

Strip-map Phase Gradient Autofocus

M. P. Hayes*, H. J. Callow, and P. T. Gough

Acoustics Research Group,

Dept. of Electrical and Computer Engineering,

University of Canterbury,

Private Bag 4800, Christchurch, NZ

Abstract

This paper describes a generalisation of the Phase Gradient Autofocus (PGA) algorithm that allows strip-map operation. A standard autofocus technique, PGA, uses prominent points within the target scene to estimate the point spread function of the system. PGA was developed for tomographic mode spotlight synthetic aperture radar (SAR) but has limited applicability for side-scan synthetic aperture operation. We show how it can be generalised to work with strip-map geometries and relate our new method to the previous PGA extension to strip-map systems.

Strip-map PGA (SPGA) works by employing the (image domain) 2D phase correction technique. Using the scale transform of the imaging system, the 2D phase errors are related to aperture position errors. Robust phase error estimates are obtained by utilising range redundancies.

We also present a modified Phase Curvature Autofocus (PCA) algorithm applicable to strip-map operation that also uses the scale transform. Preliminary results of the two algorithms are presented showing a similar performance, and a discussion of the limitations of the algorithms is given.

1 Introduction

Synthetic aperture sonar (SAS) provides high-resolution seafloor imaging by coherently combining the echoes received by a transducer as it moves along transmitting pings. Synthetic apertures can be achieved that are much longer than physical apertures, especially at low frequencies, and that can be tailored to achieve a constant along track resolution with range.

The disadvantage with synthesising long apertures is that it is difficult to constrain (or measure) the path followed by the sonar to the accuracy required for coherent processing. If the unknown sway is of the order of a tenth of a wavelength or more, the reconstructed image is severely blurred. Constraining the motion of a tow-body with the use of a fixed guide or rail has been popular in the past to prove the concept of synthetic aperture sonar but is limited for practical use. High-drag nose towed towfish have been employed to successfully reduce the sway and Autonomous Underwater Vehicles (AUVs) have been proposed as more stable SAS platforms.

There are two major approaches to solving the blurring caused by the unknown sonar motion. The first is to instrument the sonar platform with an extremely accurate inertial navigation system (INS) to monitor the path followed by the sonar. These systems are expensive (but getting cheaper) and so the second approach is to try to estimate the sonar motion from the recorded echo data. The techniques that implement this are usually referred to as “micro-navigation” or “autofocus” techniques. Both data-driven micro-navigation and autofocus techniques attempt to solve essentially the same problem, that of estimating the path of the sonar from the data (and correcting the output images, possibly in iterative fashion).

Many autofocus techniques for synthetic aperture sonar and radar have been proposed in the literature. Essentially these either utilise strong point-like scatterers or the ensemble of weak scatterers that produce the seafloor reverberation.

The most popular technique for SAR autofocus is Phase Gradient Autofocus (PGA) [Jakowatz et al. 1996; Wahl et al. 1994a] but this is restricted in application to spotlight operation rather than to strip-map operation. In spotlight operation, the sonar beam is focused on the region of interest whereas the more general strip-map configuration images different portions of the scene as viewed from the sonar. PGA has been applied to the strip-map case by mosaicing the results from small sections of the aperture (mPGA) [Bonifant 1999].

A more promising variation of PGA applicable to strip-map operation is PCA (Phase Curvature Autofocus) [Wahl et al. 1994b]. However, as shown in this paper, PCA cannot estimate local linear phase errors.

This paper presents a version similar to PCA that allows phase gradients to be estimated rather than phase curvatures. This algorithm is applicable to strip-map operation hence we call this algorithm Strip-Map Phase Gradient Autofocus (SPGA). We start by recapping the effects of motion errors on synthetic aperture imagery, then present a PCA algorithm that encompasses 2D blurring. This is then followed by how to adapt PCA to estimate phase gradients and we then illustrate its performance compared with PCA. Finally, we discuss the limitations of SPGA, draw our conclusions, and present recommendations for future research.

