Matched Doppler Processing

Stephen Searle

Thesis submitted for partial requirements for the degree of
Master of Applied Science

Dept. of Electrical and Electronic Engineering
Faculty of Engineering
University of Adelaide
South Australia
August 1997

Contents

List of Figures

List of Tables

Abstract

This thesis describes a method for estimating the ranges and speed of a nearfield narrowband sound source as it moves past a receiver. This is applied to analysis of source motion in the underwater acoustic environment.

The technique correlates or "matches" the receiver output with a large set of replicas. A replica is generated for each combination of possible candidate source ranges and speeds. These replicas are generated according to a simple propagation model which takes into account the doppler effects inherent in the motion of the theoretical source.

Particular emphasis is placed upon the treatment of signal phase. This is carefully modelled in the replica by considering the differential time–lag and DFT cell distortion induced by the doppler effect. The matching functions exploit this phase information by processing it coherently.

The matched processing algorithm is tested with both simulated data and real data from a sonobuoy. Performance bounds of the estimators are derived and discussed in context with the simulation results. It is shown that the phase–coherent matching functions satisfy these bounds. The robustness of the technique to variance in the source trajectory is investigated through further simulation.

The technique is extended to perform tracking of source range, azimuth, speed and heading with the use of multiple receivers. A matching function which processes phase coherently across space as well as time is presented. This is compared to a spatially–incoherent formulation via extensive simulation.

Publications

- S.J. Searle & D.A. Gray, "Environmental Source Tracking and the Doppler Shift", presented at ISSPA–96, August 26–28, 1996.
- S.J. Searle & D.A. Gray, "Matched Doppler Processing for Estimating Moving Source Parameters", submitted to the Journal of the Acoustical Society of America, April 1997.

Statement of Originality

This work contains no material which has been accepted for the award of any other degree or diploma in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text.

I give consent to this copy of my thesis, when deposited in the University Library, being available for loan and photocopying.

SIGNED:

DATE:

Acknowledgements

The author wishes to thank the following people and organisations:

- Prof. Barry Quinn, for providing the opportunity and initial motivation to undertake postgraduate study.
- Prof. Doug Gray, for his "expert and highly competent supervision" and his unwaning enthusiasm for the research.
- The various other students supervised by Doug Gray for their discussions and support. Particular thanks go to John Legg for much useful discussion about the Cramer–Rao lower bound.
- The Co–Operative Resarch Centre for Robust and Adaptive Systems for providing funding to support this study.
- The Co–Operative Research Centre for Sensor Signal and Information Processing for providing resources with which to pursue research.
- Various members of Maritime Operations Division, DSTO Salisbury, including Dr. John Riley, Dr. Lesley Kelly, Dr Shane Tonissen; for useful discussions and access to real data.
- Dr. Alan Bolton from MRD, DSTO Salisbury for the interest shown in this work and some insightful discussion.

- Dr. Jane Perkins (MOD) and Dr. Michael Greening for support provided within the workplace and for being a willing audience on which to practice oratory skills.

- Finally, Dr. Peter Kootsookos who, unbeknownst to himself, taught the author more than he ever wanted to know about LaTeXby leaving some source code on the MOD machines.

The author wishes to acknowledge the funding of the activities of the two aforementioned Co–Operative Research Centres by the Australian Government under the Co–Operative Research Centre Program.

Chapter 1

Introduction

1.1 Analysis of Source Motion

Determination of the track parameters of a moving signal source is a problem that arises in many areas of signal processing. In sonar, Target Motion Analysis (TMA) is a traditional passive method for inferring the parameters which describe a moving source. Using only estimates of the target's bearing, or bearing and frequency, TMA endeavours to find the target trajectory that best fits the observed time history of bearings and frequencies [33]. Methods exist for tracking both single and multiple targets which can be either broadband or narrowband sources. Usually the target trajectory is assumed to be straight line and constant speed. TMA can be regarded as a nonlinear state estimation problem. Kalman filters and batch–style iterative algorithms have both been used extensively in TMA [17, 19].

Observability is an issue in TMA. The quality of estimates produced by TMA are highly dependent on the source–observer geometry and measurement noise levels [17]. Another important issue in TMA is that of observer motion. The motion of a target cannot be properly observed via bearings–only measurements until the observation platform itself undergoes motion [18]. This prohibits the use of a stationary arrangement of receivers (e.g. a field of sonobuoys) from performing bearings–only TMA. Source mo-

tion is usually assumed to consist of constant speed straight line segments, due to ship motion considerations. However LeCadre and Gauvrit [34] have considered optimisation of observer trajectory in bearings–only TMA, demonstrating the interest of manoeuvre diversity.

Source motion information may also be inferred from received frequency alone, as this will be altered from the emitted frequency according to the well–known doppler effect. Using high–resolution frequency estimates, a parametric expression for received frequency [4] or frequency and amplitude [3] is numerically optimised to yield estimates of speed, range and rest frequency of a constant–velocity, straight–line moving narrowband source.

A benefit of this approach is that the receiver need not be in motion as the expression for received frequency (and hence the optimisation surface) is dependent only on relative source/receiver movement. With the use of three or more sensors, each of which comes into contact with the narrowband source, the doppler parameters estimated at each sensor may be processed to yield estimates of the source trajectory [23, 21].

These TMA techniques depend on the ability to form accurate measurements (of signal bearing, frequency, or both). Thus they can be extremely susceptible to the presence of strong measurement noise.

Recent advances in Matched Field Processing (MFP) [1] have proposed matching techniques for estimation of source motion parameters in the underwater acoustic environment [5, 13]. MFP can be thought of as an extension to plane–wave beamforming [1] and can be traced back to Bucker [30]. It is able to localise a sound source in range, depth and bearing.

The general MFP algorithm is simple. Initially measurements are collected from a number of sensors, frequently a vertical line array. For each unknown parameter (range, depth, bearing) a finite list of candidate values must be nominated. This defines a multi–dimensional search space. For each element of this search space (that is, each combination of candidate parameter values) a propagation model is used to generate a

replica signal. This replica is a model of the signal one would expect to receive if the true system parameters were the same as the assumed parameter combination. Each replica is "matched" with the measurement by use of a correlation or cost function. The replica which produces the best match (i.e. highest correlation or lowest cost) is deemed to be correct and the parameters used in its generation are taken as estimates of the true system parameters.

MFP employs an extremely detailed signal model to form signal replicas. It is because of this that MFP is able to estimate source positions in range, azimuth and depth, unlike plane–wave beamforming which can estimate only source bearing. Alternatively if the source's true position is known then parameters of the environmental model may be estimated; this process is known as environmental inverse modelling [11]. Various formulations exist for modelling of the acoustic field. Most can be described broadly as normal mode models, ray theory models, and parabolic equation models.

Matching functions used by MFP include the Bartlett processor, which is of the same form as a sample–matrix frequency domain beamformer. MFP assumes that the source is approximately stationary during the estimation period. Hence long time samples may be used to form time–invariant covariance matrices for the matching function. Standard stationary–source MFP techniques do not scale up well to the moving–source problem:

> "The central dilemma of applying MFP to moving sources is the integration time must be short enough to ensure the source is quasistationary during estimation time, but long enough to allow a reasonable number of time segments to be used in covariance estimation. Furthermore these segments must be long enough to yield acceptable frequency resolution ... to achieve ... narrowband gain" [5].

The standard approach taken when addressing this dilemma is to model the effect of source motion on the acoustic field, thus producing more realistic replicas and removing the dependence on quasistationarity. In principle integration time can be arbitrarily long when motion has been modelled in the replica, but is limited in practice.

Several existing propagation models have been extended to include source motion effects. Jacyna, Jacobson and Clark [26] presented a ray–theory model demonstrating that source motion could be handled by using an approximate doppler frequency in place of the source frequency. In a similar vein, Neubert [27] used existing normal–mode software to compute the acoustic field generated by a moving source simply by modifying the inputs. This involved perturbing the input frequency (or the sound speeds in a multilayered environment) by a factor which accounted for doppler dispersion.

A more thorough study presented by Hawker [28] derived an expression for the time-series received at a point in the acoustic field in terms of the normal modes of a horizontally stratified ocean. Doppler effects on the timeseries were also considered, in particular as viewed through a finite bandwidth transform. The results differed from Neubert [27], incorporating the modal–dependent doppler shifts.

Matched–field tracking methods proposed by Chen and Lu [22] and Zala and Ozard [5] also employed normal–mode models of the acoustic field. In fact the Zala & Ozard study employed the Hawker model mentioned above. Replica signals of sources moving in a straight line and with constant velocity were generated via these extended propagation models and matched with receiver measurements. Matching functions included a Bartlett processor and some eigenvector–matching variants.

A specific extension of MFP which determines the track of a moving source is known as Environmental Source Tracking (EST). [2] This technique uses the output of a single receiver to form a single measurement of signal amplitude in each timeblock of a sampling interval. This measurement vector is matched with replicas via a squared–difference cost function which is minimised over the parameter space. The replicas themselves are computed according to a parabolic equation propagation model and cylindrical propagation loss is assumed. The technique is very simple and uses sparse measurements of the acoustic field compared with more traditional MFP algorithms which may use an array of receivers and measurements from several frequency bins. However this technique could successfully localise a moving narrowband source in range, speed and depth, using only amplitude measurements from a single receiver and single frequency bin.

The EST technique demonstrates that extensive sampling of the acoustic field is not necessary in order to achieve successful localisation of a moving sound source. Yet it is obvious that simple modifications, such as the inclusion of phase information, may enhance the performance of this technique considerably.

Complex propagation models such as normal–mode formulations can provide thorough models of the acoustic field generated by a moving source. However the computational load of replica generation can be very high when a detailed model is used. This is especially so when several frequency bins and several receivers are used as more information must be generated.

The matched processing method described in this thesis uses a simple model of sound propagation. As in previous studies, source motion is accounted for in the propagation model by consideration of the doppler effect. However the model employed is not a complex normal–mode or ray–theory formulation. Instead direct path propagation and cylindrical power attenuation are assumed, thereby enabling simple and efficient generation of replicas. As in the EST approach, sparse sampling of the acoustic field is performed. However phase information is included in the signal measurement. The coherent processing of this phase information is stressed as it provides far better performance than simple amplitude matching. TMA techniques which rely upon frequency estimates can perform poorly at the low frequencies which are traditionally used in passive sonar processing due to noise on the estimates. A technique based on coherent phase processing is less susceptible to this and should thus perform well on sonar applications.

1.2 Thesis Outline

This thesis describes a phase–coherent matching technique for estimation of the motion parameters of a nearfield narrowband source in the underwater acoustic environment.

In order to tackle this problem it is necessary to employ concepts of Estimation Theory, Coherent Processing, Beamforming and Matched Field Processing. Fundamentals of these concepts are discussed in the next chapter of this thesis.

This is followed by formal models of the signal, with special attention paid to signal phase. The effect of differential time–lags and doppler effect on signal measurement is discussed and incorporated into the signal model. Several cost functions for matching signal measurements with modelled replica measurements are defined.

This thesis focusses on the performance of the matched processing cost functions. It does not consider their optimisation. All analyses are performed by an exhaustive search over a discrete domain of candidate parameters.

Each of the cost functions is tested with simulated data and the results are discussed. This leads into a discussion of performance bounds exhibited by the matching technique. Constraints on estimator accuracy imposed by the discrete nature of the domain are determined. The results of simulation are found to conform to these bounds and to the Cramer–Rao lower bound. A comparison with an existing doppler–tracking technique [3] is performed.

The robustness of the technique to phase errors induced by nonlinearities in the source motion is investigated by further simulation. It is found that the technique is robust to small track wobbles provided that enough samples are made in order to fully encompass the doppler shift of the signal. In light of this result, the technique is applied to real data with some success.

Ultimately the extension of the matched doppler processing technique for use with multiple receivers is considered. A method which coherently processes phase across time and space is considered and compared to an application of the single receiver coherent–across–time approach.

1.3 Summary of Contributions

The contributions made by this study are:

- A simple model of signal phase derived through consideration of the varying time–lag experienced by signals undergoing a doppler shift.

- Use of minimal information in the matching process. In each timeblock, only one measurement of the amplitude and phase of the acoustic field is taken.

- Coherent processing of phase. This study builds on the findings of Collins et al [2] by inclusion of phase information in the cost function. It is demonstrated that phase information can be exploited to yield better performance when processed in a coherent manner.

- Coherent processing of doppler–affected data across multiple receivers.

- Application to Real Data. The coherent doppler phase matching algorithm described in this thesis is used to analyse real data from a sonobuoy.

Chapter 2

Background

This chapter introduces some concepts used later in this thesis: desirable properties of estimation procedures, matched filtering and coherent–phase processing, conventional beamforming processors and their properties and fundamental concepts involved in matched field processing.

2.1 Estimation Theory

Analysis of sonar signals for the determination of source parameters is fundamentally a problem of *estimation*. The goal is to determine the values of some unknown system parameters (i.e. source range and velocity) through the observation of a process which is dependent on these parameters (i.e. the signal) in the presence of noise. It is assumed that a source exists and that information from this source is present in this signal. The problem is not one of detection.

A good estimator is unbiased, consistent and efficient. That is to say, the expected value of an estimate $\hat{\theta}$ of an unknown value θ is that value θ (unbiased); the variance of the estimate reduces as more observations N are made, i.e.

$$\lim_{N\to\infty} E\left\{\left(\hat{\theta}-\theta\right)^2\right\} = 0$$

(consistency); and the mean–squared error of the estimator equals an optimum level such

as the Cramer–Rao Lower bound (efficiency).

Various forms of estimation exist; the difference being largely governed by the choice of error criterion. Procedures include minimum mean square error (MMSE) estimation and maximum a–posteriori (MAP) estimation. However it is the Maximum Likelihood (ML) estimator which is of particular interest to us.

Unlike the MMSE and MAP estimators, ML estimation involves the a–priori probability density function of the observation y given a parameter value θ, $p_{y|\theta}(y|\theta)$. The parameter which maximises this function is the ML estimate.

ML techniques are frequently used in signal parameter estimation problems. No assumption is made about the possibilities of observed values, which can be a problem in applying a–posteriori based estimation techniques like MAP and MMSE. If gaussian noise on observations is assumed, estimators and bounds are reasonably easily derived.

The Cramer–Rao Lower Bound (CRLB) [31, 32] determines the ultimate accuracy of any estimation procedure; it is a lower bound on the variance obtainable by an estimator when a certain level of noise is present in the observations. The bound is independent of the form of estimation used. However the CRLB is intimately related to ML estimation in that if an unbiased estimator satisfies the CRLB, it is the Maximum Likelihood estimator. Because of this, efficiency bounds for the ML estimator are readily derived. The CRLB on a parameter vector $\underline{\theta}$ is the diagonal of the inverse of the Fisher Information Matrix. Elements of this matrix are the expected products of the partial derivatives of log pdf of the observation, with respect to the unknown parameters. That is,

$$J_{ij} = E\left\{\left(\frac{\partial}{\partial\theta_i}\ln p_{y|\theta}(y|\theta)\right)\left(\frac{\partial}{\partial\theta_j}\ln p_{y|\theta}(y|\theta)\right)\right\}$$

This is readily computed, given a signal model and an assumption of gaussian noise.

2.2 The Doppler Effect and Source Tracking

The doppler effect is a well–known phenomenon which occurs when observing a signal source in motion. As the source moves closer to the point of observation, a wavefront is

emitted more closely to the preceding wavefront than if the source had been at rest. The effect of this is to shorten the signal wavelength and to the observer the frequency appears higher than the source's rest frequency. Similarly a receding source causes frequency to appear lower to an observer. As a source travels past an observer, the frequency will continuously decrease as the source transits from approaching to receding.

The received doppler–shifted frequency f_d may be expressed in terms of emitted frequency f_s as

$$f_d = \frac{1}{1 + \frac{v}{c}} f_s$$

where c is the propagation velocity of the medium and v is the instantaneous velocity of the target relative to the observation point. A positive velocity denotes a receding source; a negative velocity for an approaching source. This relative velocity term is dependent on the source's motion parameters and thus analysis of the received frequency can yield information about the source track. There are other slightly different formulations used in acoustics which account for factors like motion of the medium (e.g. wind speed). However the formula given above is appropriate for this study.

Techniques exist which can determine source range, speed, rest frequency and the time of closest approach using only a sequence of instantaneous frequency estimates of a received signal [4] or frequency and amplitude estimates [3]. These methods essentially use the frequency measurements to enumerate expressions for the maximum likelihood estimation of signal parameters. While the technique is largely analytic, an iterative process is necessary for maximisation of an objective function. Another method [21] performs a potentially computationally demanding grid search to determine the source motion parameters from doppler shift measurements. However the dimension of the grid search can be reduced by exploiting the rate of frequency change. With the use of several receivers, doppler–shifted frequency measurements can be used to determine a source's position and heading. This can be done by analysis of the frequency at each receiver to determine speed, range and time of closest approach; this information is then integrated across receivers [23].

Standard bearings–only target motion analysis (TMA) can estimate a source's range,

bearing, heading and speed via the analysis of (possibly noisy) angle–of–arrival measurements. Inclusion of frequency information and modelling the effect of source motion on this can improve performance of TMA techniques. Blanc–Benon and Bienvenu [19] demonstrate how inclusion of multipath parameters (including doppler shift information) may improve the execution time of TMA, particularly in the case of short time periods and fast targets. Rosenqvist [20] uses estimates of the doppler–shifted source frequency to avoid the need for motion of the observation platform to determine range estimates. This technique also overcomes the biases inherent in pseudo–linear estimation and attains the Cramer–Rao lower bound for reasonable levels of noise in the bearing and frequency estimates.

It is evident that doppler–shifted frequency contains information which can be exploited by signal processing techniques. This study is concerned with modelling the doppler effect for source motion analysis, but from a matched processing framework rather than TMA.

2.3 Coherent Processing

An important attribute of a narrowband/sinusoid signal (from a signal processing point of view) is its phase. This is essentially the position of the signal within one cycle of its period, in particular at the time of origin. A signal may be represented as a complex exponential quantity known as a *phasor*. The magnitude of the phasor corresponds to the signal's amplitude and the angle of the phasor corresponds to it's phase. Figure 2.1 provides an illustration.

When samples of a signal (or several signals) are added together, the phase information from one component will interfere with other components and cause the resultant power to be degraded. However if the signal components are pre–processed in a fashion such that the phasors are oriented in the same manner, the resultant signal power will re–inforce. The signals have been processed "in phase". Figure 2.2 provides an illustration of incoherent and coherent combination of phasors. If there is a–priori information

about the characteristics of a signal, a bank of phasors may be designed to compensate for the expected phase differences between the received signal components and thus process them coherently.

2.4 Matched Filtering

In applications such as radar signal detection and digital communications it is desirable to have a filter which maximises signal to noise ratio at some point in the output sequence. The continuous waveform which is to be detected, $x(t)$, is known ahead of time. It can be shown that the impulse response of the filter which achieves optimum SNR is the timereversal of this waveform $x(t-T)$ [25].

Filtering a signal $s(t)$ with the time reversal of the waveform $x(T-t)$ is equivalent to cross–correlating $s(t)$ with $x(t)$. If the received signal is a message waveform in noise, i.e.

$$s(t) = m(t) + n(t)$$

then the filter output $y(t)$ is

$$\begin{aligned} y(t) &= s(t) \otimes x(T-t) \\ &= m(t) \otimes x(T-t) + n(t) \otimes x(T-t) \end{aligned}$$

This has two components. One component due to noise, the other is the correlation of the waveform $x(t)$ with the message $m(t)$. The power of this component will be maximised when the message is the desired waveform, i.e. $m(t) = x(t)$. Nothing correlates with a sequence better than itself. Hence it makes intuitive sense that SNR will be maximised when the filter impulse response is a time reversal of the signal of interest. Such a filter is termed a *matched filter* because the impulse response of the filter *matches* that of the signal of interest [24].

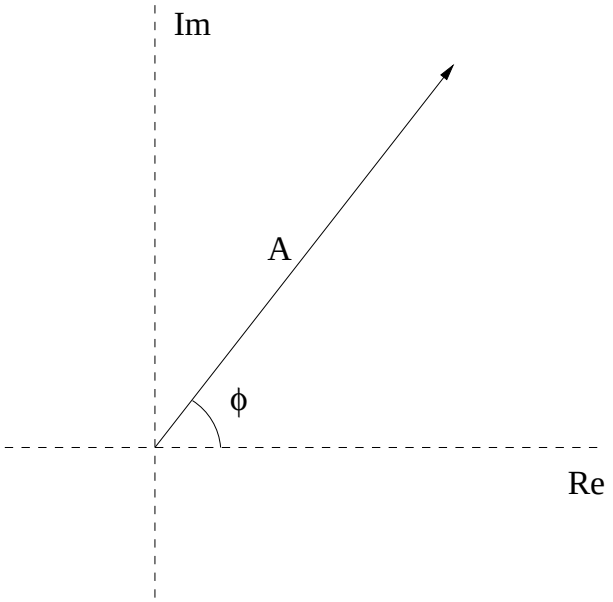

Figure 2.1: Phasor representation of signal $A \exp i(\omega t + \phi)$

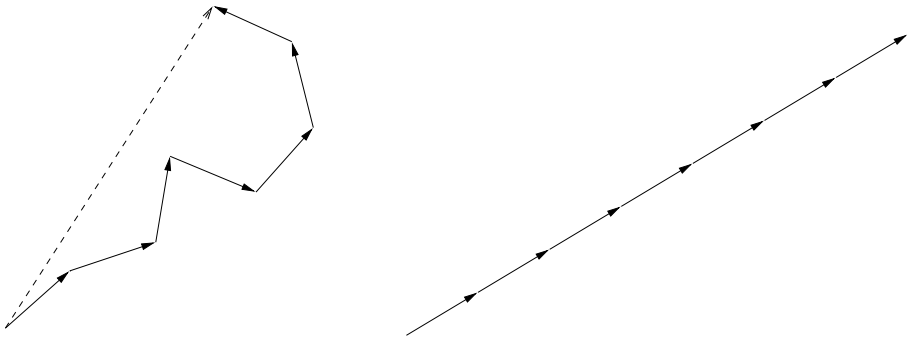

Figure 2.2: Incoherent and Coherent Combination of Signal Phasors

2.5 Beamforming

Beamforming is a generic term which may apply to a number of sensor array processing algorithms that focus an array on signals from a particular direction or location [24]. A propagating signal which is incident on an array of receivers will arrive at each sensor at a slightly different time. This time lag depends on each individual receiver's distance to the source of the signal, with respect to the other receivers in the array.

