2

Synthetic Aperture Techniques for Sonar Systems

Sérgio Rui Silva, Sérgio Cunha, Aníbal Matos and Nuno Cruz
University of Porto, Faculty of Engineering
Portugal

1. Introduction

Today a good percentage of our planet is known and well mapped. Synthetic aperture techniques used in space and airborne systems has greatly aided to obtain this information. Nevertheless our planet is mostly covered by water and the level of detail of knowledge about this segment is still very far away from that of the land segment.

Synthetic aperture is a technique that enables high resolution through the coherent processing of consecutive displaced echo data. Instead of using one static large array of transducers, it uses the along-track displacement of the sensors to synthesize a large virtual array. The resolution thus obtained is in the order of the transducer size and, most importantly, independent of the range between sensor and target. While a modern high frequency real-aperture sonar system can have a beam width below 1^o, this translates into a resolution of half a meter at a range of just 25m. A synthetic aperture system using the same transducer can obtain a resolution of about 5cm across the whole range. Moreover the transducers used for synthetic aperture can be much simpler and so a good deal less expensive. Because there is no need to have a small real aperture, the frequency employed can be considerably lower, which enables longer reach due to the better propagation of lower frequencies in water.

This potential resolution increase comes at the cost of algorithm complexity in the image formation. The sonar must also describe a movement with tight tolerances with respect to deviations from known velocity and path. Also, the platform maximum velocity is a function of the pulse repetition rate and sensor size. This limit relates to the resolution of the image that if not respected will lead to aliasing.

The most used platform for synthetic aperture sonar is the tow-fish. Good designs enable smooth motion, but the inability to use satellite navigation technology leads to expensive solutions that integrate high grade inertial navigation units and data extracted from the sonar array itself. This only works for arrays with a high count of elements that operate at the nominal or above the nominal pulse repetition frequency. The sonar can also be mounted on the hull of a ship, providing access to high precision GPS navigation that can be integrated with data from moderate cost inertial systems to further refine the navigation solution. Nevertheless a ship is seldom easy to manoeuvre and presents considerable operation and maintenance costs of the sonar system itself. An autonomous boat arises as an interesting solution for these problems. It can be used as standalone or with the support of a ship. It enables the use of GPS and inertial navigation units efficiently. Moreover, its path

Source: Advances in Sonar Technology, Book edited by: Sergio Rui Silva,
ISBN 978-3-902613-48-6, pp. 232, February 2009, I-Tech, Vienna, Austria

and velocity can be easily controlled for better motion stability. The operation and maintenance costs are low and its availability is very high. It can be used both for sporadic missions and for regular security check of harbours, river navigability assessment, infrastructure inspection, etc.
The position of the sonar must be known to a 1/8 of the wavelength for proper synthetic image formation. Traditional synthetic aperture image formation techniques assumes a straight path for the sonar motion (typically for method that operate in the frequency domain) and treat deviations from this straight path as motion errors. A newer approach uses time-domain methods (such as back-projection) that don't rely on the assumption that the sonar follows any particular path. Instead, the information obtained by the navigation system is used at each sonar sampling position to form the virtual array. Here, only the position uncertainties are considered as errors.
High frequency systems require navigation precision bellow the centimetre level. This level of precision is not feasible to be obtained with the navigation systems of today. Therefore, the image formation starts with the available navigation solution and then a global auto-focus algorithm (that searches for an optimum measurement of image quality) refines the image formation parameters to mitigate the unavailable necessary navigation precision. Instead of using redundancy in the data (that comes at a high cost), global auto-focus algorithms parameterizes the image formation process enclosing navigation errors and medium fluctuations. Because of the large number of parameters that result from this, an efficient process of focusing synthetic aperture images is necessary.

2. Overview of current systems

Active sonar systems enable underwater imaging through angular and range echo discrimination. When a single beam is used to illuminate a swath as the sonar platform moves it is said that the sonar is a side-scan. In these systems a single echo line is obtained at each time with the angular discrimination being given by the beam width. Thus a narrow beam is desirable for high angular discrimination or along-track resolution. Typical beam widths are in the order of 1^o. In these systems the along-track resolution is dependent of the range and a large array has to be used to obtain suitably low beam widths at the desired range. This type of sonar enables high area coverage speed. Alternatively, several narrow beams can be used to spatially sample the swath obtaining range and intensity information for each angular position. This is called multi-beam sonar. In this case the footprint of each beam is also dependent of the range. This type of sonar requires expensive and complex hardware to achieve a high number of sampling narrow beams. The area coverage speed is also limited by the area covered by the beams.
Synthetic aperture enables a high resolution/high area coverage binomial not possible with other sonar techniques. Instead of using a long physical array, a large virtual array is synthesised through the coherent combination of the echoes in the along track dimension of a side-scan sonar. Range independent along-track resolution is in this way obtained. Moreover the obtained along-track resolution is not influenced by the frequency of the signals employed and is in the order of the transducer physical dimensions. Lower frequency signals can thus be employed to extend the sonar range. Also because of the processing gain, the necessary transmitting power is lower when compared to its real aperture counterpart.

While it is possible to apply synthetic aperture techniques to multi-beam sonar, these have not been of great dissemination and the focus has been on obtaining high resolution sonar systems with large area coverage speed.
It is possible to obtain height estimation through the use of interferometric techniques on side-scan sonar images or synthetic aperture sonar images ([Saebo, T.O. (2007); Silva, S. et al (2008 c)]). Nevertheless multi-beam sonar obtains height measurements directly.
Synthetic aperture sonar systems can operate in strip-map or spot-light mode. In spot-light mode the beam is steered so a single zone is illuminated as the sonar platform moves, while in strip-map mode an area is sequentially sampled. Since sonar applications normally strive for area-coverage, strip-map is the more used synthetic aperture mode.
Synthetic aperture techniques are of common use in radar systems (typically know as SAR, [Tomiyasu, K. (1978)]). Here the signal propagation velocity through the medium is much higher, the wavelength long in comparison to radar platform motion uncertainties and medium phase fluctuations. Moreover the scene width is short when compared to the centre range which enables the use of simplifying approximation in the image formation algorithms. Also the bandwidth to centre frequency ratio is small. This means that SAR algorithms are not suitable for direct application in sonar data.
Medium stability for synthetic aperture sonar system is a major issue, but has been proved to be not restraining factor for the application of this technique ([Gough, P. T. et al (1989)]).
Most active synthetic aperture sonar systems are still in the research and development stage ([Douglas, B. L.; Lee, H. (1993); Neudorfer, M. et al (1996); Sammelmann, G. S. et al (1997); Nelson, M. A. (1998); Chatillon, J. et al (1999)]) but this technology is starting to emerge as a commercially advantageous technology ([Hansen, R.E. et al (2005); Putney, A. et al (2005)]). Nevertheless most of these systems rely on complex and costly underwater vehicle configurations with multiple receivers, making them not attractive cost wise. This is because to correctly focus a synthetic array one needs to know the position with a very high degree of precision. Since it is not possible to use a high precision navigation system, such as a GPS, in the underwater environments, these systems have to rely on positioning schemes that use sonar data itself. This usually means having to use a complex array with multiple receivers, limit the range to use higher pulse repetition frequency or greatly limit the area coverage speed. Inertial navigation systems are of hardly any utility in standalone mode and must be corrected by other navigation sources. The underwater vehicles can be simply towed by a ship or be autonomous.
Surface autonomous vehicles are easier to operate and maintain. Plus it is possible to use readily available high precision differential GPS navigation system. Other navigation data sources can also be effectively used to enhance the navigation solution. A simple transducer array can be used since there is no longer mandatory to have a multiple element system for navigation. This makes high resolution synthetic aperture sonar a cost effective possibility.

