IMAGE ENHANCEMENT IN SYNTHETIC APERTURE SONAR

R E Hansen Norwegian Defence Research Establishment, Kjeller, Norway
J Groen NATO Undersea Research Center, La Spezia, Italy
H J Callow Norwegian Defence Research Establishment, Kjeller, Norway

1 INTRODUCTION

Synthetic aperture sonar (SAS) images are fundamentally different from sidescan sonar (SSS) images in that multiple pulses are coherently combined in each pixel in the image. This coherent combination increases the along-track resolution in the image, such that SAS systems can have substantially better resolution than SSS. In addition, spatial-frequency (or wavenumber) information is retained in SAS imagery, allowing post-processing techniques not possible in SSS images.

The ability to properly classify an object in a SAS image depends on a number of different factors, as shown in Figure 1. Observation geometry given by range and elevation angle, is important for interpretation of the shadow, signal to noise ratio (SNR) and shallow water performance. Back-scattered target signals can contain elements of specular reflections, diffuse scattering, transparency, resonant scattering and multiple scattering all of which complicate target recognition. Shadow quality is affected by the sonar system and type of SAS processing. Finally, image resolution is critical in resolving or classifying the target, the target shadow and the surrounding scatterers on the seafloor.

By using physical models and signal processing techniques, processing of SAS data may be optimized to highlight specific features in the image. The shadow cast by an object can be corrected. The target itself and the surrounding area can be improved to reduce image degradation due to sub-optimum SAS processing[1]. In addition, wavenumber domain filtering can highlight or suppress specific features in a target.

In this paper, we first list some of the fundamental properties in SAS processing and SAS images. We then show techniques to enhance the target and shadow. We finally describe how these techniques can be included for target recognition.

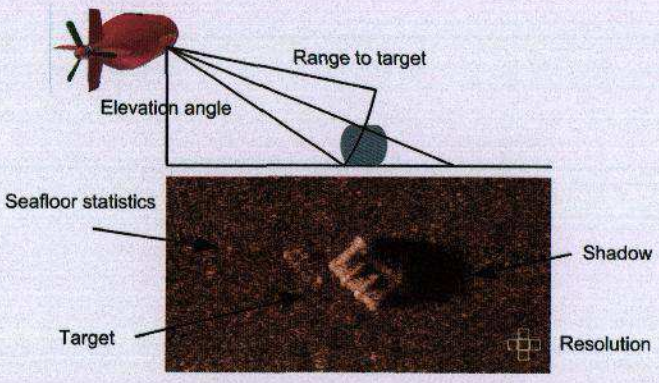

Figure 1 The image quality, and thereby the target recognition performance, is dependent on a number of different factors.

2 PROPERTIES OF SYNTHETIC APERTURE SONAR IMAGES

2.1 Resolution

Previous studies have shown that sonar image resolution is a key driver for classification performance[2]. Achieving high resolution with SSS has typically required the use of high frequencies and/or long arrays. However SSS technology struggles to deliver high resolution at long ranges due to the rapid absorption of high frequency signals[3].

In SAS, the length of the synthetic aperture increases with range providing range independent along-track resolution. Note, however, that SAS image quality is range dependent due to the range variation of the SNR and of the multipath.

A typical high frequency dynamic focused SSS has an along-track resolution of 20-30 cm at 100 m range[4]. A high resolution SAS system such as the MUSCLE[5] or the SENSOTEK[6] has a theoretical (range independent) along-track resolution better than 3 cm. A typical modern mine has a diameter of 0.5 m. This gives 2 independent pixels along-track for the SSS, and more than 15 independent pixels along-track for the high-end SAS. This increase in resolution coupled with the removal of varying resolution across the image, should support the development of higher performance automatic target recognition (ATR) techniques for SAS. The cross-track (or range) resolution is basically no different between SSS and SAS. Wideband systems with frequency modulated pulses can have a range resolution of better than 3 cm.

