Obtaining a ship's speed and direction from its Kelvin wake spectrum using stochastic matched filtering

Andreas Arnold-Bos, *Student Member, IEEE*, Arnaud Martin, *Member, IEEE*, Ali Khenchaf, *Member, IEEE*

Laboratoire E^3I^2 (EA 3876)
ENSIETA (École Nationale Supérieure des Ingénieurs des Études et Techniques de l'Armement)
2, rue François Verny
29806 Brest CEDEX 09, France
{arnoldan, martinar, khenchal}@ensieta.fr

***Abstract*—The Kelvin wake of a ship is directly linked to the ship's speed, heading and hull shape. This wake can be visible in high resolution synthetic aperture radar images or optical images. Whenever it is possible, analyzing it can provide elements to identify the ship and track its course. We propose a strategy based on the generalized Radon Transform and the Stochastic Matched Filtering where the locus of the wake signature in the 2D spectrum of the image is to be detected.**

***Index Terms*—Marine surveillance, ship wakes detection, generalized Radon transform, stochastic matched filtering**

I. INTRODUCTION

Ship wakes are a very visible feature on Synthetic Aperture Radar (SAR) images or high resolution optical images, since they stretch over kilometers. As such, they are a tell-tale signature of a ship. The wake is generally categorized into two phenomena. First, there is the turbulent ship wake, which appears as a long, dark streak behind the ship. The second phenomenon is the Kelvin wake, consisting in a system of waves living in a 39° cone (if the sea is of "nearly infinite" depth).

There is a large body of work dedicated to the detection of the turbulent wake. Generally, this boils down to performing line detection. The Radon and the Hough transform have been shown to be very suitable for this task. Various pre-processing and post-processing algorithms have been proposed to, respectively, enhance the visibility of the wake before the transformation, and to improve the thresholding of the wake signature in the transform (see [1] for a benchmark). It is also common to distinguish the wave pattern of the Kelvin wake with either airborne SAR images or high resolution optical images. Analyzing the Kelvin wake can yield an estimation of the ship's speed and heading, which can be merged later with other observations (coming, from instance, from the Doppler effect) in a fusion scheme so as to increase the robustness of the overall estimation of these parameters.

Here, we introduce a method to reliably detect the Kelvin wake based on the analysis of its 2D spectral signature. The method uses a pre-processing step based on a generalized Radon transform (GRT) of the 2-D Fourier transform of the image. The result of the transform is then thresholded using stochastic matched filtering.

II. PROBLEM STATEMENT

Theory shows that the analytic elevation ζ of the sea surface at point (x, y) can be described by a superposition of sinusoidal waves, each having a certain direction θ, amplitude, and phase $\phi(\theta)$:

$$\left\{\begin{array}{l} \zeta(x,y) = \int_{-\pi/2}^{\pi/2} A_U(\theta) \exp\phi(\theta) d\theta \\ \phi(\theta) = -ik_{U,\theta 0}(\theta)[x\cos(\theta-\theta_0) + y\sin(\theta-\theta_0)] \end{array}\right. \quad (1)$$

Here, θ_0 is the ship's heading, A_U is a (complex) amplitude sometimes called the Kochin function, which depends on the wave's direction, the hull geometry and the ship's speed U. Finally, $k_{U,\theta_0}(\theta)$ is the wavenumber of the wave travelling in direction θ in a local frame such that the x axis is aligned with the wake's axis of symmetry:

$$k_{U,\theta_0}(\theta) = \frac{k_U^0}{\cos(\theta-\theta_0)^2} \quad (2)$$

with $k_U^0 = g/U^2$ and g denoting the acceleration of gravity. This expression stems from the fact that the Kelvin wake system travels at the same speed than the ship. The 2-D Fourier Transform Z of ζ, using coordinates (κ_x, κ_y) in the Fourier plane, is:

$$Z(\kappa_x, \kappa_y) = \int_{-\pi/2}^{+\pi/2} A_U^\star(\theta)\delta(g_{U,\theta_0}(\kappa_x,\kappa_y))d\theta \quad (3)$$

with function g_{U,θ_0} defining the locus (k, θ) of the wake; this function must agree with equation (2), so:

$$g_{U,\theta_0}(\kappa_x,\kappa_y) = 0 \overset{\Delta}{\Leftrightarrow} \left\{\begin{array}{c} \theta = \arctan(\kappa_y/\kappa_x) + \theta_0 \\ k(\theta) = \frac{k_U^0}{\cos^2\theta} \end{array}\right. \quad (4)$$

Figure 1b illustrates this result by showing the Fourier Transform of figure 1a, along with the locus of the wake spectrum, which follows equation (4).

