

INTERFEROMETRIC SYNTHETIC APERTURE SONAR BATHYMETRY MAPS USING ENSEMBLES

S Steele Kraken Robotics, Dartmouth, NS, Canada
R Charron Kraken Robotics, Mount Pearl, NL, Canada

1 INTRODUCTION

Interferometric Synthetic Aperture Sonar (InSAS) can simultaneously measure seabed backscatter and relative seabed depth by utilizing two or more vertically separated receivers. InSAS systems measure the depth relative to the sonar, but true bathymetry is defined as the total water depth, so to produce bathymetry the sonar depth (estimated from pressure sensors or other similar methods) must be added to the relative seabed depth. For simplicity, in this text we will refer to the relative depth maps produced by InSAS systems as bathymetry. Accurate interferometric estimates are dependent on having high signal coherence, where even small reductions in coherence from unity will significantly impact the depth estimate¹. To produce a bathymetry map of the image scene with low vertical uncertainty, multiple estimates of the scene relative depth are required¹. There have been many different methods developed to obtain multiple interferometric estimates and they can typically be divided into two categories: those that lower image resolution (such as spatial averaging and sub-banding) and those that don't, which includes multiple interferometric pairs (three or more vertically separated arrays) and multiple frequency bands. The disadvantage of the methods that exist in the latter category is that they typically require extra hardware¹, which is not desirable as hardware space on SAS platforms is often limited. One exception to this would be multi-pass systems. However, multi-pass interferometry is still an active area of research^{2,3,4}. It should also be noted that while a multipass interferometric SAS does not require additional hardware, it does impose a significant increase in terms of survey area by requiring the same area to be imaged more than once. The method is also not ideal in highly dynamic seabed regions where temporal coherence of the SAS imagery will rapidly decrease. SAS imaging systems that cannot incorporate additional hardware face a trade-off between horizontal resolution and vertical uncertainty/resolution due to spatial averaging or sub-banding the imaging bandwidth. Centimetric resolution in all three dimensions is desirable for a variety of applications such as mine hunting⁵ or hydrographic surveys where strict requirements of resolution must be met⁶. The work presented here demonstrates a new approach that allows us to achieve centimetric resolution interferometry estimates in all three dimensions without requiring additional hardware or multiple passes. The method utilizes ensembles of observations from numerous InSAS image strips with overlapping estimates in the image scene. The method leverages a continuous SAS image mosaicing process which will be briefly described in Section 2.1. The rest of this paper will present a theoretical basis for the ensemble interferometry method and will demonstrate its effectiveness on imagery produced by a Kraken Robotics AquaPix® Miniature Interferometric Synthetic Aperture Sonar (MINSAS).

2 METHODS

2.1 Seamless SAS Processing

Traditionally, real-time SAS imagery is produced using beamformers that partition the imaging space into large rectangular blocks that extend tens of metres in the along track direction with a small fraction of spatial overlap between immediate neighbour blocks. An alternative, albeit less efficient method, is to process data on a ping-to-ping basis with the image space partitioned into the narrowest rectangular blocks that SAS processing will allow. The resulting image strips significantly overlap immediate neighbours, enabling the introduction of an image mosaicing stage capable of creating continuous

SAS images of any desired length (up to one survey leg) in the along-track direction with no visible seams. With this approach, one starts by creating memory buffers that will hold a predetermined number of backscatter returns and associated metadata. The minimum number of pings to buffer is the minimum number of pings required to form the appropriate synthetic aperture length for the desired along track resolution and across track range. The algorithm reads in one ping at a time until the buffers are full, at which point the typical SAS processing steps of implementing the Displaced Phase Centre Antenna (DPCA) technique and beamforming are commenced. However, instead of beamforming a large along-track block, we beamform a narrow image strip (Fig. 1) and save it for mosaicing later. On each ping, the buffers are updated, dropping the oldest ping and replacing it with the newest ping, and repeating the process of forming SAS image strips until the image space is fully rendered. As shown in Fig. 1, each imaging strip has significant overlap with the previous strip, leading to multiple observations of the same image pixels, which we can exploit for improved interferometric processing (Section 2.2).




