

ASMUTS – ACOUSTIC SIGNATURE MEASUREMENTS AND UNAUGMENTED TRACKING SYSTEM

September 2000

Edmond S. Sheffield
Naval Undersea Warfare Center Division
Newport, RI. 02841-1708

ABSTRACT

A real-time passive method for bearing estimation has been developed for underwater acoustic targets that capitalizes on apriori knowledge of the target's radiated acoustic signature. The sensor system is a horizontal planar array (HPA) composed of nineteen (19) hydrophones, configured as three nested hexagons, with a redundant center hydrophone. The data telemetry system supports a frequency range of 10 Hz – 1600 Hz. Data is processed with 1 Hz resolution (1590 frequencies). Azimuth uncertainty estimates and spatial associations are used to remove frequencies dominated by noise. Azimuths for potential target frequencies are averaged in bandwidths proportional to frequency, temporally and spatially associated, and correlated against expected signatures. Simulations in MATLAB with spatially correlated noise show good performance down to 6 dB signal-to-noise ratio (SNR) for an 8-tonal narrowband target.

Form SF298 Citation Data

Report Date <i>("DD MON YYYY")</i> 00092000	Report Type N/A	Dates Covered (from... to) <i>("DD MON YYYY")</i>
Title and Subtitle ASMUTS--Acoustic Signature Measurements and Unaugmented Tracking System		Contract or Grant Number
		Program Element Number
Authors Sheffield, Edmond S.		Project Number
		Task Number
		Work Unit Number
Performing Organization Name(s) and Address(es) Naval Undersea Warfare Center Division Newport, RI. 02841-1708		Performing Organization Number(s)
Sponsoring/Monitoring Agency Name(s) and Address(es)		Monitoring Agency Acronym
		Monitoring Agency Report Number(s)
Distribution/Availability Statement Approved for public release, distribution unlimited		
Supplementary Notes		
Abstract		
Subject Terms		
Document Classification unclassified		Classification of SF298 unclassified
Classification of Abstract unclassified		Limitation of Abstract unlimited
Number of Pages 12		

1.0 INTRODUCTION

The HPA is used to passively track multiple targets in azimuth using their radiated noise. The goal of the signal processing and tracking algorithms is to track targets with one degree of accuracy in azimuth at an SNR of 6 dB, as represented in Figure (1). Target signatures are known apriori as a function of both aspect and speed. The basic problem for the bearing processor is that while SNR performance at long range indicates processing in very narrow (one Hertz) bands, Doppler shift and multipath effects degrade performance of bearing estimation in such narrow frequency bands. The approach used by ASMUTS is to process with one Hertz resolution, then use spatial (azimuth and elevation) and temporal filters to identify frequencies that may be dominated by signal versus noise before frequency averaging to the target azimuth. The algorithms are designed to yield the most likely azimuth for a target as well as a metrics to determine if the target is actually present.




Figure (1). Target Spectrum and Noise Spectrum

1.1 HPA Configuration

The HPA hydrophone configuration is shown in Figure (2). For the purposes of processing, the configuration is considered as three horizontal line arrays (axes) that intersect at 60-degree angles. The spacing between adjacent hydrophones in an axis is 18.75 inches (0.476 meters), which is one-half wavelength at 1600 Hz for a sound velocity of 5000 ft/sec (1524.2 m/s).




Figure (2). The Hydrophone Configuration

1.2 The Sensor Data

Figure (3) shows the overview block diagram of the signal processing and tracking, starting with the phases for the 20 HPA hydrophones. The signals are digitized in the water, simultaneously sampled at 4096 samples/second, multiplexed into a digital data stream and transmitted over a radio frequency (RF) link to a support boat. The HPA data is extracted from the data stream and sent to a multi-processor digital signal processing (DSP) board. The DSP performs a 4096 point Fast Fourier Transform (FFT) for the simultaneous data from each of the twenty hydrophones. The magnitude of the FFT is ignored, while the phase at each one-Hertz resolution frequency is used to determine target bearings.

1.3 The Processing Concept

The raw acoustic data is processed narrowband, with one-Hertz resolution, to take advantage of as much processing gain as possible. The expected characteristics of the target and noise field are used to select and average only frequencies associated with the target at any given instant, even in the presence of multipath Doppler shifts. Target association is accomplished in azimuth, frequency, and time through several scoring systems operating as multiple hypothesis testers. The data is first operated on as a single second estimate in a one-Hertz band. Then the narrowband data is filtered and grouped into bands and examined over time. Finally, the frequency bands are associated in azimuth and correlated against expected target signatures. An operator judges the quality of the data for the target and assigns trackers.




Figure (3). Bearing Tracking Overview

2.0 DETAILED PROCESSING DESCRIPTION

The details of each step in the bearing processing shown in Figure (3) are covered in the following sub-sections.