Consider a linear array upon which a farfield source is incident from broadside (perpendicular to the line of the array). The incoming wavefronts from the farfield source are parallel and will coincide at each receiver simultaneously. Hence signals at each receiver are "in phase" and the signal components will reinforce one another upon combination of the receiver outputs. However signals from other directions will arrive at each receiver at a different time. The receiver outputs may combine incoherently due to the phase differences of the receiver outputs and the resultant signal will be degraded.

In order to receive farfield signals from a direction other than broadside, the receiver outputs must be delayed by amounts of time appropriate to the desired direction. This will cause signal components from this direction to line up in phase and reinforce the overall signal power upon combination. The array is said to be *steered* in the desired direction. The output sequence of a simple time–delay–and–sum beamformer is

$$b(t,\theta) = \sum_{k=1}^{K} x_k(t - \tau_k(\theta))$$

where $x_k(t)$ is the output of the kth receiver and $\tau_k(\theta)$ is the time delay on the kth receiver, necessary for steering in direction θ. For a linear equispaced array, this value is

$$\tau_k(\theta) = \frac{dk\sin(\theta)}{c}$$

where d is the receiver separation and $\theta = 0$ when steering at broadside.

The time domain beamformer discussed above may be dually expressed in the frequency domain using phase shifts in place of time delays. A vector of phasors (known as a steering vector) designed to compensate for phase differences in signals incident from a

particular direction is multiplied onto the spectra of the receiver outputs which are then combined to produce the beamformed output

$$\begin{aligned} B(f,\theta) &= \sum_{k=1}^{K} X_k(f) \exp\left(-i2\pi f \tau_i(\theta)\right) \\ &= \underline{v}(\theta)^H \underline{X}(f) \end{aligned} \tag{2.1}$$

where the steering vector $\underline{v}(\theta)$ is

$$\underline{v}(\theta) = \left(\exp\left(i2\pi f \tau_1(\theta)\right), \ldots \exp\left(i2\pi f \tau_K(\theta)\right)\right)^T$$

The mean power of the beamformed output is usually referred to as the output power of a conventional beamformer:

$$\begin{aligned} |B(f,\theta)|^2 &= \underline{v}^H(\theta) \underline{X}(f)\, \underline{X}(f)^H \underline{v}(\theta) \\ &= \underline{v}^H(\theta) R_X(f) \underline{v}(\theta) \end{aligned} \tag{2.2}$$

where $\underline{v}^H$ denotes the Hermitian (conjugate) transpose of vector $\underline{v}$. The quantity $R_X(f)$ is known as a spectral covariance matrix and is strictly defined as

$$R_X(f) = E\left\{\underline{X}(f)\, \underline{X}^H(f)\right\}$$

In practice this quantity is estimated, either from a single set of measurements as in equation 2.2 or as an average over several consecutive timeblocks:

$$\begin{aligned} R_X &= \frac{1}{T} \sum_{t=1}^{T} \underline{X_t}(f) \underline{X_t}(f)^H \\ &= \frac{1}{T} \left[\underline{X_1}(f) \ldots \underline{X_T}(f)\right] \left[\underline{X_1}(f) \ldots \underline{X_T}(f)\right]^H \end{aligned}$$

Averaging is commonly used because it provides better estimation of $R_X(f)$ in the presence of noise.

It is well known that the angle θ which maximises the beampattern is the maximum likelihood estimator of the bearing of a single signal in spatially uncorrelated noise. This is analogous to the maximum of a DFT power spectrum being an ML estimator of signal frequency in white noise; the beamformer is a spatial filter whereas the DFT filters in frequency. In practice, receiver outputs are filtered with a bank of steering

vectors for a wide range of steering vectors which are usually equispaced in azimuth. The output power $|B(f,\theta)|^2$ is computed for each steering vector, and the angle whose steering vector produces the largest power value is taken as an estimate of signal bearing $\hat{\theta}$. Maximisation is thus performed via an exhaustive search of candidate values from a given search domain. The resolution of the bearing estimate is determined by how coarsely this domain is spaced.

The output power of a conventional beamformer as defined by the Hermitian form in equation 2.2 is characterised by broad main lobes which afford very robust performance as a slight mismatch will generally not affect the location of the beampattern's maximum. However this wide beam, together with a sidelobe structure (i.e. areas of high correlation, not in the vicinity of the main lobe), also allow interference from signals which are not in the look direction. Variations on the conventional beamformer may be used to limit sidelobe levels or decrease main beam width. There exist other data–adaptive formulations such as the Minimum Variance Distortionless Response (MVDR) beamformer which inhibit sidelobes and have very thin main beams, but at the expense of robustness. It is important to note these characteristics because similar processors used for analysing moving sound sources inherit these properties from the beamformers.

2.6 Matched Processing

Conventional beamforming, as described above, is a very simple example of *matched processing*. The steering vector for a given direction $\underline{v}(\theta)$ is the spatial equivalent of a matched filter, which exploits coherent phase processing. It is in fact the exact set of phasors one would expect to receive if a signal was incident on the array from the look direction θ. It is a *replica* of this signal. Conjugation of the steering vector's phase in the beamformer effectively annuls the phasor differences at each receiver. Signal information from this bearing may be combined coherently, whereas signal information from other bearings will degrade. The processor used for beamforming, whether it be conventional, minimum–variance or otherwise, is essentially a correlation function which provides a

measure of how well each replica *matches* the measured data. The angle which produces the replica having best match (i.e. the peak of the beampattern) is taken as an estimate of the source's bearing $\hat{\theta}$.

2.6.1 Matched Field Processing

An overview of Matched Field Processing has been given in the introduction of this thesis (§1.1). As it's name suggests, MFP is a technique which performs matching of replica measurements in order to estimate unknown values, usually the position in range depth and bearing of a sound source. MFP has traditionally been used to process signals emanating from a stationary source, but more recently MFP techniques have been applied to estimating parameters of a moving source.

The most important difference between plane–wave beamforming and MFP is the use of a propagation model to generate realistic signal replicas for matching [1]. This propagation model may take into account several factors which reflect the physics of the environment. For example, refraction of the signal as it passes through various sound–speed layers of the ocean, reflection from the ocean floor. Numerous propagation models exist. Most can be described broadly as belonging to one of the following categories: Normal mode models, ray theory models and parabolic equation models.

A typical matching function which uses coherent phase processing is the Bartlett processor [1]. This has the same form as the conventional beamformer but is usually expressed in terms of a sample covariance matrix (e.g. equation 2.2). The Bartlett processor inherits the robustness properties of the CBF. Other data–adaptive or minimum–variance processors have also been employed for MFP, for example in [8] and [29]. These inherit the sidelobe reduction attributes of the MVDR beamformer. Some research has been done on enhancing the robustness of these minimum–variance MFP techniques [7].

The matching function describes a surface in the space of source parameters. It is the task of MFP to find the optimal point of this surface. This is usually accomplished via an exhaustive search: The matching function is evaluated for every possible combination

of candidate parameters and the optimal value located. The necessity of generating a signal model and subsequently a cost measure for each combination of parameters means that MFP is often a computationally demanding process. It is much more demanding than beamforming as the propagation models used to generate signal replicas are often complex. Furthermore the parameter space is usually multivalued and hence is usually larger than the set of candidate bearings used for beamforming.

It is important to emphasise that MFP is an estimation process, not a detection process. There is an implicit assumption that a signal is present in the receiver outputs; the role of MFP is to determine which set of parameters best account for this observed signal. Some formulations may perform detection after processing. This often takes the form of thresholding the correlation/matching function; if the correlation is suitably large, then a source is deemed to have been present. This approach works well when applied to the moving source problem, particularly because it enables the tracking of multiple targets. Each peak which exceeds the threshold may be deemed to be a source. Another such track–before–detect approach involves the integration of power from consecutive MFP outputs via a matched velocity filter and thresholding these filter outputs [16].

Chapter 3

Matched Doppler Tracking

This chapter presents a method for estimating source motion parameters by matching outputs from a single receiver. A simple model of signal amplitude and phase is determined by consideration of time–lags, which are shown to be a manifestation of the doppler shift. A number of matching functions are defined, with emphasis placed on those functions which process signal phase coherently. Simulation results which compare the various cost functions' ability to track a low–frequency sinusoid are presented and discussed. Also the matching techniques are compared with a doppler–tracking technique by means of simulation.

3.1 Problem Model

A nearfield acoustic source moves past a single receiver. It moves in a straight line and with constant velocity, emitting a narrowband signal of constant monotonic frequency, f_s Hz. Figure 3.1 presents a diagram of the source at time t'.

The problem is to determine the locus of the moving source. This may be solved by estimating its initial range, final range and range at closest approach to the receiver. These three parameters are sufficient to characterise the range history of a moving target over a time interval of known duration. Further parameters (i.e. speed) may be com-

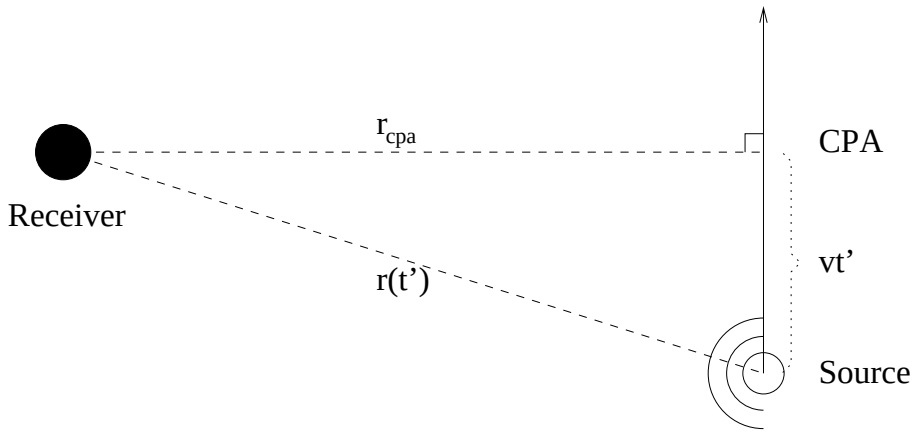

Figure 3.1: Source at time t'

puted from these. It is assumed that direct path propagation is the only mode received, which is not unreasonable for a nearfield source. However this assumption prohibits the estimation of source depth.

3.2 Signal Model

3.2.1 Effect of Time–lag on Signal Phase

At time t, the actual signal received is the signal which was emitted at some earlier time $t' < t$. The time taken in transmission, $t' - t$, may also be computed by dividing the range at time t' by the propagation speed c. If r_{cpa} is the range of the target at CPA, then

$$t - t' = \frac{r\,(t')}{c}$$

which implies

$$(c^2 - v^2)t'^2 - (2c^2 t)t' + (c^2 t^2 - r_{cpa}^2) = 0 \tag{3.1}$$

This equation has two roots t_1 and t_2. It can be shown that $t_1' < t$ and $t_2' > t$. The imposition $t' < t$ has been made so $t' = t_1'$ is chosen. The signal phase received at time t, $\phi_r(t)$, is that which was emitted at time t', $\phi_e(t')$, when the target was at range $r(t')$.

Thus

$$\phi_r(t) = \phi_e(t') = \phi_{cpa} + 2\pi f_s t' \tag{3.2}$$

where ϕ_{cpa} is the emitted phase at CPA.

When the ratio $\frac{v}{c}$ is small it can be shown that

$$t = t' + \frac{r_{cpa}}{c} + \frac{v^2 t'^2}{2cr_{cpa}} \tag{3.3}$$

which is analogous to the quadratic phase term used in coherent synthesis synthetic aperture radar processing. However for the acoustics problems considered here this formula will not be used.

3.2.2 Received Amplitude

If the power output by the target is A_0^2 at a range of 1 metre then the amplitude received at time t is

$$A(t) = \sqrt{\frac{A_o^2}{r(t')}} \tag{3.4}$$

assuming a cylindrical spreading model. Spherical spreading can be modelled by introducing an extra factor of $r(t')$ in the denominator.

3.2.3 Instantaneous Frequency

As the source travels past the receiver, it's velocity will have 2 components: A radial component $v_r(t')$, directly towards or away from the receiver, and a transverse component $v_t(t')$ orthogonal to this.

It is clear from figure 3.2 that

$$\begin{aligned} v_r(t') &= v \sin\left(\theta\left(t'\right)\right) \\ &= v\frac{vt'}{r\left(t'\right)} \\ &= \frac{v^2 t'}{c\left(t - t'\right)} \end{aligned} \tag{3.5}$$

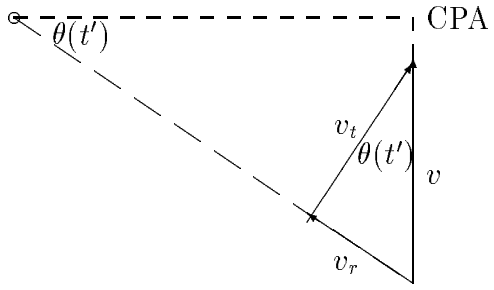

Figure 3.2: Components Of Target Velocity

The frequency received at any instant t, $f_r(t)$, will differ from the emitted frequency f_s due to the doppler effect of a sound source in motion. This doppler shift depends entirely upon the radial component of the target's velocity:

$$\begin{aligned} f_r(t) &= \frac{1}{1+\frac{v_r(t')}{c}} f_s \\ &= \frac{1}{1+\frac{v^2 t'}{c^2(t-t')}} f_s \end{aligned} \tag{3.6}$$

3.2.4 Manifestation of Time–lag as a Doppler Shift

The phase term varies in a nonlinear fashion due to the decrease in transmission time as the source moves closer to the CPA and the subsequent increase in transmission time as the source moves further away. The received frequency $f_r(t)$ can be expressed as the rate of change of phase with time. That is,

$$\begin{aligned} f_r(t) &= 2\pi \frac{d\phi_r(t)}{dt} \\ &= 2\pi \frac{d\phi_e(t')}{dt} \\ &= 2\pi f_s \frac{dt'}{dt} \end{aligned} \tag{3.7}$$

Equation 3.1 is differentiated to obtain

$$2(c^2-v^2)t'\frac{\partial t'}{\partial t} - 2c^2 t' - 2c^2 t\frac{\partial t'}{\partial t} + 2c^2 t = 0 \tag{3.8}$$

from which it can be shown that

$$\begin{aligned}\frac{\partial t'}{\partial t} &= \frac{1}{1-\frac{v^2 t'}{c^2(t'-t)}} \\ &= \frac{1}{1-\frac{v_r(t')}{c}}\end{aligned}$$

and the expression for received frequency (equation 3.7) becomes

$$f_r(t) = \frac{1}{1-\frac{v^2 t'}{c^2(t'-t)}} f_s$$

This is the exact formula for the doppler shift in frequency, as derived in equation 3.6. Thus the time–lag accurately models the effect of the doppler shift on signal phase.

3.2.5 Model of Received Signal

The expressions of received amplitude (eq. 3.4) and received phase (eq. 3.2) together form the model of the signal received at time t:

$$\begin{aligned}s_r(t) &= A(t)\exp i\phi_r(t) \\ &= \sqrt{\frac{A_0^2}{r(t')}}\exp i\left(\phi_{cpa}+2\pi f_s t'\right)\end{aligned}$$

3.2.6 Replica Measurements

The signal will be filtered in order to provide measurements of it's amplitude and phase at several points in time. The Discrete Fourier Transform (DFT) is a suitable filtration mechanism which provides narrowband gain. The f_s Hz bin of the DFT is a complex value whose magnitude is a measure of the signal's amplitude and whose angle is a measure of the phase.

However the doppler shift will cause power to leak from the f_s Hz bin. Amplitude will be degraded as a sinc function of the frequency shift and phase information will be distorted, as the received frequency $f_r(t)$ will rarely be at centre–cell. An illustration of this is provided in figure 3.3. This corruption will be negligible when the target moves

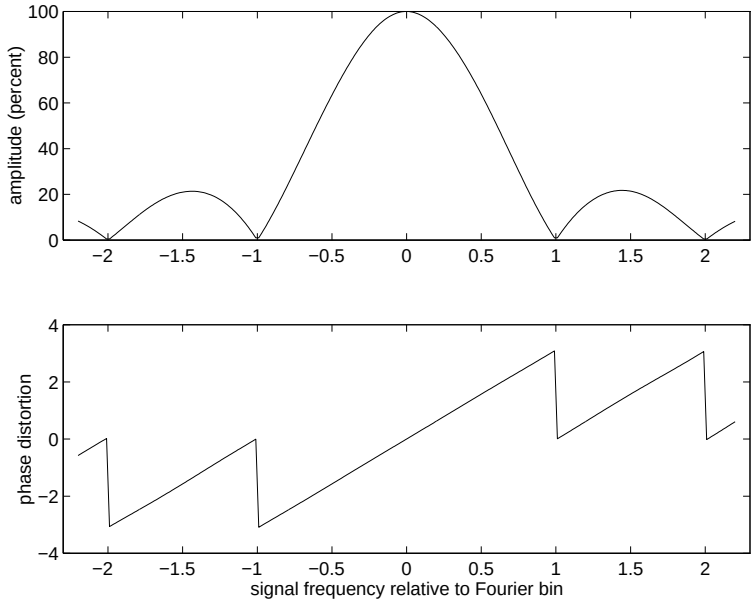

Figure 3.3: Effect of Doppler shift on Fourier bin

with a very slow speed or emits a very low frequency and in this case can be ignored [2]. However this distortion effect must be modelled in order to analyse signals which exhibit significant doppler shift.

It is assumed that the received amplitude $A(t)$ and frequency $f_r(t)$ are approximately constant over a one–second timeblock. This implies that phase varies linearly over a timeblock. If an N–point DFT is computed at the timeblock beginning at t, the f_s Hz Fourier coefficient is

$$S(f_s, t) = \sum_{n=0}^{N-1} s_r(t + nT) \exp\left(-i2\pi f_s \frac{n}{N}\right) \tag{3.9}$$

and can be shown to be given by

$$S(f_s, t) = D(t)A(t) \exp(i\phi_r(t)) \tag{3.10}$$

where

$$\begin{aligned} D(t) &= \frac{\sin \pi\delta(t)}{\sin \frac{\pi\delta(t)}{N}} \exp\left(i\pi\delta(t)\frac{N-1}{N}\right) \\ \delta(t) &= f_r(t) - f_s \end{aligned}$$

The $D(t)$ term models the power–leakage & phase distortion inherent when the sampled frequency is not at centre–cell. Replica vectors $\underline{P}$ are formed by computing the theoretical

DFT coefficients $S(f_s, t)$ in each consecutive timeblock of T seconds duration, i.e.,

$$P_n = S(f_s, t_i + (n-1)T) \tag{3.11}$$

where t_i is the initial timeblock. Similarly D_n, δ_n, A_n etc. are defined as the discretised versions of the continuous values used above. In this study it is assumed that the timeblocks are one second in duration, i.e. $T = 1$. This replica could easily be extended to comprise measurements from several frequency bins. This would enable coherent integration of information across frequency by a processor similar to those presented in [5].

3.2.7 Signal Measurement

Samples from a timeblock of length T seconds are filtered by a DFT. The complex Fourier coefficient from the bin corresponding to the known signal frequency f_s is extracted and taken as a measurement of signal amplitude and phase in that timeblock. This is performed over M consecutive timeblocks to produce a vector of complex measurements $\underline{p} = (p_1\, p_2\, \ldots\, p_M)$ where p_m is the measurement of the mth timeblock.

Replica vectors $\underline{P}$ are generated for each set of candidate parameter values according to the signal model (eq. 3.11) and then matched with the measurement vector $\underline{p}$ by means of a matching function. The process of signal measurement and matching is illustrated in figure 3.4.

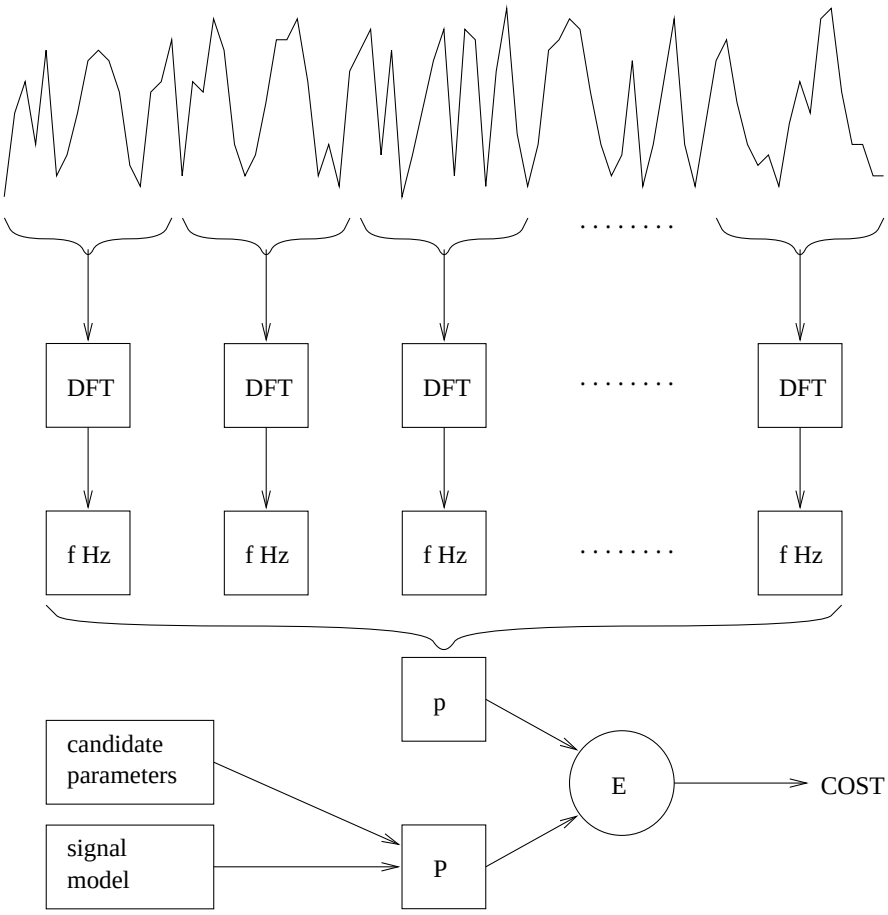

Figure 3.4: The Matching Process

3.3 Matching Functions

Source track estimation is performed by computing a replica data vector for each possible combination of parameters in the search space, via the detailed model (3.11). Each of these replicas is compared with the realisation by use of a matching function.

