IMPROVING THE QUALITY OF SYNTHETIC APERTURE SONAR IMAGES USING DATA-DRIVEN AUTOFOCUS TECHNIQUES

P. T. Gough and M. A. Miller
Acoustics Research Group, Department of Electrical and Computer Engineering,
University of Canterbury, Private Bag 4800, Christchurch, New Zealand.
gough@elec.canterbury.ac.nz

Abstract

Although originally developed for synthetic aperture radar (i.e., electromagnetic) systems, many post-reconstruction, image-improvement techniques have direct and indirect applications to synthetic aperture sonar (i.e., acoustic) systems both at the lower-frequencies (i.e., marine acoustics) and at the higher, ultrasonic frequencies (i.e., medical acoustics). The problem common to all synthetic aperture imaging systems is that the time-stamped position of the sensors needs to be known to a sub-wavelength accuracy otherwise the final images are corrupted to the extent that they become unusable and are far from what is theoretically possible which is the diffraction limit of the data collection system.

Many recently deployed synthetic aperture sonars now use a multi-hydrophone, multi-receiver array as part of the sensor platform and so can form a useful, if crude, low-resolution image from a single transmitted pulse. The ensemble of low-resolution images obtained from many pulses contains a great deal of information about how the sensor platform has moved during those pulses. A computer technique to unravel the positional histories of the sensor platform from the ensemble of low-resolution images is presented along with some simulated results derived specifically for the marine acoustics field. What is new in this proposed technique is that both horizontal angular motion (i.e., yaw around the centre of rotation of the array) as well as the more usual horizontal sideways displacement (i.e., sway perpendicular to the array) can be estimated to suprisingly high degrees of accuracy.

Introduction

Any synthetic aperture sonar (SAS) whether single receiver or multiple receiver is limited by the uncorrected random motion of the tow-fish [1, 2]. Although there are three translational displacement and three rotational degrees of freedom in a free body, only sway (side-to-side horizontal displacement from the mean tow-path) affects a single receiver SAS and only sway and yaw (horizontal angular variation around the average direction of travel) affects a multiple receiver SAS. Limiting the development in this paper to multi-receiver SAS systems, we conclude that we need to estimate sway to sub-wavelength accuracy and yaw ,depending on the length of the array, to better than about $1/10^\circ$ so as to acheive diffraction-limited imagery. So far these two critical motion parameters have been estimated by applying an autofocussing technique variously known as the "Displaced Phase Center Antenna" (DPCA), "Displaced Phase Center" (DPC) autofocus [3] or "Redundant Phase Center" (RPC) autofocus. [4]. To make the processing easier, the concept of phase centres is introduced where a phase center is a hypothetical co-located transmitter/receiver pair positioned exactly half way between the actual transmitter and actual receiver. Since we have an array of receivers, we also have an array of phase centres. The important requirement of these RPC techniques is that there **must be overlap** of the phase centres from adjacent pings.

But there is a completely different way to approach the problem of sway and yaw estimation which we have named "Displaced Ping Imaging Autofocus" (DPIA). With a multiple receiver array SAS, it is possible to compute a low angular resolution image from an object within the sector of the swath illuminated by a single pulse of the transmit beam. Provided the individual elements of the receiver array have a broad beam-width response, the low-resolution image computed from a second ping will have considerable overlap with the low-resolution image computed from the first ping. Providing that nothing moves except the SAS platform (and here we are ignoring temporal decorrelation caused by medium fluctuations), these two images have a high degree of overlap and so redundancy. Since we have a high degree of redundancy, there is a high degree of correlation in the two low-resolution images even if there is no overlap in the array of phase centers (and so little or no redundancy and correlation in the raw data). Consequently we now have a way to estimate the differential yaw and differential sway by cross correlating the two low-resolution images. Now the SAS is free to travel at the maximum allowable speed (as determined by the spatial sampling constraints) where there is an overlap of only just one half the array and little if any correlation between the echoes detected by any receiver from adjacent pings. The only requirement of this proposed technique is that there **must be overlap and correlation in the individual low-resolution images** computed from adjacent pings; a technique similar to this has been proposed before but without details in [5]. The basis for our initial simulation is a side-looking sonar with specifications close to those of the slow speed SAS recently built by Northrop-Grumman [6].

