

A GPU-Accelerated Dynamically Focused Multibeam Sidescan Beamformer

Shannon-Morgan Steele
Kraken Robotic Systems Inc.
Mount Pearl, NL, Canada
ssteele@krakenrobotics.com

Jeremy Dillon
Kraken Robotic Systems Inc.
Mount Pearl, NL, Canada
jdillon@krakenrobotics.com

Abstract—Sidescan sonar is a popular tool for seabed imaging. To produce a well focused image, corrections for the sonar bistatic geometry and near-field must be implemented during beamforming. To be computationally efficient enough to be processed in real-time, sidescan beamforming is typically performed in the frequency domain. However, the focusing corrections require dynamic focusing, meaning the corrections are range dependent and thus more suitable for time domain beamforming (TDBF). TDBF is computationally expensive and has previously been considered not well suited for real-time applications. We demonstrate how the highly parallel nature of Graphical Processing Units (GPUs) can be exploited to produce well focused sidescan images in real-time. Experimental results from Kraken’s AquaPix interferometric synthetic aperture sonar will be presented.

Index Terms—Sidescan, Sonar, Seafloor Imaging, Beamforming, Signal Processing

I. INTRODUCTION

Seabed imaging is a vital component for a variety of applications including geologic and habitat mapping, hydrography, pipeline surveys, and target detection. The significant propagation ranges of acoustic waves make them an excellent tool for imaging the seabed. Sidescan sonar is a well established acoustic seabed imaging technique. Sidescan sonars are typically mounted on the side of a towfish, remotely operated vehicle (ROV), or autonomous underwater vehicle (AUV), with a sonar mounted on both the port and starboard side of the vehicle. Sidescan sonar transmits and receives on a narrow strip of the seafloor at broadside (perpendicular to the direction of travel). The sidescan beam has a constant angular aperture, meaning the along-track resolution degrades with range.

Synthetic aperture sonar (SAS) is well suited to seabed survey operations because it can achieve rapid area coverage rates at constant, high (centimeter-scale) resolutions, with enhanced SNR [1]. SAS achieves these gains by coherently integrating successive pings as a small physical array (the sidescan sonar) moves to synthesize a larger array. Although some sidescan sonars are able to achieve constant resolution, they do so at the cost of significantly reduced resolution in the near range.

As observed in Figure 1, the Kraken Robotic Systems Inc. (Kraken) SASVIEW software displays a real aperture sidescan image for each ping, and a 50 m along track SAS processed image tile (for both port and starboard) in real-time. SAS processing requires platform motion to be estimated with sub-

wavelength accuracy. Such accuracy is not achievable with commercially available navigation systems. Instead, SAS relies on data driven micronavigation techniques such as DPCA to produce well focused images. Operating in adverse conditions such as high vehicle instability, strong multipath, or severe noise interference, can make SAS processing challenging. In cases where SAS processing fails, the software will produce a 50 m sidescan image tile in lieu of the SAS tile. In such instances, it is vital that the software provides a well focused sidescan image in real-time. Producing such an image requires the implementation of dynamic focusing in the sidescan beamformer to correct for the range dependency of the array geometry and near-field corrections.

Due to the trade-offs between beamformer robustness and computational load, dynamically focusing a sidescan image in real-time is non-trivial. In this paper, we demonstrate the advantage of time domain beamforming (TDBF) over Fourier domain beamforming (FDBF) for dynamic focusing. We then establish how the time domain beamforming is applied in real-time using GPU acceleration. The new beamformer is tested on real acoustic returns collected by two different Kraken AquaPix SAS sensors: the Interferometric Synthetic Aperture Sonar (InSAS) and the Miniature Synthetic Aperture Sonar (MINSAS). Differences between the unfocused and focused sidescan real-time processing speeds and resulting imagery are analysed.

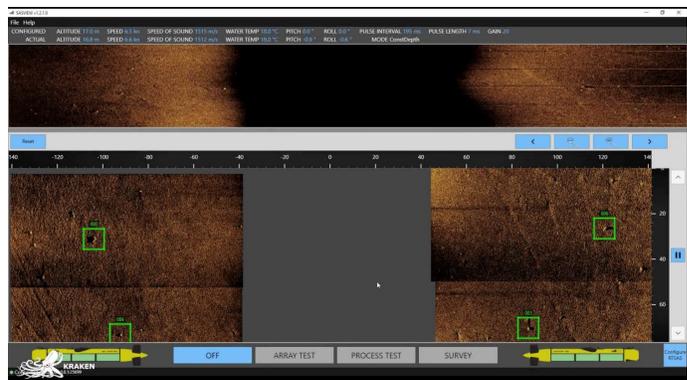


Fig. 1. Screenshot of Kraken’s SASVIEW software. Top image is real-time sidescan images. Bottom is SAS processed image tiles.