3.3.1 EST Cost Function

The cost function employed in [2] was

$$E_{EST} = \sum_{m=1}^{M} \left(|\tilde{p}_m| - \left|\tilde{P}_m\right| \right)^2 \tag{3.12}$$

where $\underline{\tilde{p}}$ is a normalised vector of received amplitude and phase:

$$\underline{\tilde{p}} = (\tilde{p}_1, \ldots, \tilde{p}_M) = \frac{\left(\sqrt{r_1}p_1, \ldots, \sqrt{r_M}p_M\right)}{\sqrt{\sum_{m=1}^{M} r_m \left|p_m\right|^2}}$$

where p_n is the f_s Hz Fourier coefficient from the nth timeblock of Fourier transformed data and r_n is the replica range at the nth timeblock. $\underline{\tilde{P}}$ is a replica vector as defined by (3.11), normalised in the same fashion as $\underline{\tilde{p}}$. The purpose of normalisation is to remove the cylindrical spreading factor of $\sqrt{r}$ so that data from around CPA does not dominate [2]. This cost value is minimised over the search space to find estimates of the unknown range parameters.

3.3.2 EST Cost Function with Phase Information

Note that the moduli of $\tilde{p}_m$ and $\tilde{P}_m$ are taken within the EST cost function (eq. 3.12), effectively discarding phase information. This was done because of difficulty in the measurement of phase [2]. However including phase information may significantly increase the performance of the technique so a modified version of the EST function which leaves the phase intact is considered.

$$\begin{aligned} E_{A\&P} &= \sum_{m=1}^{M} \left| \left(\tilde{p}_m - \tilde{P}_m \exp i\hat{\phi}_c \right) \right|^2 \\ &= \left| \left(\underline{\tilde{p}} - \underline{\tilde{P}} \exp i\hat{\phi}_c \right)^H \left(\underline{\tilde{p}} - \underline{\tilde{P}} \exp i\hat{\phi}_c \right) \right|^2 \end{aligned} \tag{3.13}$$

Normalisation of data/replicas is performed as before. However this process will only affect amplitude; phase information is not perturbed by normalisation. The term $\exp i\hat{\phi}_c$ includes an estimate of the complex phase difference between the replica vector $\underline{P}$ and data $\underline{p}$, reflecting the uncertainty due to the unknown absolute phase of the transmitted signal. This is estimated as

$$\hat{\phi}_c = \angle \sum_{k=1}^{K} \exp i\left(\angle p_k - \angle P_k\right)$$

Inclusion of this term removes the dependence of the cost function on relative phase, allowing us to ignore it in the optimisation process.

3.3.3 CBF–style Correlation Function

An alternative to minimising the least–squares difference between realisation and replica is to maximise a correlation function such as this:

$$E_{CBF} = \left|\underline{\tilde{p}}^H \underline{\tilde{P}}\right|^2 \tag{3.14}$$

This function will be maximised when the phase components of the replica match the phase in the realisation, thus allowing signal power to be summed coherently. This process is analagous to conventional beamforming for finding source direction with an array of receivers. Thus we refer to the cost function as the CBF matching function. This is equivalent to the Bartlett processor, which is often used as a robust correlation function in MFP. Note that there is no need for phase compensation; the output of the CBF is not dependent on the phase difference between $\underline{p}$ and $\underline{P}$.

3.3.4 Phase Cost Function

It is worth considering a cost function based entirely on phase in order to consider the relative benefits of phase information compared to amplitude.

$$E_\phi = \sum_{m=1}^{M} \left|\left(\angle p_m - \angle\left(P_m \exp i\hat{\phi}_c\right)\right)\right|^2 \tag{3.15}$$

Note that normalisation is not necessary as amplitude information is discarded by this cost function. However compensation for an initial phase offset in the same manner as the A&P cost function (eq. 3.13) will be necessary.

This cost function is simple, it avoids complex arithmetic as it extracts the phase value from the complex exponential component of the signal model. However the angle–finding operation $\angle$ maps all complex values onto the domain $(0, 2\pi)$. Phase–wrapping problems will occur near the boundaries of this domain. A small phase error will cause the phase measurement to jump by 2π if the boundaries of the interval are exceeded. This in turn will cause a large mismatch in the cost function. The possibility exists for this to happen in each timeblock of the sampling interval and thus this cost function is expected to be extremely non–robust in the presence of noise. However phase information can vary quite rapidly even when there is little amplitude change, for example when the source has low velocity. So this crude phase–only cost function may still have benefits over the amplitudes–only approach (eq. 3.12).

3.3.5 Ambiguity Surfaces

The value of a cost function may be presented graphically as a function of two parameters as an *ambiguity surface*, which is an enumeration of the cost value for each combination of parameters in the 2–dimensional search space. Obviously the search space used here has 3 parameters (initial, final and CPA range), so a two–dimensional ambiguity surface represents a "slice" of the parameter space.

As defined by Tolstoy [1], the ambiguity surface is an output function of the data parameters and thus incorporates noise. It is not to be confused with the ambiguity function employed in radar which is defined in the absence of noise.

Typical ambiguity surfaces generated with a -20 dB signal over a search space of initial and final ranges (representing a slice of the parameter space at true CPA) are shown in figures (3.5 – 3.8). The well–behaved structure of the EST, CBF and A&P cost functions suggests that these techniques maybe amenable to a hill–climbing optimisation

technique, thereby reducing the computation required for an exhaustive search of the domain. However the ambiguity surfaces of the phase–only cost function exhibits noise in it's structure and may not succumb to hill–climbing; an exhaustive search is thus unavoidable. In this example r_i was 810 m and r_f was 770 m. Note that the amplitude–only EST was the only cost function for which the ambiguity surface was not maximised at the correct location.

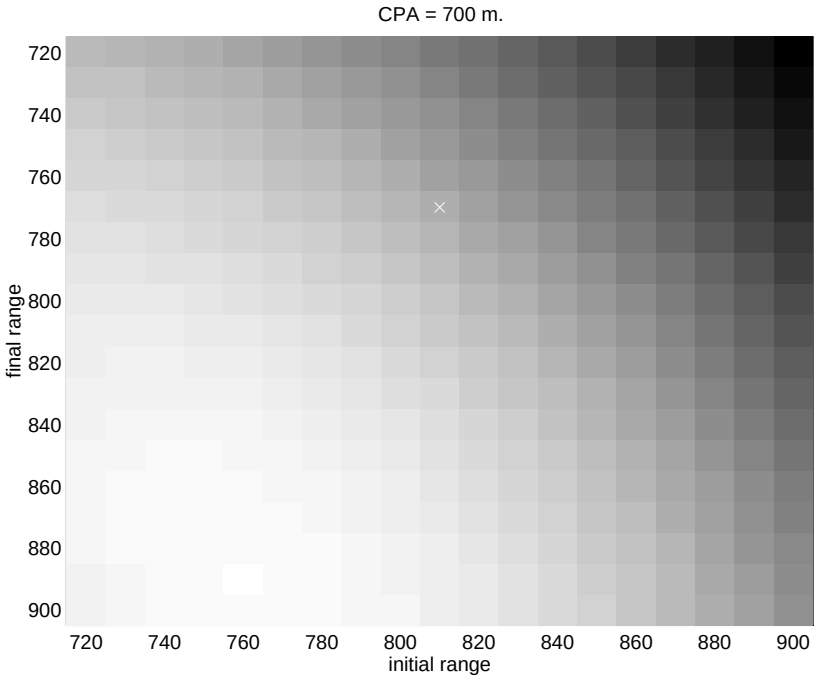

Figure 3.5: EST Cost Function Ambiguity Surface

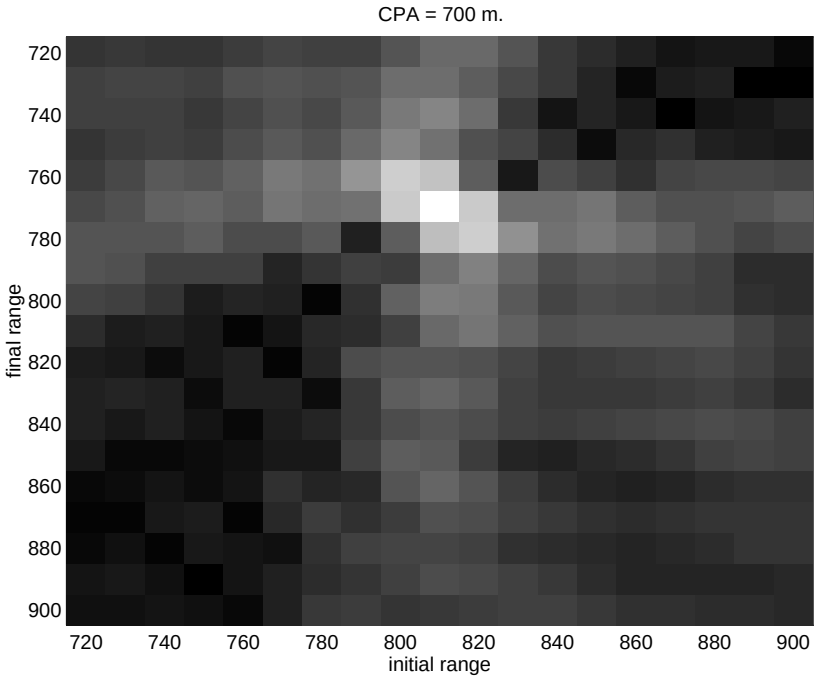

Figure 3.6: Amp & Phase Cost Function Ambiguity Surface

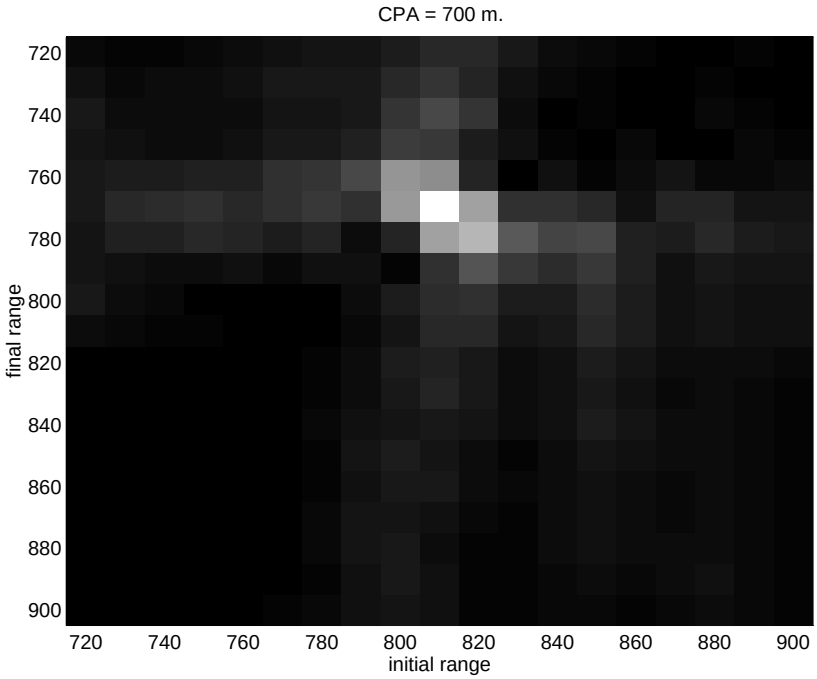

Figure 3.7: CBF Ambiguity Surface

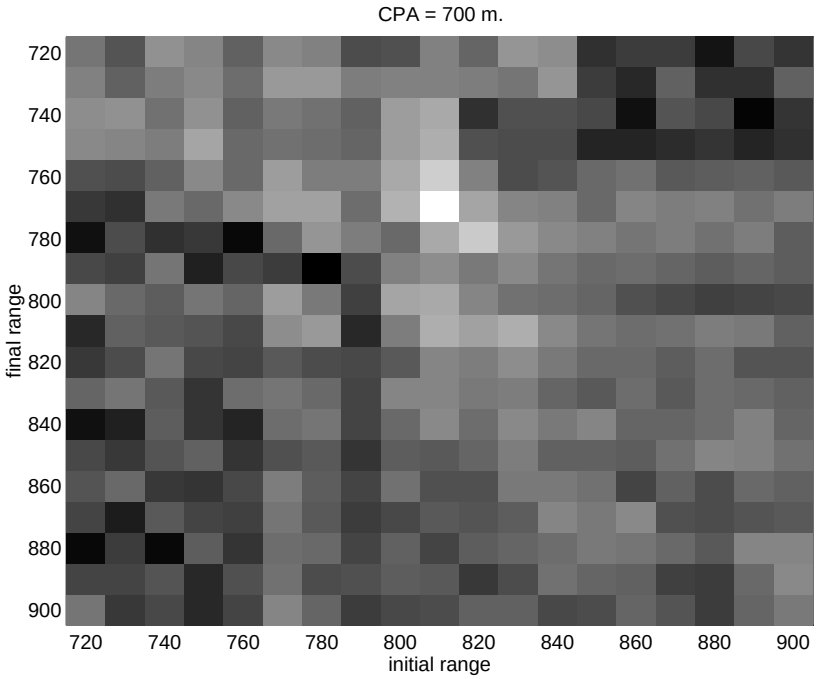

Figure 3.8: Phase Only Ambiguity Surface

3.3.6 Similarity of A&P and CBF Cost Functions

Consider the form of the A&P cost function (eq. 3.13). This may be expanded as follows:

$$\begin{aligned} E_{A\&P} &= \sum_{k=1}^{K} \left| \tilde{p}_k - \tilde{P}_k \exp(i\hat{\phi}_c) \right|^2 \\ &= \sum_{k=1}^{K} \left(\tilde{p}_k - \tilde{P}_k \exp(i\hat{\phi}_c) \right) \left(\tilde{p}_k^* - \tilde{P}_k^* \exp(-i\hat{\phi}_c) \right) \\ &= \sum_{k=1}^{K} \tilde{p}_k \tilde{p}_k^* - 2\Re \left\{ \sum_{k=1}^{K} \tilde{p}_k^* \tilde{P}_k \exp(i\hat{\phi}_c) \right\} + \sum_{k=1}^{K} \tilde{P}_k \tilde{P}_k^* \end{aligned}$$

The first term $\sum_{k=1}^{K} \tilde{p}_k \tilde{p}_k^*$ is a function of the measurement and is thus constant. The third term $\sum_{k=1}^{K} \tilde{P}_k \tilde{P}_k^*$ will also be constant due to the normalisation process which removes any range dependencies. This means that the minimum cost value with respect to some parameter combination $\underline{\theta}$ and hence replica vector $\underline{P}$ becomes

$$\begin{aligned} \min_{\underline{\theta}} E_{A\&P} &= \max_{\underline{\theta}} \Re \left\{ \sum_{k=1}^{K} \tilde{p}_k^* \tilde{P}_k \exp(i\hat{\phi}_c) \right\} \\ &= \max_{\underline{\theta}} \Re \left\{ \exp(i\hat{\phi}_c) \sum_{k=1}^{K} \tilde{p}_k^* \tilde{P}_k \right\} \end{aligned} \tag{3.16}$$

In this form it is evident that the cost function is indeed treating phase coherently. Furthermore at this point the similarity with the CBF cost function (eq. 3.14) is apparent. The CBF takes the amplitude of the complex summation, whereas the A&P cost function takes the real part. The two cost functions will have the same optimal point in a noise–free environment if and only if the phase difference estimate is exact, i.e. $\hat{\phi}_c = \phi_c$. In this case the phase of the complex sum will be nullified and the $\Re$ operation is thus the same as taking the absolute value.

Consider the case in which the replica is formed from the correct parameters and there is some noise on the measurements. Assuming once more that $\hat{\phi}_c = \phi_c$ the maximum value of (eq. 3.16) becomes

$$\begin{aligned} \max &= \Re \left\{ \exp(i\hat{\phi}_c) \sum_{k=1}^{K} \left(A_k \exp(-i\phi_k) \exp(-i\phi_c) + n_k \right) A_k \exp(i\phi_k) \right\} \\ &= \sum_{k=1}^{K} A_k^2 + \Re \left\{ \exp(i\phi_c) \sum_{k=1}^{K} n_k A_k \exp(i\phi_k) \right\} \end{aligned} \tag{3.17}$$

Compare this to the CBF cost value evaluated in the presence of noise:

$$\begin{aligned} E_{CBF} &= \left| \sum_{k=1}^{K} \tilde{p}_k^* \tilde{P}_k \right|^2 \\ &= \left| \sum_{k=1}^{K} \left(A_k \exp(-i\phi_k) \exp(-i\phi_c) + n_k \right) A_k \exp(i\phi_k) \right|^2 \\ &= \left| \sum_{k=1}^{K} A_k^2 \exp(-i\phi_c) + n_k A_k \exp(i\phi_k) \right|^2 \end{aligned}$$

It is clear that when there is no noise ($n_k = 0 \, \forall k$), this reduces to the square of the A&P form (eq. 3.17). However the presence of noise will clearly affect the two cost functions differently. This effect may be observed in the ambiguity surfaces of the two functions (figs. 3.6 & 3.7). Both maximise at the same location but there is greater evidence of sidelobes in the A&P ambiguity surface.

Currently $\hat{\phi}_c$ must be computed for every replica when using the A&P cost function. However if the phase of the signal is somehow known *a priori* it may be used for every replica and performance will improve. This is because the actual signal phase will act as a penalty when used with incorrect replica; it will introduce a complex phase into the cost function which will cause the $\Re$ operation to decimate the result of the complex summation. An ambiguity surface generated by using the exact signal phase with each replica is presented in figure 3.9. Compare this to the ambiguity surface of figure 3.6 where phase is estimated for each individual replica. The phase term obviously acts as a penalty for some of the incorrect parameter combinations, but with others it does not. The resulting ambiguity surface is obviously far too rough to be processed with a gradient–descent technique, but processing with exhaustive search will yield slightly better performance than with the estimated–phase version because of this penalty effect.

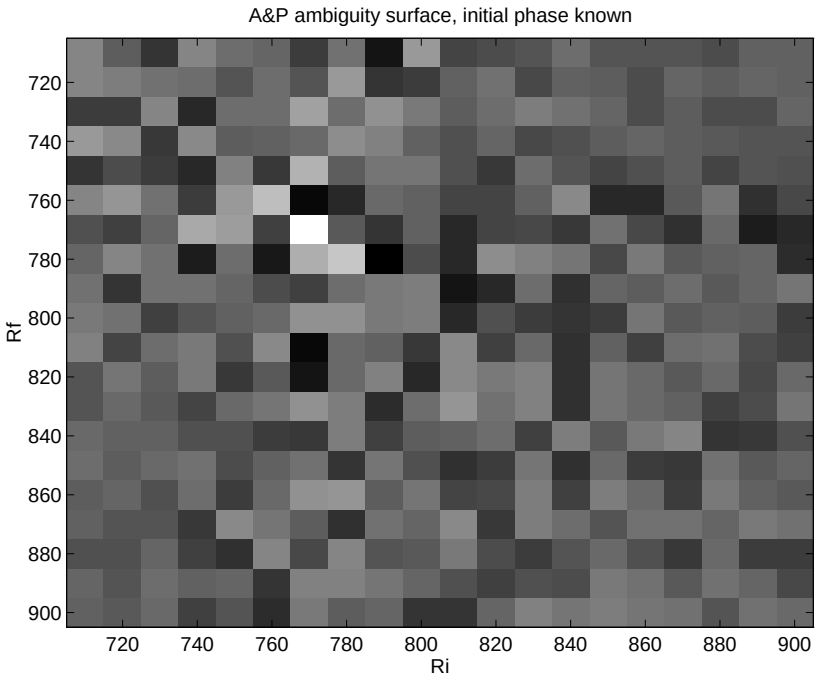

Figure 3.9: A&P Cost Function, Exact Phase Known.

3.4 Simulations

Simulated data from a narrowband source moving past a single receiver with constant velocity is generated. The four cost functions proposed in the previous section are used to estimate the range parameters of this target via matched processing.

3.4.1 Data Generation

A simple method is to simulate the output of the Fourier transform at t_n by computing an amplitude and phase as described by (3.11) and corrupting it with complex gaussian noise $\underline{z}$, i.e:

$$p_n = D_n A_n \exp i\phi_{r_n} + z_n \tag{3.18}$$

where $\phi_{r_n} = \phi_{cpa} + 2\pi f_s t'_n$, and ϕ_{cpa} is chosen from a uniform distribution on $(0, 2\pi)$. Another method involves simulating the received timeseries from a moving target and computing the relevant DFT coefficient at each timeblock. Both methods were used as appropriate.

3.4.2 Simulation 1

The simulated target had initial range of 770 m, moved to within 700 m at CPA and had final range of 770 m. It emitted a 50 Hz signal. Amplitude and phase measurements were simulated for every second of a 2 minute time interval. Target velocity was thus approximately 5.3 m/s, which is typical in sonar applications. A reasonably small time interval is chosen for 2 reasons; firstly to keep the computational load of the exhaustive search algorithm to a manageable level, seondly because over this interval the doppler frequency track is almost linear (see fig. 3.14). Thus bounds on the estimators are easy to derive by using the quadratic phase approximation as discussed in §4.1. The simulated measurements were processed with the EST cost function (3.12), the revised amplitude and phase cost function (3.13), the conventional beamformer (3.14), phase–only cost function (3.15), and a doppler tracker described in [3]. The search space comprised 2000

replica. These replicas were generated with r_{cpa} values at each hundred–metre interval between 500 and 900 m; and initial and final ranges at 10–metre intervals between $r_{cpa}+10$ and $r_{cpa}+200$ metres. That is, different r_i and r_f were used for each r_{cpa}. The CPA search domain is discretised more coarsely than the initial and final range domain, in order to sample a wide range of CPA values whilst still keeping the search space at a manageable size. It shall be demonstrated that the matched processors are poor estimators of CPA for this class of problem even with this coarse sampling level. Signal–to–noise ratios between -50 and 20 dB were considered. SNR is defined as the signal–to–noise ratio in the timeseries when the target is at CPA, i.e. before pre–processing with Fourier transforms. It is worth noting that the 1024–point FFT used in the analysis of simulated data provided a gain of approximately 30 dB. 100 simulations were performed at each SNR. Sample bias and standard deviation of the estimates were computed and are presented in tables 3.1 and 3.2. Some results are also presented in figures 3.10 – 3.13.