2 Motion effects

The sonar motion-affected pulse-compressed echo data $\tilde{ss}(t, y)$ can be related to the motion-free echo data $ss(t, y)$ using

$$\tilde{ss}(t, y) = \mathcal{M}(ss(t, y); X(y)) \quad (1)$$

where $X(y)$ is the unknown position error as a function of the along-track position, y , t is the time delay related to range, and \mathcal{M} is an operator describing the motion induced transformation. The goal of autofocus algorithms is to estimate $X(y)$ from $\tilde{ss}(t, y)$ and to recover $ss(t, y)$.

The proposed autofocus algorithms operate on regions of the distorted reconstructed image $\tilde{ff}(x, y)$,

$$\tilde{ff}(x, y) = \mathcal{A}\{\tilde{ss}(t, y)\}, \quad (2)$$

where $\mathcal{A}\{\}$ is an operator that represents the azimuth compression performed by the image reconstruction algorithm.

In the temporal Fourier domain, the pulse-compressed signal may be described using a 2D convolution of the target distribution ff with a range-variant point spread function H ,

$$\tilde{sS}(y, f) = |P(f)|^2 \int \int_{-\infty}^{\infty} ff(x_1, y_1) H(X(y) - x_1, y - y_1, f) dx_1 dy_1, \quad (3)$$

*e-mail:m.hayes@elec.canterbury.ac.nz

where the point spread function is the squared free space Green's function weighted by the beam pattern of the sonar

$$H(X(y) - x_1, y - y_1, f) = \frac{\exp\left(-j4\pi(f/c)\sqrt{(X(y) - x_1)^2 + (y - y_1)^2}\right)}{(4\pi)^2 \left((X(y) - x_1)^2 + (y - y_1)^2\right)} \times B^2 \left(\frac{(f/c)(y - y_1)D_y}{\sqrt{(X(y) - x_1)^2 + (y - y_1)^2}} \right). \quad (4)$$

To keep the mathematics manageable we have assumed 3D scattering from point targets in the z plane of the sonar and thus we can ignore elevation beam pattern effects. There are many other assumptions implicit in Eq. (4), primarily that the sonar does not move during the interval echoes are collected so that there is no temporal Doppler shift.

The point spread function Eq. (4) can be simplified by assuming that the sway is small compared to the across-track target position, i.e., $|X(y)| \ll x_1$, so

$$H(X(y) - x_1, y - y_1, f) \approx \exp\left(-j4\pi(f/c)X(y)\frac{x_1}{\sqrt{x_1^2 + (y - y_1)^2}}\right) \times H(x_1, y - y_1, f). \quad (5)$$

Then using the substitution

$$\cos \theta(y) = \frac{x_1}{\sqrt{x_1^2 + (y - y_1)^2}}, \quad (6)$$

where $\theta(y)$ is the angle from the sonar to a target at (x_1, y_1) , Eq. (5) can be written as

$$H(X(y) - x_1, y - y_1, f) \approx \exp(-j4\pi(f/c)X(y)\cos \theta(y)) \times H(x_1, y - y_1, f). \quad (7)$$

In narrowbeam applications $\cos \theta(y) \approx 1$, in which case

$$H(X(y) - x_1, y - y_1, f) \approx \exp(-j4\pi(f/c)X(y))H(x_1, y - y_1, f). \quad (8)$$

This is the simple timing error approximation. In the narrowband case this further simplifies to

$$H(X(y) - x_1, y - y_1, f) \approx \exp(-j4\pi(f/c)X(y))H(x_1, y - y_1, f) \quad (9)$$

where f_c is the centre frequency.

Using the simple timing error approximation with Eq. (3), the pulse-compressed echo measured with a motion $X(y)$ can be related to the motion free pulse-compressed echo,

$$\tilde{sS}(y, f) \approx sS(y, f) \exp(-j4\pi(f/c)X(y)). \quad (10)$$

Describing the effect of a motion error on the reconstructed image is not straightforward since it depends on the form of $X(y)$. In the simple case of a small constant sonar sway ($X(y) = X$), there is no obvious effect on the reconstructed image; it is simply displaced in range. When there is a small constant linear sway ($X(y) = \sigma y$), the reconstructed image is rotated,

$$\tilde{ff}(x, y) \approx ff(x - \sigma y, y + \sigma x), \quad (11)$$

which for very small linear sway gradients σ can be approximated by an along-track shear in y . Finally, higher order motion errors cause blurring.