II. MULTIBEAM SIDESCAN SONAR

Traditional sidescan forms a single beam during a receive cycle. Some modern sidescans digitally steer the receiver to form multiple beams on the seabed. Forming multiple beams allows the sonar platform to move faster without causing gaps in seabed coverage in the near range. In addition to multiple beams, many modern sidescan sonars produce constant along track resolution images using a dynamic aperture. This means that as time progresses in the received signal, the image formation process uses a larger portion of the physical aperture. The consequence of this method means that the entire image will have the minimum achievable angular resolution, rather than the maximum. Alternatively, we have chosen to form many beams over a swath width that covers the entire along track distance travelled between each ping ($\sim \pm 2$ degrees or less). This allows us to beamform onto a grid to maximize the resolution and produce images with pixels having the same dimensions in the along and across track directions (3×3 cm). Of the many methods available for digital steering, the simplest and most common form is the delay-and-sum beamformer, which can be implemented in either the time or frequency domain.

A. Time Domain Beamforming

One of the simplest forms of delay-and-sum beamforming is TDBF. In a conventional TDBF, the beamformer output $b(t)$ is given by (1), where N denotes the number of receivers in the array, x_n is the signal output for receiver n , w_n represents the weighting applied to each receiver, and τ_n denotes the time delays required for each receiver to steer the beam in the desired direction. For the case of broadside (traditional sidescan), there is no steering required and thus $\tau_n = 0$ for all receivers.

$$b(t) = \frac{1}{N} \sum_{n=1}^N w_n x_n(t - \tau_n) \quad (1)$$

On reception, the incoming wave front propagates in an arc. The further the receiver is from the target, the larger the diameter of the ring. Thus, when a target is far away from the receiver the wave front has a low radius of curvature and becomes approximately planar. In this far-field (Fraunhofer zone), the propagating signal is approximated as a plane wave. For a steered sidescan geometry, the delays τ_n of each receiver become proportional to the projection of the sensor position vector r_n , relative to a reference point [2]. Assuming the sound speed c and receiver positions d_n relative to the receive centre are known, (2) can be used to compute the steering delay for the desired steering angle θ [3]. Each receiver's signal must be interpolated to the calculated steering delay before summation. Even with simple linear interpolation, this step becomes quite time consuming when there are many receivers used to form the sidescan image.

$$\tau_n = \frac{d_n}{c} \sin \theta \quad (2)$$



Fig. 2. KATFISH, Kraken's actively controlled towfish, equipped with the MINSAS180.

B. Frequency Domain Beamforming

Delay and sum beamforming can also be performed in the frequency domain, based on the fact that a delay in the time-domain corresponds to a frequency-domain phase shift [4]. The Fourier transform of each received signal produces the frequency spectrum for each receiver, $X_n(f)$, where f is the frequency. In the frequency domain, each individual receiver spectrum can be corrected with its respective steering delay τ_n via (3). The corrected spectra are then transformed back to the time domain using an inverse Fourier transform.

$$b(t) = \mathcal{F}^{-1} \left[\frac{1}{N} \sum_{n=1}^N w_n X_n(f) e^{-i2\pi f \tau_n} \right] \quad (3)$$

The steering delays can be calculated in the same manner as the TDBF. The advantage of the FDBF is that it is far more computationally efficient, as it does not require interpolation. This computational efficiency makes the FDBF one of the most commonly used sidescan beamforming methods.