3. Problem statement

Synthetic aperture techniques can be used to obtain centimetre level resolution in the along-track direction. Nevertheless to obtain this level of resolution in the cross-track direction, one needs to use large bandwidths and thus high frequency signals. This makes the position accuracy issue even more problematic for synthetic aperture sonar system since the necessary accuracy is directly related to the wavelength of the signal used.

Further more a sonar platform is subject to several undesirable motions that negatively affect the synthetic aperture performance. Normally a synthetic aperture sonar platform would move through a straight path with constant velocity. But this is seldom the case. Low frequency oscillations around the desired path and high frequency motion in heave, sway and surge directions adversely affect the sonar platform. This is due to water surface motion, currents and platform instability.
In the case of a tow-fish, there is little or no control on its motion. For an autonomous under water vehicle it is difficult to maintain a path reference since there are no reliable navigation sources.
An autonomous boat offers automatic position and velocity following capabilities through the use of a highly available differential GPS system (DGPS-RTK). Furthermore, the sonar sampling position and attitude is known to a high precision level which can than be integrated in the image formation algorithms.
The synthetic aperture sonar is then thus only limited by the unknown portion of the platform motion and medium phase fluctuations. Typical DGPS-TRK systems can operate with an error in the centimetre level. This mean that a synthetic aperture image with a wavelength of about 8 cm can be directly formed using only this position estimation. Furthermore, in this way, the obtained images can be easily and accurately integrated with other geographic information systems. Nevertheless, shorter wavelengths require auto focus procedures to be applied to the data for successful image formation.
Traditional synthetic aperture image formation algorithms treat deviations from a linear path as motion errors. They do not integrate well large deviation from a theoretical linear constant velocity path, and so new algorithms had to be developed to cope with this information.
Auto-focus algorithms are essential to successfully produce high-resolution synthetic aperture images. Their role is to mitigate phase fluctuations due to medium stability and more importantly reduce the navigation precision requirements to acceptable levels.
Most auto-focus algorithms require either large pulse repetition frequencies (PRF), several receiver transducers or both. In practical terms is not always possible to use high pulse repetition frequencies due to range ambiguity. With several receiver transducers the PRF can be lower, but the system cost is considerably higher not only due to the transducer array itself but also because of the data acquisition and signal processing system increased complexity.

4. System description

The complete system setup ([Silva, S. (2007a)]) is constituted by a control base station, a surface craft and the sonar itself (Fig. 1). Because of its size and modular construction, this setup is easily portable and has low deployment time. The boat it self is modular and easy to assemble in site without the need of any special tools.
The sonar platform is an autonomous boat (Fig. 2). This is a small catamaran like craft, proving high direction stability, smooth operation and several hours of unmanned operation, which can be command manually from the base station or fulfil a pre-defined mission plan. It was built using commonly available components to lower cost and simplify maintenance. It has two independent thrusters for longitudinal and angular motion that provide high manoeuvrability at low speeds and a maximum speed of 2 m/s. Its size is

suitable for, together with the navigation system, executing profiles and other manoeuvres with sub-meter accuracy. The boat carries three GPS receivers, for position and attitude calculation, a digital compass and a inertial navigation system. The data from the digital compass and inertial system is integrated with the GPS data to provide a better navigation solution. It also embodies an on-board computer for system control, as well as for acquisition and storage of data.

The communication to the boat is obtained using an ethernet radio link (Wi-Fi). The boat and base station antennas were studied as to minimize the effect of the water reflective surface and maximize radio reach which is in the kilometre range.

The sonar system is carried by the boat as payload and transducers are placed at the front of the vessel, rigidly coupled to the boats structure.

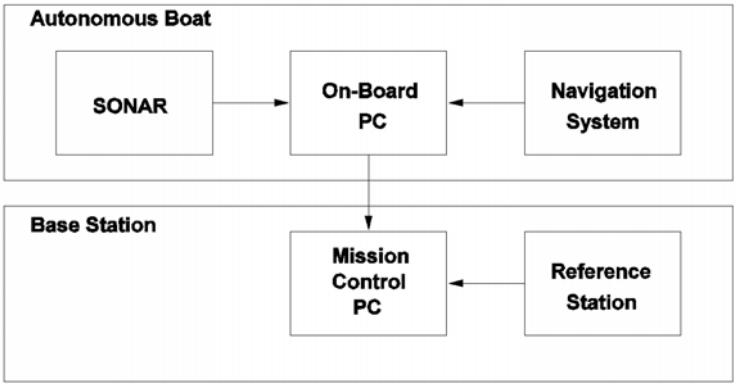

Fig. 1. Sonar system overview.

Fig. 2. Autonomous boat based synthetic aperture sonar.

Fig. 3. Autonomous boat and base station before a mission.

Compared to a solution based on a tow-fish, this sonar platform was mainly chosen because of the possibility to use a GPS positioning system. The velocity and position following is of great importance too. Low frequency errors such as deviations from the desired linear path or inconstant mean velocity are in this way limited. Other surface crafts could be used, but it would still have to provide some degree of motion control.
Having the sonar placed at the surface of the water, limits its use to shallow water realms such as rivers, lakes, dams harbours, etc, because of the maximum possible range. These are nevertheless an important part of the synthetic sonar possible application scenarios. It also makes the sonar more vulnerable to undesirable heave, roll and pitch motion that has to be measured by the boats navigation system and integrated in the sonar image formation algorithms.

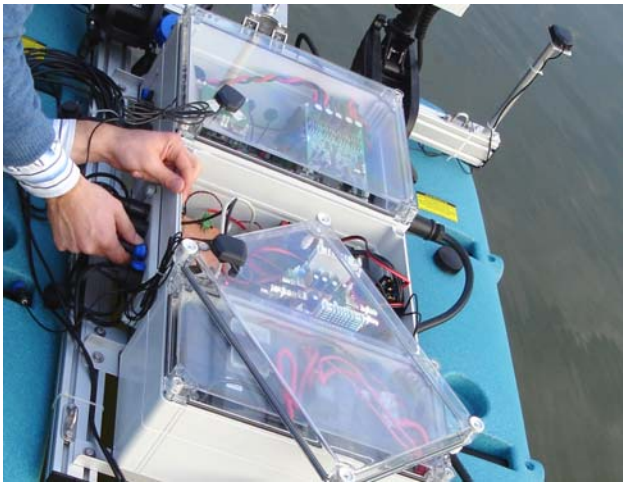

Fig. 4. Autonomous boat in preparation for a mission.