3 Modified PCA algorithm

The standard PCA algorithm [Wahl et al. 1994b] developed for SAR neglects across-track image blurring and simply uses across-track image slices. SAS has large beamwidths and the across-track blurring is more significant. Here we present a modification of PCA that utilises across-track blurring by considering regions around strong scattering targets. The key is how the motion estimates are sensibly combined since in a strip-map system the targets are only in view for a small portion of the sonar aperture.

In this algorithm we isolate patches of the reconstructed image around strong scatterers and relate phase errors back to the sonar motion using the scale transform [Soumekh 1999; Hawkins 1996]

$$f_y = \frac{y - y_m}{x_m} f_x, \quad (12)$$

derived using the relation

$$\tan \theta = \frac{y - y_m}{x_m} = \frac{f_y}{f_x}. \quad (13)$$

Starting with the pulse-compressed echo signal $ss(y, t)$, the modified PCA algorithm is:

Step 1: Fourier transform pulse-compressed echo data into the temporal frequency domain.

$$\tilde{sS}(y, f) = \int_{-\infty}^{\infty} \tilde{ss}(y, t) \exp(-j2\pi ft) dt. \quad (14)$$

Step 2: Correct for known sway $X(y)$, using Eq. (10).

Step 3: Reconstruct image $\tilde{ff}(x, y)$ from $\tilde{sS}(y, f)$, say using the wavenumber algorithm.

Step 4: Find prominent points in $\tilde{ff}(x, y)$ and associate coordinates (x_m, y_m) . Note that x_m and y_m do not need to be found to sub-pixel precision.

Step 5: Create masks $W_m(x, y)$ centred on (x_m, y_m) , where

$$W_m(x, y) = \text{rect}\left(\frac{x - x_m}{W_x}\right) \text{rect}\left(\frac{y - y_m}{W_y}\right). \quad (15)$$

The sizes of these masks is a critical aspect of the algorithm. They need to be large enough to encompass the 2D blurring of a point target but small enough to not include more than one strong target.

Step 6: Form images containing regions of interest by multiplying $\tilde{ff}(x, y)$ by the masks $W_m(x, y)$,

$$\tilde{ff}_m(x, y) = \tilde{ff}(x, y)W_m(x, y). \quad (16)$$

Step 7: 2D Fourier transform $\tilde{ff}_m(x, y)$ to $\tilde{FF}_m(f_x, f_y)$.

Step 8: Apply scale transform to form $\chi_m(f_x, y)$ from $\tilde{FF}_m(f_x, f_y)$,

$$\chi_m(f_x, y) = \tilde{FF}_m\left(f_x \frac{y_m - y}{x_m}, f_x\right). \quad (17)$$

Note that this requires an interpolation with sampled data.

Step 9: At this point we have M estimates of the phase error $\Phi_m(f_x, y)$ where

$$\Phi_m(f_x, y) = \text{phase} \{ \chi_m(f_x, y) \}. \quad (18)$$

These phase errors are related to the residual sway by

$$\Phi_m(f_x, y) = 2\pi f_x X_i(y), \quad (19)$$

which in the narrowbeam and narrowband case simplifies to

$$\Phi_m(f_x, y) \approx 4\pi \frac{f_c}{c} X_i(y). \quad (20)$$

These phase error estimates cannot be directly averaged due to errors in the estimates of x_m and y_m . An error in x_m primarily generates a localised phase offset as well as a small hyperbolic along-track error.

An error in estimating y_m produces an along-track linear phase error. Moreover, y_m is the point of closest approach but not the true centre position of each target. Thus there is a local linear phase error and it is necessary to average the phase error curvatures since these are independent of y_m . The phase error curvatures are simply the second derivative of the along track phase errors. They can be computed using successive finite differences with a sheared product

$$\Delta^2 \Phi_m(f_x, y) = \text{phase} \{ \psi_m(f_x, y) \}, \quad (21)$$

where

$$\psi_m(f_x, y) = \chi_m(f_x, y + 2\Delta y) \chi_m^*(f_x, y + \Delta y) \chi_m(f_x, y). \quad (22)$$

Step 10: Average the phase error curvatures with an amplitude weighting using

$$\Delta^2 \varphi(y) = \text{phase} \{ \Psi(y) \} \quad (23)$$

where

$$\Psi(y) = \int_{-\infty}^{\infty} \sum_{m=1}^M \psi_m(f_x, y) df_x. \quad (24)$$