2.2 Aspect Dependence

The aspect-angle variation an object is observed in, is larger for SAS than SSS. SSS typically has an aspect angle of less than 1°. Conversely, SAS have much wider beamwidths. The MUSCLE 300 kHz SAS has 5.7° beamwidth, while the SENSOTEK system at 100 kHz can have a beamwith larger than 40°. This significant difference in beamwidth between SSS and SAS will potentially affect performance and should therefore be investigated.

There are a number of different features in a SAS image (see Figure 1) that are affected by aspect variation during the collection of the synthetic aperture data:

1. Directional specular reflections (glints). A widebeam system will be able to cover more specular glints. Note that for directional scatterers, the angular spread for which the sonar receives energy from the scatterer, can be the limiting factor.
2. Complicated targets may have features, resonant scatterers and multiple reflections (cavities) that are aspect dependent.
3. Target shadow changes with aspect. Multi-aspect SAS imagery gives the opportunity to produce different, independent images each containing target shadows cast from different directions. This is equivalent to different projections of the silhouette of the target, which can be of benefit to target recognition.
4. The shadow moves according to geometry as shown in Figure 6. This results in shadow fill-in. This effect can be overcome by using shadow enhancement (see section 4).
5. The environment might be aspect dependent. If the seafloor slopes towards shore, multi-path effects become aspect dependent.

There can be an object and geometry dependent limit in the coherent gain and resolution in SAS processing. If this limit is passed, multi-aspect imagery[7] should be considered. The full length synthetic aperture can be divided into subapertures such that individual multi-aspect images can be produced. This is illustrated in Figure 2. This technique is of particular interest for widebeam SAS systems, such as the SENSOTEK SAS. Whilst this provides multi-aspect information, each subaperture image will have lower along-track resolution than the full aperture image. Note that aspect dependence should not be seen as a disadvantage, but rather as an opportunity to provide additional information on the target of interest.

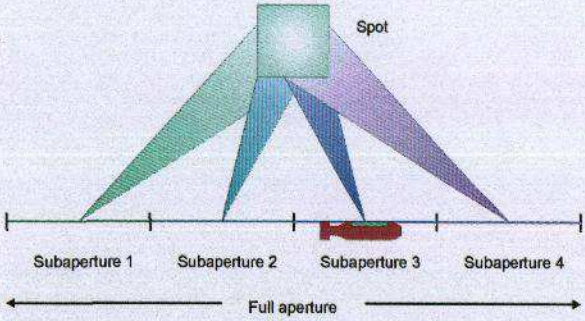

Figure 2 Each target in a SAS image is observed from a range of aspect angles.

2.3 Frequency Dependence

There is an important second order frequency dependence in SAS. For fixed along-track resolution, a higher frequency SAS will have a shorter synthetic aperture than a lower frequency SAS. For a fixed number of receivers in the sonar array, this implies that higher frequency SAS systems have less aspect angle coverage than lower frequency SAS systems. Another significant difference is the vehicle travel time. Lower frequency SAS systems will take a longer time to collect the data in the synthetic aperture build up than higher frequency SAS systems. This might be of significance in navigation accuracy and also in time varying, noise-limited environments.

SAS and SSS have the same frequency dependence regarding the acoustic properties of reflection, transmission, scattering, and absorption. The decoupling of frequency and along-track resolution in SAS gives more flexibility in the sonar design. In SAS, the carrier frequency can be chosen by other criteria than resolution, i.e. absorption. The bandwidth of the system can be used to provide range resolution and/or acoustic properties. It should be noted that the bandwidth also influences the robustness of SAS processing. With a narrow band system, ambiguities are known to trouble correlation techniques for sonar track estimation. Note that it is the bandwidth that provides resolution, while it is the relative bandwidth (bandwidth divided by centre frequency) that is of importance in mapping the acoustic properties.