The base station is constituted by a portable PC that runs the boat and sonar control software, a GPS reference station and a high speed digital radio link to the boat (Fig. 3). The purpose of the base station is to control the mission and provide real-time visualization of the sonar data.
The synthetic aperture sonar system is it self constituted by two major components: the transducer array (Fig. 5) and the digital signal processing system for signal generation and acquisition (Fig. 6).
The floating platform transports the acoustic transducer array, placed beneath the waterline. The transducer array is formed by two displaced sets of 4 transducers separated by 0.5 m. Although only one transducer is necessary for transmission and reception of the sonar echoes, the vertical arrangement enables interferometric height mapping. The horizontal set of transducers is for future use and will enable image enhancement through the use of micro-navigation techniques. Each individual set can be regulated for an angle suitable to illumination of the swath at the expected underwater surface depth.
The transducers operate at a centre frequency of 200 kHz, corresponding to a wavelength of 0.75 cm. As appropriate for synthetic aperture operations, their real aperture is large (approximately 18 degrees), but have a strong front-to-back lobe attenuation ratio which is fundamental to minimize the reflections on the near water surface. The effective transducer diameter is 5 cm, which allows for synthetic images with this order of magnitude of resolution in the along-track direction.
The usable bandwidth of the transducers (and of the signals employed) is explored thought the use of amplitude and phase compensation to obtain the highest possible range resolution from the system, thus enabling the use of bandwidths of 20 to 40 kHz with the current configuration.
The sonar signal acquisition and generation system is made up of four principal components: the digital processing and control system, the power amplifier, the low noise amplifier, and a GPS receiver for time reference (Fig. 7).
This system is tailored for a high resolution interferometric synthetic aperture sonar and so special attention was paid to the different analogue components in terms of bandwidth, frequency amplitude response and phase linearity.

Fig. 5. Transducer array and support system.

The power-amplifier is a linear amplifier which can have a bandwidth as high as 10 MHz, and has a continuous output power rate of 50 W (RMS). A trade-of was made between output power and bandwidth. Because we are interested in using our system in shallow waters, an output power of 50W was found to be adequate for this 200 kHz sonar system. In exchange it was possible to build a power-amplifier with very flat amplitude and linear phase response in addition to low distortion (THD < -60 dB) in the spectral band of interest. The power amplifier only operates through the transmission period: this saves energy and lowers the electric noise levels during echo reception. Traditionally, switch amplifiers are use to drive the transducers. A linear amplifier design of this type enables the use of amplitude modulated signals and typically provides lower noise and spectral interference. The transmitted signal can be of any kind: linear chirp, logarithmic chirp or pseudo-random sequence. It can also have any suitable windowing function applied. At this moment, a linear chirp of 30 kHz bandwidth is being used for system tests. The system is prepared to output a pulse rate between 1 to 15 Hz. Because we are interested in shallow water surveys, short ranges enable higher pulse rates. Nevertheless, because of the broad beam-width of the transducer and its effective diameter, a pulse rate of 15 Hz still imposes very low moving speeds to the autonomous boat (0,16 m/s) that in turn has to couple with higher precision trajectory tracking constrains.
Each receiving channel has a low noise amplifier and a controllable gain amplifier to handle the high dynamic range of echo signals. The pre-amplifier system has a bandwidth of 10 MHz, a noise figure of 2 dB and a 50-115 dB gain range. It can also recover rapidly from overdrive caused by the transmitting pulse. To reduced aliasing, a linear phase analogue filter is used in combination with a high sample rate A/D (up to 40 MSamples/s). This anti-aliasing filter is interchangeable and can have a cut-off frequency of, for example, 650 kHz or 2,3 MHz. Both the D/A and A/D have 12 bit resolution and are capable of 200 MSamples/s and 40 MSamples/s respectively. The 12 bit resolution of the A/D (effective 72 db SNR) was found to be high enough to cope with the front-end performance that is dominated by the transducers noise.
The sonar signal processing system uses a direct to digital system architecture. This means that only the front-end elements of the sonar, like the power amplifier and the low-noise amplifier, use analogue electronics and all other functions such as signal generation, frequency down/up-conversion, filtering and demodulation are performed in the digital domain.

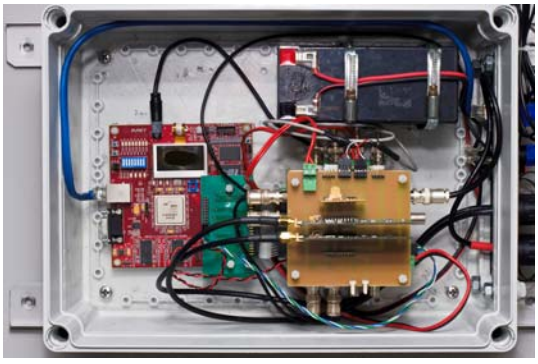

Fig. 6. Sonar signal generation and acquisition system.

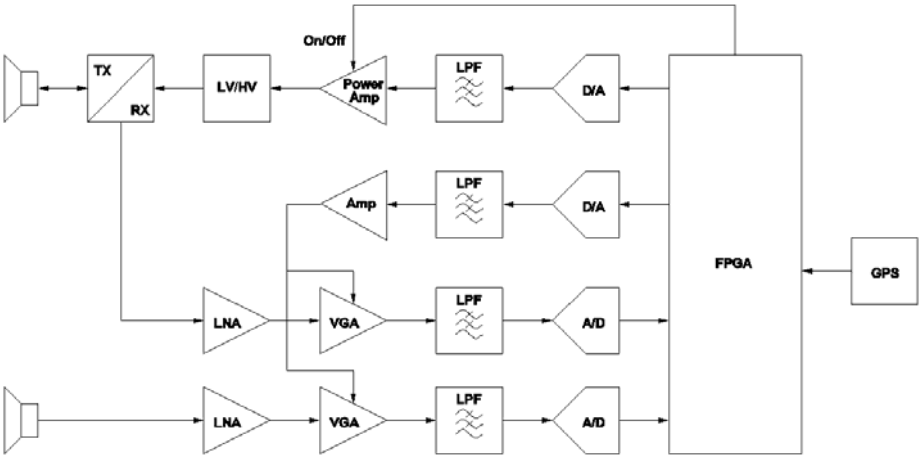

Fig. 7. Sonar schematic system overview.

This greatly simplifies the electronic hardware at the cost of more complex digital signal operations and hardware. But these in turn are powerfully performed in an field programmed array system (FPGA), providing a low cost solution.

A FPGA is a device that enables the implementation of customized digital circuits. This device can perform several tasks in parallel and at high speed. The system complexity stands, therefore, in the digital domain, enabling more flexible and higher quality signal acquisition and processing through this implementation. The use of this technology results in a low power consumption system that fits a small box, compatible with the autonomous boat both in size and energy consumption.

The FPGA system is responsible for the generation of complex acoustic signals, for controlling the transmitting power amplifiers and the adaptive gain low noise receiving amplifiers, for demodulating and for match filtering the received signals with the transmitted waveform (Fig. 8).