3.4.3 Simulation 2

100 simulations of the target described in [3] were performed. This target emitted a high frequency (300 Hz) so the doppler shift was much more pronounced than in the previous simulation. Also the sampling time was much greater (10 minutes) and encompassed the entire doppler shift, not just a fragment about the CPA. Source speed was 5 m/s, and it was 500 m from the receiver at CPA. SNR at CPA was -15 dB. The signal was analysed with the A&P cost function (3.13) and the doppler–tracking technique in [3]. Range estimates computed by the A&P were converted into estimates of time at closest approach, $\hat{t}_{cpa}$, and velocity $\hat{v}$ for comparison with the doppler tracker. Results are presented in table 3.3.

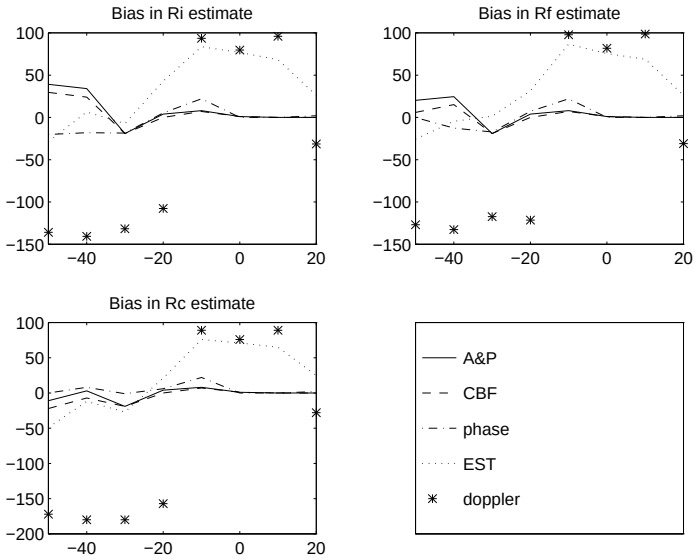

Figure 3.10: Bias in Range Estimates

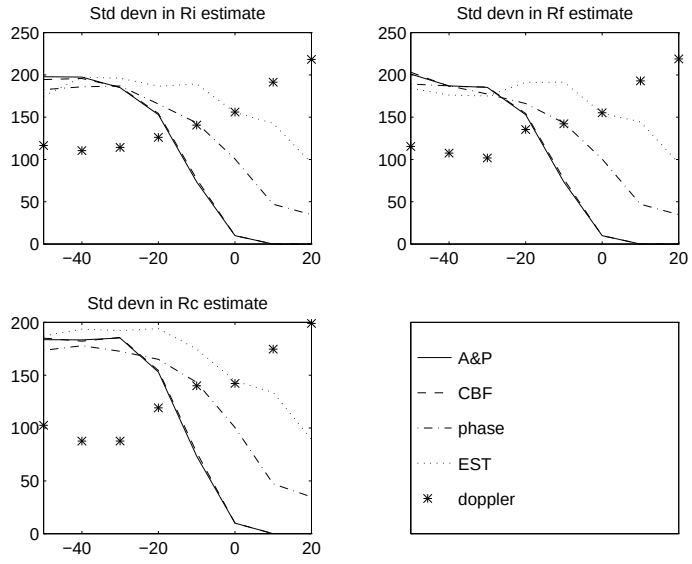

Figure 3.11: Standard Deviation in Range Estimates

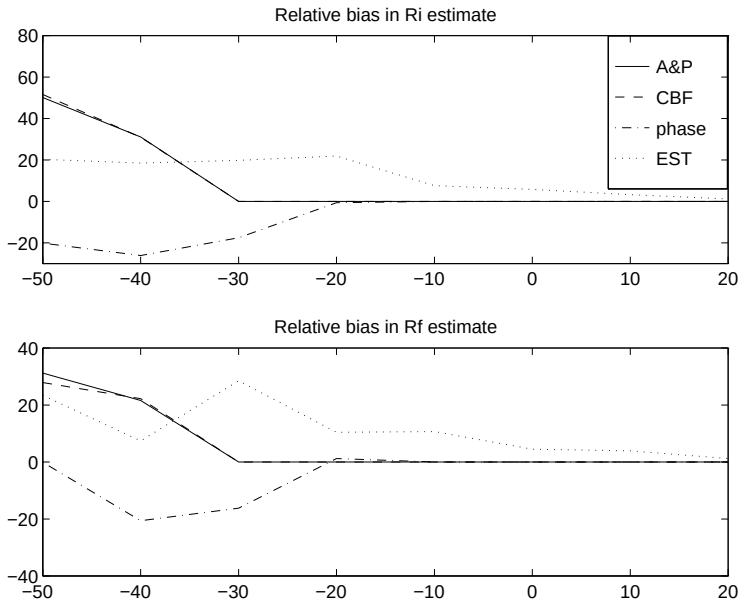

Figure 3.12: Bias in $\hat{r}_i$ & $\hat{r}_f$, relative to $\hat{r}_{cpa}$

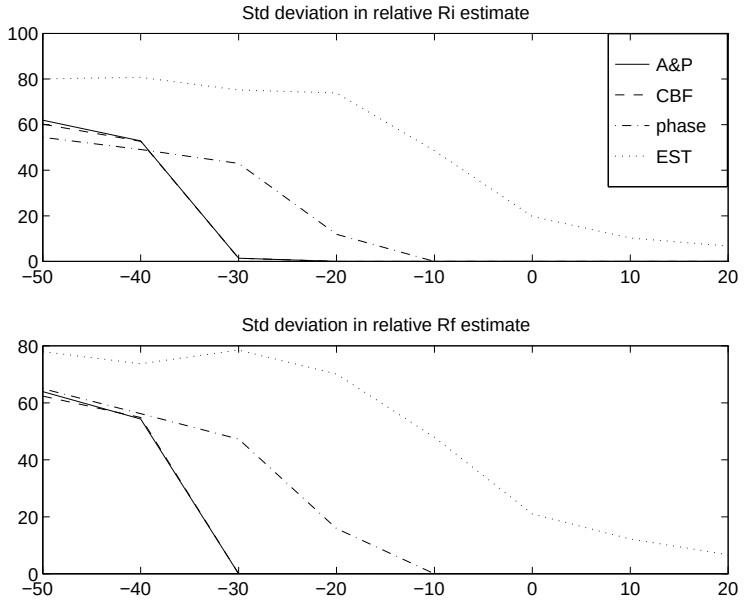

Figure 3.13: Standard deviation of $\hat{r}_i$ & $\hat{r}_f$, relative to $\hat{r}_{cpa}$

Table 3.1: Bias & std. deviation in all range estimates

Estimates of Initial Range

	A & P		CBF		ϕ only		EST		doppler	
SNR	bias	σ	bias	σ	bias	σ	bias	σ	bias	σ
-50	39.1	197.8	29.6	194.5	-20.1	182.6	-28.7	174.9	-135.9	116.5
-40	34.1	197.4	24.0	195.8	-18.1	185.9	6.5	197.3	-140.6	110.5
-30	-19.0	184.9	-19.0	185.5	-18.5	186.9	-7.2	196.2	-131.7	114.3
-20	4.0	153.0	0.0	154.4	5.4	165.5	42.9	186.6	-107.7	126.2
-10	8.0	73.4	7.0	76.9	22.0	143.3	83.6	189.3	93.5	140.7
0	1.0	10.0	1.0	10.0	0.0	100.5	76.8	155.7	79.6	156.0
10	0.0	0.0	0.0	0.0	0.0	47.1	68.3	142.7	95.9	191.3
20	0.0	0.0	0.0	0.0	2.0	34.8	26.2	95.6	-31.4	218.3

Estimates of Final Range

	A & P		CBF		ϕ only		EST		doppler	
SNR	bias	σ	bias	σ	bias	σ	bias	σ	bias	σ
-50	20.2	200.9	5.9	203.1	-0.1	189.1	-25.6	184.3	-126.8	115.3
-40	24.6	187.0	15.2	186.5	-12.6	187.0	-4.6	176.1	-132.6	107.5
-30	-19.0	185.2	-19.0	185.7	-17.2	177.4	1.5	175.1	-117.4	101.7
-20	4.0	153.0	0.0	154.4	7.2	166.2	31.4	191.1	-121.4	135.4
-10	8.0	73.4	7.0	76.9	22.0	143.3	86.7	191.5	97.7	142.2
0	1.0	10.0	1.0	10.0	0.0	100.5	75.5	154.6	81.8	155.3
10	0.0	0.0	0.0	0.0	0.0	47.1	68.9	144.4	98.6	192.9
20	0.0	0.0	0.0	0.0	2.0	34.8	26.2	95.6	-30.9	218.8

Estimates of CPA Range

	A & P		CBF		ϕ only		EST		doppler	
SNR	bias	σ	bias	σ	bias	σ	bias	σ	bias	σ
-50	20.2	200.9	5.9	203.1	-0.1	189.1	-25.6	184.3	-126.8	115.3
-40	24.6	187.0	15.2	186.5	-12.6	187.0	-4.6	176.1	-132.6	107.5
-30	-19.0	185.2	-19.0	185.7	-17.2	177.4	1.5	175.1	-117.4	101.7
-20	4.0	153.0	0.0	154.4	7.2	166.2	31.4	191.1	-121.4	135.4
-10	8.0	73.4	7.0	76.9	22.0	143.3	86.7	191.5	97.7	142.2
0	1.0	10.0	1.0	10.0	0.0	100.5	75.5	154.6	81.8	155.3
10	0.0	0.0	0.0	0.0	0.0	47.1	68.9	144.4	98.6	192.9
20	0.0	0.0	0.0	0.0	2.0	34.8	26.2	95.6	-30.9	218.8

Table 3.2: Bias in $\hat{r}_i$ & $\hat{r}_f$, relative to $\hat{r}_{cpa}$.

$\hat{r}_i - \hat{r}_{cpa}$

	A & P		CBF		ϕ only		EST	
SNR	bias	σ	bias	σ	bias	σ	bias	σ
-50	50.1	61.9	51.6	60.2	-20.1	54.4	20.3	80.1
-40	31.1	52.9	31.0	52.7	-26.1	49.1	18.5	80.7
-30	0.0	1.4	0.0	1.4	-17.5	43.0	19.8	75.2
-20	0.0	0.0	0.0	0.0	-0.6	11.9	21.9	73.9
-10	0.0	0.0	0.0	0.0	0.0	0.0	7.6	48.5
0	0.0	0.0	0.0	0.0	0.0	0.0	5.8	19.8
10	0.0	0.0	0.0	0.0	0.0	0.0	3.3	10.3
20	0.0	0.0	0.0	0.0	0.0	0.0	1.2	6.7

$\hat{r}_f - \hat{r}_{cpa}$

	A & P		CBF		ϕ only		EST	
SNR	bias	σ	bias	σ	bias	σ	bias	σ
-50	31.2	63.9	27.9	62.4	-0.1	65.0	23.4	78.0
-40	21.6	54.3	22.2	54.9	-20.6	56.2	7.4	73.7
-30	0.0	0.0	0.0	0.0	-16.2	47.3	28.5	78.5
-20	0.0	0.0	0.0	0.0	1.2	15.9	10.4	70.1
-10	0.0	0.0	0.0	0.0	0.0	0.0	10.7	47.9
0	0.0	0.0	0.0	0.0	0.0	0.0	4.5	21.0
10	0.0	0.0	0.0	0.0	0.0	0.0	3.9	12.2
20	0.0	0.0	0.0	0.0	0.0	0.0	1.2	6.7

Table 3.3: Bias & Standard Deviation for Simulation 2

Bias			
method	t_{cpa}	v	r_{cpa}
doppler	-0.51	0.02	5.01
A & P	0	0.001	0.3

Standard Deviation			
method	t_{cpa}	v	r_{cpa}
doppler	0.17	0.01	2.40
A & P	0	0.002	0.92

3.4.4 Discussion

The A&P and CBF were the best performed of the cost functions, achieving invariant unbiased estimates above 0 dB (table 3.1). That is, the technique estimated the same values for the source's range parameters in each of the 100 simulations (invariant), and the values it estimated were the correct parameter values (unbiased). Furthermore the estimates of range relative to the CPA (i.e. $\hat{r}_i - \hat{r}_{cpa}$, $\hat{r}_f - \hat{r}_{cpa}$) are unbiased at SNR as low as -30 dB (table 3.2, figs 3.12–3.13). This suggests that errors in range estimates are solely due to poor estimation of the CPA.

Both A&P and CBF formulations perform better than the original amplitude–only formulation and the phase–only method. This demonstrates the importance of exploiting both amplitude and phasor information in the matching function. Neither the amplitude–only cost function nor the phase–only cost function was able to consistently estimate the correct target ranges. However the phase–only function exhibited relatively unbiased estimation at SNRs over 0 dB (table 3.1), and achieved perfect estimation of r_i and r_f relative to $\hat{r}_{cpa}$ at -10 dB (table 3.2); the amplitude–only cost function could not achieve this level of performance at 20 dB. This demonstrates the advantage of considering raw phase information rather than amplitudes in the matching process. This is believed to be the case because signal amplitude does not vary much over a short time interval, especially when the target is not speedy. There is not enough differentiation between replicas. Conversely signal phase varies significantly over even short time intervals and even a raw measurement on $(0, 2\pi)$ will suffice, provided that SNR is not low. The 2–minute sampling interval does not encompass the entire doppler track; it incorporates only the near–linear portion about the CPA. The doppler tracks for targets with different values of r_{cpa} but similar values of $r_i - r_{cpa}$ and $r_f - r_{cpa}$ are almost indistinguishable (figure 3.14). The simulation was re–done at the SNR of -20 dB but with $r_{cpa} = 500$ and $r_i = r_f = 1170$. The velocity of this target was much greater (aprox 17.6 m/s), and a larger portion of the doppler shift was thus present over the 2–minute sampling interval. The doppler tracks of replicas with different values of r_{cpa} are much more noticeable in this case (figure 3.15) and the processors are much more robust to

mismatch, as demonstrated by simulation results: Both the A&P and CBF matching functions yielded perfect estimation of r_{cpa} in each of 100 simulations at -20 dB. Thus the amplitude and phase matching technique is robust to mismatch of CPA range only when a significant portion of the doppler track is present in the signal. This means that longer sampling times will be required for good performance when frequency is low and velocity is small, as is often the case in sonar applications.

The doppler–tracking technique [3] also performed relatively poorly on simulation 1. This technique employs a high resolution frequency estimator to obtain a doppler track for analysis. It is possible that noise in the high resolution frequency estimates may have been sufficiently large to mask the small doppler effect at this low frequency. At higher frequencies the doppler shift is more pronounced and thus the doppler–tracking technique is better suited to tracking higher frequency sources.

The results of simulation 1 show that the A&P cost function performs reasonably when analysing a low frequency -15 dB signal. Simulation 2 demonstrates that it performs near flawlessly when a greater number of timeblocks is used, despite the higher frequency. No doubt the effect of the bigger doppler shift and longer integration time is to make two replicas with similar parameters more distinguishable, hence the better performance at this problem. In fact the A&P marginally outperformed the doppler–tracking technique on simulation 2, which exhibited a small amount of variance in its estimates. It is worth noting that the search domain had to be discretised very finely (2 metres for CPA range and 5 metres for initial and final ranges) in order to detect the variance of A&P estimates. Search spaces similar to those used in simulation 1 are simply too coarsely discretised to detect small variances. This effect is discussed in §4.2.

The A&P matched processor required a much larger computational load than the doppler tracker, due to the size of the data set (10 minutes). Execution of the algorithm was slowed to such an extent that it's use would be prohibitive. With the use of a better optimisation technique it may be possible to reduce the time taken by the A&P, allowing it to run comparably to the doppler tracker with large data sets. Another alternative would be to employ a doppler tracker to obtain initial estimates of the source range

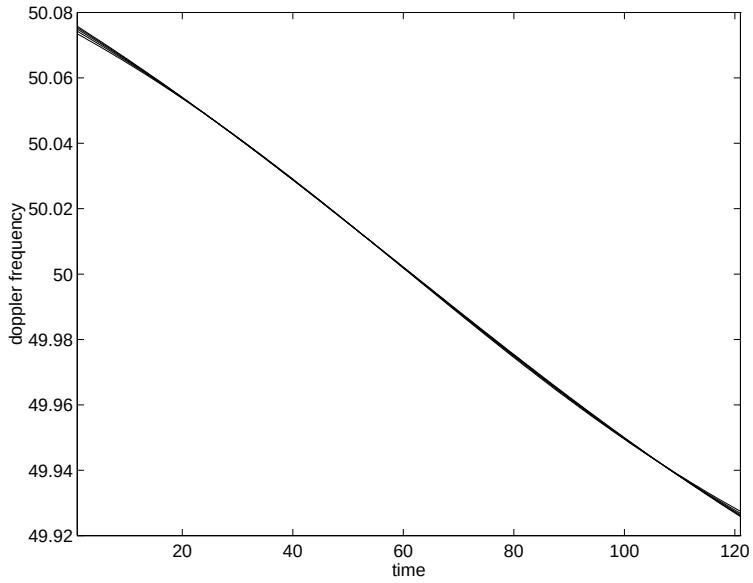

Figure 3.14: Doppler tracks of 5 replicas with different r_{cpa}

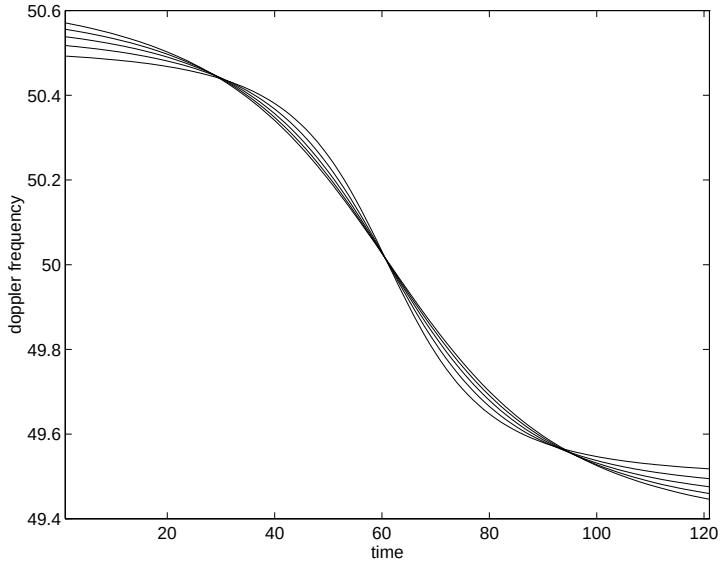

Figure 3.15: Doppler tracks of 5 replicas with different r_{cpa}, high velocity

parameters, thus allowing us to choose a small domain of candidate parameters and hence decrease computational load.

3.5 Summary

This chapter has presented a method for estimation of source motion parameters via matched processing. A simple signal model has been presented which incorporating cylindrical power loss, received (doppler) phase as a function of time–lags, and FFT distortion caused by the doppler shift in frequency. This model is used to generate sparse replicas of the acoustic field; having only one measurement of signal amplitude and phase per timeblock. Several matching functions have been defined and applied to analysing simulated data with replicas generated according to the presented model. Emphasis is given to the coherent processing of signal phase by these cost functions. The simulation results demonstrate that coherent phase matching performs far better than formulations which perform incoherent phase matching or amplitudes–only processing. The phase–matching approach is also demonstrably better than a frequency–based doppler–tracking technique at low frequency. It is believed this is because of the small doppler effect at low frequency. The two techniques both perform well at high frequency although the doppler–tracker has the advantage of being less computationally intensive.

Chapter 4

Bounds on Estimator Variance

The matched doppler processing technique is evaluated as an estimation process. The Cramer–Rao lower bound for estimator variance is derived for the simulations in the preceding chapter. The simulation results are compared to these bounds in order to determine estimator efficiency. The fact that the technique performs a search over a discrete parameter space means that the variance is limited by the size of the domain and quantisation noise will be present in the estimates. Expressions for these domain–enforced limits are determined and presented here.

4.1 The Cramer–Rao Lower Bound on Range Estimates

It is evident from figure 3.14 that the doppler frequency of the signal is nearly linear over the short two minute sampling interval and the amplitude of the signal over this period is nearly constant. The signal phase may therefore be approximated by a unitary amplitude quadratic phase signal, i.e. a chirp.

$$s(t) = \exp\left(i2\pi\left(\theta_0 + \theta_1 t + \theta_2 t^2\right)\right)$$

By considering the diagonal of the inverse Fisher information matrix for the poly-

nomial phase coefficients $\underline{\theta} = (\theta_0\, \theta_1\, \theta_2)$, the Cramer–Rao lower bounds on estimates of these coefficients can be shown to be

$$\begin{aligned}
\sigma^2_{\theta_0} &= C\left(3N^2 + 3N + 2\right) \\
\sigma^2_{\theta_1} &= C\frac{4\,(2N+1)\,(8N+11)}{(N+1)\,(N+2)} \\
\sigma^2_{\theta_2} &= C\frac{60}{(N+1)\,(N+2)}
\end{aligned}$$

where

$$C = \frac{3\sigma_n^2}{8\pi^2 N\,(N-1)\,(N-2)}$$

and σ_n^2 is the noise variance and N is the number of timeblocks sampled.

The bounds on the quadratic phase coefficients may be used to compute bounds for estimators of initial, final and CPA range. These bounds are computed numerically by multiplication of the inverse Fisher matrix by a Jacobian comprised of partial derivatives of the range parameters r_i, r_f, r_{cpa} with respect to the phase coefficients $\underline{\theta}$. An outline of this process is presented in the appendix (§A).

It must be noted that the bounds on variance of range estimates depend on the actual values of the range parameters. Furthermore the bounds decrease in a linear fashion as noise variance σ_n^2 (and hence SNR) decreases and the bounds vary inversely with the sample size N. Therefore the estimators can satisfy the performance criterion of consistency: the minimum possible error decreases as the number of observations increase.