Step 11: Double integrate the averaged phase error curvatures

$$\begin{aligned} \varphi(y) &= \varphi(y - \Delta y) + \Delta\varphi(y - \Delta y), \\ \Delta\varphi(y) &= \Delta\varphi(y - \Delta y) + \Delta^2\varphi(y - \Delta y). \end{aligned} \quad (25)$$

Step 12: Scale the phase error estimate by the wavelength of the centre frequency and divide by 4π to estimate the residual sway

$$\hat{X}_i(y) = \frac{\varphi(y)}{4\pi f_c / c}. \quad (26)$$

Step 13: Iterate algorithm until the residual sway is smaller than some threshold. The total estimated sway is the sum of the estimated residual sways from each iteration

$$\hat{X}(y) = \sum_{i=1}^I \hat{X}_i(y). \quad (27)$$

While this algorithm has been described using continuous variables for convenience, in practice it is implemented using sampled variables. Furthermore, to reduce computation, offset variables are used such as the complex envelope.

4 SPGA algorithm

The disadvantage of PCA is that it does not estimate local linear sways. Thus if there is not at least a point target per synthetic aperture length, i.e., aperture coverage is insufficient, then the reconstructed image can be disjoint.

SPGA overcomes this problem by estimating the local linear sway. It is based on the observation that a local linear sway causes a displacement between the position of closest approach and the brightest point in the echo history [Gough et al. 2000]. This effect is illustrated in Figure (1) and Figure (2) which show a single point target imaged with and without a linear sway. Note that in both cases the reconstructed point is at the position of closest approach (i.e. where the phase of the echo history is minimum) and that when there is a linear sway, the point of closest approach is not the same as when the target is in the centre of the sonar beam pattern (i.e., where the echo history is brightest).

The effect of a linear sway in the spatial frequency (wavenumber) domain can be found from a 2D Fourier transform of Eq. (11) to be a rotation

$$\widetilde{FF}(f_x, f_y) \approx FF(f_x - \sigma f_y, f_y + \sigma f_x). \quad (28)$$

Again, for small sway gradients σ , this is simply a shearing of the spatial frequency domain. Thus by estimating this spectral shift, σ can be estimated.

Step 1: Fourier transform pulse-compressed echo data into the temporal frequency domain.

$$s\tilde{S}(y, f) = \int_{-\infty}^{\infty} \tilde{s}s(y, t) \exp(-j2\pi ft) dt. \quad (29)$$

Step 2: Correct for known sway $X(y)$, using Eq. (10).

Step 3: Reconstruct image $\widetilde{ff}(x, y)$ from $s\tilde{S}(y, f)$, say using the wavenumber algorithm.

Step 4: Find prominent points in $\widetilde{ff}(x, y)$ and associate coordinates (x_m, y_m) . Note that x_m and y_m do not need to be found to sub-pixel precision.

Step 5: Create masks $W_m(x, y)$ centred on (x_m, y_m) , where

$$W_m(x, y) = \text{rect} \left(\frac{x - x_m}{W_x} \right) \text{rect} \left(\frac{y - y_m}{W_y} \right). \quad (30)$$

Step 6: Form images containing regions of interest by multiplying $\widetilde{ff}(x, y)$ by the masks $W_m(x, y)$,

$$\widetilde{ff}_m(x, y) = \widetilde{ff}(x, y) W_m(x, y). \quad (31)$$

Step 7: 2D Fourier transform $\widetilde{ff}_m(x, y)$ to $\widetilde{FF}_m(f_x, f_y)$.