3 TARGET ENHANCEMENT

Target features may be enhanced through forming SAS imagery in different ways. These techniques have a large potential to be exploited for improving target classification. We demonstrate just one of many possibilities allowed by the coherent nature of SAS imagery.

Figure 3 shows a SAS image from SENSOTEK SAS on HUGIN AUV. The target scene contains an unknown man made object of approximately 2 m in length on a soft-mud seafloor. To the right side of the object is an image artefact, most likely a side or grating lobe of the object response itself. When targets are strong and directional, any sonar may suffer from response sidelobes. Due to their proportionally larger angular coverage, SAS systems are more likely to pick up strong-directional responses and suffer resulting side-lobe structures.

To investigate the angular dependency from the front face of the object of interest, we take the 2D Fourier transform of the image giving the 2D wavenumber spectrum shown in Figure 4. A number of strong periodic features are apparent. These features correspond to specular reflections and give clues as to object structure. The alignment of object glints suggests some a periodic structure along its length.

Figure 3 Man-made object of interest (2 m x 0.3 m) on a soft-mud seafloor. The image size is 40 m x 25 m and the dynamic range is 50 dB. The range to the target is 75 m.

For comparison, the right panel in Figure 4 shows the 2D Fourier transform of the magnitude image. We see a strong directional component in the back-scatter image but information regarding target structure has been lost.

An extension of this analysis allows for glint removal and adaptive filtering in SAS imagery. Figure 5 shows a SAS image of a ship-wreck lying on the seafloor at 180 m depth outside of Horten, Norway. The data was collected by SENSOTEK SAS on HUGIN AUV. The image contains a number of interesting angular dependent features. A large resonance is seen when the imaging angle is perpendicular to the wrecks alignment. This is most likely due to internal reflections within the wreck itself. The resonance, although being an important information source, might cause difficulty in interpreting the traditional backscatter response (in particular filling the shadow region of the wreck with backscatter). Band-stop filtering in wavenumber domain allows for the removal of the specific feature and the target scene enhancement is noted in Figure 5.

The SAS imagery in Figure 5 also shows a number of features which may be caused through imperfect processing. These are particularly apparent in comparing the glint-component with the higher-resolution imagery. The resolution of the (narrow-beam) glint component appears to be better than in the original imagery. This may be due to a number processing induced factors such as inaccurate navigation. As noted earlier the appearance of defocus also could be due to a target dependent resolution limit.

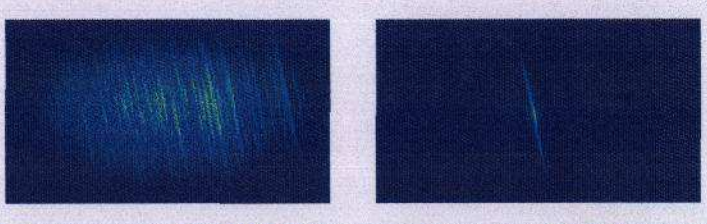

Figure 4 Coherent spectra (left) and magnitude image (right) spectra of the scene shown in Figure 3. The y-axis represents the cross-track wavenumber (related to frequency) and the x-axis represents the along-track wavenumber (linked to angle of arrival). Note the large number of strongly angular dependent scattering lines in the coherent spectrum.

Figure 5 Glint removal through wavenumber filtering. Left: original image. Center: glint removed image. Right: glint component. All images are peak normalized and show 60 dB dynamic range. The glint component corresponds partly to internal reflections inside the wheel-house of the wreck.

Finally, the shadow region in the SAS imagery is somewhat confused. The wreck hull is partly transparent at the frequencies we use. When combined with the multi-aspect illumination, a large amount of energy leaks into the shadow zone. Said another way, on average it is possible to see most of the sea-floor behind the wreck. This feature of SAS imagery is addressed with the shadow enhancement technique outlined in the next section.