III. DYNAMIC FOCUSING

A. Near-field Correction

For the plane wave approximation used in the previous section to be considered appropriate, the far-field condition must be met. To meet the far-field condition, the wave front disparity δ_r across the aperture due to wave front curvature must be less than $\frac{1}{8}$ of a wavelength λ , often termed the Fraunhofer Zone [5]. The wave front disparity can be obtained from the range R to the location of interest and the length L of the receive aperture (array) using

$$\delta_r(t) = R(t) \left[\sec \left(\tan^{-1} \frac{L}{2R(t)} \right) - 1 \right]. \quad (4)$$

The Fraunhofer far-field requirement can also be expressed as a separation requirement between target and receiver

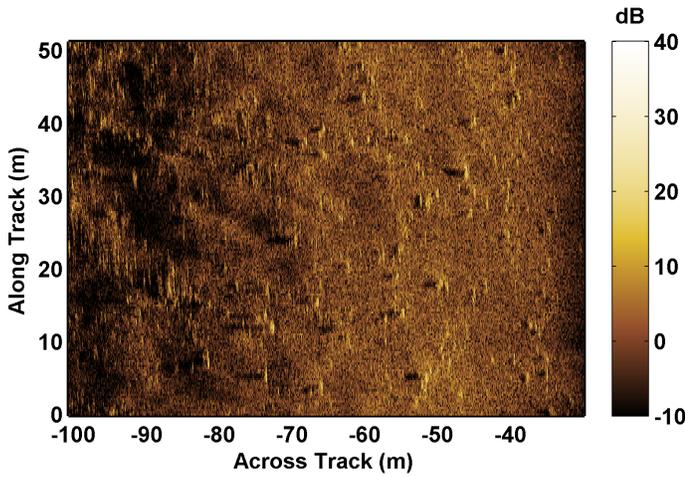


Fig. 3. Baltic Sea unfocused sidescan image.

$$R \geq \frac{4(\frac{L}{2})^2}{\lambda}. \quad (5)$$

For many sidescan sonars, almost the entire imaging range of interest is in the near-field. For example the KATFISH, Kraken’s actively controlled towfish (Fig. 2), is equipped with a MINSAS180, which has a receive aperture length of about 1.6 m. The MINSAS centre frequency wavelength is 4.45 mm, and thus the MINSAS180 operates in the near-field until about 569 m.

The near-field correction is simple to include in TDBF, where an additional delay $\tau_{df}(t)$ can be applied in (1). The δ_r changes for each time or range step in (4) and thus the image must be dynamically focused with increasing slant range. In the time domain, appropriate delays can be applied to each individual element and range of interest using

$$\tau_{df}(t) = \frac{R(t)}{c} \left[\sec \left(\tan^{-1} \frac{d_n}{R(t)} \right) - 1 \right]. \quad (6)$$

For the TDBF, dynamically focusing the image with range is simple. This is not the case for the FDBF as we no longer have a time or range component of the data once in the frequency domain. To apply range dependent focusing in the frequency domain, we would need to utilize block FDBF. To perform block FDBF, the data must be divided into separate blocks of data based on range and (3) is applied to each data block separately with a different τ_{df} correction. In order to get reasonable dynamic focusing range resolution, we would need to divide the data into many blocks, thus eliminating the speed advantage of the FDBF.

B. Bistatic Geometry Correction

Sidescan sonar often has a bistatic configuration, meaning there is a separation between the transmitter and each receive element. Many sidescan sonar processing algorithms will approximate the bistatic configuration as monostatic by either assuming co-location (the transmitter and receiver in the same

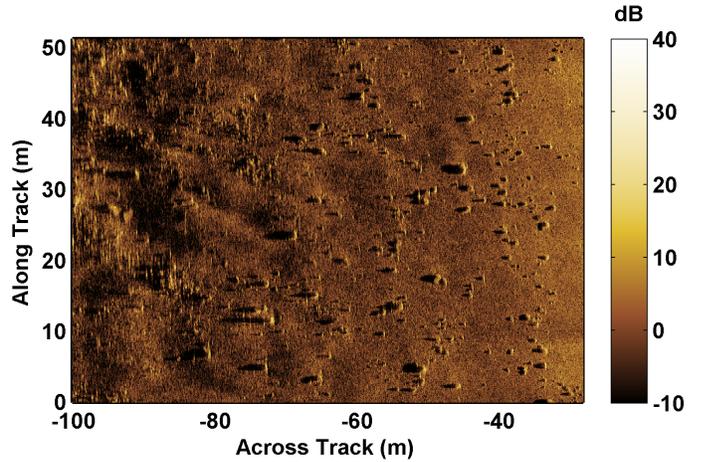


Fig. 4. Baltic Sea dynamically focused sidescan image.

location) or using the phase centre approximation (PCA). PCA assumes that transmission and reception occur from the phase center $C = (T_x + R_c)/2$, where T_x is the transmission location and R_c is the location of the receiver center [6]. Many sidescan processing algorithms also further simplify the bistatic geometry by using the “stop-and-hop” model [7]. This model assumes the sonar does not move between signal transmission and reception. However, the platform is continuously moving, and this motion will cause the distance between the transmit location and each receiver location to change over the receive cycle.