The multi-receiver sonar system model has H re-

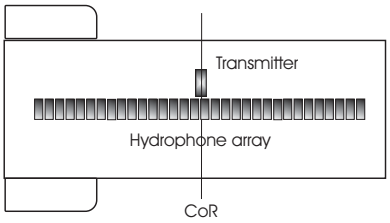

Figure 1: Hydrophone array details.

ceiving hydrophone elements all aligned in the direction of travel where each hydrophone has an integer index $h \in [0, H-1]$. The position of the h^{th} hydrophone element relative to the centre of rotation around the sonar's long axis is $d = (h - H/2)\Delta d$ with Δd the separation of the hydrophone element centres; see fig 1 for a sketch of the towfish geometry. The data collected by this new multi-hydrophone sonar is now

$$ss_p(t,d) = \int_x \int_y ff(x,y) \cdot a(t, x, y-u-d) \odot_t \wp\left(t - \frac{\sqrt{x^2+(y-u)^2}}{c} - \frac{\sqrt{x^2+(y-u-d)^2}}{c}\right) \cdot dx\, dy. \quad (1)$$

where the location of the sonar along the aperture is related to the ping index p by $u = p\Delta u$. Since the sonar now has an array of hydrophones, the received data needs correcting for yaw as well as sway.

Yaw estimates using DPIA

We have a number of options in computing the individual low-resolution images from the data collected from each ping. However since images from adjacent pings are to be cross-correlated to estimate sway (a cross-track distance in x) and yaw (an angle θ), the most appropriate image is in fact $\widehat{ff}_p(x, \theta_p)$ and $\widehat{ff}_{p+1}(x, \theta_{p+1}) \approx \widehat{ff}_{p+1}(x, \theta_p)$ where the origin of x, θ coordinates lies at the centre of rotation of the towfish at ping p. Although a perfect matched filter can be determined for each point in x and θ, there is no really efficient way to process blocks of data in one inverse-filtering operation.

A more efficient process is based on individual ping-centred, polar images where the range to a target is given by $r = \sqrt{x^2 + (y-u)^2}$. Now the low resolution image is given by $\widehat{ff}_p(r, \theta)$ which can be computed with Fourier-based techniques directly from the raw data. There is some range migration in off-axis reflectors between the two ping images, however, this can be accomodated by a small decrease in range resolution if this causes any problems. Taking the data from a single ping, $ss_p(t, d)$, we can reconstruct a limited-resolution, ping-centred polar image $\widehat{ff}_p(r, \theta)$ using any convenient azimuth compression algorithm that results in a geometrically correct image in ping-centred polar r, θ coordinates.

A simple Fourier-based transformation can be used to compute the individual ping-centred polar image $\widehat{ff}_p(r, \theta)$ from $ss_p(t, d)$. The basis for this is

$$\begin{aligned} \widehat{ff}_p(r,\theta) &\approx \mathcal{R}_{t,\,k_d}\{\mathcal{F}_d\{ss_p(t,d)\}\} \\ &= \mathcal{R}_{t,k_d}\{sS_p(t,\,k_d)\} \end{aligned} \quad (2)$$

where the $\mathcal{R}\{\cdot\}$ transformation is a coordinate mapping given by $t \Rightarrow 2r/c$, and $k_d \Rightarrow k_0 \sin\theta$ centred on ping p. Note that no interpolation is required for $\mathcal{R}\{\cdot\}$ mapping from t and k_d coordinates into r and θ coordinates but it does rely on the narrowband approximation of $k_d \approx k_0 \sin\theta$.