4 SHADOW ENHANCEMENT

Shadow is a projection of the silhouette of an object, cast onto the seafloor behind the object, as illustrated in Figure 1. A problem with any SAS system is the eventual blend of shadow and echo, or geometrical shadow fill-in, due to the motion of the illuminating source during the collection of the data in the synthetic aperture. This is illustrated in Figure 6. This leads to a loss in edge definition between shadow and echo from the surrounding seafloor and possible loss of performance for classification based on shadow shape. The problem of geometrical shadow fill-in is most dominant for wide beam SAS systems.

The shadow can be considered as a target moving in opposite direction to the vehicle, as illustrated in Figure 6. This motion can be compensated for by a technique called Fixed Focus Shadow Enhancement (FFSE)[8,9], such that the shadow becomes sharp, and the surrounding seafloor becomes defocused and smeared. Thus shadows from objects other than the object of interest will smear. This secondary effect of FFSE is of particular interest in the case of a seafloor with a topography that causes shadows. An example of interest in mine hunting is a target on a seafloor with sand ripples. We demonstrate the statistical improvement of shadow contour estimation accuracy with FFSE via simulation of 100 SAS images of a target on a sand ripple seafloor.

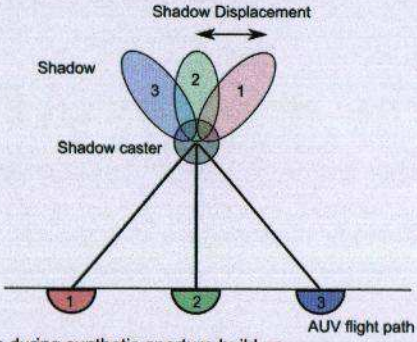

Figure 6 Shadow motion during synthetic aperture build up.

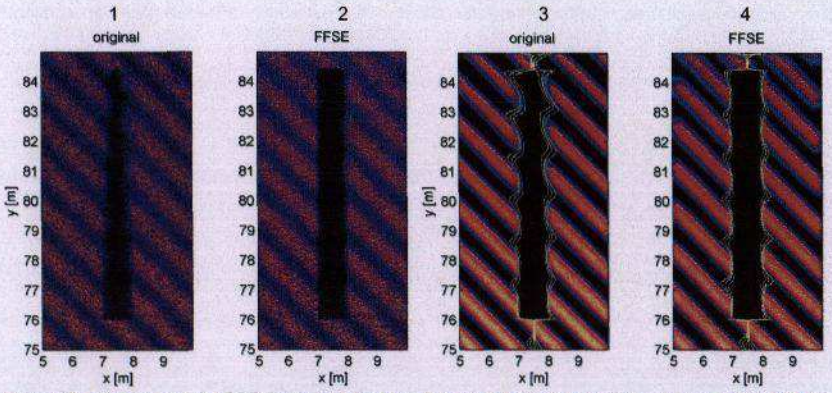

Figure 7 Left image pair: SAS image of simulated target shadow superimposed on a sand ripple region. The color scale used has a range of 40 dB. Panel 1: Original SAS image; Panel 2: FFSE image. Right image pair: Average of the total statistical population generated with 100 SAS images. The mean and standard deviation of the estimated shadow contours are overlaid in white and yellow respectively. Panel 3: Original SAS image; Panel 4: FFSE image.

The simulations involve a 1 m x 1 m plate on the seafloor that is facing the sensor. The sonar settings are those of MUSCLE, operated at an altitude of 10 m and with a frequency band of 270-330 kHz. The target is at a range of 76 m and the synthetic aperture length is 17 m, which enables sufficient data collection to cover the system's horizontal beamwidth of 5°. The bottom consists of 2 cm high and 1 m long sand ripples with a wavelength of 1 m and an orientation of 45° relative to the plate. Panels 1 and 2 of Figure 7 show SAS images from a single simulation. The target echoes are omitted for convenience. Reverberation is simulated as a large number of scatterers on the seafloor, which become insonified according to the shadow of the object and the shadow of the ripples. The scatterers are randomly distributed on the imaged patch of the seafloor and effectively create K-distributed reverberation. Note that in this simulation, we have omitted the sand ripple vertical variation effect on the target shadow itself. This effect is visible in Figure 7 by the far shadow end being exactly horizontal.