4.2 Bounds Imposed by the Discrete Nature of the Domain

4.2.1 Variance on a Discrete Uniformly Spaced Interval

The discrete and finite nature of the search space used in practice imposes bounds on the variance with which parameters can be estimated. It can be shown that values which are

uniformly sampled from a length L interval of N equispaced points will have a variance equal to

$$\sigma_u^2 = \frac{(N+1)}{(N-1)}\frac{L^2}{12}$$

However a greater variance may be observed if the sampling is biased towards either end of the interval and only the most extreme elements are ever selected. This is likely to occur at low SNR. In this case, $N = 2$ and thus

$$\sigma_{max}^2 = \frac{L^2}{4}$$

4.2.2 Quantisation Noise

The resolution with which estimates of the source parameters can be made is determined by how finely the search space is sampled. Assuming that the optimum value will be associated with its closest point on the search domain, the estimates will be in error by up to $\frac{\Delta}{2}$ units, where

$$\Delta = \frac{L}{N-1}$$

This "quantisation noise" on the estimates will have a variance equal to that of a uniformly distributed random value on the interval $(-\frac{\Delta}{2}, \frac{\Delta}{2})$, which can be shown to be

$$\sigma_q^2 = \frac{L^2}{12(N-1)^2}$$

4.3 Actual Bounds on Simulations

The r_{cpa} domain for simulation 1 had 5 elements, equally spaced 100 metres apart; the domain for r_i and r_f was not equally spaced because only certain values were allowed for each value of r_{cpa}, however the relative range domain (for $r_i - r_{cpa}$ and $r_f - r_{cpa}$) contained 20 elements, equally spaced 100 metres apart.

The CRLBs for range estimates from a unitary amplitude chirp signal (dashed line) are compared with observed variance using the phase–only (3.15) cost function (solid line) in figure 4.1. The variance bounds imposed by the domain discretisation are

also shown as dotted lines. The uppermost of these lines is the maximum variance σ^2_{max}. Beneath this is the variance on a discrete uniformly spaced interval σ^2_u. The lowest dotted line corresponds to the quantisation noise σ^2_q. Observe that for SNRs below -10 dB, $\hat{r}_{cpa}$ exhibits more variance than a uniform random sample of the domain; the optimal point of the ambiguity surface varies too wildly to be located within the sample space.

In the simulations care was taken to ensure that the set of true target parameters was always an element of the search domain, so quantisation noise will not have the same effect as it would under real conditions. In spite of this the phase–only estimates are very poor; the CRLB is not approached by the estimator.

Further simulation was performed with the phase–only cost function, in which the true values of the source parameters were not exact sample points on the domain, but instead were randomly selected from a neighbourhood about a sample point. Results show that even at high SNR the phase–only estimator performed as badly as a random guess (figure 4.2). This bad performance is no doubt due to the extremely noisy nature of the phase–only ambiguity surface as is evident from (fig 3.8). Even a small deviation from a sampling point is enough to cause the peak of the ambiguity function to be missed. Thus the phase–only cost function cannot be satisfactorily used for estimation of source parameters.

Fortunately the Amplitude–and–Phase cost function (3.13) is better behaved under conditions of quantisation noise. Further simulation was performed using this cost function and a domain which did not necessarily include the exact true parameter values. Variance of the estimates is displayed in figure 4.3. This was generated with a domain coarseness of 100m.

By discretising the r_{cpa} domain more finely, more accurate estimates may be formed. This is evident in figure 4.4 which was done with a domain sampled at every 10m. The finer discretisation allows the estimator variance to approach the CRLB at higher values of SNR. The dashed line representing the CRLB is plotted here both raw and with the quantisation noise variance added. Notice how the sum of CRLB and quantisation noise curves away from the plain CRLB at approx. 0 dB. Notice also that the range estimates

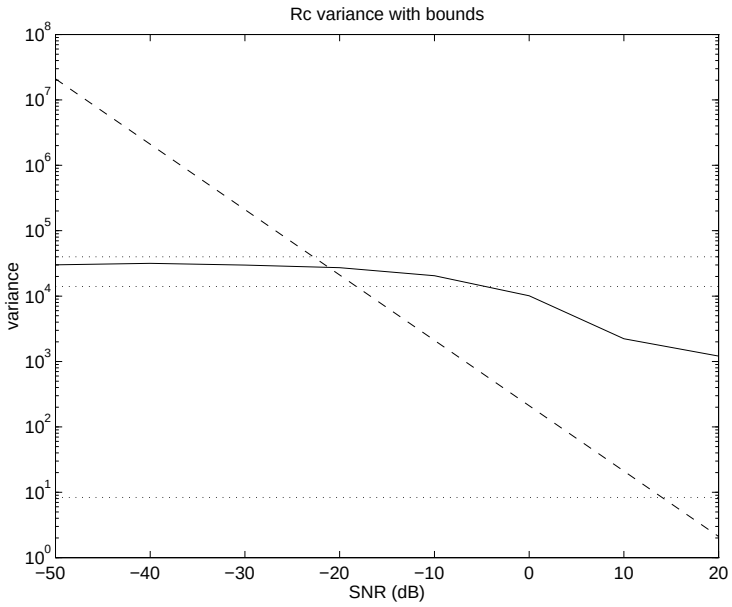

Figure 4.1: Bounds on r_{cpa} estimator variance, ϕ–only

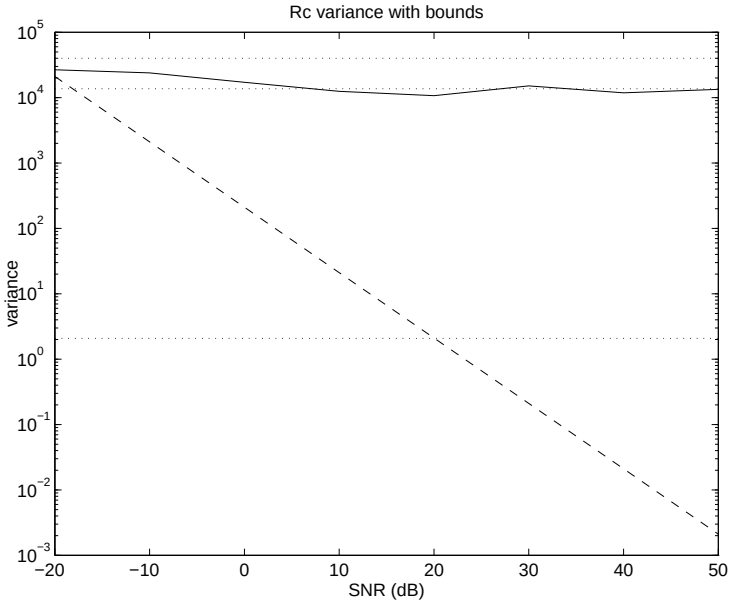

Figure 4.2: Bounds on r_{cpa} estimator variance, ϕ–only with quantisation errors

closely adhere to this bound until the quantisation noise becomes the dominant component at around 30 dB. This demonstrates that both CRLB and quantisation noise affect the ability to estimate source range parameters with accuracy, and that the performance bounds are approached by the coherent phase matching technique. Furthermore equations for the CRLB indicate that it decreases as the number of samples increases. The phase–coherent technique is thus efficient and consistent.

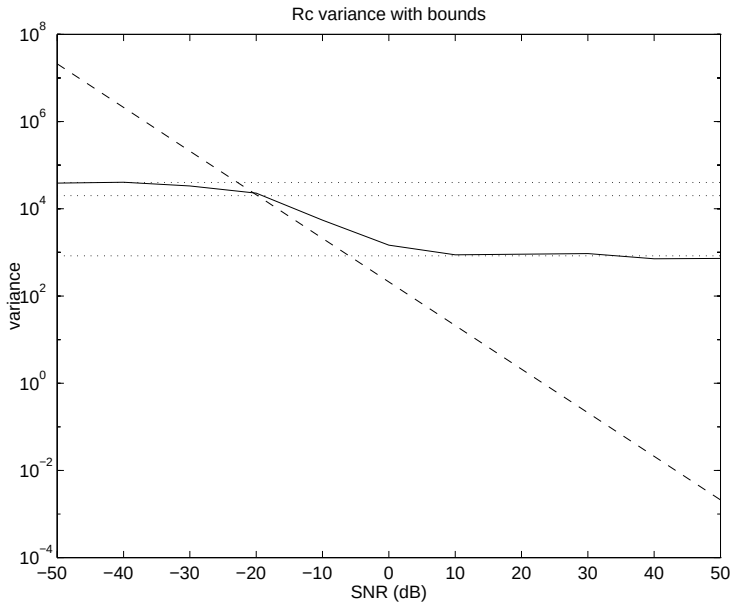

Figure 4.3: Bounds on r_{cpa} estimator variance, A&P

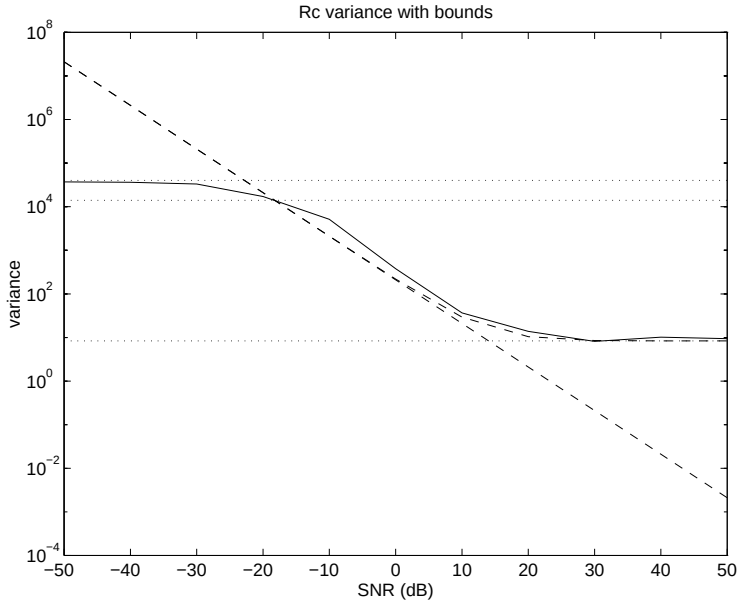

Figure 4.4: Bounds on r_{cpa} estimator variance, A&P, fine domain

Chapter 5

Robustness Issues

Previous chapters demonstrated that matched processing with the A&P cost function performs reasonably well when tracking a simulated target which moves in a straight line. However this will not be the case in reality. No craft will move in a perfectly straight line; variations in the craft's trajectory will be induced by ocean currents and turbulence. Furthermore the receiver will not necessarily remain in a fixed position. In the case of a sonobuoy, it is not uncommon for these receivers to drift from the position at which they are deployed. The ability of the matched doppler processing technique to perform in conditions of drift must be determined.

5.1 Effect of Range Perturbation on Signal Phase

The amount of phase corruption caused by a shift in range is determined by the signal wavelength λ. A maximum phase error of π radians will occur when the source deviates from it's locus by $\frac{\lambda}{2}$ metres. The phase error is cyclic on the interval $(-\pi, \pi)$, becoming zero when the range deviates by a whole number of wavelengths. In terms of signal frequency,

$$\xi \text{ metres} \rightarrow 2\pi\frac{f}{c}\xi \text{ radians} = 360\frac{f}{c}\xi \text{ degrees}$$

Thus signals of higher frequency will experience greater disruption of phase due to source wobble.

5.2 Model for Range Perturbation

A model of the range perturbations from a straight–line locus for a source having given range parameters (initial, final and CPA range) is required. Such a model must take into account the fact that the range deviance at time t_n depends somewhat on the deviance at time t_{n-1}. In other words, the target's position at one timeblock depends on the target's position in the previous timeblock. Furthermore, the expectation of the range perturbations in each timeblock should be zero. This is required to ensure that the target wobbles about it's locus; were it not the case, the cumulative effect of the perturbation would cause the track's actual parameters to differ significantly from it's conceptual parameters.

Thus the range perturbation is modelled as a first–order autoregressive process

$$\xi_n = \alpha \xi_{n-1} + N(0, \sigma^2)$$

where ξ_n is the range perturbation at time t'_n, $N(0, \sigma^2)$ is a normal random deviate of 0 mean and specified variance and α a number on $(0, 1)$. It is obvious from this equation that ξ_n depends upon ξ_{n-1} and it can be shown that $\mathrm{E}\{\xi_n\} = 0$, satisfying the criteria for a model of track wobble. Furthermore the variance of an AR(1) process can be shown to be

$$E\left[\xi_n^2\right] = \frac{\sigma^2}{1 - \alpha^2}$$

and thus the range variance generated by a particular AR model may be computed.

5.3 Signal Model with Range Perturbations

Any such deviations in the relative range between source and receiver will cause perturbations to the received amplitude and phase of the signal; most notably to the phase.

The assumption that the signal we receive at time t is the signal which was emitted at time t' is maintained. However the range of the target at time t' will have an extra term to account for deviance from a straight line trajectory. The received amplitude is thus

$$A(t_n) = \sqrt{\frac{A_o^2}{r(t'_n) + \xi_n}} \tag{5.1}$$

The signal emitted at time t'_n must travel an extra ξ_n metres to reach the receiver, hence the signal will be retarded by a further $\frac{\xi_n}{c}$ seconds. The signal phase received is actually the phase which was emitted at time $t'_n - \frac{\xi_n}{c}$, i.e.

$$\phi_r(t_n) = \phi_e\left(t'_n - \frac{\xi_n}{c}\right) = \phi_{cpa} + 2\pi f_s\left(t'_n - \frac{\xi_n}{c}\right) \tag{5.2}$$

A signal can be simulated from this modified signal model in the same manner as (3.18), using (5.1) and (5.2) in place of the original amplitude and phase equations (3.4) and (3.2).

5.4 Robustness Simulations

The simulated target had initial and final range of 770 metres and had a range of 700 metres at CPA. It emitted a 50 Hz. signal during the 2 minute sampling interval. No noise was added to the simulated signal, so any error in the range estimates will be wholly due to source wobble. Several combinations of AR model parameters were used to generate the target wobble. 100 simulations were performed for each combination and processed with the A&P cost function (3.13). Sample bias and standard deviation were computed and are presented in table 5.1.

The simulation results demonstrate that even a small amount of wobble in the source track is enough to produce variance in the estimates of range at CPA, $\hat{r}_{cpa}$, and hence variance in initial and final range. However the estimates of initial and final range relative to range at CPA (i.e. $\hat{r}_i - \hat{r}_{cpa}$ and $\hat{r}_f - \hat{r}_{cpa}$) are relatively unaffected by track wobble. Simulation suggests that these estimates will be invariant and unbiased if the standard deviation of the track wobble, σ_ξ, is less than about 3 metres.

The conclusion is that estimates of initial and final range relative to $\hat{r}_{cpa}$ are robust; however estimation of r_{cpa} itself is extremely sensitive to track deviation when analysing this class of signal. This is not surprising; it occurs because of the great similarity between replicas which differ only in r_{cpa} as discussed in §3.4.4.

The doppler effect is more pronounced at higher frequency and replicas are thus more easily distinguishable. However this benefit is countered by the fact that at high frequency, larger phase errors are produced by deviations in source range. For example, an AR process which has $\sigma^2 = 1$ and $\alpha = 0.1$ produces a sequence of range offsets which has standard deviation $\sigma_\xi = 1.0$. The phase standard deviation σ_ϕ associated with a 50 Hz signal is 12.1 degrees. However at 180 Hz, the associated phase standard deviation σ_ϕ is 43.4 degrees. At 300 Hz the deviation in phase is 72.3 degrees. Some investigative simulations were performed at $f_s = 180$, and the results obtained were no better than the results of the 50 Hz simulations.

Since a larger frequency cannot be reliably used, a longer sampling time is required so that more doppler information is included in the replicas and thus robust estimation of r_{cpa} may be achieved for this class of problem.

The simulation was re–done for the same target, but this time a slightly larger sampling interval (5 minutes) was used. The results are presented in table 5.2. The improvement in performance is quite noticeable; in fact the technique is fairly robust to range perturbation when σ_ξ, the standard deviation of the AR process used to model its wobble, is below 2 metres. This supports the theory that better robustness to track variation is achieved when longer sampling times are used.

Table 5.1: Bias & std deviation in estimates of wobbly track

ξ parameters				r_i		r_f		r_{cpa}		$r_i - r_{cpa}$		$r_f - r_{cpa}$	
σ^2	α	σ_ξ	σ_ϕ	bias	σ	bias	σ	bias	σ	bias	σ	bias	σ
0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
1	0.1	1.0	12.1	1.0	94.8	1.0	94.8	1.0	94.8	0.0	0.0	0.0	0.0
1	0.5	1.2	13.9	-32.0	139.2	-32.0	139.2	-32.0	139.2	0.0	0.0	0.0	0.0
3	0.1	1.7	20.9	-2.0	142.1	-2.0	142.1	-2.0	142.1	0.0	0.0	0.0	0.0
3	0.5	2.0	24.0	33.0	162.1	33.0	162.1	33.0	162.1	0.0	0.0	0.0	0.0
1	0.9	2.3	27.5	21.0	186.0	21.1	186.1	21.0	186.0	0.0	2.0	0.1	1.7
5	0.5	2.6	31.0	-22.0	175.6	-22.0	175.6	-22.0	175.6	0.0	0.0	0.0	0.0
3	0.9	4.0	47.7	12.3	190.1	12.3	191.8	12.0	191.4	0.3	6.7	0.3	6.9
5	0.9	5.1	61.6	-22.8	182.3	-25.2	187.3	-25.0	186.1	2.2	13.8	-0.2	12.7

Table 5.2: Bias & std deviation in estimates of wobbly track, longer sample

ξ parameters				r_i		r_f		r_{cpa}		$r_i - r_{cpa}$		$r_f - r_{cpa}$	
σ^2	α	σ_ξ	σ_ϕ	bias	σ	bias	σ	bias	σ	bias	σ	bias	σ
1	0.1	1.0	12.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
1	0.5	1.2	13.9	1.0	10.0	1.0	10.0	1.0	10.0	0.0	0.0	0.0	0.0
3	0.1	1.7	20.9	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
3	0.5	2.0	24.0	0.0	14.2	0.0	14.2	0.0	14.2	0.0	0.0	0.0	0.0
1	0.9	2.3	27.5	11.0	51.0	11.0	51.0	11.0	51.0	0.0	0.0	0.0	0.0
5	0.5	2.6	31.0	2.0	14.1	2.0	14.1	2.0	14.1	0.0	0.0	0.0	0.0
3	0.9	4.0	47.7	14.2	93.5	15.2	93.1	15.0	92.5	-0.8	4.2	0.2	4.5
5	0.9	5.1	61.6	12.4	120.1	11.3	123.0	12.0	120.8	0.4	15.9	-0.7	11.1

5.5 Summary

The matched doppler processing technique is robust to perturbations in the source track provided that enough of the signal is sampled in order to enable clear differentiation between replicas. In other words, the sampling period must be long enough to encompass a significant portion of the doppler shift.

The speed of the source in these simulations was approximately 6.4 m/s. It is reasonable to asume that a real source travelling at this speed will exhibit less than 2 metres deviation in it's track. It is not uncommon for an underwater source to have approximately constant course and speed [23] and thus any wobbles are expected to be of an order smaller than the source's velocity. Hence, provided that the assumption of linear motion is met, the technique should be applicable to real data.

Chapter 6

Application to Real Data

Real data from a sonobuoy was procured, in which a target moved past while emitting two low–frequency sinusoids and a high–frequency pilot tone. This data was processed with the A&P cost function (3.13). It was known that range at CPA was approximately 700 m and that the craft was moving with a speed of roughly 5 knots (aprox. 2.5 m/s). Therefore the r_{cpa} domain was restricted to below 1000m and initial and final ranges between 25 and 400 metres greater than r_{cpa} were considered.

The robustness simulations in the previous chapter determined that 5 minutes of data was sufficient to determine CPA range in the presence of source wobble at low SNR and low frequency. The set of real data was much longer but unfortunately the computational demands of the matching process prohibited the use of a segment longer than this. The point of closest approach to the sonobuoy was determined by examining a high–resolution lofargram of the data and locating the time sample at which the doppler shift occurred. A 5 minute segment about the CPA was chosen for analysis.

Figure 6.1 shows the amplitude and phase measured from the real data in one of the low–frequency bins. The dashed lines correspond to the amplitude and phase of the replica which produced the lowest cost when processed with the A&P cost function. This replica was generated at a CPA of 800m and velocity of approximately 2.7 m/s, which are realistic estimates of the true CPA and velocity.

The amplitude and phase from the other low–frequency bin are shown in figure 6.2. Notice that the replica amplitude is much stronger than the measured amplitude. This occurs because the signal power was assumed to be higher than it actually is. However this amplitude difference factors out of the cost function and has no effect. What really matters is the shape of the amplitude curve, which is a manifestation of source range and the doppler effect. Similar estimates of the craft's range and velocity were obtained, despite the lower SNR evident at this frequency.

The amplitude and phase from the high–frequency pilot tone were also analysed with the matched processor. This yielded the somewhat different CPA estimate of 400 m and a craft velocity of 1.7 m/s. The processors are expected to be far less robust to track wobble and sensor movement at high frequency because large phase variations can occur with extremely small range perturbations. Real data is expected to be subject to such perturbations. This may be a possible cause of the processor's poor estimation of the target range at this frequency.

A cursory examination of the high–frequency bin's amplitude and phase plot (fig. 6.3) reveals that amplitude varied wildly in the latter half of the sampling interval. This "noise spike" may also account for the poor performance. Despite this the technique does appear to have detected CPA quite well, as the selected replica appears to exhibit a phase transition near the centre of the sampling interval.