2.1 Once Per Second Bearing Estimation

This step is accomplished in the DSP. The input is the raw phases for the hydrophones at each one-Hertz resolution frequency bin. The outputs are estimates of the azimuth, elevation, and uncertainties in the estimates for each frequency bin. The specific steps are as follows:

2.1.1 FFT of Simultaneously Sampled Hydrophones

- Identify “bad” hydrophones by wide-band variance comparisons between hydrophones

2.1.2. Unwrap Phases Using 2π Multiples and Wavelength Versus Spacing.

- Account for previously identified “bad” hydrophones
- Establish maximum frequency without spatial aliasing for each axis

2.1.3. Compute the Trace Wavenumber (k_t) for each axis as the Linear Regression of the Phases.

- k_t is the Δ phase (radians/ meter) along the axis at each frequency
- Compute the standard error in the regression line estimate and trace wavenumber.

2.1.4. Compute Solid Angles for the Three Axes (S_1, S_2, S_3) from their Trace Wavenumbers

- Identify non-acoustic frequencies ($k_t > k$) per axis where possible.

$$S_{\text{axis}} = \cos^{-1}[k_t / k] = \cos^{-1}[k_t \cdot c / 2\pi f] \quad (1)$$

where k_t = Observed Trace Wave number (rad/m)

k = Acoustic Wave Number (rad/m) at the HPA depth

f = Frequency (Hz)

c = Sound Velocity (m/sec) at the HPA depth

2.1.5 Compute the Azimuth from the Three solid Angles

Three solid angles provide three possible solutions for azimuth and elevation, using axes A-B, B-C, or A-C. The solutions using A-B and A-C are averaged, ignoring the BC solution, which is completely dependent on the other two. Figure (4) shows an example of the intersection of the three solid angle cones for the HPA with a unit sphere projected onto the horizontal plane. The dashed lines show the intersection of the solid angle cones with the horizontal plane. The solid lines perpendicular to each axis represents the intersection of the solid angle cone and the unit sphere, projected onto the horizontal plane.

The intersection of the solid lines for axes A-B and A-C are computed in XY coordinates relative to the center of the unit sphere and axis-A as the X-axis. The X coordinates and Y coordinates of the two solutions are averaged. The azimuth is the inverse tangent of the average Y divided by the average X. The actual point of intersection lies on the 3-D unit sphere, but the elevation is the inverse cosine of the distance from the point of XY intersection to the origin.




Figure (4). Top View of Solid Angle Intersections Projected with Horizontal Plane

2.1.6 Compute the Azimuth and Elevation Uncertainty

This process relies on a simplified representation of the array as four hydrophones in two complementary equilateral triads, as shown in Figure (5).




Figure (5). Simplified Array Representation

The phases for the four representative hydrophones are assigned as the average slope with error based on the regression line and the standard errors in phase for each axis as shown in Equation (2). Zero is used as the reference phase of the center hydrophone in the simplified array. The azimuth is then computed from the phases using a simple formula for an equilateral triangle. The result is compared to the azimuth computed from the regression line and the difference between the two is the uncertainty estimate.

$$\begin{aligned}
 \varnothing_1 &= 0 \\
 \varnothing_2 &= k_{tA} L - (\sigma_A / \sqrt{n_A}) = 1.43k_{tA} - (\sigma_A / 2) \\
 \varnothing_3 &= k_{tB} L + (\sigma_B / \sqrt{n_B}) = 1.43k_{tB} - (\sigma_B / 2) \\
 \varnothing_4 &= k_{tC} L + (\sigma_C / \sqrt{n_C}) = 1.43k_{tC} - (\sigma_C / 2)
 \end{aligned} \tag{2}$$

Where $\varnothing_n$ is the phase on the n_{th} representative hydrophone.

k_{tA} is the slope of the regression line (Δ phase / meter) for Axis-A.

σ_A is the standard error about the regression line for Axis-A.

The equations for azimuth based on the relative phases of the equilateral triangles 1-2-3, and 1-2-4 of the representative array are given in equation (3), below.

$$\begin{aligned}
 AZ_{123} &= \text{TAN}^{-1} \left[\frac{\varnothing_2 + 2 \varnothing_{23}}{\sqrt{3} \varnothing_2} \right] \\
 AZ_{124} &= 0 - \text{TAN}^{-1} \left[\frac{\varnothing_2 + 2 \varnothing_{24}}{\sqrt{3} \varnothing_2} \right]
 \end{aligned} \tag{3}$$

$$\begin{aligned}
 \text{Where: } \varnothing_{23} &= \varnothing_3 - \varnothing_2 \\
 \varnothing_{24} &= \varnothing_4 - \varnothing_2
 \end{aligned}$$

The azimuth error at zero elevation ($AZERR_0$) is taken as the RMS of the error for the two triads in the representative array as shown in Equation (4), below.

$$AZERR_0 = \sqrt{((AZ_{123} - AZ)^2 + (AZ_{124} - AZ)^2) / 2} \tag{4}$$

The azimuth error at zero elevation is converted to azimuth error at the observed elevation estimate ($AZERR_{EL}$) and elevation error at the observed elevation (estimate) by the simple relationships listed in Equation (5).

