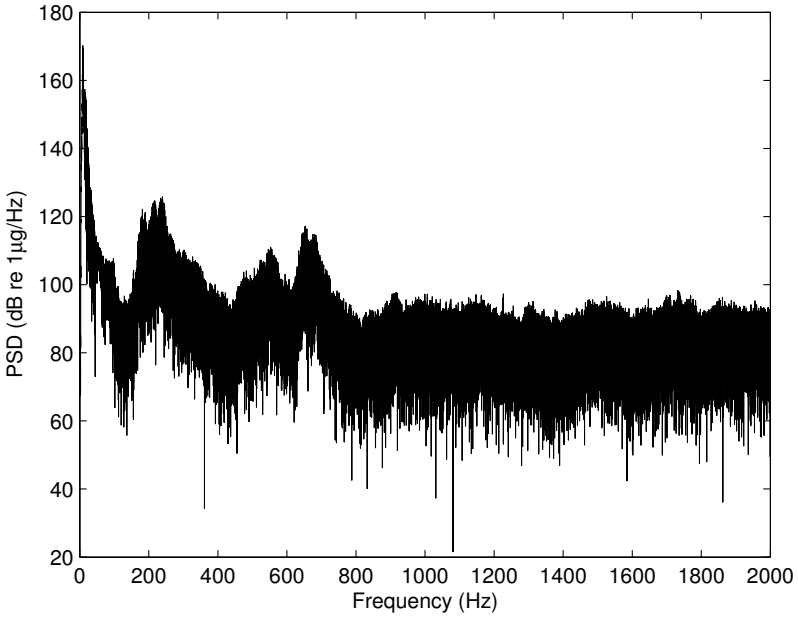

Figure 45: Power spectral density of one minute of velocity channel data taken from the August 2012 dataset.

Spectrograms of pressure and velocity channels are shown below in Figures 46 and 47.

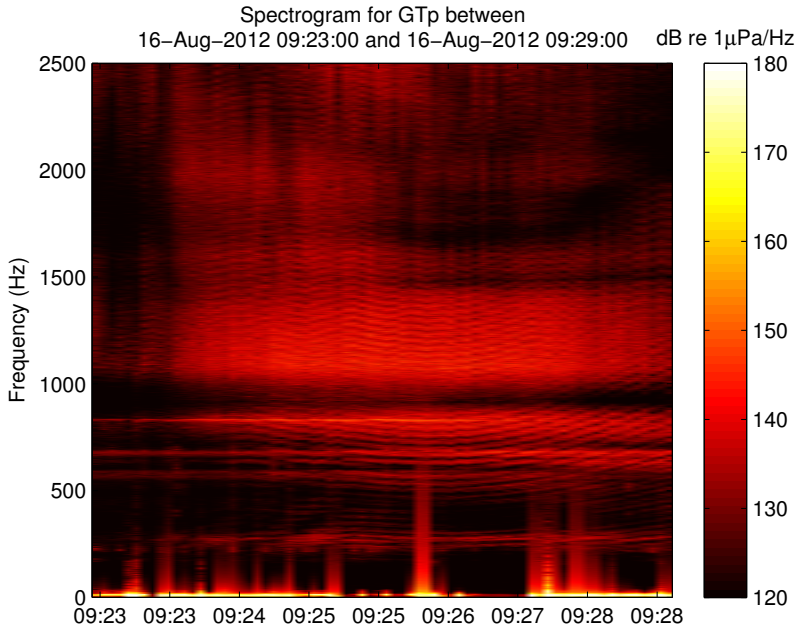

Figure 46: Spectrogram of August 2012 pressure channel data using 10 second windows.

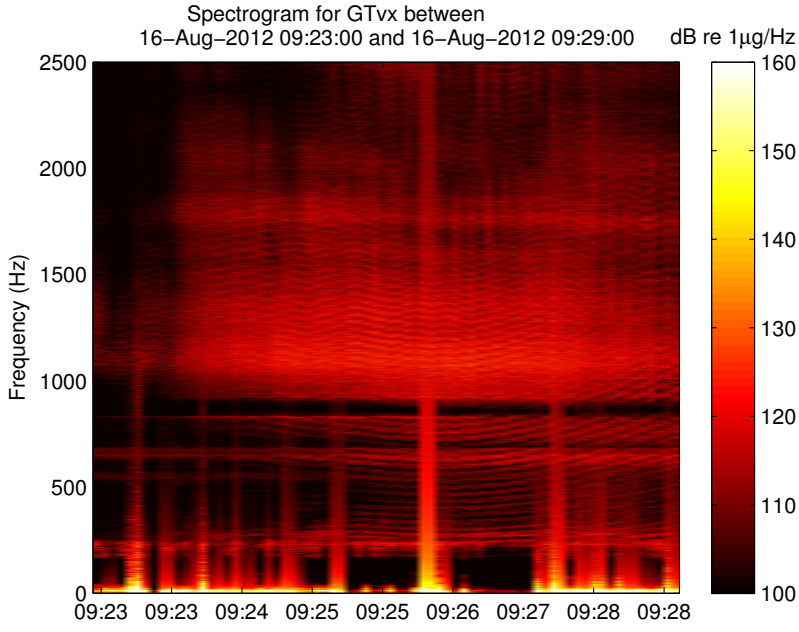

Figure 47: Spectrogram of August 2012 velocity channel data using 10 second windows.