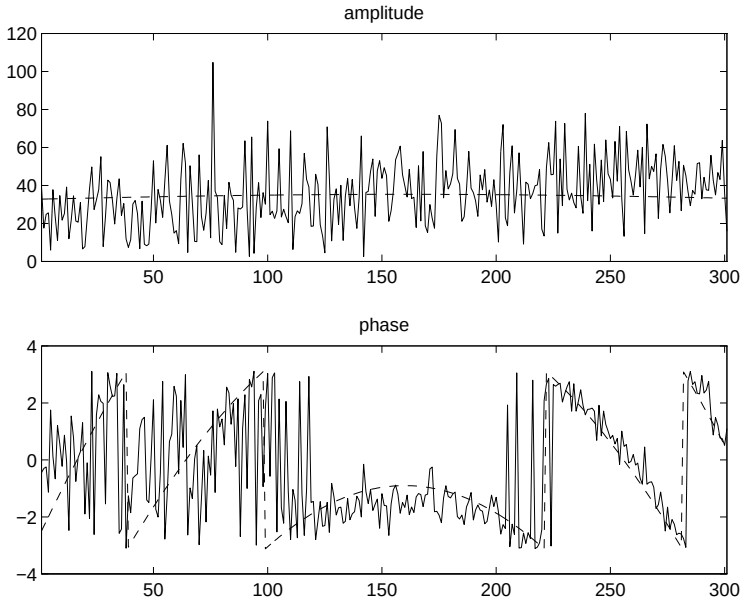

Figure 6.1: Measured Amp & Phase with replica, low freq 1

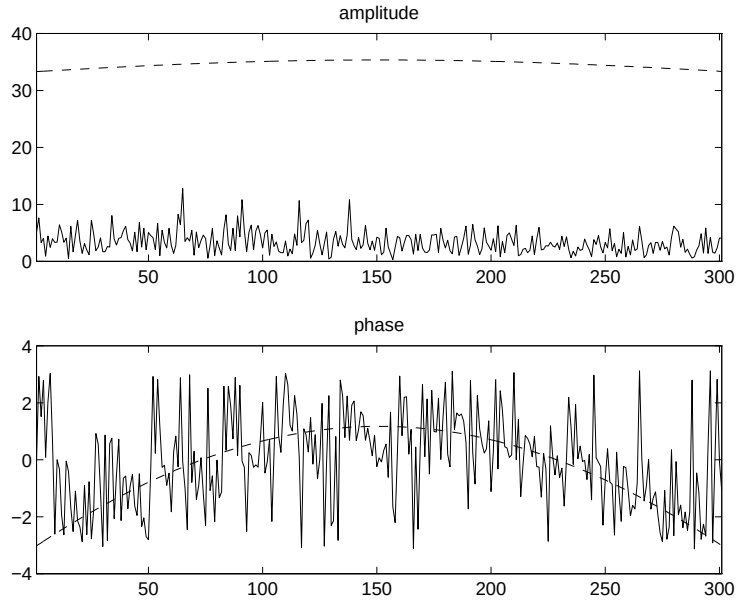

Figure 6.2: Measured Amp & Phase with replica, low freq 2

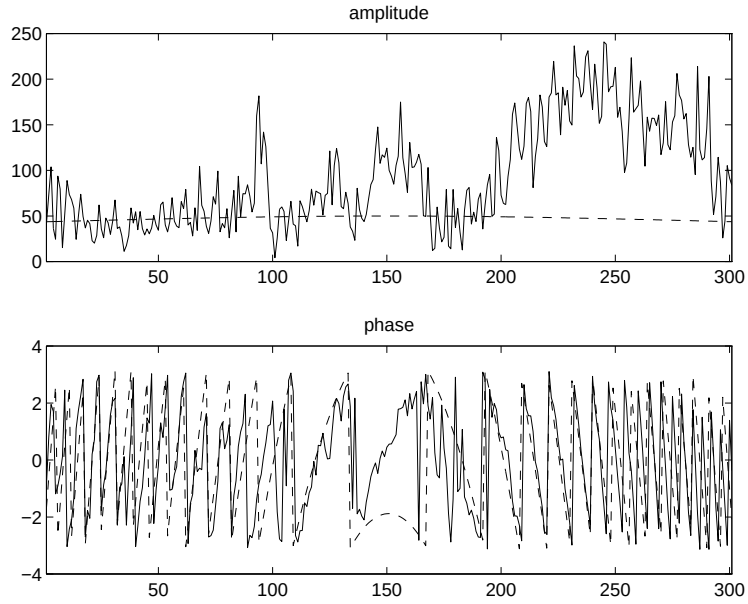

Figure 6.3: Measured Amp & Phase with replica, hi freq

Chapter 7

Multiple Receivers

It is desirable to extend the matched doppler processing algorithm to function with multiple receivers for several reasons. Firstly is the motivation of better performance in noise. If the outputs of N receivers are combined in such a fashion that signal components are combined in phase, the effective signal strength increases by a factor of N. Thus more receivers should enable the technique to function better at low SNR.

Secondly using more than one omnidirectional receiver enables us to locate a source in bearing as well as range. Furthermore it will be possible to determine a target's heading as well as it's speed.

Thirdly it may be possible to perform the technique in the event of a receiver failure. Multiple receiver TMA can function adequately even when the signal drops out at one or more receivers. This robustness to receiver failure can be achieved only with the use of several receivers.

Fourthly it will also be possible to focus on one particular target of interest when there are other sound sources present in the environment.

However the advantages gained by extending the technique to use multiple receivers does involve a trade–off. As more receivers are considered, the computational load required to process them will increase. Furthermore the ability to pinpoint the source in two dimensions will require a search space of greater dimension, resulting in a combina-

torial increase of the processing required to enumerate all cost values.

Another problem which needs to be addressed is the fact that the doppler shift at each receiver will be unique; that is, each receiver experiences a doppler shift different from each of the other receivers. This prohibits a simple time–delay–and–sum approach, as is the case in conventional beamforming for a stationary source, because time–lags alone are not enough to account for the difference in signal phase at each receiver.

7.1 Geometry of the Problem

In the single receiver case expressions for values like the time of CPA, t_{cpa}, were defined in terms of source ranges and speed. However since one of the benefits of using multiple receivers is the ability to determine source positions and headings, it is worth expressing the problem model in terms of cartesian positional information. This enables evaluation of expressions in terms of actual source position and heading, rather than ranges and speed.

For a source with hypothesised initial position $\underline{S} = (S_x, S_y)$ and two–dimensional velocity $\underline{V} = (V_x, V_y)$, the range and time of the CPA to the kth receiver located at $\underline{H}^{(k)} = (H_x^{(k)}, H_y^{(k)})$ can be computed as

$$t_{cpa}^{(k)} = \frac{(H_x^{(k)} - S_x)V_x + (H_y^{(k)} - S_y)V_y}{v^2} \tag{7.1}$$

$$= \frac{(\underline{H}^{(k)} - \underline{S})^T \underline{V}}{v^2} \tag{7.2}$$

$$r_{cpa}^{(k)} = \frac{\left|(S_x - H_x^{(k)})V_y - (S_y - H_y^{(k)})V_x\right|}{v^2} \tag{7.3}$$

where

$$v = \|\underline{V}\| = \sqrt{V_x^2 + V_y^2} \tag{7.4}$$

The emission time at this receiver, $t^{(k)\prime}$, may be computed as before (3.1). Taking the smaller root of this quadratic gives

$$t^{(k)\prime} = \frac{c^2 t + \vec{\underline{HS}}^T \underline{V} - \sqrt{c^2 \left\|\vec{\underline{HS}} + t\underline{V}\right\|^2 + \left(\underline{V}^T \vec{\underline{HS}}\right)^2 - v \left\|\vec{\underline{HS}}\right\|^2}}{c^2 - v^2}$$

where

$$\underline{\vec{HS}} = \underline{S} - \underline{H}^{(k)}$$

Likewise the doppler frequency received at this receiver at time t can be computed from the doppler equation (3.6)

$$f_r^{(k)}(t) = \frac{1}{1 + \frac{v^2 t^{(k)\prime}}{c^2(t - t^{(k)\prime})}} f_s$$

and the associated phase is thus

$$\phi_r^{(k)}(t) = \phi_{cpa} + 2\pi f_s t^{(k)\prime}$$

The amplitude and phase distortion inherent in the f_s Hz bin of a DFT performed on a signal shifted $\delta^{(k)}(t) = f_r^{(k)}(t) - f_s$ Hz from centre cell at time t is

$$D^{(k)}(t) = \frac{\sin \pi \delta^{(k)}(t)}{\sin \frac{\pi \delta^{(k)}(t)}{N}} \exp\left(\pi \delta^{(k)}(t) \frac{N-1}{N}\right)$$

The signal measurements (DFT values) are as before (eq. 3.10), except now we must model the signal at several receiver locations instead of just one. If an N–point DFT is computed at the timeblock beginning at time t at the kth receiver then the f_s Hz Fourier coefficient is

$$S^{(k)}(f_s, t) = D^{(k)}(t) A^{(k)}(t) \exp(i\phi_r^{(k)}(t))$$

Samples of this form the signal replica vectors for each receiver:

$$P_n^{(k)} = S^{(k)}(f_s, t_i + (n-1)T)$$

7.2 Spatially Incoherent Method

The obvious straightforward method of using multiple receivers for matched tracking is to process each receiver separately and sum the individual cost values. That is,

$$\begin{aligned} E_{inc} &= \sum_{k=1}^{K} E_{CBF}(k) \\ &= \sum_{k=1}^{K} \left| \underline{\tilde{p}}^{(k)H} \underline{\tilde{P}}^{(k)} \right|^2 \\ &= \sum_{k=1}^{K} \left| \sum_{n=1}^{N} \tilde{p}_n^{(k)*} \tilde{P}_n^{(k)} \right|^2 \end{aligned} \tag{7.5}$$

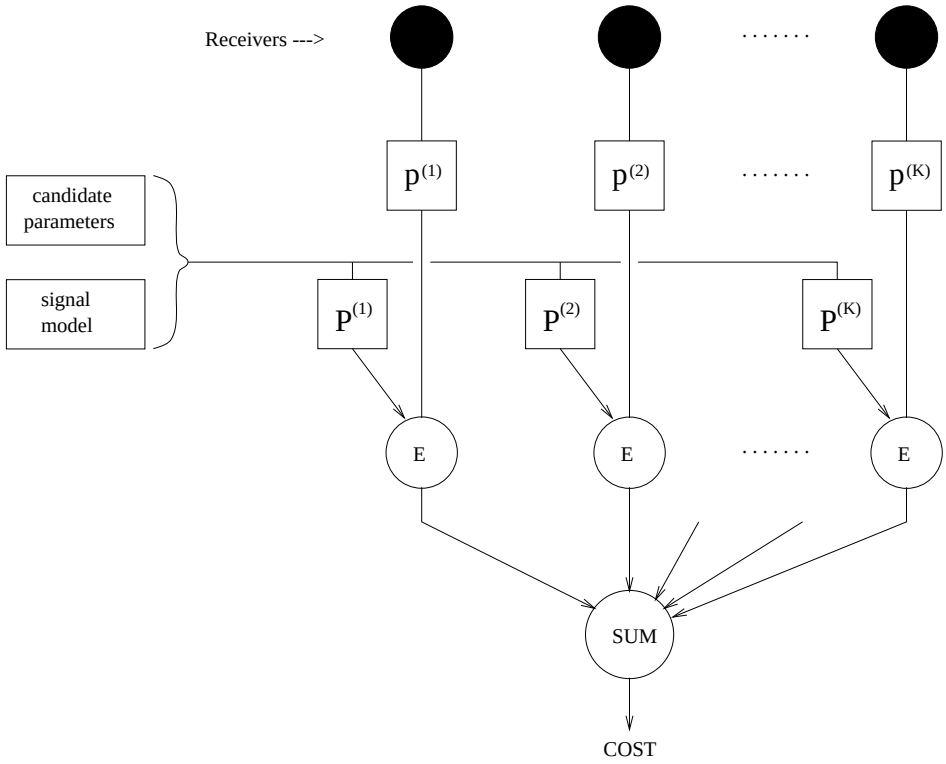

Figure 7.1: Spatially Incoherent Matched Processing with K Receivers

where $E_{CBF}(k)$ is the CBF cost value derived from the kth receiver (eq. 3.14). Note that any arbitrary cost function could be used. The CBF function is chosen due to it's good performance in analysis of simulated data and it's form obviates the coherent processing of phase. Each cost value $E_{CBF}(k)$ is a function of the source's ranges from that single receiver k. It is not a unique function of the source's two dimensional positions. However the sum of the cost values *is* a function of the craft's actual positions because the set of ranges to each receiver defines the source position unambiguously when there are 3 or more receivers arranged in a non–linear fashion. Two candidate positions having the same range to one receiver may have different ranges to another differently placed receiver.

This method has the advantage that it is very simple; it is easily implemented as an application of the single–receiver cost functions described in the previous chapters. A diagram of the process is given in figure 7.1.

A simple modification of this involves a set of weights $\underline{\alpha} = (\alpha_1, \alpha_2, \ldots \alpha_K)$, one for

each receiver such that

$$E = \sum_{k=1}^{K} \alpha_k E_{CBF}(k)$$

Receivers with a higher weighting would have greater effect on the cost value. Presumably one would wish to ascribe sensors closer to the source a greater weighting as the signal–to–noise ratio will be greater at these sensors due to less power attenuation. However such a weighting is already present in the signal model; received signal amplitude is modelled as varying inversely with the distance from the source. This is also reflected in the replicas. The signal model ensures that the closest receivers will get the better weighting. Furthermore this weighting evolves over time to reflect the signal strength as the target move toward or away from a receiver.

Currently the normalisation process removes range dependencies from the data and thus dispenses with the intrinsic receiver weighting. It would be necessary to use un–normalised data and replica vectors $\underline{p}$ and $\underline{P}$ in order to exploit the natural weighting mechanism described above. Receiver weighting has not been used in this study, but it is worthwhile to note the possibility.

It is also worth noting another variant on the spatially incoherent cost function. By reversing the order of the summation, it is possible to process phase coherently across space but not time:

$$E = \sum_{n=1}^{N} \left| \sum_{k=1}^{K} \tilde{p}_n^{(k)*} \tilde{P}_n^{(k)} \right|^2$$

In practice the number of receivers is usually far less than the number of time samples. Because of this, the gain afforded by this matching function is expected to be less than that of the function discussed above. Therefore this variant is not considered any further.

7.3 Spatially Coherent Method

Each receiver will experience a different doppler effect as the source moves past, by virtue of the fact that each receiver has different ranges to CPA and times of CPA. The incoherent combination of cost values described above blithely ignores the doppler–phase

relationships between the receivers as it treats each receiver as a stand–alone; information from the several receivers is combined only once it has been condensed into a set of cost values, none of which contains any relative phase information.

The immediate problem is to coherently combine phase information from the receivers in order to exploit the relative doppler–phase differences and thus improve performance.

In conventional time–delay–and–sum beamforming, the signal from a stationary source at each receiver exhibits a different time–lag from the source due to the varying distances between each receiver and the source. Coherent combination of the received signals is achieved by compensating for the time–lag at each receiver so that the signals line up in phase. Signal power is thus amplified when the receiver outputs are correctly delayed and combined.

The matched tracking problem is analogous in that each receiver experiences a different time of CPA, but a simple time–shift is not sufficient to line up the signals in phase. This is because the different range at CPA of each receiver causes the doppler shift to vary in duration between receivers. The closer a receiver is to the source, the shorter is the time interval over which the doppler shift takes place. Since the signal phase is dependent on the doppler shift, it is clear that simply translating the reciever outputs to a common time of CPA will not be sufficient to enable coherent integration. The doppler effect must be accounted for at each receiver before the signals may be combined.

The single receiver Matched Doppler Processor can be thought of as an extension of matched filtering. Similarly the multiple receiver algorithm be thought of as matched filtering, only in this case the replicas are matched in time and space, not merely in time. For a given source trajectory, a set of replica vectors may be generated which coherently combine the signal phase from this locus, incident at several spatial locations and several consecutive time instants. The signal model presented in chapter 3 and reformulated in terms of position rather than range in §7.1 is suitable for generating such replicas. What remains is to extend the existing correlation functions in order to process phase coherently across space as well as time.

Ideally phase information will be processed coherently across time and space. The replica vectors as defined previously are already a function of both time and space; replica phase depends on both emission time t' and range r. Because of this, we may filter the received signals in space and time by a simple modification to the spatially–incoherent cost function (eq. 7.5):

$$\begin{aligned} E_{co} &= \left| \sum_{k=1}^{K} \sum_{n=1}^{N} \tilde{p}_n^{(k)*} \tilde{P}_n^{(k)} \right|^2 \\ &= \left| \sum_{k=1}^{K} \underline{\tilde{p}}^{(k)H} \underline{\tilde{P}}^{(k)} \right|^2 \end{aligned} \tag{7.6}$$

The updated cost function above leaves phase intact while processing across receievers and thus any residual phase information will be degraded. However the correlation value for a perfectly matched replica will not be altered (disregarding noise effects) because phase will combine coherently across space for this replica in either case.

Consider a signal in the absence of noise

$$p_n^{(k)} = A_n^{(k)} \exp i \left(2\pi f_s t_n^{(k)\prime} + \phi_c \right)$$

and the replica of this signal

$$P_n^{(k)} = A_n^{(k)} \exp i 2\pi f_s t_n^{(k)\prime}$$

Then the values of the spatially incoherent cost function E_{inc} and the spatially coherent cost function E_{co} are

$$\begin{aligned} E_{inc} &= \sum_{k=1}^{K} \left| \sum_{n=1}^{N} {A_n^{(k)}}^2 \exp -i\phi_c \right|^2 \\ &= \sum_{k=1}^{K} \left| \sum_{n=1}^{N} {A_n^{(k)}}^2 \right|^2 \end{aligned}$$

$$\begin{aligned} E_{co} &= \left| \sum_{k=1}^{K} \sum_{n=1}^{N} {A_n^{(k)}}^2 \exp -i\phi_c \right|^2 \\ &= \left| \sum_{k=1}^{K} \sum_{n=1}^{N} {A_n^{(k)}}^2 \right|^2 \end{aligned}$$

Any phase difference between the signal components is cancelled by the replica, allowing signal components to combine in phase. The residual phase term $\exp -i\phi_c$ is removed by the modulus.

When the replica is not properly matched to the noiseless signal, i.e.

$$P_n^{(k)} = \bar{A}_n^{(k)} \exp i2\pi f_s \bar{t}_n^{(k)\prime}$$

and the cost function values become

$$E_{inc} = \sum_{k=1}^{K} \left| \sum_{n=1}^{N} A_n^{(k)} \bar{A}_n^{(k)} \exp i2\pi f_s \left(\bar{t}_n^{(k)} - t_n^{(k)} \right) \right|^2$$

$$E_{co} = \left| \sum_{k=1}^{K} \sum_{n=1}^{N} A_n^{(k)} \bar{A}_n^{(k)} \exp i2\pi f_s \left(\bar{t}_n^{(k)} - t_n^{(k)} \right) \right|^2$$

The terms of the innermost summations (i.e. $\sum_n$) will be out of phase due to residual phase terms $\exp i2\pi \left(\bar{t}_n^{(k)} - t_n^{(k)} \right)$. These terms will combine incoherently to form a complex value $\Sigma_n(k)$ with amplitude $B(k)$ and phase $\psi(k)$:

$$\Sigma_n(k) = B(k) \exp i\psi(k)$$

As this value was formed via incoherent combination,

$$B(k) \leq \sum_{n=1}^{N} {A_n^{(k)}}^2$$

and the cost values become

$$\begin{aligned} E_{inc} &= \sum_{k=1}^{K} |B(k) \exp i\psi(k)|^2 \\ &= \sum_{k=1}^{K} B(k)^2 \end{aligned}$$

$$E_{co} = \left| \sum_{k=1}^{K} B(k) \exp i\psi(k) \right|^2$$

The spatially incoherent matching function discards any residual phase when summing across receivers. However the spatially coherent function leaves phase intact, taking modulus only after the summation. The residual spatial phase information which

is present when the replica is mismatched will cause degradation of amplitude in the summation, since

$$\left|\sum_{k=1}^{K} B(k) \exp i\phi(k)\right| \leq \left|\sum_{k=1}^{K} B(k)\right|$$

The spatially coherent function is thus expected to have greater sidelobe suppression ability than the spatially coherent function. This may be a problem when attempting to estimate the parameters of a source when there are other sources present in the environment; particularly if these extra sources have a range history which is similar to that of the target of interest.

7.4 Simulations

7.4.1 Simulation Details

A simulation scenario of a source moving past the origin on a cartesian reference plane was considered. The source's initial position was at point (-200m,600m) with respect to the origin and it moved past the origin at a speed of 5 m/s at an angle of 2 degrees to the x–axis. It emitted a 27 Hz sinusoidal signal.

Initial simulations were performed assuming the sonobuoy field was a linear array of receivers, equally spaced along the x–axis and centred about the origin. Two array configurations were considered. The first of these was an array with a fixed aperture of 50 metres, but with a variable number (between 2 and 10) receivers. The second configuration was an array not of fixed aperture but with a fixed number of elements (4), having variable separation. Separations between 5 and 500 metres were considered. Signal–to–noise ratios between 0 dB and -50 dB were considered. 100 individual simulations were performed for each combination of SNR and array configuration.

The search space comprised 750 elements, having initial x and y positions within 100 metres of their true values, speeds between 4 and 6 m/s, and headings between 10 and -10 degrees from the x–axis. A plot of all candidate trajectories in the search space is given in figure 7.2, shown in relation to a 6–element array. It was necessary

to choose a restricted search domain in order to keep the computational load of the simulations at a tolerable level. Thus a small neighbourhood about the known position of the simulated source in the parameter space was selected. In reality the search domain may be constrained by the use of a–priori information. For example, the approximate positions of the source could be determined by beamforming and a search domain defined about this neighbourhood.

The simulated multiple receiver outputs were processed with both the spatially–incoherent and spatially–coherent matched processing algorithms described previously. Results for the spatially–incoherent technique are presented in tables 7.1 – 7.4. Results for the spatially–coherent technique were almost identical to these and thus have not been presented here. The results of these simulations are discussed in §7.4.2.

Further simulation was performed using a scenario designed to highlight the difference between the spatially coherent and spatially incoherent processing methods. The receiver array comprised 4 elements arranged in the shape of a 5–by–5 metre square. A non–linear shape was chosen in order that the array should posess better focussing properties; a linear array exhibits inherent ambiguity in that it is impossible to differentiate between sources from either side of the array. The small receiver separation was chosen in order to limit the variability across the receivers and thus enhance the effect of sidelobes.

Two sources were simulated. The first source was initially at a position of (-300,600) metres relative to the array centre and moved parallel to the x–axis at a speed of 5 metres per second. The second source was a 90–degree rotation of the first source: It moved from (600,300) with a speed of 5 metres per second, parallel to the y–axis. Both sources emitted a 50 Hz sinusoidal signal.