To achieve a diffraction limited image, the bistatic geometry must be accounted for during beamforming. Similar to the near-field correction, bistatic geometry correction is range dependent and must be applied to each receive element individually. Thus the bistatic correction is also most easily applied with TDBF. As discussed previously, TDBF is rarely used in sidescan sonar due to the costly nature of the required interpolations. In the next section we will discuss how time domain backprojection, an implementation of TDBF, can produce focused images in real-time with GPU acceleration.

C. GPU Acceleration

The most intuitive method to implement both the near-field and bistatic corrections is time domain backprojection because it can handle an arbitrary array geometry [8]. In time domain backprojection, we compute the range from the transmitter to the pixel and the pixel to each receiver for each pixel position in the image. The received signals are then interpolated to the corresponding range and the value found is summed across all receive channels. Despite its simplicity, time domain backprojection is rarely used in sidescan sonar because the number of required operations is proportional to N^3 [9], where N is the number of receivers in the aperture.

The highly parallel structure of a GPU makes it significantly more efficient than a central processing unit (CPU). Kraken produces SAS imagery and bathymetry in real-time through hardware acceleration on embedded NVIDIA GPUs [10].

TABLE I
IMAGE ACQUISITION PARAMETERS AND PROCESSING TIMES FOR SAMPLE DATA SETS

	Vehicle Altitude (m)	Along Track Swath (m)		# of data samples	Data Acquisition Time (s)	Processing Time (s)	
		Min	Max			Unfocused	Focused
Baltic Sea	10	32	100	1.20×10^8	32.68	0.52	0.67
Bedford Basin	15	42	145	7.70×10^7	16.67	0.60	0.71
Narragansett Bay	18	34	194	1.13×10^8	31.25	1.27	1.40

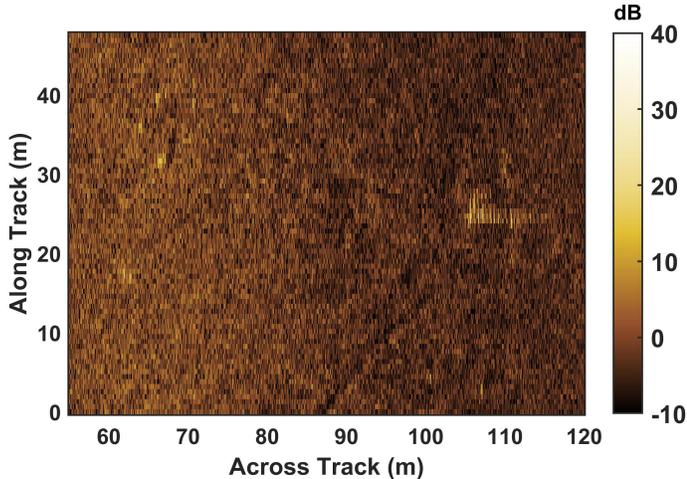


Fig. 5. Bedford Basin sidescan image produced without focusing.

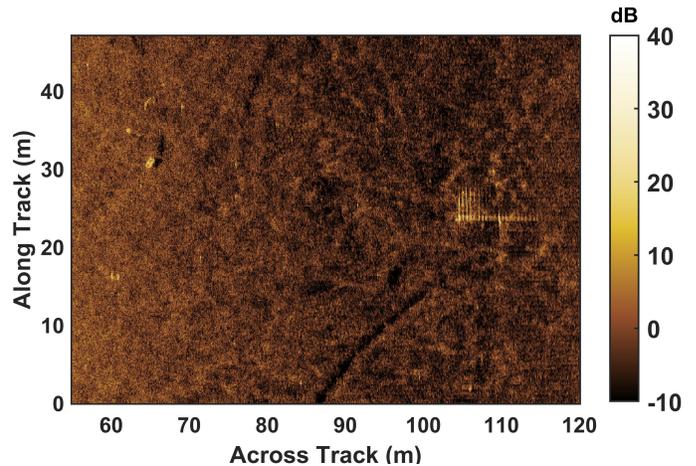


Fig. 6. Bedford Basin sidescan image produced with dynamic focusing.