A.2 March 2014 Dataset

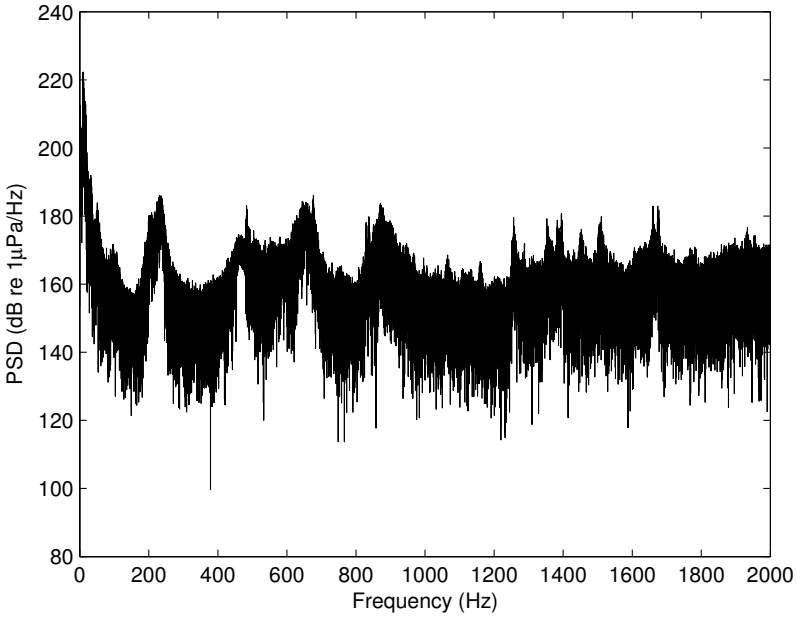

Figure 48: Power spectral density of one minute of pressure channel data taken from the March 2014 dataset.

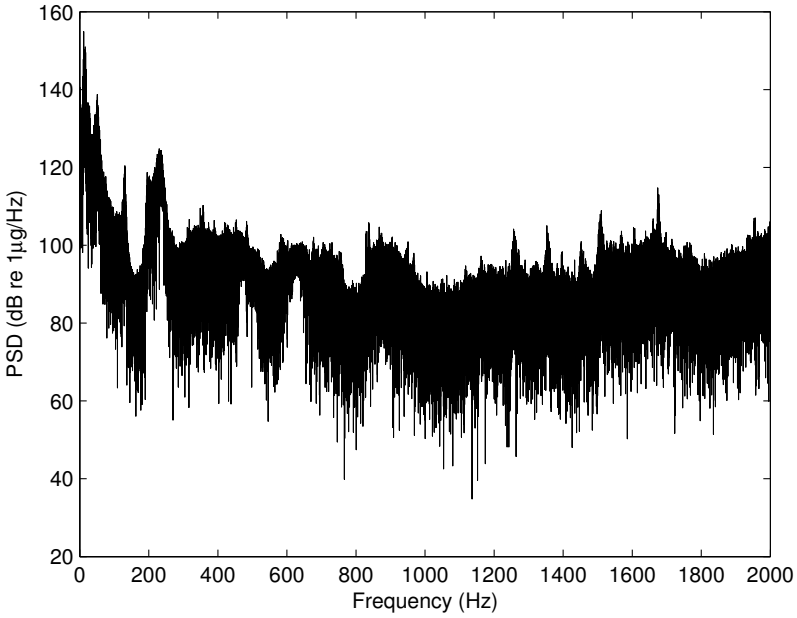

Figure 49: Power spectral density of one minute of velocity channel data taken from the March 2014 dataset.

Spectrograms of pressure and velocity channels are shown below in Figures 50 and 51.

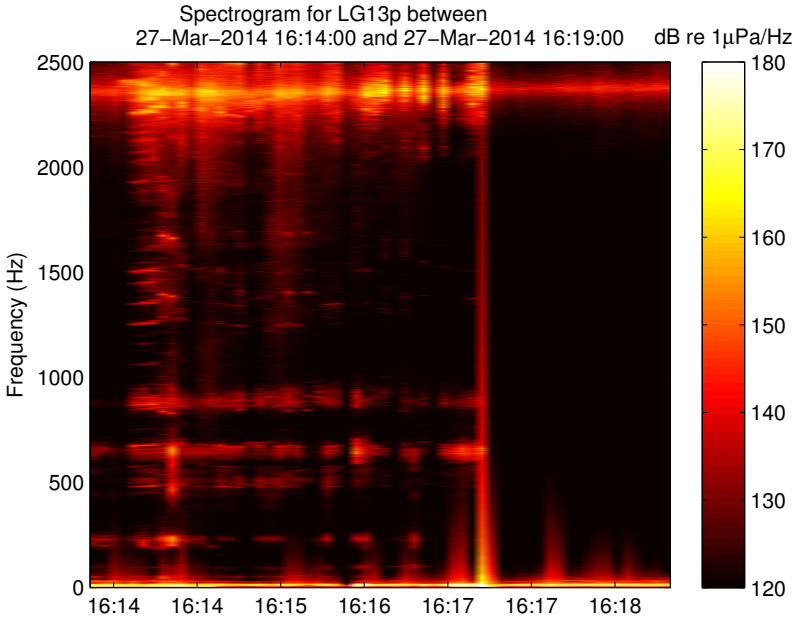

Figure 50: Spectrogram of March 2014 pressure channel data using 10 second windows.