Figure 1: Sketch demonstrating the Seamless SAS imaging concept. The gridded region represents the image scene. Synthetic apertures (SA) 1 and 2 are the same length but SA 2 drops the earliest ping from SA 1 and replaces it with the most recent ping. The imaging area of SA 1, strip 1 (dark grey), has significant overlap with imaging area of SA 2, strip 2 (black dashed lines).

Once the SAS strips are completed a mosaic of the entire imaging area can be assembled. One can start by superposing a grid onto the image scene. The vehicle Inertial Navigation System (INS) data defines the ground area represented by each image strip. As mentioned previously, there is significant redundancy between the strips in that multiple strips can be used to populate the same ground plane pixel. There are a variety of approaches available for deciding which strip contributes to an individual ground plane pixel. For instance, one can use local aspect angle, or illumination metrics, or quality metric associated with the micronavigation process to guide the decision. Our interest here is not so much the image mosaic details, rather our focus is the interferometric process, wherein we can now utilize repeated observations of the same ground plane pixel to compute an ensemble estimate of the seabed bathymetry.

2.2 Interferometric Processing

Relative seabed depth can be calculated from the phase-difference map (interferogram), which can be estimated with the 2D zero-lag complex cross-correlation, γ , of co-registered signals a and b , from interferometric receiver number one and two, respectively⁷,

$$\gamma = \frac{E\{ab^*\}}{\sqrt{E\{|a|^2\}E\{|b|^2\}}}. \quad (1)$$

The argument of γ is the maximum likelihood estimate of the phase difference $\theta = \arg\{\gamma\}$ and the degree of coherence is $\mu = |\gamma|$. The expected values can be obtained through ensemble averages of multiple observations for each pixel⁸; however, synthetic aperture systems typically only observe each pixel once and thus most literature assumes the observations are stationary and the ensemble averages can be replaced with spatial averages over rectangular windows of $N \times M$ pixels surrounding the pixel of interest. The SAS processing method of Section 2.1 provides multiple observations of

each ground plane pixel and while the number of repeat observations is not sufficient to eliminate the need for spatial averaging, the spatial and ensemble averaging methods can be combined such that the number of pixels required for spatial averaging is reduced. The combined spatial and ensemble coherence is estimated as

$$\hat{\gamma} = \frac{\sum_{k=1}^K \sum_{n=1}^N \sum_{m=1}^M a_k[m, n] b_k^*[m, n]}{\sqrt{\sum_{k=1}^K \sum_{n=1}^N \sum_{m=1}^M |a_k[m, n]|^2 |b_k^*[m, n]|^2}}, \quad (2)$$

where K is the total number of repeated SAS observations of a given pixel by the degree of overlap afforded by the particulars of the imaging method (Section 2.1), and a_k and b_k are the k th observations of a given pixel from interferometric receiver number one and two, respectively. Since the samples are correlated, the effective number of observations is likely less than the number of observations used in Eq. 2. As described in Section 2.3, we will investigate the effective number of looks. From Eq. 2, one can choose between reducing vertical uncertainty or improving the resolution of the bathymetry estimate. The vertical uncertainty is effectively the variance of the interferogram (σ_θ^2) converted to depth. Converting the standard deviation of the time delay estimate⁷ to variance, the interferogram variance can be estimated as

$$\sigma_\theta^2 = \frac{3}{\pi^2 B^2 L} \left[\frac{1}{\text{SNR}} + \frac{1}{2\text{SNR}^2} \right], \quad (3)$$

where B is the bandwidth, L is the number of samples in the estimate, SNR is the signal to noise ratio, and f_c is the centre frequency of the signal. The SNR is derived from the degree or strength of coherence (μ) as $\text{SNR} = \frac{\mu}{1-\mu}$. The variance of the interferogram can be converted to relative depth variance (vertical uncertainty, σ_z^2) with

$$\sigma_z^2 = \frac{\lambda}{2\pi} \frac{r}{D} \frac{\cos(\phi + \phi_0)}{\cos \phi} \sigma_\theta, \quad (4)$$

where λ is the signal wavelength, r is the slant range to the point of interest, D is the interferometric vertical baseline, ϕ is the depression angle from the sonar to point of interest on the seabed, and ϕ_0 is the array tilt relative to vertical.