Step 8: Determine the average along-track spatial energy density spectrum using

$$\beta_m(f_y) = \int_{-\infty}^{\infty} \left| \widetilde{FF}_m(f_x, f_y) \right|^2 df_x. \quad (32)$$

Step 9: Calculate the centroid of the along-track spatial energy density spectrum using

$$\Delta f_{y_m} = \frac{\int_{-\infty}^{\infty} f_y \beta_m(f_y) df_y}{\int_{-\infty}^{\infty} \beta_m(f_y) df_y}. \quad (33)$$

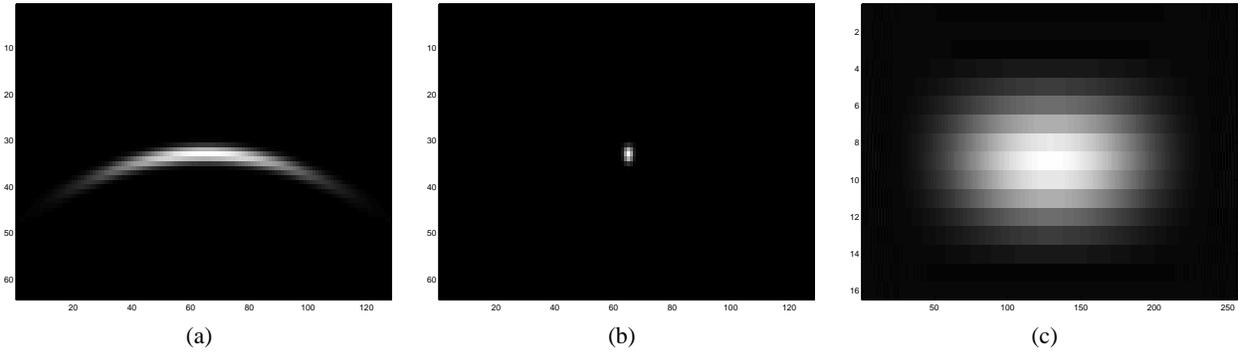


Figure 1: Images of a single prominent point with no sway: (a) magnitude of pulse-compressed echo history $ss(y,t)$, (b) magnitude of reconstructed image $ff(x,y)$, (c) real part of spatial Fourier transform of reconstructed image $FF(x,y)$.

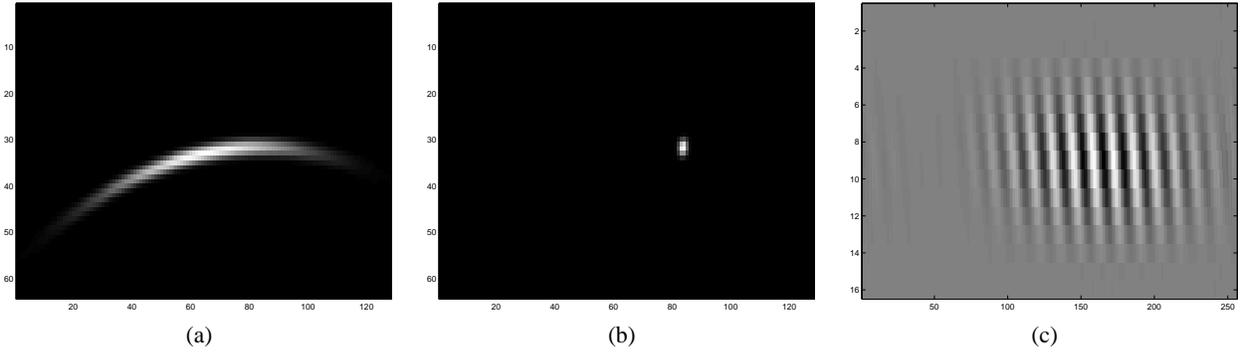


Figure 2: Images of a single prominent point with a linear sway ($X(y) = 0.05y$): (a) magnitude of pulse-compressed echo history $\tilde{ss}(y,t)$, (b) magnitude of reconstructed image $ff(x,y)$, (c) real part of spatial Fourier transform of reconstructed image $\tilde{FF}(f_x, f_y)$.

Step 10: Estimate the local linear sway gradient,

$$\sigma_m = \frac{\Delta f_{y_m} c}{2f_c}. \quad (34)$$

Step 11: Estimate the true target position (x'_m, y'_m)

$$x'_m = x_m, \quad (35)$$

$$y'_m = y_m - \sigma_m x_m. \quad (36)$$

Step 12: Phase correct spatial spectrum for estimated true target position

$$\begin{aligned} \tilde{FF}'_m(f_x, f_y) &= \tilde{FF}_m(f_x, f_y) \\ &\times \exp(-j2\pi [f_x(x'_m - x_m) + f_y(y'_m - y_m)]). \end{aligned} \quad (37)$$