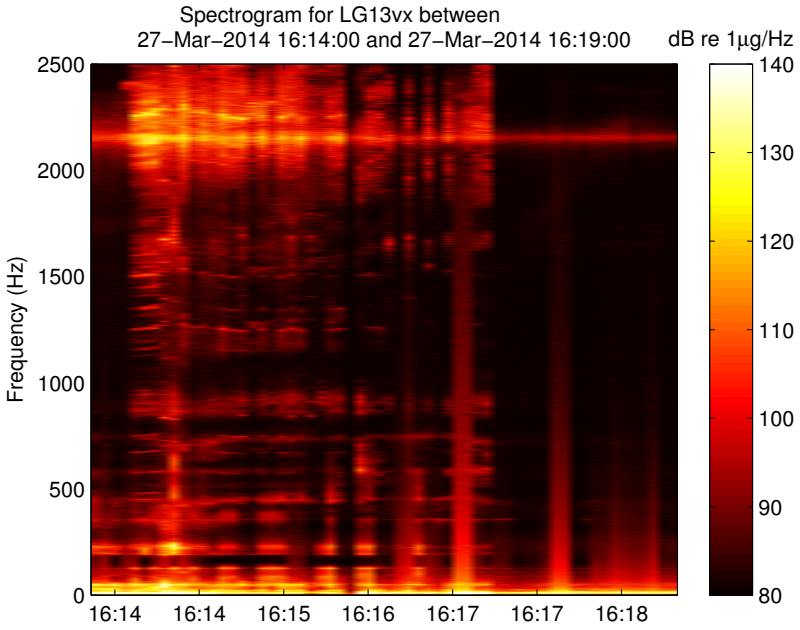

Figure 51: Spectrogram of March 2014 velocity channel data using 10 second windows.

APPENDIX B

DATA PREPROCESSING

This appendix describes the data processing steps that are implemented before beamforming and localization.

B.1 Downsampling

Raw data were sampled at a rate f_s but the band of interest lies at frequencies much lower. The data is downsampled by a factor of $q = 8$ using the built-in MATLAB function "resample". This function first lowpass filters the data using an 8th order Chebyshev Type I IIR filter with a cutoff frequency of $0.8f_s/q$ and passband ripple of 0.05 dB to avoid aliasing, followed by selecting every q–th point in the data.

B.2 Rotations

The vector sensor data are rotated from their individual sensor frames into a common reference frame. Since the telemetry data is typically sampled more slowly than the acoustic record and tends to change slowly, the nearest telemetry data point to any acoustic sample is used for the rotation (i.e. nearest–neighbor interpolation).

Telemetry data can come from the glider's internal measurements of yaw, pitch and roll angles or from the vector-sensor's gravity and heading vectors. It was determined that the vector sensor telemetry was not reliable for either at–sea experiment, and that the glider's internal measurements were more accurate.

Rotation itself is performed by multiplying the acoustic velocity vector at each acoustic sample time by a rotation matrix which rotates the vector from the local sensor coordinate system into the geodetic ENU coordinate system. This involves knowledge of the sensor roll angle β (see Figure 4); the glider's yaw, pitch, and roll; and the magnetic declination (see Figure 3).

B.3 Whitening

Whitening the data channels normalizes the magnitude of a signal's spectra within a band of interest. Taking the FFT of the data, the complex amplitude at each frequency is divided by its magnitude (plus machine epsilon to avoid division by zero). The result is also windowed in frequency with a Tukey window (r=0.1) which is essentially rectangular but with slightly rounded edges.

Velocity channels whitened in this way exhibit constant amplitudes relative to each other, which is undesirable in incoherent processing where the relative amplitudes of the velocity is what helps determine the direction of a source. Thus, for velocity channels, the method of whitening normalizes the velocity components by the same normalization factor, rather than performing whitening independently on each channel. The normalization factor, $V(\omega)$, is given by

$$V(\omega) = \sqrt{v_x(\omega)v_x^*(\omega) + v_y(\omega)v_y^*(\omega) + v_z(\omega)v_z^*(\omega)} \tag{17}$$

where the frequency components of the velocity vector are $v_x(\omega)$, $v_y(\omega)$, $v_z(\omega)$, and the superscript $*$ denotes complex conjugation. The normalization factor is used to determine the whitened spectrum for the i–th velocity channel, $\tilde{v}_i(\omega)$, given by

$$\tilde{v}_i(\omega) = \frac{v_i(\omega)}{V(\omega) + \epsilon} \tag{18}$$

where $\epsilon \approx 2 \times 10^{-16}$.