It is anticipated that for this ripple case, three effects contribute to better estimation of the shadow contour. Two known defocus effects are the sharper shadow edge and a more Gaussian reverberation distribution[9]. The additional effect we show here is shadow fill-in of ripples at ranges other than the target range. By simulation it is possible to reveal the extent and sensitivity of this third effect.

The result of one simulation is shown in the left image pair in Figure 7. The element level data are processed with original SAS processing (panel 1) and with FFSE (panel 2). Three improvements are clearly visible: the shadow edge is much sharper for FFSE, the defocusing of the reverberation is observable beyond a range of about 78 m, and the ripples blur. Near the plate the ripples are the same, but their shadows are visually affected at longer ranges.

The contour estimation method for the simulations is relatively simple and is based on the average reverberation level (RL). The method starts exactly in the middle of the shadow and sets the right and left bound of the shadow as the first pixel with value higher than RL-3 dB. No 2D information is used. For the statistical analysis, we repeated this scenario with the accompanying shadow contour estimation 100 times. The result is shown in the right image pair in Figure 7, where the average images are shown with the average and standard deviation of the estimated contour. We see that the standard deviation, or shadow estimation accuracy, of FFSE is indeed much better.

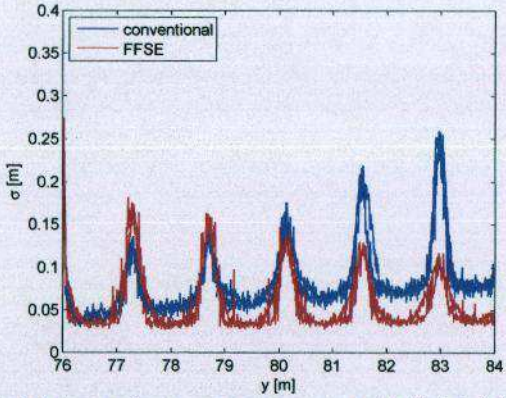

Figure 8 Estimated shadow contour accuracy versus range calculated statistically; two red curves for FFSE and two blue curves for conventional SAS.

Figure 8 shows the standard deviation on the shadow contour estimate. We see that the accuracy is improved by a factor 2.5 at the far end of the shadow. For the same case, but with a flat seafloor, the improvement in terms of accuracy was slightly better than a factor two. The fact that FFSE is marginally worse for ranges close to the target range was found to be artificial, caused by shadow fill-in by sidelobes. To conclude, it can be stated that the shadow fill-in effect is significant, when comparing the ripple and flat bottom case.

These results are preliminary, consisting of only one ripple case. However, they demonstrate another important property of FFSE: target shadow and non-target shadow is easier to seperate by applying FFSE. The result is expected to depend on the ripple characteristics such as height, wavelength and orientation. The ripple shadow fill-in effect is expected to be of less gain for FFSE when the sonar is sailing perpendicular or parallel to the ripples. For the former there is no shadow anyway and for the latter the defocus is expected to fill-in the shadow to a lesser extent. In general, the results here are applicable to a wider variety of cases such as rocks in the vicinity of the target and bottom topography other than ripples.

5 TARGET RECOGNITION STRATEGY

Target recognition (automatic or manual) on SAS images can be achieved in a variety of ways. With reference to Figure 9, we propose the following method to incorporate SAS image enhancement in the target recognition: First, the sonar data are processed to high resolution SAS images and used as input to the detection algorithm. Then, a small area around each detection is reprocessed to provide multi-aspect imagery, target enhancement and shadow enhancement. These results are subsequently fed into the classification algorithm.