Step 13: Apply scale transform to form $\chi_m(f_x, y)$ from $\tilde{FF}'_m(f_x, f_y)$,

$$\chi_m(f_x, y) = \tilde{FF}'_m\left(f_x \frac{y_m - y}{x_m}, f_x\right). \quad (38)$$

Step 14: Calculate the phase error gradient using a sheared product

$$\Delta\Phi_m(f_x, y) = \text{phase}\{\psi_m(f_x, y)\}, \quad (39)$$

where

$$\psi_m(f_x, y) = \chi_m(f_x, y + \Delta y) \chi_m^*(f_x, y). \quad (40)$$

Step 15: Average the phase error gradients with an amplitude weighting using

$$\Delta^2\varphi(y) = \text{phase}\{\Psi(y)\} \quad (41)$$

where

$$\Psi(y) = \int_{-\infty}^{\infty} \sum_{m=1}^M \psi_m(f_x, y) df_x. \quad (42)$$

Step 16: Integrate the averaged phase error gradients

$$\varphi(y) = \varphi(y - \Delta y) + \Delta\varphi(y - \Delta y). \quad (43)$$

Step 17: Scale the phase error estimate by the wavelength of the centre frequency and divide by 4π to estimate the residual sway

$$\hat{X}_i(y) = \frac{\varphi(y)}{4\pi f_c / c}. \quad (44)$$

Step 18: Iterate algorithm until the residual sway is smaller than some threshold. The total estimated sway is the sum of the estimated residual sways from each iteration

$$\hat{X}(y) = \sum_{i=1}^I \hat{X}_i(y). \quad (45)$$

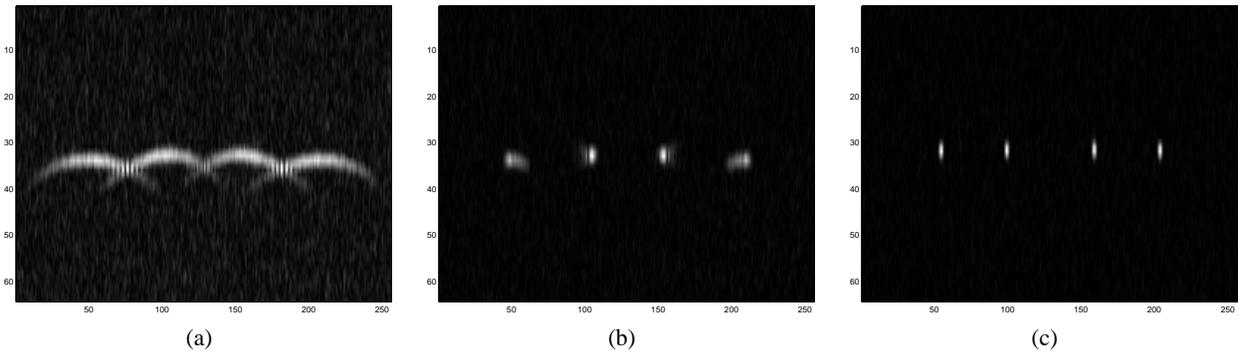


Figure 3: Images of four prominent points with a two cycle sinusoidal sway: (a) magnitude of pulse-compressed echo history $\tilde{s}(y,t)$, (b) magnitude of reconstructed image $\tilde{f}(x,y)$, (c) magnitude of reconstructed image $ff(x,y)$ after motion compensation estimated by SPGA has been applied.

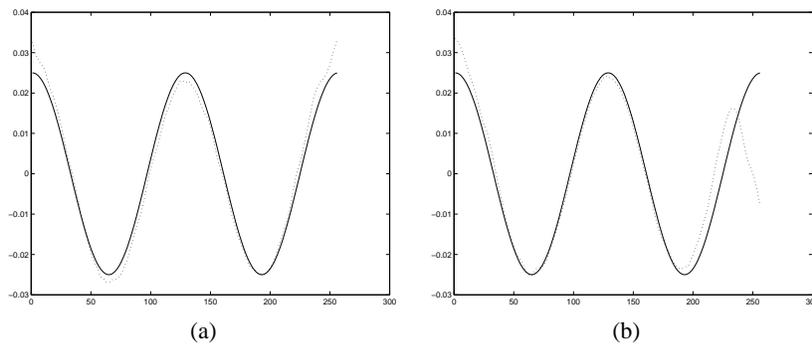


Figure 4: Sway (metres) estimated by (a) modified PCA and (b) SPGA. The actual sway in both cases is denoted by the solid line.