B.4 Clipping

Clipping takes any data further than a specified value from the mean and replaces the data point with the appropriate positive or negative value from the mean.

The clipping process used in experiments takes the standard deviation of the data within a snapshot and clips any values beyond three times the standard deviation.

B.5 Filtering

Filtering without whitening is done after clipping to remove high-frequency peaks introduced by the clipping process.

The filtering is identical to the filtering with whitening, however no amplitude normalization is performed. The signal is merely windowed in the frequency domain by a Tukey window with $r = 0.1$, then returned to the time domain. To remove any possible filter transients, the time–domain signal is windowed in time with another Tukey window, this time with $r = 0.01$.

APPENDIX C

PROCESSING FLOW

This appendix describes the interaction of functions and code within the repository to a basic extent. Shown below is a legend for the blocks used in the following flowcharts.

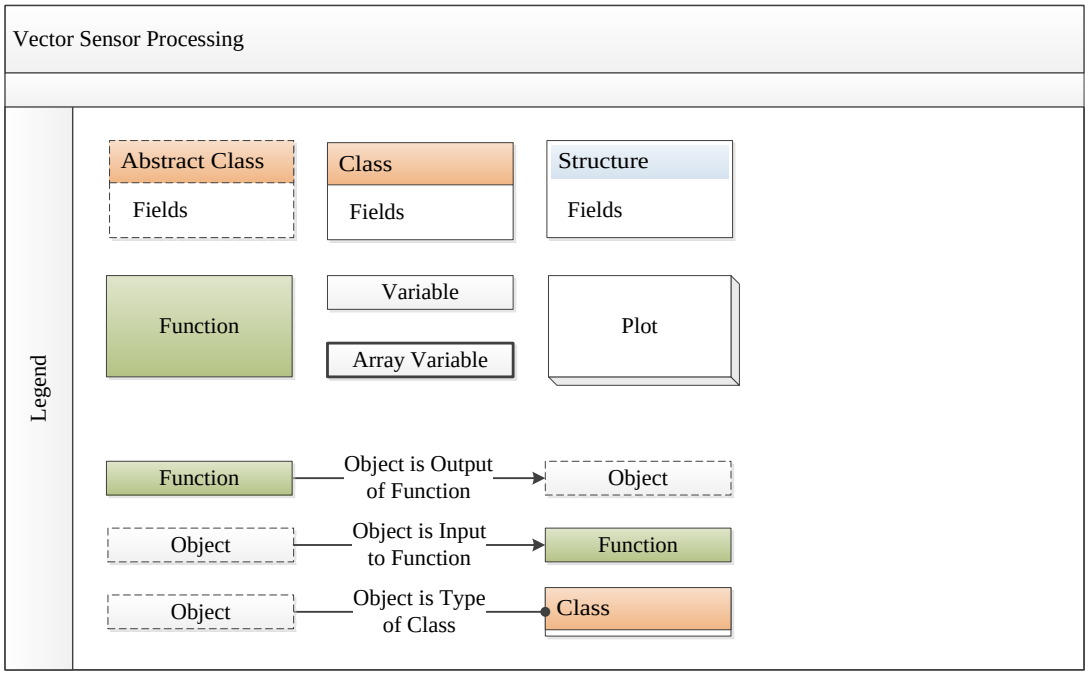

All functions pertaining to data loading and beamforming use some basic functionality provided by the Resource Manager (see below). The Resource Manager is responsible for maintaining information on entities (an abstract class), such as their name, display color and location. Both the boat and glider object are entities, however the glider adds pressure and three channels of velocity data as well as a structure for user–specific data. The locations of an entity are managed by the location class, which

stores the entity's reference location (x0, y0, z0), which in this case was the latitude, longitude and altitude of the reference coordinate system. The position of the entity over time is stored in channels x, y, z; the orientation of the entity (if required) is stored in yaw, pitch and roll. All channels mentioned so far are a separate class called "channel", which stores a descriptive name of the data contained within, as well as two arrays, one for time values and the other for the data values. All of these lower–level classes are used to store and manage the data of the gliders and boat, and allows the beamforming methods to operate on a consistent object.

To fill the data contained within an entity, there exist specialized classes which extend the abstract class "resource". The resource class stores a resource's name, its start time, and its end time. Each individual resource class is able to take a range of times as input and return the correct data from a resource file. The "resource mult." class is a resource itself, and contains a list of resources. When asked for data for a particular time window, the multiple resource class is able to concatenate data from any relevant sub–resources. Lastly, a class called a data record maintains a list of all resources within a data record (for example, all glider telemetry files). For a given time period, it is able to return the correct resource for which data exists during that time period.