But what if there is a change in towfish yaw between the two images? Clearly each image is displaced in angle relative to each other but not in range. Thus the actual angular change required to bring the images into coincidence can be estimated. If two adjacent ping-centred and distorted complex images are cross-correlated in θ, an estimate of the differential yaw can be obtained from the envelope of the cross-correlation function. Because the rotational centre of the towfish for ping $(p+1)$ and ping p are displaced by the distance moved between pings, the origin of each image is slightly different and the cross-correlation has a slight range-dependent error. By using a slightly different matched filters for each ping or non-symetrical zero padding appropriately in the Fourier domain, we apply

$$cc_p(r,\theta) = |\widetilde{ff}_{p+1}(r,\theta_p) \star_\theta \widetilde{ff}^*_p(r,\theta_p)| \quad (3)$$

and now the origins of coordinates for the two ping-centred, polar images are the same. The maximum value of $cc_p(r, \theta)$ for each value of r identifies the shift in θ—call this $\theta_p(r)$—required to bring that range of the two ping-centred images into coincidence. Now $\theta_p(r)$ is the **differential** yaw of the sonar platform between ping $(p+1)$ and ping p at range r and averaging $\theta_p(r)$ over r gives the best estimate of the differential yaw. It is also common to cross correlate the envelopes of the two images although zero padding may be needed in the Fourier domain to accomodate bandwith expansion by the non-linear process. This amplitude only correlation is given by

$$cc_p(r,\theta) = |\widetilde{ff}_{p+1}(r,\theta_p)| \star_\theta |\widetilde{ff}^*_p(r,\theta_p)| \quad (4)$$

and although more robust doesnt have the accuracy of the full complex correlation. To test the efficacy of the proposed algorithms, we need to estimate the yaw and sway using the proposed technique and then compare it in some mean-squared-error sense to the actual yaw and sway injected into the simulated data collection process. Firstly we simulate a 32 hydrophone receiver SAS with the layout of the hydrophone elements illustrated in fig.1 The centre of rotation (CoR) is considered to pass through the centre of the hydrophone

array and the middle of the transmitting projector. The SAS travels at $3.5\,\mathrm{ms}^{-1}$ radiating a 10° beam of sound at a centre frequency of 100kHz with a pulse bandwidth of 20kHz and repeating every 0.29 s. The tow-velocity of the SAS is such that there is no phase centre overlap and the leading-most phase centre of one ping is adjacent to the trailing-most phase centre of the next ping. The received data from each of the 32 hydrophone receivers is well over-sampled, demodulated to complex baseband and decimated to a sampling frequency of 30kHz. The first image shown in fig. 2. This displays the magnitude of the raw data (unprocessed other than demodulated to complex base-band and pulse compressed in range) from just one hydrophone receiver detected from a rectangular block object surrounded by point reflectors and immersed in random uncorrelated noise. The effect of random sway on the unprocessed image is quite noticable. Since the actual sway and yaw used to create the raw data are known, they can be corrected for in the image reconstruction algorithm (in this case a modified form of the wavenumber algorithm) and the diffraction-limited image intensity available for visual comparison.

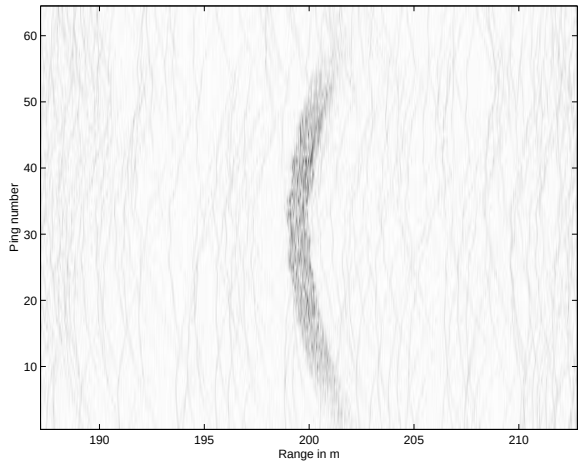

Figure 2: Unprocessed raw data with non-zero yaw and sway included

Now for the test of the yaw and sway estimation procedure. The data from all 32 receivers for any two sequential pings is applied to the single ping imaging algorithm and the yaw differences estimated from the θ cross-correlation of the intensity of the two images. This is combined with the sway differences estimated from the r cross-correlation of the same two images. Now the difficult step is to estimate the sway and yaw from the measured sway differences and measured yaw differences since there is no unique way of separating constant yaw from linear sway. Once this procedure has been completed, we can compare the actual sway and yaw that was inserted into the data collection simulation with the estimated sway and yaw determined from the unprocessed raw data alone.