2.3 Effective Number of Looks

Due to lack of independence between samples, SAS interferometry estimates often have an effective number of looks less than the nominal number of looks. This applies to both spatial and ensemble averaging techniques. The complex nature of determining sample independence makes it difficult to determine analytically. We follow a procedure utilized in synthetic aperture radar (SAR) to estimate the effective number of looks⁹. This procedure consists of comparing a model of the probability density function (pdf) of the interferometric phase to the interferometric phase pdf derived from the interferogram statistics using the Chi-square goodness-of-fit test¹⁰. The multi-look phase difference pdf is parameterized by the number of looks and the coherence level¹, with the pdf trending towards a Dirac-delta function as the number of looks or coherence is increased. The multilook pdf for L independent looks can be expressed as¹¹

$$\text{pdf}(\theta, L) = \frac{\Gamma\left(L + \frac{1}{2}\right) \left(1 - |\gamma|^2\right)^L |\gamma| \cos(\theta - \theta_0)}{2\sqrt{\pi} \Gamma(L) \left(1 - |\gamma|^2 \cos^2(\theta - \theta_0)\right)^{L+\frac{1}{2}}} + \frac{\left(1 - |\gamma|^2\right)^L}{2\pi} {}_2F_1\left(L, 1; \frac{1}{2}; |\gamma|^2 \cos^2(|\theta - \theta_0|)\right), \quad (5)$$

where, θ_0 is the true phase difference, θ is the estimated phase difference, $\Gamma(\cdot)$ is the gamma function, and ${}_2F_1$ is the hypergeometric function. Assuming the coherence is fixed, we can generate pdfs over a large range of number of looks and compare them to the binned phase distribution from a patch of seabed. The Chi-square goodness-of-fit test can then be used to determine if there is a significant difference between the observed and modeled distribution. Here, we use a significance

level of 0.05, and thus on average we can expect a well fitting distribution to have a failure rate up to 20 %. The Chi-square goodness-of-fit test was computed over hundreds of square seabed regions of 3 m² obtained from a generally featureless seabed image with an along-track length of 294 m. While the original imaging region has a maximum across-track range of 160 m, we discarded seabed patches beyond approximately 122 m range across-track to avoid issues with the signal degrading at long range impacting the phase estimates (and thus biasing the estimate of the effective number of looks). The pdf model assumes a single coherence value. To accommodate for this in our real data that has variable coherence values, we will filter our phase estimates to only use coherence values between 0.96 and 0.98, giving an average coherence of about 0.97. These coherence values were chosen as the majority of the pixels are within this range. Our analysis of the ensemble method effective number of looks is complicated by the non-independence of the spatial averaging technique. To help control for this, we performed the number of looks analysis with the spatial averaging technique alone ($K = 1$) and the combined spatial and ensemble averaging technique ($K = 3$).

3 RESULTS AND DISCUSSION

We demonstrate our results on MINSAS data collected from the NOAA Okeanos Explorer Vessel on a region featuring a shipwreck (Fig. 2, top) with many vertical structures well above the seabed and lots of clutter (small individual objects) surrounding the wreck. This sample image was chosen because it allows us to visually assess the performance over rapidly varying height estimates and on clutter that require centimetric resolution to observe. For estimating the effective number of looks, we have chosen a large region of flat featureless seabed near the wreck. Typically, the MINSAS bathymetry is averaged over a local neighbourhood of $N = 8$ and $M = 16$ samples in the along and across track directions, respectively (MINSAS resolution is 3 cm along track and 1.5 cm across track), yielding 24 cm horizontal resolution in both directions. The processing method described in Section 2.1 provides $K = 3$ observations for each ground pixel, which we supplement by averaging over at least $N = 2$ pixels in the along track direction and $M = 4$ pixels in the across track direction. With this technique we expect we can maintain similar vertical uncertainty while improving the bathymetry spatial resolution up to 6 cm. Comparing the middle (24 cm resolution) and bottom (6 cm resolution) bathymetry maps in Fig. 2, the bathymetric resolution improvement is obvious; the wreck is better focused and there are numerous objects that can be observed in the high resolution bathymetry map that cannot be observed in the low resolution bathymetry map.