The first task of the FPGA is therefore to control the digital-to-analog (D/A) and analog-to-digital interface (A/D). At each transmission pulse trigger, the FPGA reads the signal wave form from memory (that is stored in base-band, decimated form), converts the signal to pass-band (frequency up conversion), interpolates and filters the signal and finally supplies it to the D/A.

During the receiver stage, the samples are read from the A/D, converted to base-band (frequency down-conversion), filtered, decimated and finally placed in a FIFO memory.

Each transmit pulse is time-stamped using a real-time clock implemented in the FPGA system that is corrected using the time information and pulse-per-second trigger from a GPS receiver. This enables precise correlation with the navigation data.

The processor embedded in the FPGA bridges the low level hardware and the control PC, providing an ethernet access to the sonar.

The results are then supplied to an embedded computer for storage and acoustic image computation. The shore base station accesses this data through a high-speed digital radio-link making available the surveyed data.

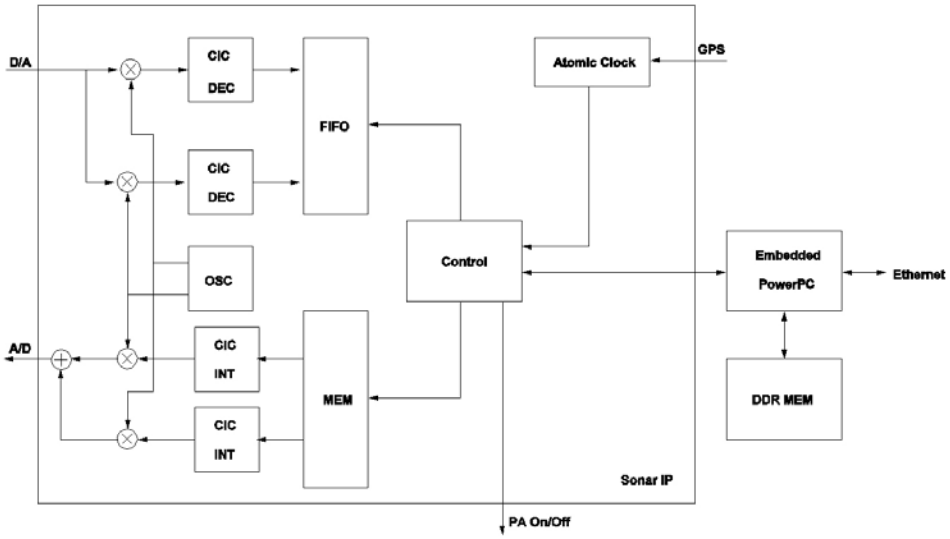

Fig. 8. FGPA system detail.

5. Navigation system

The accurate target area image formation process depends on the precise knowledge of the sensor position as it traverses the aperture. One of the major strengths of using a surface vessel is the use of satellite systems for navigation. Since the used wavelengths are in the order of centimetres, this is the level of relative accuracy that is required for the vessel along the length of the synthetic aperture. This is possible with GPS navigation if operated in differential mode using carrier phase measurements. The long term errors of a CP-DPGS solution for a short baseline to the reference station is in the order of the decimetre; the short term errors, due to noise in the carrier phase tracking loops, is in the order of the centimetre. This error can further be smoothed out through integration of the GPS solution with inertial measurements. For height computation through In-SAS processing of pairs of images, precise attitude data has to be obtained so to precisely transform the lever between two sensors into target position estimate. For this purpose inertial sensor play a crucial role: through blending with the GPS data, it supplies pitch and roll estimates with levels of accuracy in the order of arcminutes; the heading error estimate is slightly poorer, but has

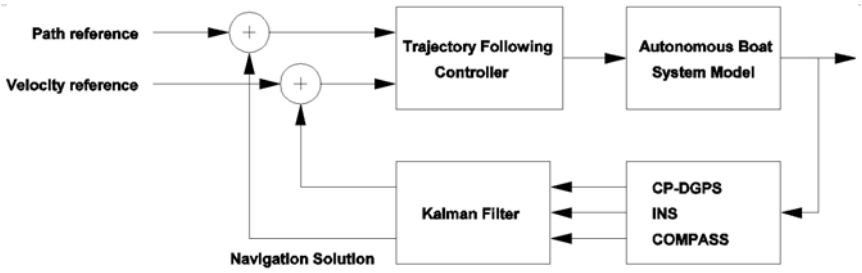

Fig. 9. Navigation system diagram.

less significant impact in the In-SAS processing algorithms. The Navigation subsystem is basically a Kalman filter mixing the dynamic behaviour of the vessel through the inertial sensor with the GPS reading. The major states are positioning, velocity and altitude (alignment) errors, but also includes inertial sensor biases. The output is a complete navigation solution, including position velocity and attitude information at a high data rate.

6. Autonomous boat control

To achieve the best results, the boat should follow the predefined paths at constant speed over ground and with minimal roll and pitch motions. Furthermore, it is desirable that the intended boat trajectories are described with minimal deviation to allow multiple sweeps in close parallel lines which are mandatory for dual passage interferometric height mapping.
To accomplish these goals, a control system that automatically drives the vehicle along user specified trajectories at given speeds was developed. The control system is organized in two independent control loops, one for the velocity and the other for the horizontal position of the vehicle. The velocity loop determines the common mode actuation of the boat while the horizontal plane loop determines the differential mode actuation. These values are then combined to produce the commands for the starboard and port thrusters.
The velocity loop is based on a proportional plus integral controller that assures that the velocity of the boat is, in steady state, the one defined by the user. The controller parameters where tuned to assure a smooth motion by rejecting high frequency noise from the navigation sensors. Different controller parameters can be used to obtain the best behaviour at different velocity ranges.
The horizontal plane loop implements a line tracking algorithm. It is composed by an outer controller that computes a heading reference based on the cross track error (distance of the boat to the desired straight line) and an inner controller that drives the vehicle heading to the given reference ([Cruz, N. et al (2007)]). This two stage control loop assures zero cross track error in steady state regardless of the water current and the tuning of the controller parameters took into account a dynamic model of the boat.
The whole system has already been tested in operational scenarios with great success. In particular, it has been observed that the vehicle describes intended trajectories with average cross tracking errors below 10 cm. For illustration purposes, Fig. 10 presents a histogram of the cross tracking error for a 50 m straight line described by the autonomous boat while collecting SAS data.

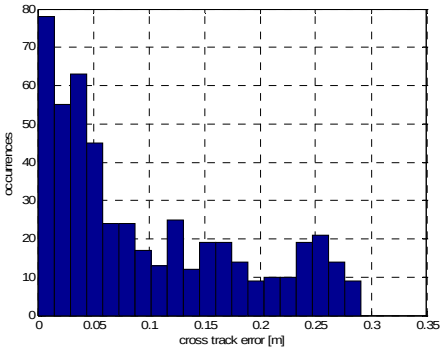

Fig. 10. Line tracking performance.

In Fig. 10 it is illustrated the path following ability of the autonomous boat, where it was programmed to execute two parallel profiles with a distance of 1 m of each other.