SNRs between -40 and -25 dB were considered for both sources. 100 simulations were performed for each combination of SNRs. 2 search domains were employed: The first similar to the domain depicted in figure 7.2 except that headings between -8 and 8 degrees were considered. The second search domain used was twice the size of the first, comprising this domain plus a 90–degree rotation of it, thus encompassing the second source also. This is depicted in figure 7.3. Results of the simulations performed in this

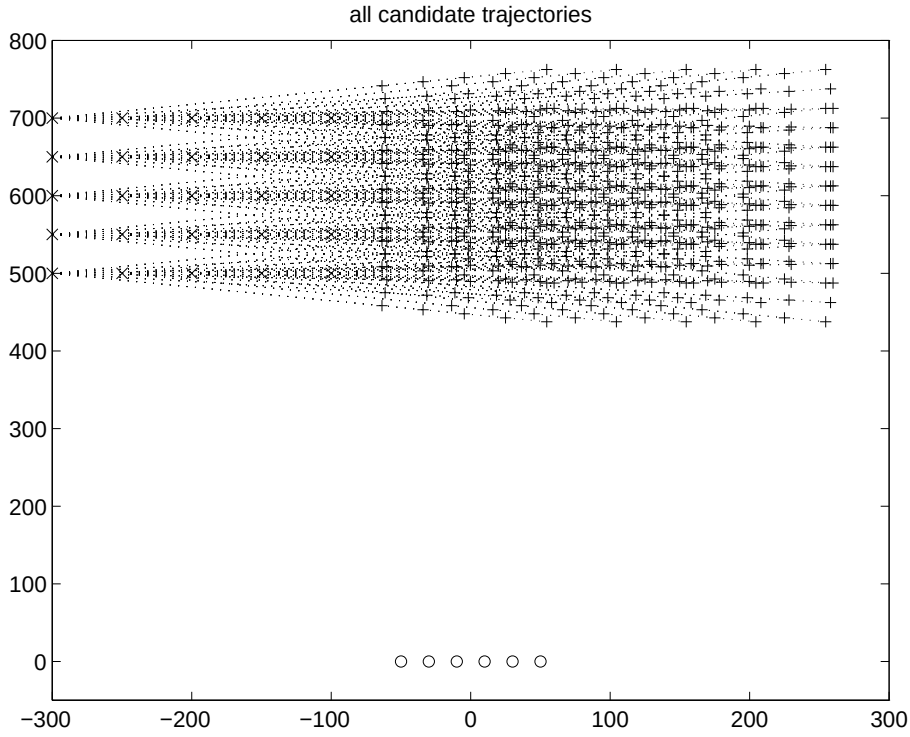

Figure 7.2: All Candidate Trajectories for Simulated Data Analysis

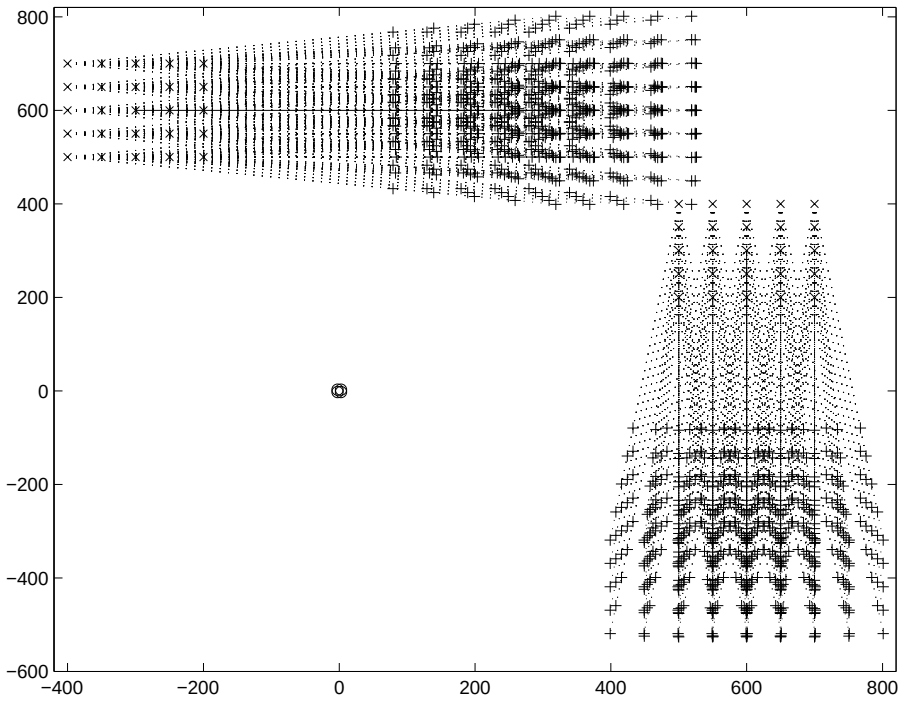

Figure 7.3: All Candidate Trajectories for Simulated Data Analysis with Two Sources

two–source scenario are presented in tables 7.5 – 7.8.

Table 7.1: Bias, Varying Number Of Receivers

SNR	# rec	Initial Position		Speed	Heading
-50	2	6.5	-11.5	0.1	1.2
-45	2	8.5	-16.5	0.0	1.2
-40	2	9.0	-21.0	-0.1	0.4
-35	2	5.0	-12.5	-0.1	-0.6
-30	2	3.0	-6.0	-0.0	-0.2
-25	2	1.0	-2.0	-0.0	-0.1
-20	2	0.0	0.0	0.0	0.0
-50	4	2.5	-14.5	0.1	2.5
-45	4	3.5	-25.5	-0.1	1.3
-40	4	3.5	-21.5	-0.2	0.0
-35	4	9.5	-18.0	-0.1	-0.8
-30	4	6.5	-13.0	-0.1	-0.5
-25	4	0.0	0.0	0.0	0.0
-20	4	0.0	0.0	0.0	0.0
-50	6	0.5	-11.5	0.1	2.1
-45	6	-1.0	-21.5	-0.2	0.5
-40	6	10.0	-21.5	-0.1	-0.5
-35	6	7.5	-15.0	-0.1	-0.6
-30	6	1.0	-2.0	-0.0	-0.1
-25	6	0.0	0.0	0.0	0.0
-20	6	0.0	0.0	0.0	0.0
-50	8	-3.5	-23.5	0.1	1.4
-45	8	1.5	-18.0	-0.1	0.6
-40	8	9.5	-21.5	-0.1	-0.7
-35	8	5.5	-11.0	-0.1	-0.4
-30	8	1.5	-3.0	-0.0	-0.1
-25	8	0.0	0.0	0.0	0.0
-20	8	0.0	0.0	0.0	0.0
-50	10	-20.0	-14.0	-0.1	1.9
-45	10	-4.0	-16.5	-0.1	0.5
-40	10	11.0	-20.0	-0.1	-0.8
-35	10	3.5	-7.0	-0.0	-0.3
-30	10	0.0	0.0	0.0	0.0
-25	10	0.0	0.0	0.0	0.0
-20	10	0.0	0.0	0.0	0.0

Table 7.2: Standard Deviation, Varying Number Of Receivers

SNR	# rec	Initial Position		Speed	Heading
-50	2	74.1	74.5	0.7	6.4
-45	2	72.8	63.6	0.7	6.2
-40	2	50.9	62.8	0.5	5.0
-35	2	29.7	46.8	0.3	2.1
-30	2	11.9	23.9	0.1	1.0
-25	2	7.0	14.1	0.1	0.6
-20	2	0.0	0.0	0.0	0.0
-50	4	80.2	67.5	0.8	7.1
-45	4	67.1	68.7	0.7	5.7
-40	4	45.6	60.0	0.4	4.1
-35	4	23.2	41.1	0.2	1.8
-30	4	16.9	33.8	0.2	1.4
-25	4	0.0	0.0	0.0	0.0
-20	4	0.0	0.0	0.0	0.0
-50	6	70.5	73.1	0.7	6.8
-45	6	60.7	66.0	0.6	5.5
-40	6	36.9	54.2	0.3	3.0
-35	6	17.9	35.9	0.2	1.4
-30	6	7.0	14.1	0.1	0.6
-25	6	0.0	0.0	0.0	0.0
-20	6	0.0	0.0	0.0	0.0
-50	8	70.4	65.3	0.7	6.3
-45	8	52.4	56.2	0.5	4.8
-40	8	27.2	43.4	0.3	1.9
-35	8	15.7	31.4	0.2	1.3
-30	8	8.6	17.1	0.1	0.7
-25	8	0.0	0.0	0.0	0.0
-20	8	0.0	0.0	0.0	0.0
-50	10	71.1	75.2	0.7	6.8
-45	10	49.6	56.4	0.4	4.5
-40	10	22.0	42.0	0.2	1.8
-35	10	12.8	25.6	0.1	1.0
-30	10	0.0	0.0	0.0	0.0
-25	10	0.0	0.0	0.0	0.0
-20	10	0.0	0.0	0.0	0.0

Table 7.3: Bias, Varying Hydrophone Separation

SNR	separation	Initial Position		Speed	Heading
-50	5	-4.5	-8.5	0.1	1.3
-45	5	3.0	-20.5	-0.0	1.6
-40	5	9.5	-27.0	-0.2	-0.1
-35	5	16.0	-33.0	-0.2	-1.3
-30	5	8.0	-16.0	-0.1	-0.6
-25	5	1.5	-3.0	-0.0	-0.1
-20	5	0.0	0.0	0.0	0.0
-50	10	-8.0	-18.5	0.1	1.2
-45	10	-3.5	-18.5	-0.0	2.0
-40	10	5.0	-22.0	-0.1	-0.2
-35	10	14.0	-27.0	-0.1	-1.1
-30	10	4.5	-9.0	-0.0	-0.4
-25	10	1.0	-2.0	-0.0	-0.1
-20	10	0.0	0.0	0.0	0.0
-50	50	1.5	-8.5	0.1	1.6
-45	50	-3.0	-15.0	-0.1	-0.6
-40	50	2.5	-9.5	-0.1	-0.1
-35	50	2.0	-3.0	-0.0	-0.1
-30	50	0.0	0.0	0.0	0.0
-25	50	0.0	0.0	0.0	0.0
-20	50	0.0	0.0	0.0	0.0
-50	100	1.5	5.5	-0.1	2.5
-45	100	7.5	-3.0	0.1	0.8
-40	100	1.0	5.5	-0.0	0.1
-35	100	0.0	0.0	0.0	0.0
-30	100	0.0	0.0	0.0	0.0
-25	100	0.0	0.0	0.0	0.0
-20	100	0.0	0.0	0.0	0.0
-50	500	1.0	0.5	0.0	1.6
-45	500	-3.5	7.5	0.1	1.6
-40	500	2.5	8.0	0.1	-0.2
-35	500	0.0	0.0	0.0	0.0
-30	500	0.0	0.0	0.0	0.0
-25	500	0.0	0.0	0.0	0.0
-20	500	0.0	0.0	0.0	0.0

Table 7.4: Standard Deviation, Varying Hydrophone Separation

SNR	separation	Initial Position		Speed	Heading
-50	5	72.2	72.5	0.8	6.9
-45	5	69.2	69.7	0.6	6.0
-40	5	41.9	51.9	0.4	4.3
-35	5	24.5	47.3	0.2	1.9
-30	5	18.4	36.8	0.2	1.5
-25	5	8.6	17.1	0.1	0.7
-20	5	0.0	0.0	0.0	0.0
-50	10	76.1	68.0	0.8	6.7
-45	10	63.7	68.0	0.7	6.7
-40	10	46.9	55.7	0.4	3.7
-35	10	22.6	44.6	0.2	1.8
-30	10	14.4	28.8	0.1	1.2
-25	10	7.0	14.1	0.1	0.6
-20	10	0.0	0.0	0.0	0.0
-50	50	66.5	72.8	0.7	6.4
-45	50	61.5	62.2	0.6	5.0
-40	50	35.8	45.3	0.5	2.9
-35	50	9.8	17.1	0.1	0.7
-30	50	0.0	0.0	0.0	0.0
-25	50	0.0	0.0	0.0	0.0
-20	50	0.0	0.0	0.0	0.0
-50	100	65.3	70.7	0.7	6.2
-45	100	57.5	57.7	0.7	5.2
-40	100	23.5	31.7	0.3	1.7
-35	100	0.0	0.0	0.0	0.0
-30	100	0.0	0.0	0.0	0.0
-25	100	0.0	0.0	0.0	0.0
-20	100	0.0	0.0	0.0	0.0
-50	500	70.7	65.3	0.7	6.7
-45	500	64.1	57.0	0.6	6.6
-40	500	33.6	38.7	0.3	3.1
-35	500	0.0	0.0	0.0	0.0
-30	500	0.0	0.0	0.0	0.0
-25	500	0.0	0.0	0.0	0.0
-20	500	0.0	0.0	0.0	0.0

Table 7.5: 2 Sources In Domain, Spatially–Coherent cost function
Estimates of parameters of first source:

Bias

SNR 1	SNR 2	Initial Position		Speed	Heading
-40	-40	460.5	-163.0	-0.1	-45.9
-35	-40	77.0	-42.5	-0.1	-7.1
-30	-40	15.0	-15.0	-0.1	-1.2
-25	-40	0.0	0.0	0.0	0.0
-40	-35	799.5	-268.0	-0.0	-79.8
-35	-35	405.5	-146.0	-0.1	-40.2
-30	-35	35.0	-23.0	-0.1	-3.2
-25	-35	0.0	0.0	0.0	0.0
-40	-30	900.0	-300.0	0.0	-90.0
-35	-30	900.0	-300.0	0.0	-90.0
-30	-30	419.0	-149.0	-0.1	-41.6
-25	-30	4.0	-4.0	-0.0	-0.3
-40	-25	900.0	-300.0	0.0	-90.0
-35	-25	900.0	-300.0	0.0	-90.0
-30	-25	900.0	-300.0	0.0	-90.0
-25	-25	413.0	-143.0	-0.0	-41.1

Standard Deviation

SNR 1	SNR 2	Initial Position		Speed	Heading
-40	-40	430.7	143.5	0.2	43.7
-35	-40	195.1	76.0	0.2	19.5
-30	-40	35.9	35.9	0.2	2.9
-25	-40	0.0	0.0	0.0	0.0
-40	-35	276.0	91.7	0.1	27.9
-35	-35	433.1	139.4	0.2	43.6
-30	-35	129.8	54.8	0.2	12.8
-25	-35	0.0	0.0	0.0	0.0
-40	-30	0.0	0.0	0.0	0.0
-35	-30	0.0	0.0	0.0	0.0
-30	-30	438.5	141.1	0.2	44.1
-25	-30	19.7	19.7	0.1	1.6
-40	-25	0.0	0.0	0.0	0.0
-35	-25	0.0	0.0	0.0	0.0
-30	-25	0.0	0.0	0.0	0.0
-25	-25	443.5	145.1	0.1	44.5

Table 7.6: 2 Sources, Only 1 In Domain, Spatially–Incoherent cost function
Estimates of parameters of first source:

Bias

SNR 1	SNR 2	Initial Position		Speed	Heading
-40	-40	399.0	-147.5	-0.1	-40.2
-35	-40	363.5	-131.5	-0.1	-36.9
-30	-40	156.0	-54.0	-0.0	-15.5
-25	-40	27.0	-9.0	0.0	-2.7
-40	-35	523.5	-179.5	-0.0	-52.7
-35	-35	451.0	-154.5	-0.0	-45.4
-30	-35	253.0	-85.0	-0.0	-25.3
-25	-35	27.0	-9.0	0.0	-2.7
-40	-30	783.0	-261.0	0.0	-78.3
-35	-30	639.0	-213.0	0.0	-63.9
-30	-30	441.0	-147.0	0.0	-44.1
-25	-30	108.0	-36.0	0.0	-10.8
-40	-25	882.0	-294.0	0.0	-88.2
-35	-25	882.0	-294.0	0.0	-88.2
-30	-25	792.0	-264.0	0.0	-79.2
-25	-25	522.0	-174.0	0.0	-52.2

Standard Deviation

SNR 1	SNR 2	Initial Position		Speed	Heading
-40	-40	437.5	150.8	0.2	44.6
-35	-40	434.3	149.9	0.2	44.3
-30	-40	338.8	113.2	0.1	33.9
-25	-40	154.3	51.4	0.0	15.4
-40	-35	444.4	152.3	0.1	44.8
-35	-35	447.8	151.9	0.1	45.2
-30	-35	405.6	135.1	0.0	40.6
-25	-35	154.3	51.4	0.0	15.4
-40	-30	304.2	101.4	0.0	30.4
-35	-30	410.4	136.8	0.0	41.0
-30	-30	452.2	150.7	0.0	45.2
-25	-30	293.9	98.0	0.0	29.4
-40	-25	126.6	42.2	0.0	12.7
-35	-25	126.6	42.2	0.0	12.7
-30	-25	293.9	98.0	0.0	29.4
-25	-25	446.4	148.8	0.0	44.6

Table 7.7: 2 Sources In Domain, Spatially–Incoherent cost function
Estimates of parameters of first source:
Bias

SNR 1	SNR 2	Initial Position		Speed	Heading
-40	-40	44.0	-32.5	-0.2	-3.5
-35	-40	36.0	-31.0	-0.2	-2.9
-30	-40	15.0	-15.0	-0.1	-1.2
-25	-40	0.0	0.0	0.0	0.0
-40	-35	60.5	-49.0	-0.3	-4.8
-35	-35	34.5	-33.0	-0.2	-2.8
-30	-35	18.0	-18.0	-0.1	-1.4
-25	-35	0.0	0.0	0.0	0.0
-40	-30	75.0	-73.5	-0.4	-6.0
-35	-30	46.0	-46.0	-0.2	-3.7
-30	-30	34.0	-34.0	-0.2	-2.7
-25	-30	4.0	-4.0	-0.0	-0.3
-40	-25	95.0	-95.0	-0.5	-7.6
-35	-25	84.0	-84.0	-0.4	-6.7
-30	-25	71.0	-71.0	-0.4	-5.7
-25	-25	16.0	-16.0	-0.1	-1.3

Standard Deviation

SNR 1	SNR 2	Initial Position		Speed	Heading
-40	-40	49.4	54.8	0.2	4.0
-35	-40	47.2	49.1	0.2	3.8
-30	-40	35.9	35.9	0.2	2.9
-25	-40	0.0	0.0	0.0	0.0
-40	-35	48.4	52.2	0.2	3.9
-35	-35	47.5	48.8	0.2	3.8
-30	-35	38.6	38.6	0.2	3.1
-25	-35	0.0	0.0	0.0	0.0
-40	-30	43.5	43.5	0.2	3.5
-35	-30	50.1	50.1	0.3	4.0
-30	-30	47.6	47.6	0.2	3.8
-25	-30	19.7	19.7	0.1	1.6
-40	-25	21.9	21.9	0.1	1.8
-35	-25	36.8	36.8	0.2	2.9
-30	-25	45.6	45.6	0.2	3.6
-25	-25	36.8	36.8	0.2	2.9

Table 7.8: 2 Sources, Only 1 In Domain, Spatially–Coherent cost function
Estimates of parameters of first source:

Bias

SNR 1	SNR 2	Initial Position		Speed	Heading
-40	-40	18.5	-13.0	-0.1	-1.5
-35	-40	17.0	-14.5	-0.1	-1.4
-30	-40	5.0	-5.5	-0.0	-0.4
-25	-40	0.0	0.0	0.0	0.0
-40	-35	8.0	-4.5	-0.0	-0.6
-35	-35	7.0	-7.0	-0.0	-0.6
-30	-35	1.0	-1.0	-0.0	-0.1
-25	-35	0.0	0.0	0.0	0.0
-40	-30	3.0	-3.0	-0.0	-0.2
-35	-30	1.0	-1.0	-0.0	-0.1
-30	-30	1.0	-1.0	-0.0	-0.1
-25	-30	0.0	0.0	0.0	0.0
-40	-25	0.0	0.0	0.0	0.0
-35	-25	0.0	0.0	0.0	0.0
-30	-25	0.0	0.0	0.0	0.0
-25	-25	0.0	0.0	0.0	0.0

Standard Deviation

SNR 1	SNR 2	Initial Position		Speed	Heading
-40	-40	40.0	41.8	0.2	3.2
-35	-40	37.1	37.8	0.2	3.0
-30	-40	21.9	22.4	0.1	1.8
-25	-40	0.0	0.0	0.0	0.0
-40	-35	27.3	27.6	0.1	2.2
-35	-35	25.6	26.6	0.1	2.1
-30	-35	10.0	10.0	0.0	0.8
-25	-35	0.0	0.0	0.0	0.0
-40	-30	17.1	17.1	0.1	1.4
-35	-30	10.0	10.0	0.0	0.8
-30	-30	10.0	10.0	0.0	0.8
-25	-30	0.0	0.0	0.0	0.0
-40	-25	0.0	0.0	0.0	0.0
-35	-25	0.0	0.0	0.0	0.0
-30	-25	0.0	0.0	0.0	0.0
-25	-25	0.0	0.0	0.0	0.0

7.4.2 Discussion

It is clear from the single–source simulations that estimation of a moving source's position, speed and heading can be achieved with the use of multiple receivers. Invariant unbiased estimation of these parameters is achieved at -20 dB using a 2–element array with an aperture of 50 metres. As more receivers are added to the fixed–aperture array the SNR at which invariant unbiased estimation can be achieved drops, as shown by tables 7.1 and 7.2. An SNR as low as -30 dB can be successfully processed when the array has 10 elements.

It also appears that the aperture of the receiver array has an effect on the technique's performance (tables 7.3 and 7.4). Four receivers separated by 5 metres attain invariant unbiased estimation of all parameters at -20 dB. However as the receiver separation is increased, the SNR at which invariant performance is achieved decreases. A separation of 100 metres enables accurate estimation at -35 dB. This effect occurs no doubt due to the fact that a wider aperture enhances variability across the array. Several receivers close together will receive much the same signal. The doppler frequency will be nearly identical at each receiver. However receivers having large separation may experience far different signals, possibly experiencing both extremes of the doppler shift at the same time.

However it is worth noting that there does appear to be a maximum limit on aperture beyond which no performance gain will be made. A separation of 500 metres conferred no greater performance than a separation of 100 metres. In fact the 500 metre separation produced more variance in estimates at -40 dB than did the 100 metre separation. A possible explanation for this is signal attenuation: Hydrophones separated by 5 metres receive roughly the same signal power, but hydrophones seperated by 500 metres will not all experience a strong signal. The SNR is defined at the centre of the array; receivers which are at a distance from the array centre will experience a signficantly lower SNR.

Consider the results of the two–source simulations, using a large search domain which encompasses both sources (table 7.5). The spatially–coherent processor consistently

estimated the position, speed and heading of the first source when it's SNR was -25 dB and the SNR of the second source was -35 or -40 dB. It exhibited a very small amount of variance when second source's SNR was -30 dB, or the first source had an SNR of -30 dB and the second -40 dB. This demonstrates that when the first source is of reasonable strength and is stronger than the second, the spatially–coherent processor can determine the source's motion parameters. This is further supported by the observation that when the second source's SNR is stronger than that of the first, the spatially–coherent processor consistently picks an initial position with a bias of (900,-300) metres and a heading bias of 270 degrees. In this case the processor has estimated the parameters of the second source, because it is the stronger. When both sources have the same SNR the processor exhibits large variance in its estimates. This is because the processor cannot distinguish between the two sources and will select one or the other depending on random effects like noise. This is further evident when examining the biases when the SNRs are identical: The biases are approximately halfway between the values of the first source and those of the second.