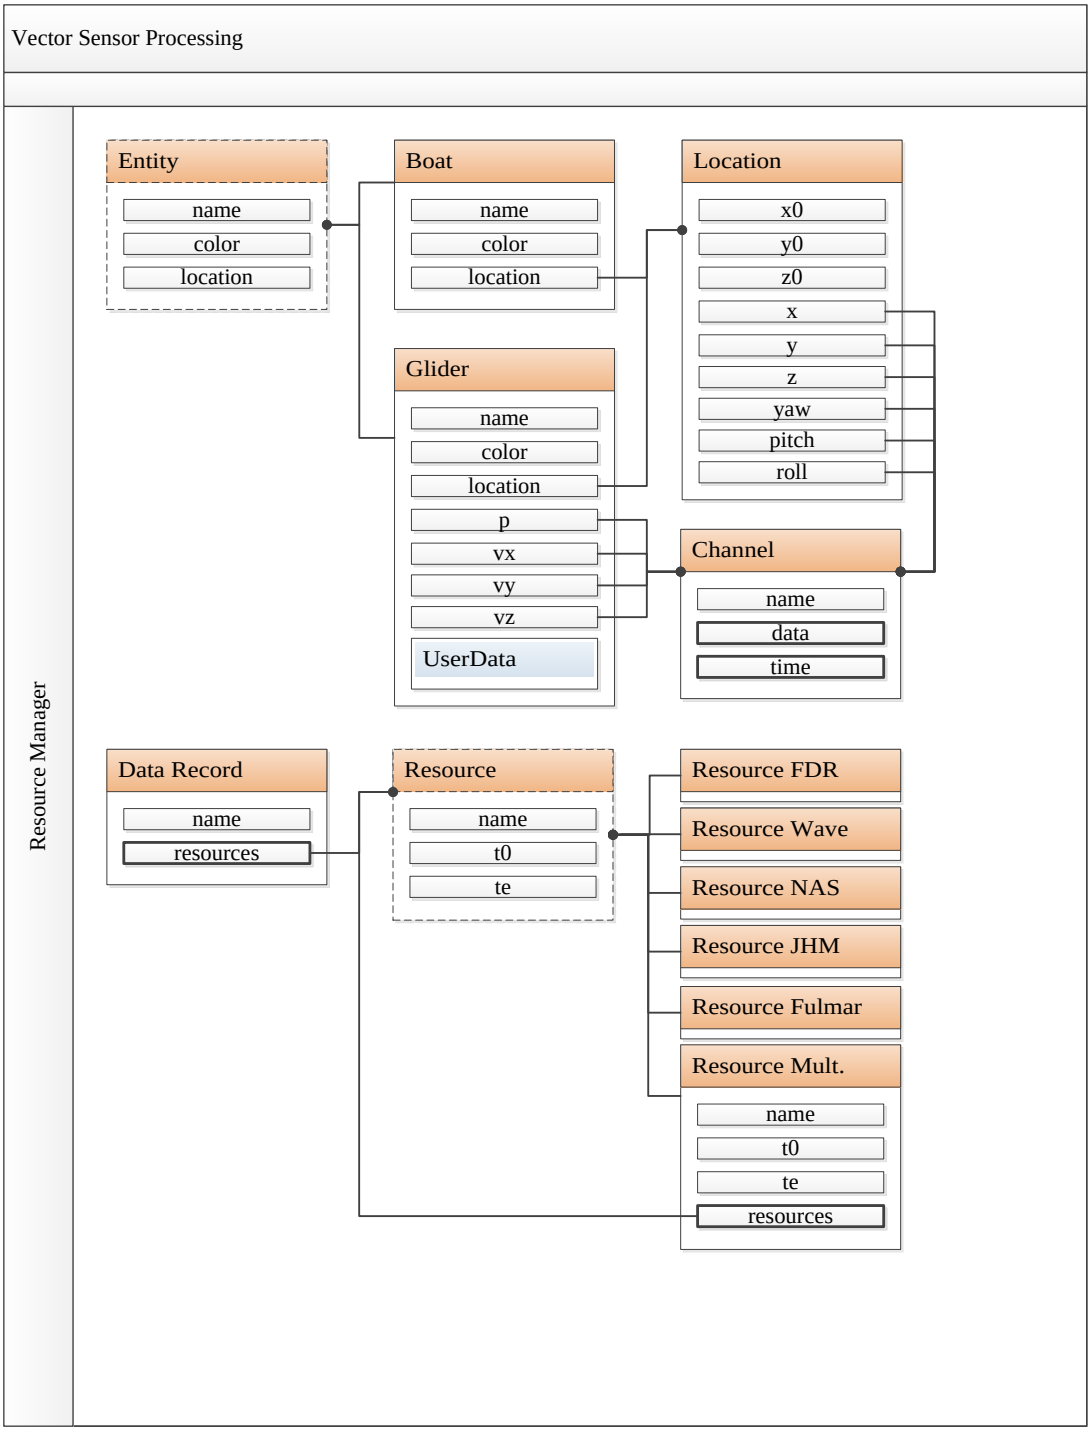

Functionality from the Resource Manager is used in the Data Loader (see below), which is responsible for creating the glider and boat objects with all their data loaded. From a resource definition filename (defsfilename), the function "loadResources"

maps filenames in certain folders to their proper resource type, packaging them into the acoustic data records, non–acoustic vector sensor data records, glider telemetry records, and research vessel records. Along with any definitions specific to the dataset, the data records are put into a resource structure called "R". At this point, all data pertaining to a specific dataset is contained within the resource variable "R". To then load data from a given dataset, its resource variable "R" along with a start time, end time, type of data to load, indices of gliders to load and an optional method of rotation are passed to "loadData", which returns an array of glider objects (g) and source objects (s). The resulting source and glider objects can be passed to plotting functions or the beamforming functions for processing, as their data is already been loaded and stored for the times requested.