5 Results

Figure (3) shows the effect of a mild sinusoidal sway on four point targets at the same range and how this can be estimated using SPGA and removed to improve the image quality. The sway has a peak-peak amplitude of 0.05 m (one wavelength of the centre frequency) with two cycles over the displayed region. White Gaussian noise with a power of 0 dB compared to the maximum echo power has been added to the echo to model sea-noise and receiver noise but not coherent speckle from seafloor reverberation. Figure (3)(a) shows the magnitude of the pulse compressed echo signal; the pulse compression has improved the signal to noise ratio by approximately 20 dB. After synthetic aperture image reconstruction, the signal to noise ratio is further improved by the image is not diffraction limited due to the effect of the uncompensated sway. Finally, after using SPGA to estimate the sway and applying motion compensation, a close to diffraction limited image can be seen in Figure (3)(c).

The sway estimated by the modified PCA and SPGA algorithms is shown along with the actual sway in Figure (4). SPGA has performed better in the central portion of the aperture but poorer at the edges. To overcome the poor response at the edge may require zero padding of the image data to avoid edge effects. Obviously, many more trials are required to determine the relative merits of the two algorithms.

6 Discussion

In the embodiment described here, SPGA has a number of possible limitations. The most severe limitations are caused by SPGA's requirement for a statistically homogeneous scene. The main limitations are:

Window size — SPGA windows individual targets to ensure that the phase signature of only one scatterer at a time is used. If a number of targets exist within the window, the path estimation gives incorrect results. As the algorithm iterates, the along-track window size decreases reducing the likelihood of multiple targets with a single window. Selecting many targets at different locations alleviates some of the symptoms.

A related problem exists if the window size chosen does not fully encompass the target blurring. The algorithm is then not able to estimate high frequency blurring components leading to low frequency path estimates (see [Warner et al. 2000]). SPGA, like PGA, starts with a wide window that decreases over time to allow estimation of the low frequency errors after the high frequency ones.

Point scatterer assumption — In the derivation of SPGA the image scene was assumed to consist of a number of randomly placed point scatterers. The algorithm exploits the randomly varying scene phase structure to estimate the constant phase structure caused by the sonar sway. Any statistically homogeneous scene is suitable for SPGA operation, where the selected targets have zero average Fourier phase (the same as

the Knox-Thompson and zero bispectral-phase priors used in astronomical imaging).

The selection of a strongly scattering target that has significant Fourier phase (often a strongly scattering extended target) leads to estimate biasing. The bias may be reduced by selecting unrelated point-like targets or point-like targets that have a high contrast [Chan and Yeo 1998]. Selecting many point-like targets is a useful way of ensuring the algorithm's input scene is statistically homogeneous.

Window separation — SPGA requires the input targets patches to be uncorrelated—to ensure statistical homogeneity. Ideally the targets selected should be well separated [Chan and Yeo 1998]. SPGA requires that the majority of energy in the patch comes from the patch. If under-sampling of the aperture leads to grating lobes, these must not be selected.

Clutter information — The information contained in the sea-floor clutter is useful for autofocus. SPGA currently ignores this source of information. However, we expect that SPGA is able to operate on clutter regions with a different patch selection policy. Clutter regions (on average) satisfy the requirement for a statistically homogeneous scene (if not the point-scatterer assumption) and so are able to be used directly.

Non Lambertian scatterers — SPGA assumes that the energy scattered from a target is the same from all incidence angles. Because SPGA relies on beam symmetry for the linear slope estimation targets that have non Lambertian scattering profiles can cause incorrect results.

Squinted systems — SPGA will not operate on systems that image with a squinted geometry (see above). However, if the squint angle is known, this can be compensated.

Bulk sway estimation — Approximations made in SPGA's derivation require that the sway is small—due to the scale transform assuming no shift in the stationary phase point. SPGA should not be used for bulk sway estimation (use Shear Average Autofocus instead [Johnson et al. 1995; Callow et al. 2001]).