GPUs contain hundreds or thousands of processing cores operating in parallel, and thus any computations that can be completed in parallel are suitable for GPU acceleration. For example, one of the most computationally expensive steps in both SAS and sidescan processing is matched filtering (used to maximize the across track resolution and SNR). The SASVIEW software computes all matched filtering (including that done for unfocused sidescan), on the GPU. Similarly, we can accelerate the sidescan dynamic focusing on the GPU because time domain backprojection is an inherently parallel process. By computing the backprojection on the GPU, the dynamic focusing can be completed in real-time on consumer grade personal computers, laptops, and low power embedded processors equipped with suitable GPUs, such as the Jetson series of NVIDIA processors.

IV. RESULTS

In this section, we validate the GPU-accelerated time domain backprojection method for focusing a multibeam sidescan sonar using sample data from both the AquaPix MINSAS and InSAS. The InSAS consists of two vertically separated arrays of transducers, with each array containing a unique dual-row transducer for multipath suppression. Vertically separated arrays provide bathymetry in addition to co-registered seabed imagery. The MINSAS is a miniaturized version of the InSAS, for use on medium and small sized vehicles, featuring two vertically separated single rows of transducers. Both the InSAS and MINSAS have modular configurations enabling the

selection of different array lengths to accommodate varying area coverage rates. The MINSAS is currently available in four increasing length configurations (MINSAS60, 120, 180, and 240). The MINSAS operates at a centre frequency of 337 kHz with a bandwidth of 40 kHz, whereas the dual rows of the InSAS operate in distinct bands at 240 and 337 kHz, each with a bandwidth of 40 kHz.

We have selected three different real sidescan images to test the new GPU-accelerated multibeam sidescan beamformer. The first two were collected with the MINSAS and the third was collected with the InSAS. In our analysis, we evaluate the focused sidescan images and compare the computation time of the focused and unfocused images. A summary of the image acquisition parameters and corresponding processing times for each data set can be found in Table I. Note that vehicle altitude refers to the vehicle's height above the seabed. The unfocused processing times include GPU acceleration for the matched filtering, but not the beamforming (because the beamforming is just a simple sum of the signal from all receivers). All computations were done on a consumer grade laptop equipped with a NVIDIA GeForce GTX 1050 Ti GPU.

A. Baltic Sea

In 2017, Kraken participated in a sea trial conducted in the Baltic Sea with the MINSAS120. This data set includes rocks of a variety of sizes distributed evenly from near to far range, making it an ideal data set for testing the range dependency of the focusing correction. Even at near range, the unfocused sidescan image appears to be poorly focused (blurry), with the

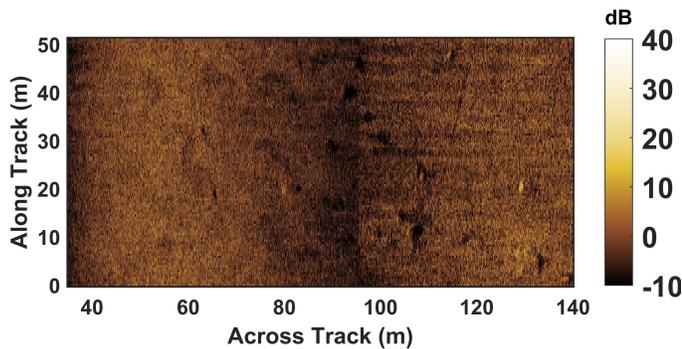


Fig. 7. Narragansett Bay sidescan image produced without focusing.

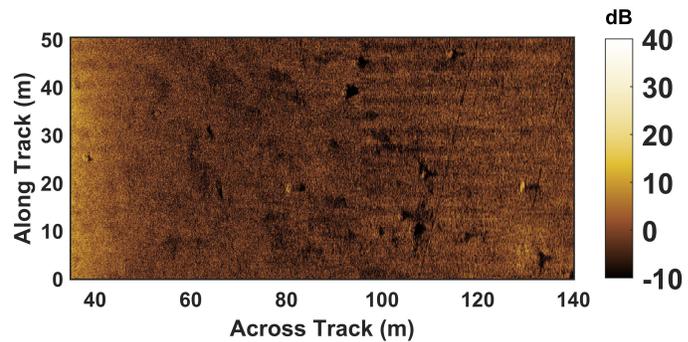


Fig. 8. Narragansett Bay sidescan image produced with dynamic focusing.

blurriness increasing with range (Fig. 3). By applying dynamic focusing we are able to produce a well focused image until the far range (Fig. 4). We are able to achieve these gains by adding just 0.15 s to the sidescan processing time (Table I).