2.1.7 Compute the RMS Azimuth and Elevation Errors for the observed elevation

$$\begin{aligned} \text{AZERR}_{\text{EL}} &= \text{AZERR}(\text{elev}=0) / \cos(\text{EL}) \\ \text{ELERR}_{\text{EL}} &= \text{AZERR}(\text{elev}=0) / \sin(\text{EL}) \end{aligned} \quad (5)$$

2.2 Spatial & Quality Filter

The tracking range operates as four planar arrays deployed in 200-600 meters of water as a rectangle, with an HPA in each corner. The targets are confined to the interior of the rectangle, maintaining at least one-kilometer “safety zone” about each HPA. The defined “boundary” for a given HPA is the 90-degree sector in azimuth and less than 45-degrees elevation. The data for each HPA is examined and frequencies that define targets that are most likely outside of the range are labeled as “bad”. A target’s azimuth and elevation with the estimated uncertainty defines a range of values that must lie, at least partially, inside the range boundaries.

For example: The range has boundaries of 0-90 degrees in azimuth relative to the HPA. A certain frequency has an azimuth estimate of 120 degrees with azimuth uncertainty of 20 degrees, so the target is most likely somewhere in azimuth between 100 and 140 degrees, which is entirely outside the range boundary. That frequency is considered “bad”. Another frequency also has an azimuth estimate of 120 degrees but with azimuth uncertainty of 35 degrees, so that target is likely somewhere between 85 and 155 degrees, which includes azimuths inside the range boundary. That frequency would be “good”, even though the error in the azimuth estimate is larger than the first frequency.

A second filter labels frequencies as “bad” that have excessive uncertainty in the azimuth estimate. The goal is to track with one degree of bearing accuracy using as few as four seconds of averaging and as few as four frequencies. This implies that the uncertainty in a single second’s data for a single frequency should be less than about four degrees (with assumptions of independent normally distributed estimates over time and frequency). The threshold value used to identify “bad” frequencies is typically assigned between 8 and 12 degrees of azimuth uncertainty, depending on the expected signature. The elevation error is not used as a filter because it is related to the azimuth uncertainty.

2.3 Narrowband Frequency Association

The frequency spectrum from 10 Hertz to 1600 Hertz is divided into 400 frequency bands. Each band has a bandwidth of 1.5 % of the lowest frequency in the band and 50% overlap with the two adjacent bands. There are never less than two narrowband frequencies assigned to a band. Doppler shifts for moving targets are accommodated by the proportional bandwidth before correlating the observed frequencies with the stored replica frequency sets. The 50% overlap insures that any specific frequency is always present in at least one band regardless of Doppler shift.

The Narrowband Frequency Association Algorithm is designed to selectively average the azimuths for the narrowband frequencies in the band on a second to second basis. The narrowband azimuths are weighted inversely proportional to the uncertainty estimates of the azimuths. Zero weight is assigned for any narrowband frequencies labeled as “bad” in previous steps, any frequency not associated in azimuth with the “best” frequency in the band, or any frequency with azimuth uncertainty more than twice that of the “best” frequency in the band.

The first step is to select one frequency as the most likely to be associated with any acoustic target. To do that, a scoring system is used to rate each narrowband frequency in the band by assigning points for low elevation, low azimuth error, and agreement in azimuth with the band output for each of the previous three seconds. The azimuth agreement with the band output is measured in degrees and is weighted against time so that the most recent second is the most significant. Initially zero points are assigned for history comparison when examining the first second of data, so the selection of “best” frequency in the band depends mostly on low azimuth uncertainty. This is equivalent to picking out the frequency with the highest SNR in the band. Once several seconds have been examined, consistency within the four seconds is the driving factor.

Once each narrowband frequency is “scored” and the “best” frequency is selected, the score of that frequency is compared to a minimum score threshold. If the score does not exceed the threshold, then the band is labeled as “bad”. Every narrowband frequency in the band is tested against the best frequency to determine if the azimuths (with uncertainty) overlap. If they do, and if the azimuth uncertainty of the other frequency is less than twice the azimuth uncertainty of the best frequency, then the azimuths of the frequencies will be averaged together and the combined uncertainty is reduced. The individual frequency azimuths are weighted inversely proportional to the azimuth uncertainty in the band average.