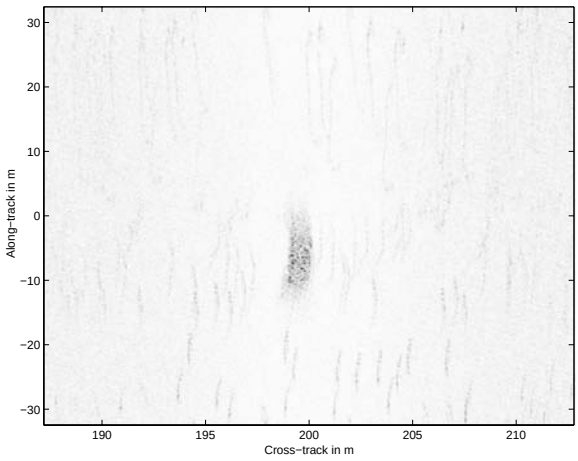

Figure 3: Processed image showing effect of uncorrected yaw and sway

Constant yaw vs linearly increasing sway

Consider the situation where a towfish has constant yaw for all pings and no other motion errors. Simple geometry suggests that, with a redefinition of the direction of travel, more or less the same data would be generated by a situation with zero yaw and linearly increasing sway. Since the same data set would be generated by an ensemble of different scenarios, it means that the inverse problem of determining the actual motion from a single data set is ambiguous. In fact the same data set could be generated with an infinite number of different combinations of constant yaw and linearly increasing sway.

We deal with the ambiguity by assuming that all yaw-offset can be corrected as linearly increasing sway. So before passing the yaw-corrected data onto the sway estimation algorithm, the yaw estimates derived entirely from integrating the yaw differences are set to a zero mean. Given no *apriori* information, it is reasonable to assume that the yaw value is over the span of the aperture is zero-mean. Thus, setting the mean of the yaw estimates to zero lessens the average amount of constant yaw to be detected as linear sway in the sway estimation algorithm since it minimises the variance in the difference between the true mean and the mean of the estimates. Once the raw data has been corrected for yaw (to within an unknown constant), the corrected data is passed to the sway estimator. Since a constant yaw is equivalent to a linearly increasing sway, the unknown constant yaw is detectable in the sway differential as a non-zero trend.

Finally it remains to be seen just how effective the DPIA algorithm is in estimating the yaw and sway and how the image is improved using these estimates. Comparing the magnitude of the reconstructed image

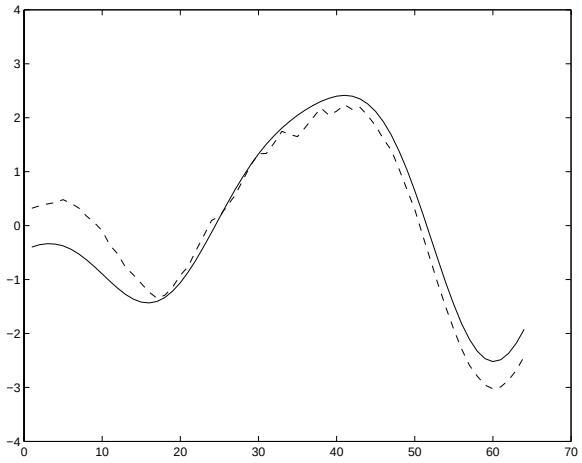

Figure 4: Yaw estimates in degrees (dashed line) compared to actual yaw (solid line) both with means removed

shown in fig. 5 and it can been inferred that the image is very much closer to the diffraction-limit than the image using no correction at all, fig. 3. It is also clear that a strong cultural (*i.e.*, man-made) target appears to bias the results so that images of areas including such targets are not as good as images of targets fields comprising of nothing more than a selection of almost-equal intensity point reflectors.