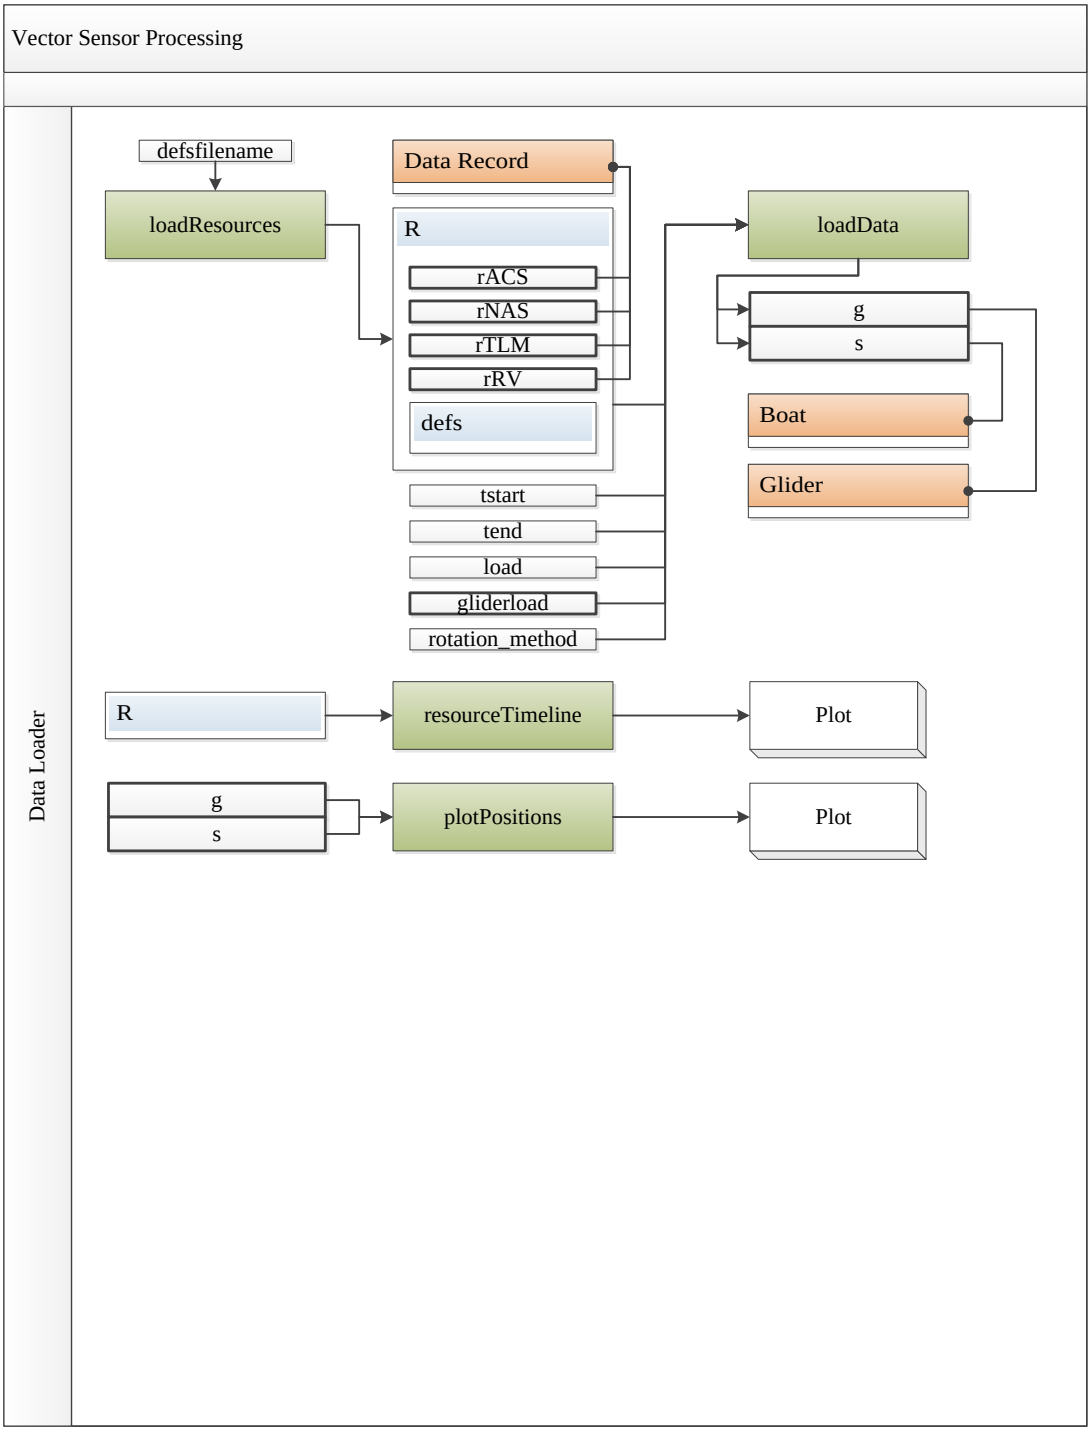

For any given snapshot, the glider and source data can be loaded using the Data Loader. To perform processing, the gliders can first be optionally passed to a pre-processing step to condition their data. Afterward, the gliders are passed to the