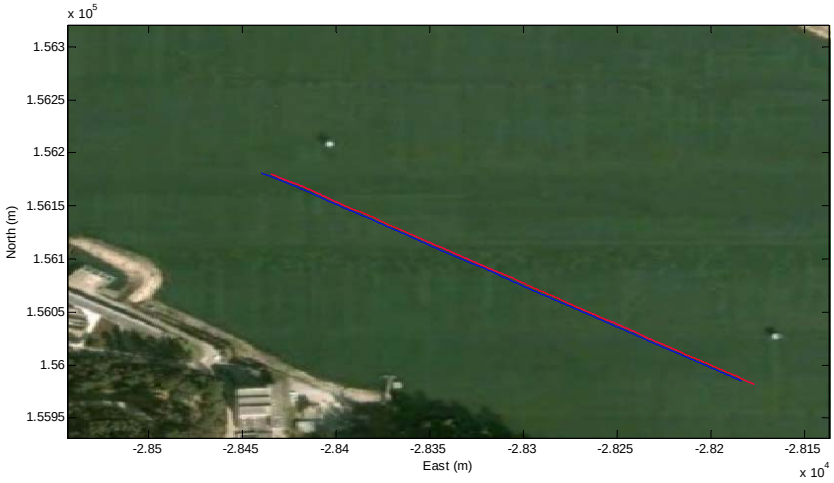

Fig. 11. Example path.

Fig. 12 shows the boat velocity graph, measured by the navigation system, during the execution of a profile.

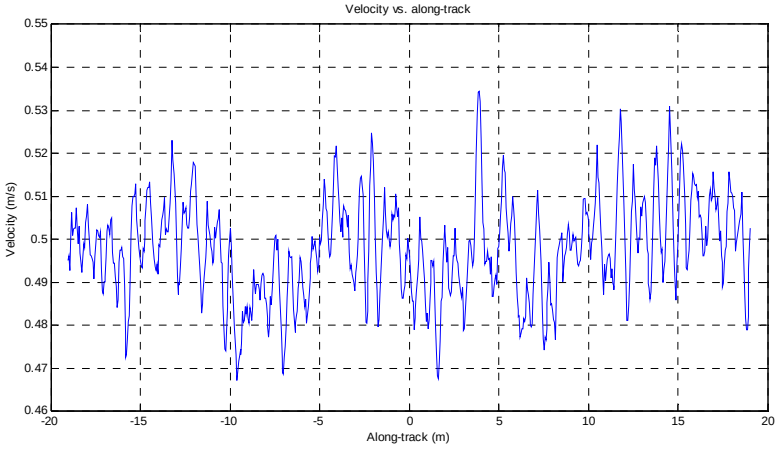

Fig. 12. Autonomous boat velocity during a mission.

7. Synthetic aperture sonar system model

Consider a platform that moves trough a nominal straight path. Since the boat is programmed to follow straight lines at a constant speed, this is a good assumption. The sonar transducers are rigidly coupled to the sonar structure, through which the acoustic signals are sent and their respective echoes are received (Fig. 13).

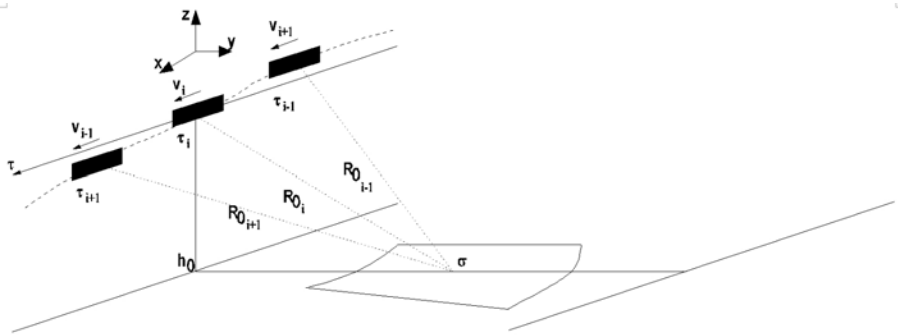

Fig. 13. Autonomous boat based SAS model.

Each echo contains indistinctive information of an area corresponding to the radiation pattern of the transducers. The transducers are adjusted so that a swath of about 50 m is at the boat left broadside, which is the strip-map configuration.

At each along-track sampling position, the echoes caused by the swath reflectivity are received and recorded along the track.

The echo ($\mathrm{ee}(\tau, \mathrm{t})$) data is formed trough the reflection of the transmitted pulse ($p_m(t)$) on the swath ($\mathrm{ff}(\mathrm{x},\mathrm{y}_s)$):

$$ee(\tau,t) = \iint ff(x,y_s)\, p_m\left(t - \frac{2}{c}\sqrt{(v\tau - x)^2 + y_s^{\,2}}\right) dxdy \tag{1.1}$$

Transducer beam pattern effects and signal propagation loss have not been included in the equation for simplicity but should be accounted for in the image reconstruction process.

The echo data is converter to base-band though multiplication with the signal carrier frequency:

$$ee_b(\tau,t) = ee(\tau,t)e^{-j2\pi f_0 t} \tag{1.2}$$

Let $p_b(t)$ be the baseband transmitted signal, the cross-track pulse compressed echo image will be:

$$ss_b(\tau,t) = ee_b(\tau,t) * p_b(t) \tag{1.3}$$

And thus resulting in the following equation:

$$ss_b(t,\tau) = \iint ff(x,y_s)\, p_c\left(t - t_v(\tau)\right) e^{-j2\pi f_0 t_v(\tau)} dxdy \tag{1.4}$$

Where $p_c(t)$ is the autocorrelation of the transmitted baseband signal and $t_v(\tau)$ is the time-of-flight and is given by:

$$t_v(\tau) = \frac{2}{c}\sqrt{\left(x_p(\tau) - x_0\right)^2 + \left(y_p(\tau) - y_0\right)^2 + \left(z_p(\tau) - z_0\right)^2} \tag{1.5}$$

The point-target has coordinates (x_0, y_0, z_0) and $(x_p(\tau), y_p(\tau), z_p(\tau))$ are the coordinates of the sonar platform position at each along-track sampling instant τ. In the absence of motion errors, $y_p(\tau)$ and $z_p(\tau)$ will be zero and $x_p(\tau) = v\tau$, where v is the platform velocity.

The task to retrieve the estimated image ($\hat{ff}(x, y_s)$) is done through the inversion of this model, coherently combining the along-track samples to form an image:

$$\hat{ff}(x, y_s) = \iint ss_b(\tau, t)\left(t - t_v(\tau)\right) e^{j2\pi f_0 t_v(\tau)} d\tau dt \tag{1.6}$$

The sonar platform seldom dislocates through a straight line, and so the sonar model must account for the irregular along-track sampling positions. Therefore the polar coordinate model must be abandoned and a broader expression for the echo travel time must contain the estimated position of the transducers and a rough estimation of the bottom height (a flat bottom is enough for a first approximation). Moreover, different transducer geometries mean that the travel time will not be exactly the simple double of the distance to reach the target, but the distance from the transmitting transducer to the target and back to the receiver which is given by:

$$t_v(\tau) = \frac{1}{c}\left(\left\|T_{TX}^{'}(\tau) - X_\sigma\right\| + \left\|X_\sigma - T_{RX}^{'}(\tau)\right\|\right) \tag{1.7}$$

Attitude variations must also be accounted for in the calculation of the transducer positions. Knowing the arm between the boat navigation centre (reference for the navigation system) and each transducer, and also the roll, pitch and yaw angles, the positions of the transducers can be correctly calculated. The transducers attitude also influences the swath footprint. During image synthesis this information should be used to calculate the correct weights of the combined echoes.