In remote sensing, only the image $I(x, y)$ of the surface as acquired by a camera or a radar is known; $\zeta(x, y)$ is not. With SAR, the spectrum of I can be linked to the spectrum of ζ by a transfer function $\mathcal{H}$ [2]; the same is valid for optical pictures (provided that there are no perspective distortion). Function $\mathcal{H}$ accounts, for instance, for the position of the light or the radar in the scene, and links the angles of incidence of light, to the

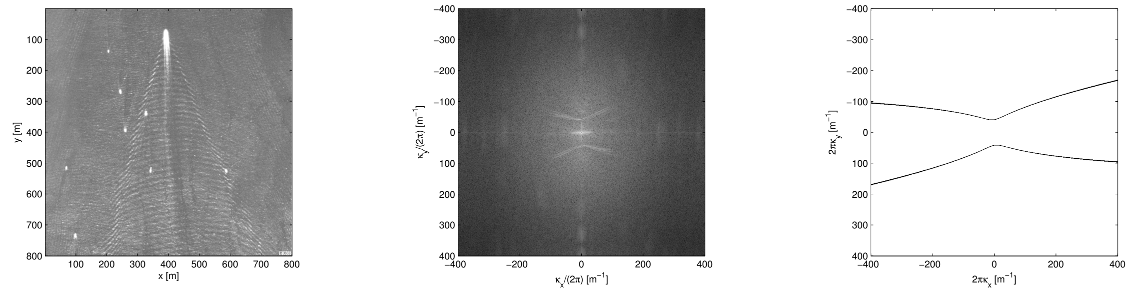

(a) Original picture (b) 2D Fourier Transform (c) Locus of wake spectrum

Fig. 1: An optical picture of a wake (source: USGS) and its locus in the Fourier plane.

local reflectivity. If $\mathcal{H}$ is assumed to be nearly linear, then the locus of the wake in the spectrum of I will be essentially the same than in Z.

The problem is then to detect this locus reliably, since the wake signal will be lost in the spectral components coming from other waves. The only approach that we are aware of, consists in detecting couples of points (P_{1i}, P_{2i}) on the locus (by thresholding the Fourier transform) and then using them to determine an estimation (U_i, θ_0^i), the final estimation of U and θ_0 being obtained by averaging the U_i's and the θ_0^i's [3]. We believe this method to be unreliable since only a few falsely detected couples of points can corrupt the average. Even if a RANSAC scheme [4] could be used to reject the outliers, it is much better to improve the detection in the first place.

III. Presentation of the algorithm

In this section, we detail the steps used in our algorithm to detect the locus of wake spectrum. The algorithm is divided into three steps: first, a preprocessing on the Fourier transform of the image, then a Generalized Radon Transform is used, and finally, the transform is enhanced using Stochastic Matched Filtering prior to a thresholding.

A. Pre-processing

First, the local average of image I is subtracted; we worked with a window being ten pixels wide. Then the magnitude of the 2D Fourier Transform is computed. The contrast of the image is then enhanced by using a high-pass filter. We write $F(\kappa_x, \kappa_y)$ the result of this processing.

B. The Generalized Radon Transform

To detect the wake, two things are desirable: first, to increase the signal to noise ratio (SNR) of the wake signature; second, to reduce the surface of the signature (*i.e.* to concentrate the signal into a single spike), which is easier to detect. Thus, the main idea is to compute the sum of the image spectrum along all possible wake loci. If a wake is indeed present, then the sum on the locus of the wake will have a much higher value than the sum for other loci, and the wake can be detected by thresholding the sums. This is in fact nothing more than the Generalized Radon Transform (GRT) of the spectrum defined for a family of curves of parameters (U, θ_0) having $g_{U,\theta_0}(\kappa_x, \kappa_y) = 0$ as implicit equation. The transform R of F for a wake locus of parameters (U, θ_0), will be:

$$R(k_U, \theta_0) = \iint_{-\infty}^{+\infty} F(\kappa_x, \kappa_y) \delta \left(g_{U,\theta_0}(\kappa_x, \kappa_y) \right) d\kappa_x d\kappa_y \quad (5)$$