When using the spatially–incoherent processor (table 7.6), estimator variance is again highest when the SNRs of the two sources are equal. Once again, the estimator cannot distinguish between the sources. Sources of differing SNR produce less variance than sources with the same SNR, as with the spatially–coherent processor, but the results are biased in favour of the stronger source. This bias is greater when the SNR difference is larger. This indicates that the spatially–incoherent processor requires large difference in SNR in order to robustly estimate the motion parameters of the stronger source. There is some small variance on speed estimates when both SNRs are very low, as is to be expected at low SNR, but otherwise speed estimation is performed consistently. This is no doubt because both sources have the same speed of 5 metres per second and is further evidence that the spatially–incoherent processor will correctly estimate the parameters of one of the sources, provided that the SNRs are not both very low.

Consider now the same two–source scenario but employing a search domain which encompasses only the first source (table 7.7). The results using the spatially–coherent

processor are comparable to those of the full–domain simulations, when the first source is of greater strength than the second: The parameters of the first source are accurately estimated when the SNR is -25 dB and the second source SNR is -35 or -40 dB and there is a very small amount of variance in other combinations for which the first source is dominant. However performance is poor when the second source is dominant.

At first glance it appears that the spatially–incoherent processor (table 7.8) has performed better than the coherent version on this restricted domain problem. Performance is similarly good when the main source is of sufficiently high SNR; estimates are consistently accurate. However it appears that estimates of the main source's parameters are also perfectly estimated when it's SNR is as low as -40, if the SNR of the second source is -25 dB. This is obviously not the case; it has been demonstrated that -40 dB is too weak to estimate signal motion parameters in previous simulations. This behaviour occurs because the second source is quite similar to the first source, in that they have the same speed and very similar range histories. The cost values for these sources will thus be quite similar; the second source is at a sidelobe of the first and leakage from this sidelobe has occurred.

The spatially–coherent processor has not been "fooled" by a stronger second source, because the relative phase information between the receivers will force signals arriving from this second source to combine incoherently and be degraded. The extra coherent processing has provided greater sidelobe suppression.

In conclusion, the benefit of using the spatially–coherent over the spatially–incoherent processor is obvious. The spatially–coherent processor offers far greater robustness to relative signal strength than the spatially–incoherent version. Furthermore, the spatially–incoherent processor is much less robust to interference from strong sources outside the search domain. The spatially–coherent version provides better suppression of interference from sidelobes.

7.4.3 Summary and Extensions

The results of simulations suggest that the matched doppler processing technique can localise a moving source in heading and bearing as well as range and speed when multiple receivers are used. The technique performs better at low SNR when more receivers are used. Better performance can also be achieved when there is significant separation between the receivers, although there appears to be a limit on this separation.

The matched doppler processing technique is capable of determining the motion parameters of the stronger source when there are two similar sources in the environment. A spatially–coherent processor is better able to distinguish between two similar sources when the difference in signal power is small. A spatially–incoherent processor exhibits sidelobe leakage which can be a problem when one source is not in the search space.

The simulations involving two sources suggest that when the technique does err, it selects the weaker source rather than some arbitrary point on the domain. This suggests that there are two "peaks" in the ambiguity surface; one peak corresponding to each source. If one of the sources is stronger, then it's respective peak will be greater and thus will be picked in preference to the weaker source. When the sources are of similar power, the selection is dependent on the random effects of noise.

The fact that the ambiguity surface is bimodal when two sources are present suggests that the technique could be employed for tracking multiple sources. Instead of finding the maximum correlation across the entire search domain, the optimisation process must instead find local maxima on the search domain and ascribe each of these to a source. Sidelobe suppression will be important in order that sidelobes are not detected as local maxima and reported as true sources. This suggests that only a spatially–coherent processor should be used for computing cost values. It is likely that thresholding of the cost values will be performed in order to detect true sources in a similar fashion to the track–before–detect algorithm suggested by Chan et al [16].

Chapter 8

Conclusion

8.1 Summary

This thesis has presented a method for estimation of motion parameters of a nearfield narrowband source in the underwater acoustic environment.

Initially the specific problem of estimating the ranges of a source undergoing straight–line motion past a single receiver was addressed. A simple signal model involving doppler–affected signal phase was presented. This model was used for generating signal replicas for matching via a bank of cost functions: One using only amplitude information, a similar form which also exploited signal phase, one having the form of the conventional beamformer, and another which used only raw phase. The similarity of the amplitude and phase cost function to the conventional beamformer cost function was investigated. Extensive simulation using these cost functions determined that the amplitude and phase and conventional beamformer cost functions were the best performed of the cost functions, consistently estimating signal ranges at low SNR and outperforming an existing doppler–tracking technique for tracking a high–frequency signal.

Bounds on the performance of the matched doppler processing algorithm were derived. The Cramer–Rao Lower Bound on estimator variance for the particular scenario under analysis was computed. Further bounds on variance imposed by the discrete na-

ture of the search domain were determined. The variance on estimates observed in the simulations of the preceding chapter were compared to these bounds. It was found that the amplitude–and–phase estimates adhered to the Cramer–Rao Lower Bound.

Chapter five examined the effects of track variation on the performance of the matched doppler processing algorithm. A model for track wobble was proposed and used to generate simulated data which was analysed with the amplitude–and–phase cost function. It was found that the technique was robust to a certain amount of deviation in the source locus and that the amount of variance required to fault the technique was too large to realistically occur.

The technique was applied to real data from a sonobuoy which was sampling a signal from a nearfield source. Reasonable estimates of the craft's speed and CPA range were obtained with the matched doppler processing technique when analysing low–frequency bins. However the results obtained from the high–frequency bin were not realistic. This was explained in terms of greater phase variations at higher frequency and evidence of a "noise spike" in the data.

An extension of the single–receiver technique for exploiting multiple receivers was presented. Two processing functions were proposed: One spatially–incoherent funtion which was essentially a summation of individual CBF cost function values from each receiver, and a spatially–coherent function which processed phase coherently across receivers as well as time. Analysis of simulated data revealed that both techniques could accurately localise a moving source in range, speed, bearing and heading. Performance improved as more receivers were used, or as the receiver field's aperture increased. Analysis of simulated data comprising two similar sources demonstrated that the spatially–coherent technique was better able to determine the parameters of the stronger source and that the spatially–incoherent function exhibited sidelobe effects.

8.2 Extensions

This study has demonstrated that matched processing is a viable algorithm for estimation of moving source parameters. However it has only scratched the surface.

This study has avoided use of complex propagation models. In this study a simple straight–line propagation model was used. This was sufficient for a nearfield source. However it is possible that consideration of a few other additional modes of propagation, even simple ones such as surface and bottom bounce, may improve performance of the technique.

One potential extension of this work would be to implement the matched doppler processing algorithm on a parallel computer architecture. The bulk of the time spent by this algorithm is in the serial processing of the replica. Much time could be saved if the replica could be matched simultaneously, as may be possible on a parallel machine.

The technique presented here is quite limited in that the source locus is assumed to be constant. In reality a source of interest will not necessarily adhere to a fixed trajectory. For example, an ocean vessel may undergo a change in course during the sampling period. The behaviour of the technique under these conditions has not been investigated. Presumably the cost value of the "correct" replica would be fine up to the point in time where the craft diverges from it's original trajectory, but as further time samples are used the cost value would degrade. It may be necessary to devise a dynamic programming routine which controls the matching process. It would detect when a cost value begins to degrade and would re–initialise the search at this point in time, possibly also redefining the search domain to concentrate on small deviations from the original estimated trajectory.

The robustness of the technique to small amounts of variance in the source's straight–line locus was determined in chapter 5. Another issue which has not been addressed is that of robustness to variance in source frequency. The real data analysed in this study was produced by a ship towing a projector emitting known fixed frequencies and thus variance was not an issue. However in most practical applications of the technique the

signals emitted by a source of opportunity will vary, for example as the craft changes gear. Robustness to this kind of signal perturbation must be investigated.

One important area which has not been addressed by this study is that of optimisation of the cost function. In all simulations and analyses the optimal combination of parameters was determined by an exhaustive search of a finite parameter space. This has proven to be quite expensive in terms of computational overhead. Unless the sampling interval and the size of the parameter space are kept small the demand of the matching process becomes prohibitive. In a real application it would be desirable to have no constraints on the search space, partly so that a large domain may be searched, and partly to reduce quantisation noise by sampling the domain as finely as is required. It would also be desirable to sample for an arbitrarily long interval, as it has been demonstrated that greater robustness can be achieved with longer sampling intervals. Furthermore it would be desirable to track a source for as long as it's signal is being received. For these reasons it will be necessary to dispense with exhaustive search as an optimisation procedure. Some of the cost functions exhibit a reasonably smooth, well–behaved ambiguity surface (e.g. CBF). It is hypothesised that such cost functions may be adequately processed with a numerical gradient–descent technique. Such numerical methods have an added bonus in that they can determine optimal points on a function to a desired level of accuracy, thereby avoiding quantisation noise, but they may be susceptible to sidelobes and local optima. Recent research in seismic signal processing (for example, [14, 15]) has reported some success using a domain search based on Genetic Algorithms. This is essentially a directed random search of a discrete domain. It does not usually succumb to local optima and can function even when the search terrain is rough. Only a fraction of elements of the search domain are evaluated, thereby avoiding the massive processing required by an exhaustive search. However a Genetic Algorithm estimate will still suffer the effects of quantisation noise.

Appendix A

Derivation of Cramer–Rao Lower Bounds

Peleg & Porat [12] present a method for computing the Cramer-Rao bounds on the amplitude and phase terms of a signal with constant but unknown amplitude and phase which is polynomial of an arbitrary order. When using the phase–only cost function, amplitude information is not used and may be treated as unity. Over the short sampling durations used in the simulations, amplitude varies very little and thus CRLBs for constant–amplitude signals are a good approximation to CRLBs for the simulated signals in this study. The signal is approximately a linear chirp

$$s(t) = \exp i2\pi \left(\theta_0 + \theta_1 t + \theta_2 t^2\right)$$

and thus phase is quadratic.

It can be shown that the entries of the Fisher Information Matrix for the three polynomial phase coefficients $\underline{\theta} = [\theta_0\, \theta_1\, \theta_2]$ is

$$J_{\theta_{k,l}} = \frac{8\pi^2}{\sigma_n^2} T^{k+l-2} S_{k+l-2}$$

where

$$S_m = \sum_{n=1}^{N} n^m$$

and σ_n^2 is noise variance, N is the number of timeblocks in the signal and the timeblock length is one second. The Fisher Information Matrix J_θ for the case where $T = 1$ is thus the hermitian matrix

$$\begin{aligned} J_\theta &= \frac{8\pi^2}{\sigma_n^2}\begin{bmatrix} S_0 & S_1 & S_2 \\ S_1 & S_2 & S_3 \\ S_2 & S_3 & S_4 \end{bmatrix} \\ &= \frac{8\pi^2}{\sigma_n^2}\begin{bmatrix} N & \frac{N(N+1)}{2} & \frac{N(N+1)(2N+1)}{6} \\ \frac{N(N+1)}{2} & \frac{N(N+1)(2N+1)}{6} & \frac{N^2(N+1)^2}{4} \\ \frac{N(N+1)(2N+1)}{6} & \frac{N^2(N+1)^2}{4} & \frac{N(N+1)(2N+1)(3N^2+3N-1)}{30} \end{bmatrix} \end{aligned}$$

It follows that the inverse of the Fisher Information Matrix for $\underline{\theta}$ is

$$J_\theta^{-1} = C\begin{bmatrix} 3N^2+3N+2 & -12N-6 & 10 \\ -12N-6 & \frac{4(2N+1)(8N+11)}{(N+1)(N+2)} & \frac{-60}{N+2} \\ 10 & \frac{-60}{N+2} & \frac{60}{(N+1)(N+2)} \end{bmatrix}$$

where

$$C = \frac{3\sigma_n^2}{8\pi^2 N(N-1)(N-2)}$$

The Cramer-Rao bounds on the polynomial phase coefficients $\underline{\theta}$ are

$$\sigma_{\underline{\theta}}^2 = \text{diag}\left(J_\theta^{-1}\right)$$

Target parameters r_{cpa}, v, and the timeblock at which CPA occurs t_{cpa} can be expressed as functions of $\underline{\theta} = [\theta_0, \theta_1, \theta_2]$ as follows:

$$\begin{aligned} t_{cpa} &= \frac{f-\theta_1}{2\theta_2} \\ r_{cpa} &= \frac{c}{f}\left(\frac{\theta_1^2 - f^2}{4\theta_2} - \theta_0 + \phi_0\right) \\ v &= \sqrt{\frac{-c^2}{2f^2}\left[\theta_1^2 - 4\theta_2\left(\theta_0 - \phi_0\right)\right] + \frac{c^2}{2}} \end{aligned}$$

The inverse of the Fisher information matrix for these parameters, J_{trv}^{-1}, may be expressed as a function of $J_{\underline{\theta}}^{-1}$

$$J_{trv}^{-1} = D J_{\underline{\theta}}^{-1} D^T$$

where D is the Jacobian

$$\begin{aligned} D &= \begin{bmatrix} \frac{\partial t_{cpa}}{\partial \theta_0} & \frac{\partial t_{cpa}}{\partial \theta_1} & \frac{\partial t_{cpa}}{\partial \theta_2} \\ \frac{\partial r_{cpa}}{\partial \theta_0} & \frac{\partial r_{cpa}}{\partial \theta_1} & \frac{\partial r_{cpa}}{\partial \theta_2} \\ \frac{\partial v}{\partial \theta_0} & \frac{\partial v}{\partial \theta_1} & \frac{\partial v}{\partial \theta_2} \end{bmatrix} \\ &= \begin{bmatrix} 0 & \frac{-1}{2\theta_2} & \frac{\theta_1 - f}{2\theta_2^2} \\ \frac{-c}{f} & \frac{c\theta_1}{2f\theta_2} & \frac{c(f^2-\theta_1^2)}{4f\theta_2^2} \\ \frac{\theta_2 c^2}{vf^2} & \frac{-\theta_1 c^2}{2vf^2} & \frac{\theta_0 c^2}{vf^2} \end{bmatrix} \end{aligned}$$

The Cramer-Rao bound on the parameters t_{cpa}, r_{cpa} and v are the diagonal values of the matrix J_{trv}^{-1}; $\sigma_{r_{cpa}}^2$ is the bound of particular interest to us.

Initial and final range are formulated in terms of these 3 parameters as

$$\begin{aligned} r_i &= \sqrt{r_{cpa}^2 + v^2 \left(t_{cpa} - 1\right)^2} \\ r_f &= \sqrt{r_{cpa}^2 + v^2 \left(t_{cpa} - N\right)^2} \end{aligned}$$

The inverse of the Fisher matrix may be updated to account for these parameters by use of a second Jacobian D_2:

$$J_{r_{if}}^{-1} = D_2 J_{trv}^{-1} D_2^T$$

where

$$D_2 = \begin{bmatrix} \frac{\partial r_i}{\partial t_{cpa}} & \frac{\partial r_i}{\partial r_{cpa}} & \frac{\partial r_i}{\partial v} \\ \frac{\partial r_f}{\partial t_{cpa}} & \frac{\partial r_f}{\partial r_{cpa}} & \frac{\partial r_f}{\partial v} \end{bmatrix}$$

The diagonal of $J_{r_{if}}$ is the Cramer–Rao bounds on r_i and r_f. Similarly partial derivatives of relative initial and final range $r_i - r_{cpa}, r_f - r_{cpa}$ may be used to compute bounds on these values.

These values may be computed numerically for any given set of parameters to produce Cramer–Rao bounds on source motion parameters. It is worth noting that the bounds are a function of the parameters themselves. Thus targets with different tracks may be estimated with different levels of accuracy.

Bibliography

[1] A. Tolstoy; "Matched Field Processing", World Scientific 1993.

[2] M.D. Collins, L.T. Fialkowski, W.A. Kuperman, J.S. Perkins; "Environmental Source Tracking", J. Acoust. Soc. Am. **94**(6), pp. 3335–3341, 1993.

[3] B. G. Quinn; "Doppler Speed and Range Estimation using Frequency and Amplitude Estimates", J. Acoust. Soc. Am. **98**(5), pp. 2560–2566, November 1995.

[4] B. Ferguson & B. Quinn, "Application of the Short Time Fourier Transform and Wigner–Ville Distribution to the Acoustic Localisation of Aircraft", J. Acoust. Soc. Am. **96** 821–827 1994.

[5] C.A. Zala & J.M. Ozard; "Matched–Field Processing For A Moving Source", J. Acoust. Soc. Am. **92**(1), pp. 403–417, 1992.

[6] G.M. Frichter, C.L. Byrne, C. Feuillade; "Sector–Focussed Stability Methods For Robust Source Localisation", J. Acoust. Soc. Am. **88**(6), pp. 2843–2851, December 1990

[7] J.M. Ozard, G.H. Brooke & P. Brouwer; "Improving Performance For Matched Field Processing With A Minimum Variance Beamformer", J. Acoust. Soc. Am. **91**(1) pp. 141–150, January 1992.

[8] J.L. Krolik; "Matched–Field Minimum Variance Beamforming In A Random Ocean Channel", J. Acoust. Soc. Am. **92**(3) pp. 1408–1419, 1992.

[9] C.A. Zala & J.M. Ozard; "Matched–Field Processing In A Range Dependent Environment", J. Acoust. Soc. Am. **88**(2), pp. 1011–1019, August 1990.

[10] M.D. Collins, L.T. Fialkowski, W.A. Kuperman, J.S. Perkins; "The Multivalued Bartlett Processor and Source Tracking", J. Acoust. Soc. Am. **97**(1), pp. 235–241, January 1995.

[11] A. Tolstoy; "Review Of Matched Field Processing For Environmental Inverse Problems", Int. J. of Modern Physics, **3**(4), pp. 691–708, 1992.

[12] S. Peleg & B. Porat; "The Cramer–Rao Lower Bound for Signals with Constant Amplitude and Polynomial Phase", IEEE Trans. Sig. Proc. **39** (3), pp. 749–752, March 1991.

[13] S.J. Searle & D.A. Gray; "Environmental Source Tracking and the Doppler Shift", presented at ISSPA-96, August 26–28, 1996.

[14] M.K. Sen, P.L.Stoffa; "Rapid Sampling of Model Space using Genetic Algorithms: Examples from Seismic Waveform Inversion", Geophys. J. Int. **108** pp. 281–292, 1992.

[15] M. Sambridge, G. Drijkoningen; "Genetic Waveforms in Seismic Waveform Inversion", Geophys. J. Int. **109** pp. 323–342, 1992.

[16] Y.T. Chan, G.H. Niezgoda & S.P. Morton; "Passive Sonar Detection and Localisation by Matched Velocity Filtering", IEEE J. Ocean. Eng. **20**(3), pp. 179–189, 1995.

[17] S.C. Nardone & M.L. Graham; "A Closed–form Solution to Bearings–only Target Motion Analysis", IEEE J. Ocean. Eng. **22**(1) pp. 168–178 1997.

[18] J.R. Moon & C.F. Stevens; "An Approximate Linearisation Approach to Bearings–Only Tracking", IEE Colloquium on Target Tracking and Data Fusion p88, 8/1-16 1996/253.

[19] P. Blanc–Benon & G. Bienvenu; "Passive Target Motion Analysis using Multipath Differential Time–Delay and Differential Doppler Shifts", Proc. ADFS–96 p228, 190–5 1996.

[20] P.A. Rosenqvist; "Passive Doppler–Bearing Tracking using a Pseudo–Linear Estimator", IEEE J. Ocean. Eng. **20**(2), pp. 114–118, 1995.

[21] Y–T Chan & F.L. Jardine; "Target Localisation and Tracking from Doppler Shift Measurements", IEEE J. Ocean. Eng. **15**(3), pp. 251–7 1990.

[22] H–Y Chen & I–T Lu; "Localization of a Broadband Source Using a Matched–Mode Procedure in the Time–Frequency Domain", IEEE J. Ocean. Eng. **19**(2), pp. 166–174, 1994.

[23] Y.T. Chan, J.J. Towers, "Sequential Localization of a Radiating Source by Doppler–Shifted Frequency Measurements", IEEE Trans. Aero. Elec. Sys. **28**(4) pp. 1084–1090, 1992.

[24] D.H. Johnson & D.E. Dudgeon, "Array Signal Processing – Concepts and Techniques", Prentice Hall 1993.

[25] W.B. Davenport & W.L. Root, "An Introduction to the Theory of Random Signals and Noise", IEEE, 1987.

[26] G.M. Jacyna, M.J. Jacobson & J.G. Clark, "General Treatment Of Source Motion On The Total Acoustic Field With Application To An Isospeed Channel", J. Acoust. Soc. Am. **60**(4), pp. 815–824, 1976.

[27] J.A. Neubert, "The Effect Of Doppler On Long–range Sound Propagation", J. Acoust. Soc. Am. **62**(6), pp. 1404–1411, 1977.

[28] K.E. Hawker, "A Normal Mode Theory Of Acoustic Doppler Effects In The Oceanic Waveguide", J. Acoust. Soc. Am. **65**(3), pp. 675–681, 1979.

[29] C.L. Byrne, R.T. Brent, C. Feuillade & D.R. DelBalzo, "A Stable Data–adaptive Method For Matched–field Array processing in acoustic waveguides", J. Acoust. Soc. Am. **87**(6), pp. 2493–2502, 1990.

[30] H.P. Bucker, "Use of Calculated Sound Fields and Matched–Field Detection to Locate Sources In Shallow Water", J. Acoust. Soc. Am **59** pp. 368–373, 1976.

[31] H. Cramer, "Mathematical Methods of Statistics", Princeton University Press, Princeton, NJ, 1946.

[32] S.M. Kay, "Fundamentals of Statistical Signal Processing: Estimation Theory", Prentice Hall, Englewood Cliffs, NJ, 1993.

[33] J.A. Fawcett & B.H. Maranda, "A Hybrid Target Motion Analysis / Matched Field Processing Localisation Method", J. Acoust. Soc. Am **94**(3) pp. 1363–1371, 1993.

[34] J.P. LeCadre & H. Gauvrit, "Optimisation of the Observer Motion for Bearings–Only Target Motion Analysis", Proc. First Australian Data Fusion Symposium, ADFS–96, p. 228, 190–5, 1996.