The frequency domain model of the imaging system can be obtained through Fourier transformation of equation (1.4) and the Principle of Stationary Phase ([Gough, P. T. (1998)]). The image reconstruction task can, as a result, be done by multiplication of the frequency domain inverse sonar imaging model.

For calculation of the obtained along-track resolution and sampling frequency, consider that the target will be seen by the sonar during the time it is inside the aperture (3dB lobe) of the transducer, which is approximately given by:

$$\theta_{3dB} \approx \frac{\lambda}{D} \tag{1.8}$$

The array spacing from Nyquist spatial sampling and classical array theory is $\lambda/2$ (two way equivalent 2π phase shift), which means that for angles of arrival of a wave-front the inter-element phase difference must be less than 2π ([McHugh, R. et al (1998)]).

This is also true for motion errors. To correctly form a synthetic aperture the platform position must be known within 1/8 of a wavelength so the echoes can be coherently combined with negligible image deterioration ([Cutrona, L. J. (1975); Tomiyasu, K. (1978); Fornaro, G. (1999)]).

We can also consider the array spacing to be given by a pulse repetition frequency (PRF) that is at least equal the maximum Doppler shift experienced by a target. The Doppler shift f_D is related to the radial velocity v_r by:

$$f_D = \frac{2v_r}{\lambda} = \frac{2v\sin\theta(\tau)}{\lambda} \tag{1.9}$$

The maximum radial velocity is obtained at the beam edge and so the lower bound for the PRF is ([McHugh, R. (1998)]):

$$PRF \geq \frac{2v\sin\theta_{3dB}}{\lambda} \approx \frac{2v\lambda / D}{\lambda} = 2\frac{v}{D} \tag{1.10}$$

The minimum synthetic array spacing is thus:

$$d_{SA} = v\frac{1}{PRF} = \frac{D}{2} \tag{1.11}$$

The along-track resolution is independent of the range and wavelength. This results from the fact that for a transducer with a fixed length D, the synthetic aperture length DSA will be given, approximately, by:

$$D_{SA} \approx 2R_0\theta_{3dB} = 2R_0\frac{\lambda}{D} \tag{1.12}$$

Where R_0 is the distance to the center of the scene.
This than gives the classical synthetic aperture along-track resolution δ_{AT} formula:

$$\delta_{AT} \approx R_0\theta_{SA} = R_0\frac{\lambda}{D_{SA}} = R_0\frac{\lambda}{2R_0\frac{\lambda}{D}} = \frac{D}{2} \tag{1.13}$$

We see here that the phase relations that enable the synthetic array formation are tightly related to the wavelength of the signal and the effective synthetic array length. Normally these two values are interconnected due to the transducers real aperture width, but can be explored to mitigate some of the problems inherent to synthetic aperture.
The image formed in this way has a cross-track resolution of c/2BW and an along-track resolution of $D/2$ (where c is the speed of sound, BW is the transmitted signal bandwidth and D is the effective transducer diameter). More importantly, the along-track resolution is independent of the target range. To correctly synthesize an image without aliasing artefacts in the along-track dimension, it is necessary to sample the swath with an interval of $D/2$ (considering the use of only one transducer for transmission and reception). This constraints, together with the maximum PRF defined by the longest distance of interest and the along-track sampling restrictions, imposes a very speed to a sonar platform ([Cutrona, L. J. (1975); Gough, P. T. (1998)]).

8. Image formation process

The sonar acquires the data in pass-band format which is then converted to base-band and recorded. Starting with this uncompressed base-band recorded data, the first step in image

formation is cross-track pulse compression. This is also known as match filtering. This is step is necessary because using a longer transmitting pulse carries more energy than a short pulse with the same peak power which enhances the signal-to-noise ratio. The resulting cross-track resolution is not given by the duration of the transmitted pulse, but instead by its bandwidth. The task of pulse compression is done through correlation of the received data with the base-band transmitted pulse.

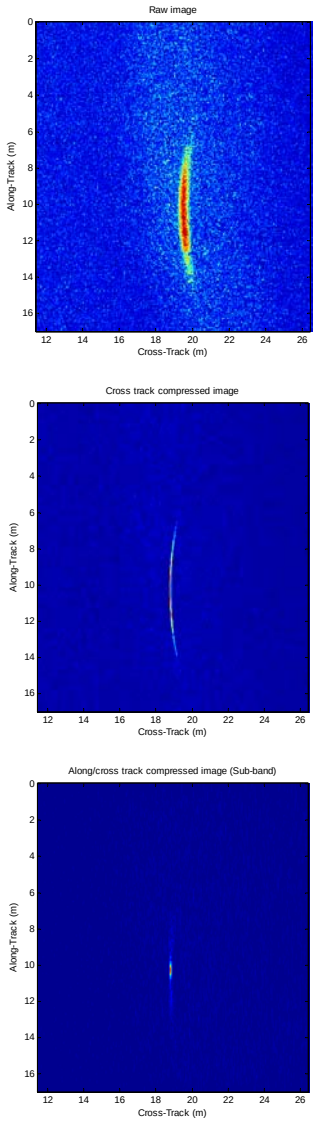

Fig. 14. Raw image, cross-track compressed image and along/cross-track compressed image.