To properly interpret the vertical uncertainty associated with the high resolution bathymetry, we need to first determine the effective number of looks the ensemble method achieves. However, our analysis is complicated by the fact that the spatial averages that we combine with the ensemble averages are likely not completely independent themselves and should be considered in our analysis. At all spatial resolutions tested, the portion of measured pdfs significantly different from that expected from the nominal number of looks (expressed in the left side image of Fig. 3 as the portion of nominal looks) does not exceed the significance threshold. Thus, we consider the spatial averaging effective number of looks to be equivalent to the nominal number of looks. When spatial and ensemble averaging are combined, we have found the effective number of looks is dependent on the spatial resolution, where lower spatial resolution results in a decrease in the effective number of looks (Fig. 3, right). This result should be expected because, as we average over a wider spatial area, our ensembles become decreasingly independent. A similar trend occurs in the spatial averaging looks analysis where, as the spatial resolution degrades, the portion of rejected pdfs increases (Fig. 3, left). Additionally, we have found that at resolutions worse than 12 cm ($N = 4$) the minimum number of rejected pdfs occur at contributions less than the number of nominal looks. The minimum occurs at around 0.75 and 0.31 independent contributions (in terms of proportion of nominal looks) for 15 cm ($N = 5$) - 21 cm ($N = 7$) and 24 cm ($N = 8$) resolutions, respectively. While no statistically significant difference between the nominal and effective spatial averaging looks was measured, it is possible that the looks used in the spatial averaging are not completely independent and may be contributing to the reduction of independent looks at lower resolutions when both spatial and ensemble averaging are applied.




Figure 2: Comparison of MINSAS intensity image [top], 24 cm resolution bathymetry ($N = 8$, $M = 16$, $K = 1$) [middle], and 6 cm resolution bathymetry ($N = 2$, $M = 4$, $K = 3$) [bottom]. Bathymetry ranges from 23 m (dark blue) to 13 m (red) depth below sonar platform. The mosaic is 144 m in the along track (x-axis) and 75 m in the across track (y-axis).




Figure 3: Effective number of looks for spatial resolutions of 6 to 24 cm with spatial averaging applied (left) and combined spatial and ensemble averaging applied (right). Here, the total number of looks in the estimate is presented as $L = N \times M \times K$.

VU (cm)	24 cm		18 cm		12 cm		6 cm	
	$K = 1$	$K = 3$	$K = 1$	$K = 3$	$K = 1$	$K = 3$	$K = 1$	$K = 3$
mean	10	10	11	10	13	11	13	9
median	3	2	4	2	6	3	10	6
max	232	311	208	305	233	157	64	75

Table 1: Comparison of shipwreck image vertical uncertainty (VU) statistics for $K = 1$ and $K = 3$ ensembles at four different spatial resolutions.

We have found that at spatial resolutions greater than approximately 15 cm ($N = 5$) the ensemble technique does not provide many additional independent samples. The minimum of the Chi-square goodness of fit test curve for 24 cm resolution ($N = 8$) is located at $K = 0.58$ and crosses the significance threshold at $K = 1.23$ (Fig. 3, right), indicating the effective number of looks is between 0.58 and 1.23 and thus no extra ensemble looks are achieved. At 21 cm resolution ($N = 7$), the effective number of looks is between 0.61 and 1.7, and thus we can expect the ensembles to contribute very little at this resolution. Similarly, at 18 cm resolution ($N = 6$), the effective number of looks is between 1.1 and 2.9, and thus one may observe a small improvement in the effective number of looks. Once the resolution improves to 15 cm ($N = 5$) or better, we see a change in the shape of the effective number of looks curve from hockey stick (relatively flat followed by steep curve upward) to a more gentle arc. This may be an indicator that the spatial resolution is high enough to allow the ensembles to contribute some independent looks. At 15 cm resolution we can expect an effective number of looks between 1.8 and 4.7, and thus we should expect to see some impact on our bathymetry estimates. For $N = 2 - 4$ the curves do not cross the significance threshold; this likely has to do with the lack of significant difference between the modeled interferometric phase pdf for the different looks of interest, rather than a property of the measured interferometric phase at these resolutions. Thus, below $N = 5$ we will consider the curve minima alone to be the effective number of looks. At 12 cm ($N = 4$) the effective number of looks is greater than two (2.2) and thus we can expect to observe a small improvement in our interferometry estimate. At 6 cm and 9 cm ($N = 2 - 3$) resolution we find the effective number of looks to be quite similar ($N = 3$ not plotted for figure visibility), with a curve minima occurring at $K = 2.9$ looks for both. This indicates that at centimetric spatial resolutions the effective number of looks is almost equivalent to the number of looks used (three).