correlation function, which generates a covariance matrix for all frequencies (C). The covariance matrix is then passed to a weighting function which performs its weighting to the elements, and returns the updated covariance matrix. This matrix, along with a vector of the ambiguity surface points to evaluate, the glider array, the speed of sound, and the processing bandlimits are passed to the beamforming function, which returns time–domain beamformer outputs for each point in the ambiguity surface. The estimation function then performs time–windowing and summation to yield a scalar quantity for each point in the ambiguity surface (e). This ambiguity surface is then passed to the search function, which estimates the position of the source and returns it in "p". Given the estimated source position and the true source position, metrics of accuracy (acc) are returned. From the ambiguity surface (e), the precision can be found as well.

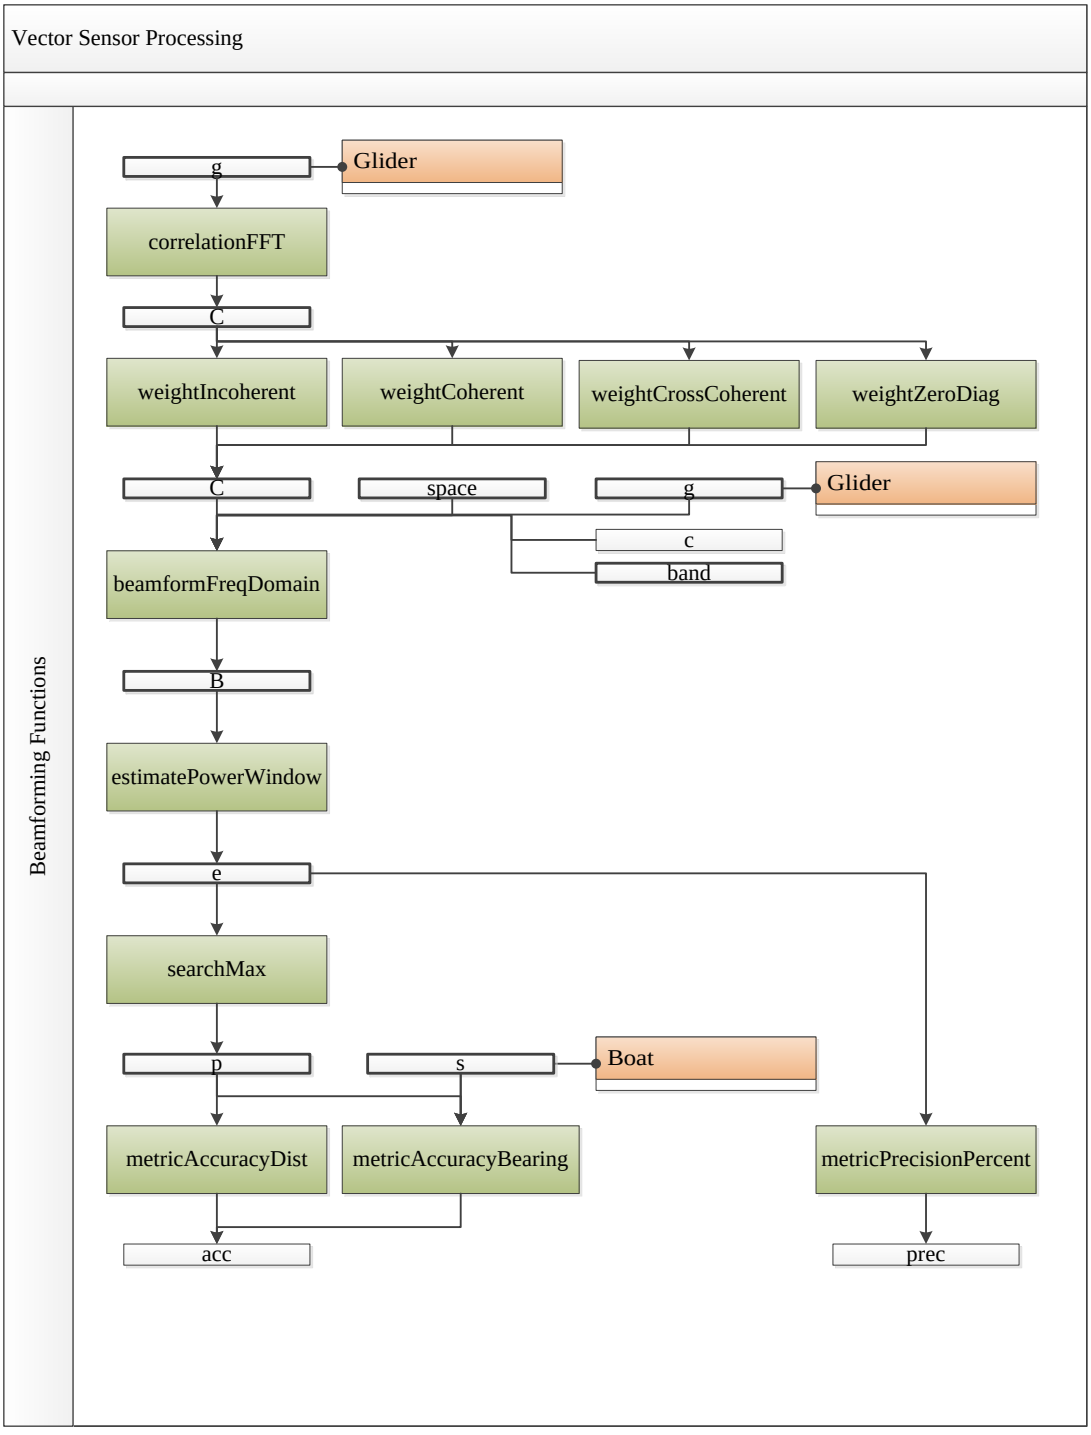
Vector Sensor Processing
Beamforming Functions
g
Glider
correlationFFT
C
weightIncoherent
weightCoherent
weightCrossCoherent
weightZeroDiag
C
space
g
Glider
c
band
beamformFreqDomain
B
estimatePowerWindow
e
searchMax
p
s
Boat
metricAccuracyDist
metricAccuracyBearing
metricPrecisionPercent
acc
prec