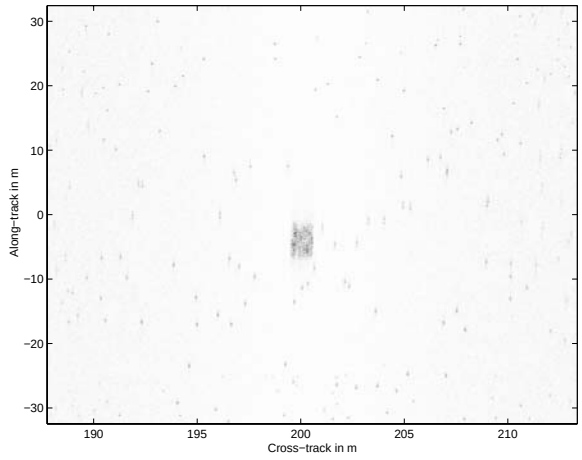

Figure 5: Reconstructed image corrected for estimated yaw and estimated sway

Conclusions

The reconstructed images created using a hydrophone array SAS are almost always impaired by unknown random yaw and sway. Consequently autofocus algorithms need to be able to estimate the random yaw and sway histories reasonably accurately. By cross-correlating two limited-resolution, ping-centred images computed using data collected from two adjacent pings, we can estimate the differential yaw and sway that occured between the two pings. With suitable integration over the full set of recorded echoes, we can estimate overall yaw histories down to 1/10° and sway histories down to a fraction of a wavelength (at least in computer simulations). Target fields that include strong targets need to be handled carefully since their presence appears to bias the motion estimates and we have also noticed this in single hydrophone SAS sway estimators based on the unprocessed echos such as shear averaging [7, 8, 9].

References

[1] P.T. Gough & D.W. Hawkins, "Imaging algorithms for a strip-map synthetic aperture sonar: Minimizing the effects of aperture errors and aperture undersampling," *Oceanic Engineering IEEE J.*, vol. 22(1), pp. 27–39, 1997.

[2] S.M. Banks & H.D. Griffiths, "Imaging and motion estimation for synthetic aperture sonar based on Fast Factorised Back-Projection," in Proceedings ECUA'2002, Gdansk, Poland, 24-27 June 2002, pp.529—534.

[3] A. Bellettini & M. Pinto, "Experimental investigation of SAS micronavigation," in Proceedings ECUA'2000, Lyon, France, July 2000, pp.445-450.

[4] J.T. Christoff, "Motion-compensated high frequency synthetic aperture sonar", *J. Acoust. Soc. Am.*, vol.8l, p.2950, 1998.

[5] J. Groen & J.C. Sabel, "Interaction between autofocussing and synthetic aperture processing: a study based on simulations and experimental rail data", in Proceedings ECUA'2002, Gdansk, Poland, 24-27 June 2002, pp.235—240.

[6] H. Keeter, "Navy awards Northrop Grumman $3.5 million synthetic aperture sonar contract," *Defense Daily*, vol. 209(3), March 2002.

[7] H.J. Callow, M.P. Hayes & P.T. Gough, "Wavenumber domain reconstruction of SAR/SAS imagery using single transmitter and multiple-receiver geometry," *Electronics Letters*, vol.38(7), pp.336–337, 2002

[8] H.J. Callow, M.P. Hayes & P.T. Gough, "Noncoherent autofocus of a single-receiver broad-band synthetic aperture sonar imagery", in Proceedings IEEE Oceans, Sept. 2001, Honolulu, Hawaii, pp.157—162.

[9] K.A. Johnson, P.T. Gough & M.P. Hayes, "A method of estimating the sub-wavelength sway of a sonar towfish", *Oceanic Engineering, IEEE J.*, vol. 20 (4), pp.258—267, 1995.