Alternatively, it can be noticed that a particular point of F located at (κ_x, κ_y) can intercept a family of curves which parameters (U, θ_0), or alternatively (k_U^0, θ_0), are defined by the following equation:

$$k_U^0(\theta_0) = \sqrt{\kappa_x^2 + \kappa_y^2} \cos^2 \left(\theta_0 - \arctan(\kappa_y / \kappa_x) \right) \quad (6)$$

These curves are sinusoids in the (k_U^0, θ_0) plane. After performing the GRT, the local mean is removed to the result so as to enhance the visibility of the spike that would be present should a wake be present in the image. We note R' the result of this operation.

C. Improving the detection with Stochastic Matched Filtering

The spikes in R' can be isolated from noise using a detection threshold, which could for example be set at $m + k\sigma$, where m is the average of R', σ its standard deviation, and k is a constant which should ideally be higher than 3 if the noise is Gaussian. The noise in R' comes from the Radon Transform of anything that is not a wake in the original image, for instance, other waves of the ocean. Noise becomes more of a problem for high sea states and small, slow ships. The position (k_U, θ_0) of the spike gives, of course, the speed and direction of the ship.

If the shape of a spike corresponding to a wake in R' were constant everywhere in the image, whatever the imaging conditions or the ship and its speed, matched filtering could have been used to increase the visibility of the spike in noise. Matched filtering is indeed the optimal linear filter for maximizing the SNR in the presence of additive stochastic white noise. Alas, this is not the case here.

Recently, Stochastic Matched Filtering (SMF) has been proposed [5]; it extends the traditional matched filtering to cases where all the instances of the signal to detect bear some resemblance but also differ in some random changes. This is exactly the situation we are in. To quickly explain this method, let us note $\mathbf{r}$ a chunk of R': this is a 1-D vector resulting from the concatenation of lines of a 2-D window of n pixels. This vector can be written:

$$\mathbf{r} = \mathbf{s} + \mathbf{n} \quad (7)$$

Here $\mathbf{s}$ is the useful signal, n is the noise. If the signal to detect typically lives in a window $\mathbf{w}$ of length n, it is possible to compute the $n \times n$ normalized covariance matrix $\mathbf{C}_w$ of windows containing useful signal. This is done over several sample windows $\mathbf{w}_1 ... \mathbf{w}_m$ $(m > n)$ of w acquired either in real images or simulated. The same can be done for windows known to contain only noise (we note $\mathbf{C}_n$ that matrix). Then matrix $\mathbf{C} = \mathbf{C}_n^{-1}\mathbf{C}_w$ of size $n \times n$ is computed. It can then be shown that filtering $\mathbf{r}$ with any eigenvector h_i of $\mathbf{C}$ associated to an eigenvalue $\lambda_i > 1$ guarantees an increase of the signal-to-noise ratio; this increase is theoretically equal to λ_i if $h_i^t C_w h_i = 1$. As such, it is possible to determine if the signal is present or absent by thresholding the result of the filtering by the h_i's.

IV. Building the Stochastic Matched Filters

In this stage, the covariances of the signal (free of any noise) and of the noise (free of any useful signal) will be estimated via simulation and by using actual images. This is often long (up to a day in our case), but it can be done offline.

First, the size of the spike to detect must be estimated. This is done via trial and error; in our case, we chose a nine pixels wide window (thus $n = 9 \times 9 = 81$). Then the covariance matrices are estimated. At least n samples must be used to estimate $\mathbf{C}_n$ and $\mathbf{C}_w$, or the $\mathbf{C}$ will be singular.

A. Estimating $\mathbf{C}_n$

The pre-processing and GRT algorithm is used on a body of real images. For each image, random windows are chosen and added to a pool of samples. The covariance matrix $\mathbf{C}_n$ of that pool can then be computed. This matrix is experimentally found to be nearly diagonal, which means the noise is spatially uncorrelated.