At this stage data filtering and frequency equalization can be applied.
The next step is synthetic aperture formation that should use the available navigation data to synthesize the virtual array and form the sonar image. Fig. 14 shows these steps in succession for an image of an artificial target placed in the river bottom for a test mission. Note that the first image has low along and cross track resolution because its unprocessed, the second image has better cross-track resolution due to pulse compression and finally the last image, which is the result of synthetic aperture processing, resembles a small point.
Synthetic aperture image formation can be done through the use of several algorithms which can be classified into frequency domain algorithms, such as the wave-number algorithm, chirp scaling algorithm or the inverse scaled Fourier transform algorithm, and time domain algorithms such as the explicit matched filter or the back-projection algorithm ([Gough, P. T. (1998); Silkaitis, J.M. et al (1995)]).
The wave-number algorithm relies on inverting the effect of the imaging system by the use of a coordinate transformation (Stolt mapping) through interpolation in the spatial-frequency domain. The compressed echo data is converted to the wavenumber domain (along/cross-track Fourier transforms), matched filtering is applied supposing a target at a reference range followed by a nonlinear coordinate transformation ([Gough, P. T. (1998)]).
The chirp-scaling algorithm avoids the burdensome non-linear interpolation by using the time scaling properties of the chirps that are applied in a sequence of multiplications and convolutions. Nevertheless the chirp scaling algorithm is limited in use to processing of uncompressed echo data obtained by the transmission of chirp signals.
An approach based on the inverse scaled Fourier transform (ISFFT) previously developed for the processing of SAR data can also be followed. This algorithm interprets the raw data spectrum as a scaled and shifted replica of the scene spectrum. This scaling can then be removed during the inverse Fourier transformation if the normal IFFT is replaced by a scaled IFFT. This scaled IFFT can be implemented by chirp multiplications in the time and frequency domain (Fig. 15). The obtained algorithm is computationally efficient and phase preserving (e.g. fit for interferometric imagery). Motion compensation can be applied to the acquired data in two levels: compensation of the known trajectory deviations and fine corrections trough reflectivity displacement, auto-focus or phase-retrieval techniques. The deviations from a supposed linear path are compensated thorough phase and range shift corrections in the echo data. Velocity variations can be regarded as sampling errors in the along-track direction, and compensated through resampling of the original data ([Fornaro, G. (1999)]).
The back-projection algorithm, on the other hand, enables perfect image reconstruction for any desired path (assuming that rough estimate of the bottom topography is known), since it does not rely on the simple time gating range corrections ([Hunter, A. J. et al (2003); Shippey, G. et al (2005); Silva, S. (2007b)]). Instead, it considers that each point in one echo is the summation of the contributions of the targets in the transducer aperture span with the same range. With this algorithm one is no longer forced to use or assume a straight line for the sonar platform displacement. The platform deviations from an ideal straight line are not treated as errors, but simply as sampling positions. In the same way, different transducers array geometries are possible without the need for any type of approximation. This class of synthetic aperture imaging algorithms, although quite computational expensive in comparison with frequency domain algorithms, lends itself very well to non-linear acquisition trajectories and, therefore, to the inclusion of known motion deviations from the expected path. To reconstruct the image each echo is spread in the image at the correct

coordinates (back-projected) using the known transducer position at the time of acquisition (Fig. 16). It is also possible to use an incoherent version of this algorithm (e.g.: that does not use phase information). But the obtained along-track resolution is considerably worse ([Foo, K.Y. et al (2003)]).

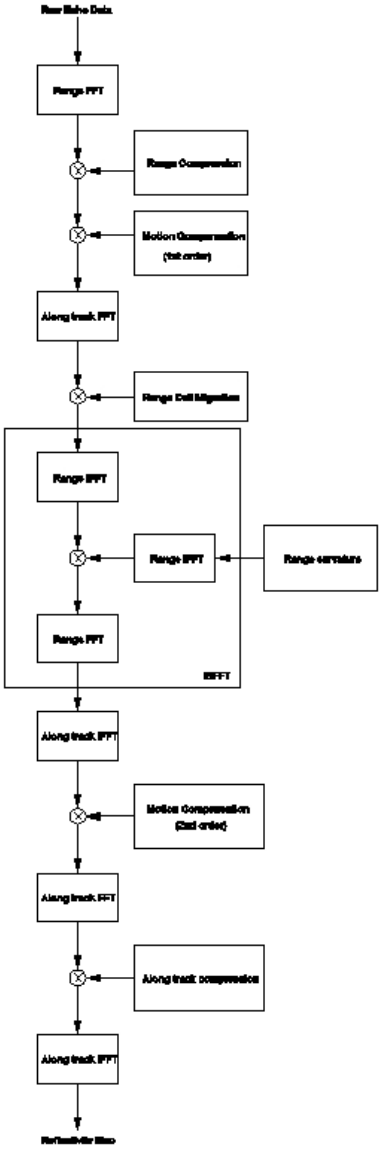

Fig. 15. ISFFT algorithm flow diagram.

The back-projection algorithm can also be implemented in matrix annotation ([Silva, S. et al (2008 a)]). The navigation information and system geometry is used to build the image

formation matrix leading to the reconstructed image. The transmitting and receiving beam patterns and the corresponding swath variation with the platform oscillation is also weighted in the matrix. This makes this algorithm well suited for high resolution sonar systems with wide swaths and large bandwidths that have the assistance from high precision navigation systems. The main advantage of this algorithm is the ease of use within an iterative global contrast optimization auto-focus algorithm ([Kundur, D. et al (1996)]). The image formation is divided into two matrixes: a fixed matrix obtained from the sonar geometrical model and navigation data (corresponds to the use of a model matching algorithm, such as the explicit matched filtering); and a matrix of complex adjustable weights that is driven by the auto-focus algorithm. This is valid under the assumption that the image formation matrix is correct at pixel level and the remaining errors are at phase level (so that the complex weight matrix can correct them).

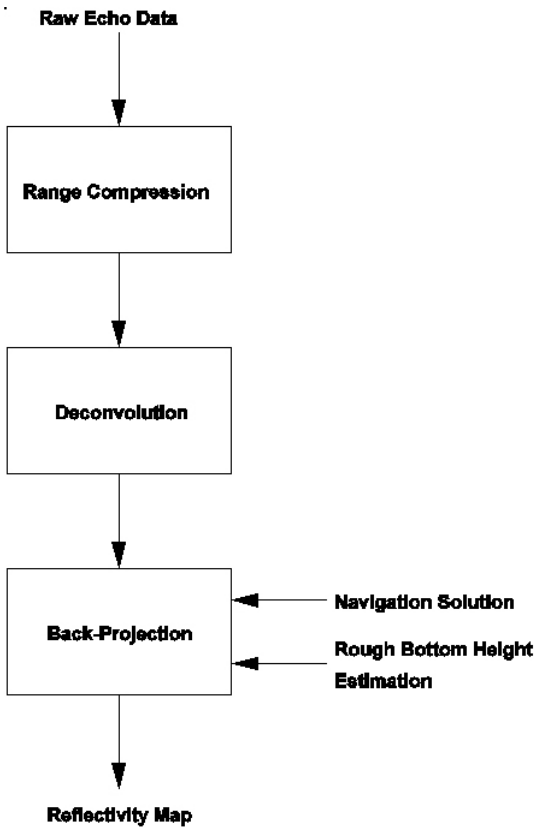

Fig. 16. Back-projection algorithm signal flow diagram.

9. Auto-focus

Since the available navigation data sources, be it DGPS or INS systems, cannot provide enough precision to enable synthetic aperture processing of high resolution (high frequency)