As predicted by our analysis of the effective number of looks, at a spatial resolution of 24 cm, using en-

VU (cm)	24 cm		18 cm		12 cm		6 cm	
	$K = 1$	$K = 3$	$K = 1$	$K = 3$	$K = 1$	$K = 3$	$K = 1$	$K = 3$
mean	4	2	5	3	7	4	11	7
median	2	1	3	2	5	3	8	5
max	170	190	144	185	134	176	63	68

Table 2: Comparison of featureless seabed image vertical uncertainty (VU) statistics for $K = 1$ and $K = 3$ ensembles at four different spatial resolutions.

sembles provides no value in terms of introducing independent looks, and thus even with ensembles the vertical uncertainty does not significantly improve (Table 1). Consistent with the effective number of looks analysis, the vertical uncertainty improves slightly (mean reduced by 1 cm, median reduced by 2 cm) for spatial resolutions ranging from 21 cm to 15 cm (18 cm spatial resolution vertical uncertainty presented in Table 1). At 6 cm spatial resolution we can achieve centimetric vertical resolution using ensembles, reducing the mean and median vertical uncertainty by 4 cm (Table 1). Since we achieve fewer independent ensemble looks as we increase the spatial averaging, the vertical uncertainty improvement is slightly less at 12 cm than at 6 cm. At 12 cm resolution the vertical uncertainty mean and median was reduced by 2 and 3 cm respectively (Table 1). To test if we can also expect to achieve centimetric resolution in all three dimensions on featureless (mostly speckle) seabeds, we computed the vertical uncertainty for the featureless seabed image used to estimate the nominal number of looks in Section 2.3. As shown in Table 2, we can achieve 7 cm vertical resolution at 6 cm spatial resolution on a featureless seabed. We observe a much more consistent relationship between the number of looks and vertical uncertainty with the featureless seabed image than the shipwreck image. This is likely caused by the dominance of rapidly varying seabeds and shadow regions in the shipwreck image, which may be biasing or limiting the achievable vertical uncertainty. For both the featureless and shipwreck images, the ensemble method appears to increase the sensitivity to outliers, causing the maximum vertical uncertainty to increase at all spatial resolutions. The disparity of maximum vertical uncertainty is the least severe at the 6 cm spatial resolution; we believe this is because reduced spatial averaging may reduce variance on rapidly varying seabeds and shadow regions. We observe evidence of this in the shipwreck vertical uncertainty images (Fig. 4): the regions around the wreck with rapid changes in bathymetry and in shadow areas are where the vertical uncertainty is the highest. The vertical uncertainty in these regions increases as the spatial resolution degrades, regardless of the application of ensemble averaging.

4 CONCLUSION

InSAS requires multiple interferometric estimates to generate accurate bathymetry maps, which is often achieved through spatial averaging. This spatial averaging sacrifices spatial resolution for vertical resolution/uncertainty. Here, we described how redundancy in SAS processing and ensemble averaging can be leveraged to produce bathymetric maps with centimetric resolution in both the horizontal and vertical directions. We demonstrated this technique is capable of producing centimetric resolution in all three dimensions on MINSAS data consisting of both a featureless (primarily speckle) image and a much more complex image featuring a large object (shipwreck). The improved bathymetry spatial resolution from 24 cm to 6 cm allowed us to detect more clutter and reduce the vertical uncertainty in regions with rapidly varying seabed heights. Our analysis showed the effective number of looks obtained from the ensemble technique are dependent on the amount of spatial averaging used. For the MINSAS, spatial averaging should be limited to $N = 5$ (15 cm) or less for the ensembles to provide an independent estimate. At centimetric resolutions we found the independent number of looks to be roughly equivalent to the nominal number of looks.