REFERENCES

Harry L. Van Trees. *Detection, Estimation and Modulation Theory, Part IV, Optimum Array Processing*. John Wiley & Sons, 2004. ISBN 978-0-471-09390-9.

A. Nehorai and E. Paldi. Acoustic vector-sensor array processing. *IEEE Transactions on Signal Processing*, 42(9):2481–2491, 1994. doi: 10.1109/78.317869.

M. Hawkes and A. Nehorai. Acoustic vector-sensor beamforming and capon direction estimation. *IEEE Transactions on Signal Processing*, 46(9):2291–2304, 1998. doi: 10.1109/78.709509.

Gerald L. D'Spain, James C. Luby, Gary R. Wilson, and Richard A. Gramann. Vector sensors and vector sensor line arrays: Comments on optimal array gain and detection. *J. Acoust. Soc. Am.*, 120(1):171, 2006. doi: 10.1121/1.2207573.

Henry Cox and Hung Lai. Simultaneous grating lobe and backlobe rejection with a line array of vector sensors. pages 1320–1323, 2009. doi: 10.1109/acssc.2009.5469917.

H. W. Chen and J. W. Zhao. Wideband MVDR beamforming for acoustic vector sensor linear array. *IEE Proceedings - Radar, Sonar and Navigation*, 151(3):158, 2004. doi: 10.1049/ip-rsn:20040651.

G. C. Lauchle, J. Wang, and M. S. Howe. Flow-induced noise on underwater pressure-vector acoustic sensors. pages 1906–1910, 2002. doi: 10.1109/oceans.2002.1191921.

M. Hawkes and A. Nehorai. Wideband source localization using a distributed acoustic vector-sensor array. *IEEE Transactions on Signal Processing*, 51(6):1479–1491, 2003. doi: 10.1109/tsp.2003.811225.

Y. Rockah and P. Schultheiss. Array shape calibration using sources in unknown

locations–part I: Far-field sources. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 35(3):286–299, 1987a. doi: 10.1109/tassp.1987.1165144.

Y. Rockah and P. Schultheiss. Array shape calibration using sources in unknown locations–part II: Near-field sources and estimator implementation. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 35(6):724–735, 1987b. doi: 10.1109/tassp.1987.1165222.

Peter M. Schultheiss. Optimum range and bearing estimation with randomly perturbed arrays. *The Journal of the Acoustical Society of America*, 68(1):167, 1980. doi: 10.1121/1.384616.

Joe Imlach and Ray Mahr. Modification of a military grade glider for coastal scientific applications. pages 1–6, 2012. doi: 10.1109/oceans.2012.6405134. URL `http://ieeexplore.ieee.org/ielx5/6387510/6404774/06405134.pdf`.

M. Hawkes and A. Nehorai. Acoustic vector-sensor correlations in ambient noise. *IEEE Journal of Oceanic Engineering*, 26(3):337–347, 2001. doi: 10.1109/48.946508.

John G. Grimes. Global positioning system standard positioning service performance standard, 2008.

J. C. Chen, R. E. Hudson, and Yao Kung. Maximum-likelihood source localization and unknown sensor location estimation for wideband signals in the near-field. *IEEE Transactions on Signal Processing*, 50(8):1843–1854, 2002. doi: 10.1109/tsp.2002.800420.

R/V John H. Martin specifications, 2011. URL `http://marineops.mlml.calstate.edu/JM-Specs`.

R/V Fulmar specifications, 2013. URL http://montereybay.noaa.gov/marineops/about/fulmar/specifications.html.