B. Estimating $\mathbf{C}_w$

We considered nine classes $H_i, i \in [1..9]$ of generic hulls: sailboat, liner, coastal patrol ship, frigate (DTMB 5415), tanker/freighter, attack submarine, and Wigley parabolic hulls of various length-to-beam ratio (4, 5 and 6). The length distribution of each ship in a given class followed a hand-tuned Gaussian distribution. The speed distribution was uniform. Each class was given a reference model s_0^i of length L_0^i. The other ships $s_j^i, j \geq 1$ in the class were scaled versions of s_0^i, having a length L_j^i, and a speed U_j^i. Speeds are taken so that wavelengths k_U^0 vary at a step of 1 m^{-1} with 50 ships per k_U^0. Figure 2 shows the distribution of ships we used for this paper. For each reference model s_0^i, the amplitude function $A_0^i(U, \theta)$ (see eq. 1) can be computed for all angles θ and a wide and finely sampled range of speeds. This is done by numerical integration over a ship hull [6]. This operation is slightly computation intensive, which is why it has only been done for the reference models. Figure 3 shows $A_U(\theta)$ for the reference

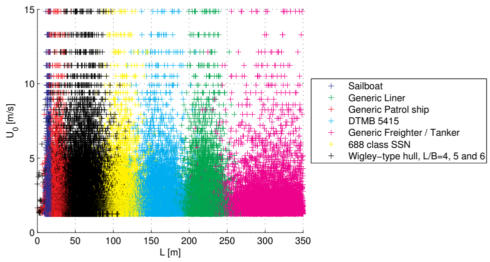

Fig. 2: Distribution of the ship samples used to build the stochastic matched filters. The legend indicates the reference ship used for each class.

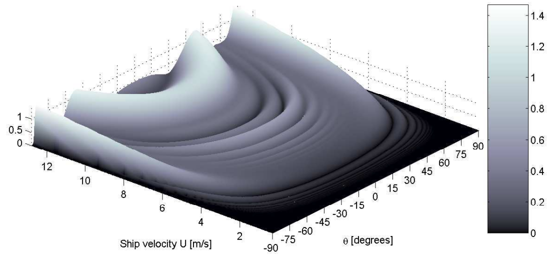

Fig. 3: Magnitude of complex function $A(\theta)$ for a SSN 688 class submarine at periscope depth.

chosen for the submarine class, with $\theta \in [-\pi/2, +\pi/2]$ and U between 0.3 and 12.8 m/s.

Then, the amplitude function $A(\theta)$ for ship s_j^i is deduced from the amplitude function of the class model s_0^i. Indeed, two ships differing only from their scale, and sharing the same Froude number $U/\sqrt{gL}$, have the same wake pattern; the ratio of the wave heights is equal to the ratio of the ship scales. Thus, by knowing U_j^i, L_j^i and L_0^i, the amplitude function for s_j^i is:

$$A_j^i(\theta) = \frac{L_j^i}{L_0^i} A_0^i \left(\sqrt{\frac{L_j^i}{L_0^j}}, \theta \right) \quad (8)$$

By knowing the amplitude function $A_j^i(\theta)$ and the locus of the wake spectrum $g_{U_j^i}(\kappa_x, \kappa_y)$, a simulated F image can be built from scratch as if it were obtained by the preprocessing explained in section III-A; then its Radon transform R' can be computed, and the window containing the spike can be extracted and added to the pool of signal samples, as in figure 4.

But there is one final catch: amplitude function $A(\theta)$ cannot be used directly, but rather $\mathcal{H}A(\theta)$ (see section II). In practice, this would add a whole new set of parameters and complicate the simulation. However, some effects of the imaging situation cannot be neglected. In particular, depending on the position of the sun (in an optical photograph), only one arm of the wake will be well visible. For instance, in figure 1a, the left arm is better seen than the right one; and indeed, only one half of the wake spectrum is clearly visible in figure 1b. Some additional diversity can be added to mimic the effect of lighting. For

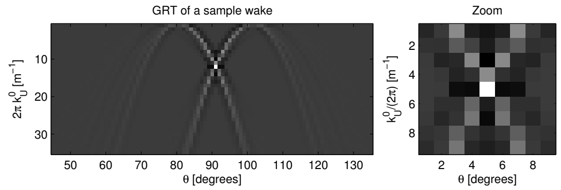

Fig. 4: Typical GRT of a noise-free wake spectrum.

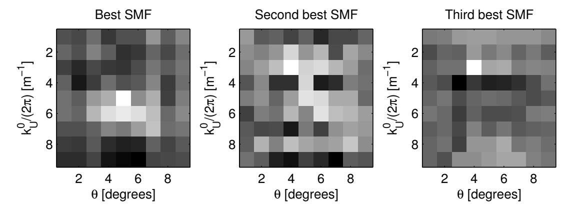

Fig. 5: The three best SMF obtained through the simulation process.