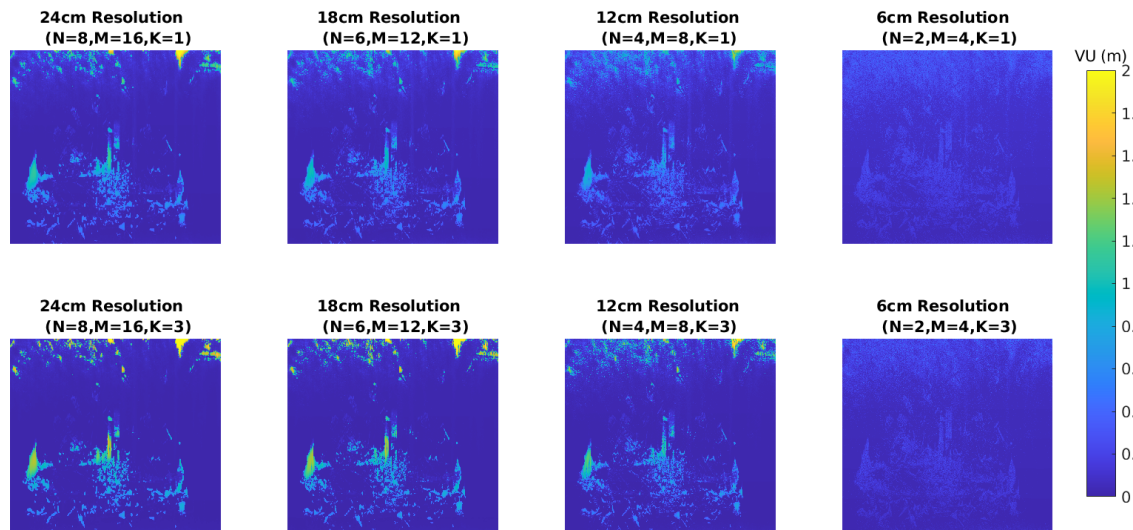


Figure 4: Vertical uncertainty without (top) and with (bottom) ensembles at 24, 18, 12, and 6 cm resolution.

5 REFERENCES

- [1] Philip J. Barclay. *Interferometric synthetic aperture sonar design and performance*. PhD thesis, University of Canterbury, Christchurch, New Zealand, 2006.
- [2] Jeremy Dillon and Vincent Myers. Baseline estimation for repeat-pass interferometric synthetic aperture sonar. In *EUSAR 2014; 10th European Conference on Synthetic Aperture Radar*, 2014.
- [3] Jeremy Dillon and Vincent Myers. Coherence estimation for repeat-pass interferometry. In *2014 Oceans - St. John's*, 2014.
- [4] Ruggero De Paulis, Claudio Maria Prati, Silvia Scirpoli, Fabio Rocca, Alessandra Tesei, P. A. Sletner, Stefano Biagini, Piero Guerrini, Francesco Gasparoni, Cosmo Carmisciano, and Marina Locritani. Sas multipass interferometry for monitoring seabed deformation using a high-frequency imaging sonar. *OCEANS 2011 IEEE - Spain*, pages 1–10, 2011.
- [5] Franck Florin, Franck Fohanno, Isabelle Quidu, and Jean-Philippe Malkasse. Synthetic aperture and 3d imaging for mine hunting sonar. *Undersea Defence Technology (UDT) Europe 2004, Nice France*, 2004.
- [6] International hydrographic organization standards for hydrographic surveys s-44 edition 6.0.0, 2020.
- [7] Torstein Olsmo Saebo. *Seafloor Depth Estimation by means of Interferometric Synthetic Aperture Sonar*. PhD thesis, University of Tromso, Tromso, Norway, 2010.
- [8] Ramon F. Hanssen. *Radar interferometry: Data Interpretation and error analysis*. Kluwer, 2001.
- [9] L.R. Joughin and D.P. Winebrenner. Effective number of looks for a multilook interferometric phase distribution. *Proceedings of IGARSS '94 - IEEE International Geoscience and Remote Sensing Symposium*, 1994.
- [10] Jay L. Devore. *Probability and statistics for engineering and the Sciences*. Brooks/Cole, Cengage Learning, 8 edition, 2012.
- [11] Richard Bamler and Philipp Hartl. Topical review: Synthetic aperture radar interferometry. *Inverse Problems*, 14:1, 07 1998.