2.4 Time Association for Bands

The purpose of this step is to select from the current second and the previous three seconds either one or two seconds of azimuths to average for each frequency band. No more than two seconds of band azimuths are ever selected. A scoring system similar to that described in sec 2.3 is used for the narrowband frequency association. The time association assigns point values to the possible pairs of seconds for both agreement in azimuth and for the weighted average azimuth uncertainties. The weighting is inversely proportional to the azimuth uncertainty for each band. If only one second, of the four, is considered “good” then it is automatically selected. If two seconds are selected, then the azimuth uncertainties must be within a factor of two, otherwise only one second will be used.

2.5 Replica Correlation of Band Groups

The output of the time association algorithm is a list of frequency bands that may contain a target along with the azimuth estimate and azimuth uncertainty estimates for each band. The next step is to determine which subsets of frequency bands may be associated with some target. To accomplish that, the frequency bands are associated in azimuth as illustrated in Figure (6).

An azimuth association “window” is scanned through the azimuths that are within the range boundaries relative to the specific HPA. At each increment, the limits of the azimuth estimate (azimuth +/- uncertainty) for each time averaged frequency band is tested to determine if it crossed into or across the association window. The list of the associated frequency bands and their respective azimuth uncertainties at each increment is stored in a table. The association window has a finite width, typically 4-8 degrees wide, and scans in increments equal to half its width. For limited numbers of targets and frequencies, it is possible to greatly reduce the increments by deterministically indexing through the possible associations based on the minimum and maximum azimuths for each band.




Figure (6). Azimuth Association

2.6 Replica Correlation of Band Groups

The purpose of this step is to determine the most likely azimuth(s) for the target considering that the frequencies observed at any instant in time may change. The current design limits the number of frequency sets that are assigned trackers to only three, so that the tracks can be displayed and evaluated by an operator in real-time.

Each target has as many as 20 frequency sets that identify depending on the aspect and speed dependence of the signature. Each frequency set is a list of up to 25 of the 400 bands, as described in sec 2.3. Typically 12-20 sets of frequencies are used to define a single target, with commonality between the frequency sets. Twenty frequency sets would cover five aspects and four speeds.

The actual signature is the source sound pressure level in decibels relative to one microPascal versus narrowband frequency. The signature must be converted to what is actually observed by the HPA processing algorithms, which are the azimuth and azimuth error for each band. To convert the target spectra in bandwidth is a simple power summation, but the conversion from source sound pressure level to azimuth uncertainty requires some effort.

MATLAB was used to run Monte-Carlo simulations of direct path target signals in a spatially coherent noise field. The resulting azimuth error versus frequency and SNR for the HPA is shown in Figure (7). For a prospective target signature, the highest level frequency band is used as a level reference of 6 dB, then all frequency bands are converted to azimuth error by the relative level and frequency through curve-fitted equations from Figure (7). The result is scaled such that frequency band with the least azimuth error is scaled to one degree of azimuth uncertainty.




Figure (7): Azimuth Error versus Frequency and SNR

As the association band is iterated through azimuth, a table of correlations between the observed frequencies and the replica frequencies is constructed. The correlation value for a frequency set is the sum of the correlations for each frequency in the set. The correlation for a single frequency is the product of the inverse of the observed and replica azimuth uncertainties. If a frequency band is present in both the association window and the replica frequency set and both have an azimuth uncertainty of one degree, then the correlation for that frequency is 1.0. If the observed and replica azimuth uncertainties are both two degrees, then the correlation for that frequency is $1/2 \times 1/2 = 1/4$.

When the association band has been scanned across all the azimuths within the range boundaries, the top three correlations are saved as the frequency set number, correlation value, azimuth, and azimuth uncertainty. The top-three correlations are maintained in a table for a moving 8-16 second time interval. A scoring system is used to rate each frequency set over the time interval to determine the best three frequency sets to use at any given time. The time window keeps the choice of the best three frequencies from changing too fast, so history and statistics can be obtained for a frequency set, while allowing for change over time.

3. RESULTS

The processes described have been implemented in MATLAB and tested on simulated acoustic data for two stationary targets composed of eight frequencies each. The first target was assigned an azimuth of 30 degrees, elevation of zero degrees, at frequencies of 100, 300, 500 ... 1500 Hz. The second target was assigned an azimuth of 60 degrees, elevation of 10 degrees, at frequencies of 200, 400, 600, ... 1400 Hz. Spatially coherent ambient noise was added at an SNR equal to 6 dB. The resulting azimuths derived for the two targets is shown in Figure (8), below. As a control, the algorithms attempt to track a third target that is not actually present in the simulated acoustic data (triangles). Other simulations have been performed with two moving targets with crossing paths, various SNR, and broadband targets. The single most significant limitation in the processing is that multiple targets must not be tracked using the same frequency bands (of the 400), otherwise the targets may be indistinguishable or averaged in azimuth, depending on the circumstances.




Figure (8). Azimuth Tracking Results