If the SAS processing is time-limited, the resolution of the streaming SAS imagery can be reduced for the detection stage. The spot reprocessing must then access the raw data to generate SAS imagery at full resolution. In non-time-limited systems, the original input images can be processed to full resolution and stored with phase information.

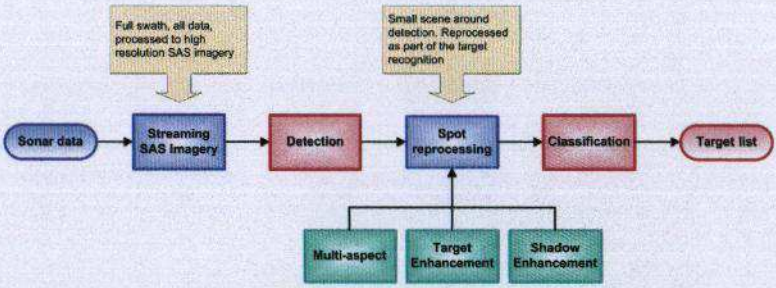

Figure 9 Proposed processing scheme for target recognition including SAS image enhancement.

6 CONCLUSIONS

SAS images are different to SSS images, not only regarding resolution. The information gathered in every pixel in a SAS image, contains information from a number of different aspect angles. The data is also gathered over a period of time up to several tens of seconds. In addition, SAS images may have reduced quality (defocus, ghost targets) due to inaccuracies in the SAS processing. In target recognition, this should be taken into consideration.

We have listed some of the differences between SAS and SSS, and some system dependencies in SAS. Different techniques can be used to enhance the information gathered in order to improve the target recognition capability. We have shown a specific technique to enhance/remove glints in targets. In addition, we have demonstrated that a secondary effect of the Fixed Focus Shadow Enhancement (FFSE) technique is that the surrounding seafloor around the shadow of interest will defocus. We confirm that this has a clear and positive effect in segmenting the shadow in difficult areas such as sand ripple regions with shadow from the ripples blending with the target shadow.

7 REFERENCES

1. R. E. Hansen, H. J. Callow, T. O. Sæbø, P. E. Hagen and B. Langli. Signal processing for the SENSOTEK interferometric SAS: Lessons learned from HUGIN AUV trials. Proc. Synthetic Aperture Sonar and Radar 2006, Lerici, Italy. (2006).
2. V. Myers and M. Pinto, Information theoretic bounds of ATR algorithm performance for sidescan sonar target classification, Proceedings of the SPIE. (2005).
3. X. Lurton, An Introduction to Underwater Acoustics, Principles and Applications, Springer-Praxis. (2002).
4. Klein 5000 product specification, Klein Associates Inc, www.l-3klein.com
5. A. Bellettini, M. A. Pinto and B. Evans, Experimental results of a 300 kHz shallow water synthetic aperture sonar, Proc. Underwater Acoustic Measurements: technologies and results, Crete, F.O.R.T.H., Heraklion. (2007).
6. R. E. Hansen, T. O. Sæbø, H. J. Callow and B. Langli, The SENSOTEK synthetic aperture sonar: Results from HUGIN AUV trials, Proc. European Conference on Underwater Acoustics, Carvoeiro, Portugal. (2006).
7. D. A. Cook, J. T. Christoff and J. E. Fernandez, Broadbeam multi-aspect synthetic aperture sonar, Proc. Oceans 2001 MTS/IEEE, Honolulu, HI, USA. (2001).
8. J. Groen, Adaptive motion compensation in sonar array processing, PhD Thesis, Delft UT/TNO, The Netherlands, www.darenet.nl/en/page/language.view/promise.page. (2006).
9. J. Groen, R. E. Hansen, V. Myers, H. J. Callow and M. A. Pinto, Measurements on shadow contour estimation accuracy in synthetic aperture sonar images, Proc. Underwater Acoustic Measurements: technologies and results, Crete, F.O.R.T.H., Heraklion. (2007).