B. Bedford Basin

In March of 2019, Kraken conducted a sea trial using the KATFISH equipped with the MINSAS180. For this trial, Kraken deployed a 0.5 m cube target on the seabed in the Bedford Basin, Halifax, Nova Scotia. The cube target is an excellent feature to test the sidescan dynamic focusing corrections on. The unfocused sidescan image appears extremely pixelated, where even the cube at ~ 105 m is difficult to pick out in the scene (Fig. 5). In the focused image, we are not only able to detect the cube, we are also able to observe the texture of the seafloor and identify small objects (such as a lobster pot at ~ 65 m) in the scene that could not be detected in the unfocused image (Fig. 6). With GPU acceleration, focusing the image only added an extra 0.11 s to the processing time (Table I).

C. Narragansett Bay

In October 2012, Kraken conducted a sea trial with the U.S. Navy's Naval Undersea Warfare Center (NUWC) Newport Division with survey data collected in Narragansett Bay, RI. During this trial, the InSAS was deployed from a REMUS 600 AUV. The data selected from this trial shows many small targets (rocks and lobster traps), distributed in the near to very far range, allowing us to test the range dependency of the focusing correction at even further ranges. In the unfocused image, we observe significant blurring of the targets at all ranges (Fig. 7). In the focused image, we observe a significant improvement in the image quality. The targets are not nearly as blurry and previously undetected targets can be identified (Fig. 8). Focusing this image only added an extra 0.13 s to the processing time of the image (Table I). Note that the image is a composite of the imagery from short range and long range rows of the dual-row sensor. In Fig. 8, the new processing software employs improved vertical beam pattern inversion, which eliminates the brightness artifact at 95 m across track.

V. CONCLUSION

In this paper, we introduced a new real-time dynamically focused multibeam sidescan beamformer. Using time domain

backprojection, both near-field and bistatic corrections can be applied to the entire image. The dynamically focused beamformer operates in real-time through GPU acceleration. Overall, the additional processing time required for dynamic focusing is negligible. In general, TDBF increases the processing time for an entire image by just over a tenth of a second, about one tenth of which is time spent putting the data on the GPU. The additional processing time needed to dynamically focus the image is undetectable by the user when performed on the GPU. In comparison, the processing time required to focus all three sets of images to the same resolution with a block FDBF, would take on the order of minutes, which means the sidescan processing would no longer be considered real-time. Future work will concentrate on further improving the image quality by incorporating corrections for platform motion during dynamic focusing.

REFERENCES

- [1] S.-M. Steele, R. Charron, J. Dillon, and D. Shea, "Shallow water survey with a miniature synthetic aperture sonar," *OCEANS MTS/IEEE Seattle*, Oct 2019.
- [2] R. Mucci, "A comparison of efficient beamforming algorithms," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 32, no. 3, pp. 548–558, 1984.
- [3] L. E. Kinsler, A. R. Frey, A. B. Coppens, and J. V. Sanders, *Fundamentals of Acoustics*, 4th ed. John Wiley & Sons, Inc., 2000.
- [4] T. Curtis, P. Kember, and B. Oconnell, "Wide band, high resolution sonar techniques," *IEE Colloquium on Underwater Applications of Image Processing*, Mar 1998.
- [5] X. Lurton, *Introduction to Underwater Acoustics: Principles and Applications*. Springer-Verlag Berlin AN, 2016.
- [6] A. Belletini and M. Pinto, "Theoretical accuracy of synthetic aperture sonar micronavigation using a displaced phase-center antenna," *IEEE Journal of Oceanic Engineering*, vol. 27, no. 4, pp. 780–789, Oct 2002.
- [7] D. W. Hawkins and P. T. Gough, "Temporal Doppler effects in SAS," *Proceedings of the Institute of Acoustics*, vol. 26, 2004.
- [8] U. Hamid, R. A. Qamar, and K. Waqas, "Performance comparison of time-domain and frequency-domain beamforming techniques for sensor array processing," *Proceedings of 11th International Bhurban Conference on Applied Sciences & Technology (IBCAST) Islamabad, Pakistan*, Jan 2014.
- [9] N. Odegaard, *Fast time domain beamforming for synthetic aperture sonar*, 2004.
- [10] J. Dillon, "Real-time interferometric SAS processing with ultra-low power consumption," *OCEANS MTS/IEEE Charleston*, Oct 2018.