Large linear sways — Large linear sways wrap the linear slope calculation stage causing biasing in the spatial Doppler centroid estimation. Calculating the linear slope using correlation techniques will ameliorate the biasing.

Platform yaw — A reliance on the spatial Doppler centroid makes the algorithm vulnerable to platform yaws. A yaw will shift the spatial Doppler centroid leading to inaccurate linear slope estimation. SPGA assumes that no yaw is present in the system.

Improvements to the algorithms could include the use of a higher order eigenvector phase derivative estimator [Jakowatz and Wahl 1993] to reduce the number of iterations and a better estimator of the local linear phase gradient. Rather than estimating centroids, it is likely that correlation with the expected beam-pattern response will be more robust when there are large local linear phase gradients or when the data is marginally sampled.

7 Conclusion

In this paper we have presented two new autofocus algorithms, both using 2-D patches around bright regions in the reconstructed image as estimated of blurred prominent points. The scale transform is used to map the 2-D Fourier transform of these patches and a

sheared product is used as the phase gradient estimator. The modified PCA algorithm averages phase curvatures while the SPGA algorithm averages phase gradients after estimating and compensating local linear phase gradients.

The results at this stage are preliminary but promising. Many more trials are required to gain an understanding of the relative performance of the algorithms. This is a difficult task since the performance of autofocus algorithms is strongly dependent on the scene, motion, and models used. Further work is also required to make the algorithms robust so that they will work reliably on real sonar data.

References

- BONIFANT, JNR., W. W. 1999. *Interferometric synthetic aperture sonar processing*. Master's thesis, Georgia Institute of Technology.
- CALLOW, H. J., HAYES, M. P., AND GOUGH, P. T. 2001. Non-coherent autofocus of single receiver, broad-band synthetic aperture sonar imagery. In *OCEANS 2001*, vol. 1, IEEE, 157–162.
- CHAN, H. C., AND YEO, T. S. 1998. Noniterative quality phase-gradient autofocus (QPGA) algorithm for spotlight SAR imagery. *IEEE Transactions on Geoscience and Remote Sensing* 36, 5 (September), 1531–1539.
- GOUGH, P. T., HAYES, M. P., AND CALLOW, H. J. 2000. Strip-map path correction using phase matching autofocus. In *ECUA 2000*, ECUA, 412–418.
- HAWKINS, D. W. 1996. *Synthetic Aperture Imaging Algorithms: with application to wide bandwidth sonar*. PhD thesis, Department of Electrical and Electronic Engineering, University of Canterbury.
- JAKOWATZ, JNR., C. V., AND WAHL, D. E. 1993. Eigenvector method for maximum-likelihood estimation of phase errors in synthetic-aperture-radar imagery. *Journal of the Optical Society of America* 10, 12 (December), 2539–2546.
- JAKOWATZ, JNR., C. V., WAHL, D. E., EICHEL, P. H., GHIGLIA, D. C., AND THOMPSON, P. A. 1996. *Spotlight-Mode Synthetic Aperture Radar: A Signal Processing Approach*. Kluwer Academic Publishers, Boston.
- JOHNSON, K. A., HAYES, M. P., AND GOUGH, P. T. 1995. Estimating sub-wavelength sway of sonar towfish. *IEEE Journal of Oceanic Engineering* 20, 4 (October), 258–267.
- SOUMEKH, M. 1999. *Synthetic Aperture Radar Signal Processing with MATLAB Algorithms*. John Wiley & Sons, Inc.
- WAHL, D. E., EICHEL, P. H., GHIGLIA, D. C., AND JAKOWATZ, JNR., C. V. 1994. Phase gradient autofocus – a robust tool for high resolution SAR phase correction. *IEEE Transactions on Aerospace and Electronic Systems* 30, 3 (July), 827–835.
- WAHL, D. E., JAKOWATZ, JNR., C. V., AND THOMPSON, P. A. 1994. New approach to strip-map SAR autofocus. In *Sixth IEEE Digital Signal Processing Workshop*, IEEE, 53–56.
- WARNER, D. W., GHIGLIA, D. C., FITZGERRELL, A., AND BEAVER, J. 2000. Two-dimensional phase gradient autofocus. *SPIE-Image Reconstruction from Incomplete Data* 4123, 162–173.