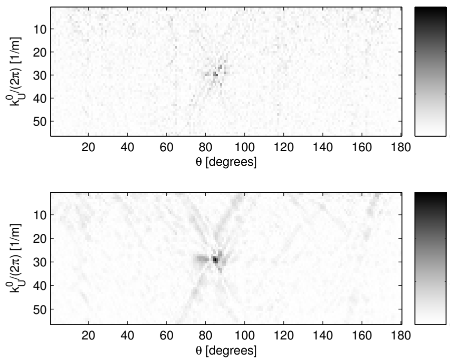

Fig. 6: Raw GRT for image from figure 1a (above) and result after using SMF (below). Images are scaled so that the standard deviation of noise is unitary. There is a 3.28 dB improvement in the SNR when SMF is used.

each ship, we chose to weight the amplitude function $A_j^i(\theta)$ by a Gaussian function of center Θ_j^i and of width σ_j^i; the directions where this function is nearly 1 represent those where the wake is clearly visible because it is well lighted. Also, the amplitude of the signal must be scaled properly to match the range of values returned by the sensor. In the end, the SMF's are obtained using eigenvector decomposition; figure 5 shows the filters associated with the three best eigenvalues. Notice that only the spike appears. When the noise level decreases, however, the filters become very similar to the sample obtained in figure 4, which agrees with intuition.

V. Results and conclusions

We made experiments on a body of 64 optical, metric resolution images courtesy of the USGS obtained on `terraserver.com`. The GRT performs well in concentrating the energy of the wake but we observed it never performed better than the human eye, in that the wake spectrum had to be visible in the Fourier transform for the algorithm to have a chance to work.

The transform was then cleaned using the SMF. In our case, the highest eigenvalue of $\mathbf{C}$ had a value of 81.35, that is 19.10 dB. Yet, the gain in SNR never exceeded 4 dB for all our images, see for instance figure 6. However, our body of "learning" images was a mix various sea states and configurations: the theoretical covariance of noise hence does not necessarily fit the noise actually present in the image since some operational compromises had been made. To try and reach the theoretical SNR improvement, it would therefore be preferable to have several banks of matched filters, one for each sea state. This involves higher simulation costs, but once this is done, the SMF has exactly the same complexity as any linear filter and can be computed efficiently by a fast Fourier transform. There is therefore no reason not to use it to increase the robustness of the signal.

Our approach is general, in that it can be used for optical pictures, but also for SAR images, provided that the Kelvin wake is visible, which means that the imaging configuration must have been designed for that task. There are, however, restrictions to the method in that the depth of the sea must be infinite (otherwise the locus has a different expression, but this could be adapted) and the ship must mot be turning (which smears the locus of the spectrum and precludes proper detection).

As for stochastic matched filtering, its originality is precisely to help improve the SNR when the signal is not known deterministically, but only through its covariance. This is often the case when thresholding transforms, such as wavelet transforms, other types of Radon transforms, etc. As such, there is probably plenty of room for exploration with this method.

Acknowledgment

This work was supported by a grant from the regional council of Brittany.

References

[1] A. Arnold-Bos, A. Khenchaf, and A. Martin, "An evaluation of current ship wake detection algorithms in SAR images," in *Caractérisation du Milieu Marin / Marine Environment Characterization, SeaTechWeek '06*, Brest, France, Nov. 2006.

[2] K. Hasselmann and S. Hasselmann, "On the nonlinear mapping of an ocean wave spectrum into a synthetic aperture radar image spectrum and its inversion," *Journal of Geophysical Research*, vol. 96, no. C6, pp. 10,713–10,729, Jun. 1991.

[3] Z. Wu, "On the estimation of a moving ship's velocity and hull geometry information from its wave spectra," Ph.D. dissertation, The University of Michigan, 1991.

[4] M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the American Mathematical Society*, vol. 24, pp. 381–395, Jun. 1981.

[5] P. Courmontagne, "An improvement of ship wake detection based on the Radon transform," *Signal Processing*, vol. 85, no. 8, pp. 1634–1654, 2005.

[6] E. O. Tuck, L. Lazauskas, and D. C. Scullen, "Sea wave pattern evaluation, part I report, primary code and test results (surface vessels)," Applied Mathematics Department, The University of Adelaide, Tech. Rep., Apr. 1999.