sonar data [Bellettini et al (2002); Wang et al (2001)], the phase errors caused by the unknown motion components and medium turbulence must be estimated to prevent image blurring.
Auto-focus algorithms exploit redundancy and or statistical properties in the echo data to estimate certain image parameters that lead to a better quality image. Therefore, the auto-focus problem can be thought as a typical system estimation problem: estimate the unknown system parameters using a random noise input. If the auto-focus algorithms estimates the real path of sonar platform they are called micronavigation algorithms [Bellettini et al (2002)] (sometimes with the aid of navigation sensors such as inertial units) otherwise they are generically designated as auto-focus algorithms. Redundant phase centre algorithm and shear average algorithm are examples of micronavigation algorithms.
Since redundancy in data is greatly explored, common auto-focusing algorithms require restrictively along-track sample rates equal or higher than the Nyquist sample rate. This imposes unpractical velocity constrains, especially for system that use few receivers (as is the case with the sonar system described here). It is not possible to obtain micro-navigation from an under-sampled swath or to perform displaced centre phase navigation with only one transducer. So, with these impairments, global auto-focus algorithms are required in sonar systems that use simple transducers arrays and under-sampled swath. The use of global auto-focus algorithm presents several advantages for synthetic aperture sonar image enhancing. They differ from other algorithm because they try to optimize a particular image metric by iteratively changing system parameters instead of trying to extract these parameters from the data. Global auto-focus algorithms can correct not only phase errors due to navigation uncertainties, but also phase errors that are due to medium fluctuations.
It is required that the synthetic aperture algorithm uses the available navigation solution to form an initial image. Starting with the available navigation solution, the errors are modelled in a suitable way. If the expected errors are small they can be modelled as phase errors for each along-track position. If the sonar platform dynamic model is known, the number of search variables can be greatly reduced by parameterizing this model ([Fortune, S. A. et al (2001)]). These parameters are weighted together with the image metric and serve as a cost function for the optimization algorithm to search the solution space (Fig. 18).
Nevertheless, these errors are hardly ever smaller than the original signal wavelength, and so create a solution surface that is difficult to search for the optimum set of parameters. However, if we have access to the raw data, by dividing the received signal bandwidth in several smaller bands and conjugate complex multiplying the pulse compressed signals obtained in each band one by the other, a new resulting signal is obtained with an effective longer wavelength corresponding to the frequency difference between the two sub-bands ([Silva, S. (2008 b)]). This longer wavelength effectively reduces the impact of phase fluctuation from the medium and platform motion uncertainties. Using this, it is possible to divide the signal bandwidth into several sub-bands and combine them in to signals with different wavelengths. At the first step, a large wavelength is used since the expected motion correction is also large. After achieving a predefined level of image quality, the auto-focus algorithm then proceeds by using a smaller wavelength and the previous estimated position parameters.
This step is repeated with decreasingly smaller wavelength and position error, until the original wavelength is used. The result is a faster progression through the solution surface, with lower probabilities of falling into local minima.

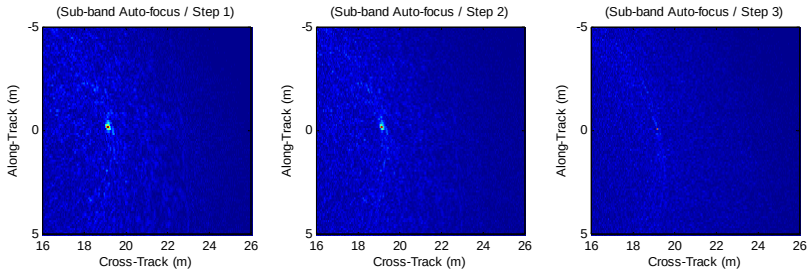

Fig. 17. Sonar image of the artificial target through the various auto-focus steps.

Fig. 17 shows an image of an artificial point target in 3 successive auto-focus steps. The algorithm starts wit a longer wavelength thus producing a low resolution image. As it progresses through the process, the target gets a sharper appearance.

For image quality metric a quadratic entropy measure can be used, which is a robust quality measure and enables fast convergence than a first order entropy measure or a simple image contrast measure. This is a measure of image sharpness. The lower the entropy measure, the sharper the image.

To calculate the quadratic entropy one needs to estimate the image information potential IP. Instead of making the assumption that the image intensity has a uniform or Gaussian distribution, the probability density function is estimated thought a Parzen window method using only the available data samples ([Liu, W. et al (2006)]):

$$IP(x) = \frac{1}{N^2} \sum_{j=1}^{N} \sum_{i=1}^{N} k_\sigma (x_j - x_i) \tag{1.14}$$

Where $k_\sigma (x - x_i)$ is the Gaussian kernel defined as:

$$k_\sigma (x - x_i) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-x_i)^2}{2\sigma^2}} \tag{1.15}$$

Because this method of estimation requires a computational intensive calculation of the sum of Gaussians, this is implemented through the Improved Fast Gaussian Transform described in [Yang, C. et al (2003)].

This auto-focus method is suitable for systems working with an under-sampled swath and few transducers. No special image features are necessary for the algorithm to converge.

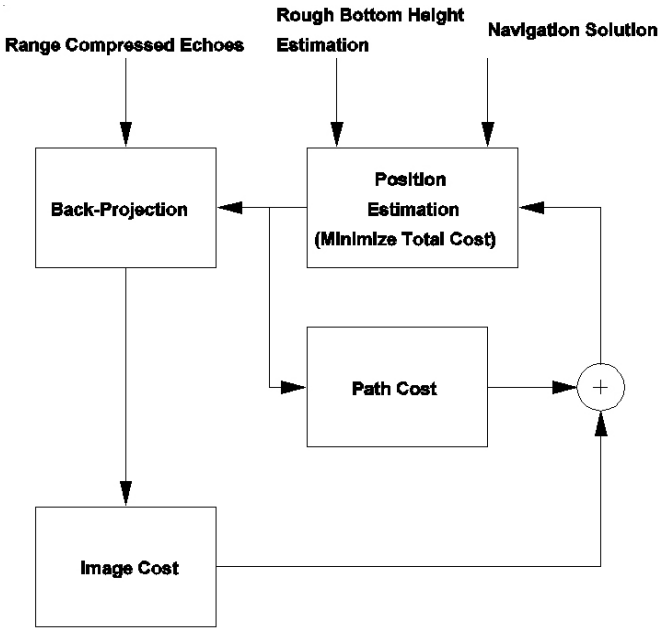

Fig. 18. Auto-focus block diagram.

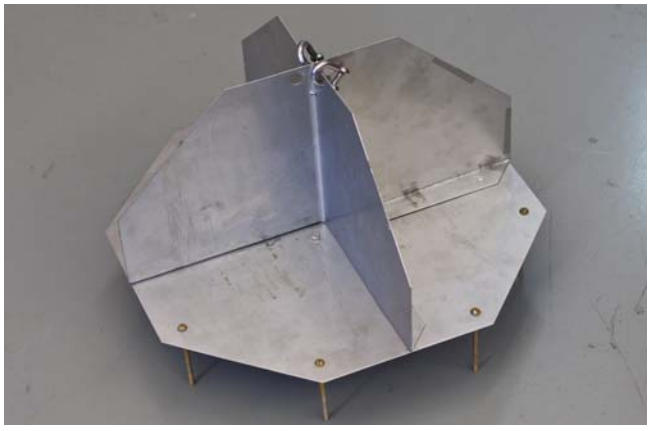

Fig. 19. Artificial target used for resolution tests.

10. Results

To test the system and access its capabilities a series of test missions were performed in the Douro river, Portugal. For the first tests an artificial target was placed in the muddy river bottom and the autonomous boat programmed to make several paths through the area.
The artificial target is a half octahedral reflector structure made of aluminium (Fig. 19). It measures 20x20x20cm, but the target response seen by the sonar should be like a point after correct image synthesis.

Fig. 20 shows one image of the artificial target obtained though the matrix implementation of the back-projection algorithm as describe previously.

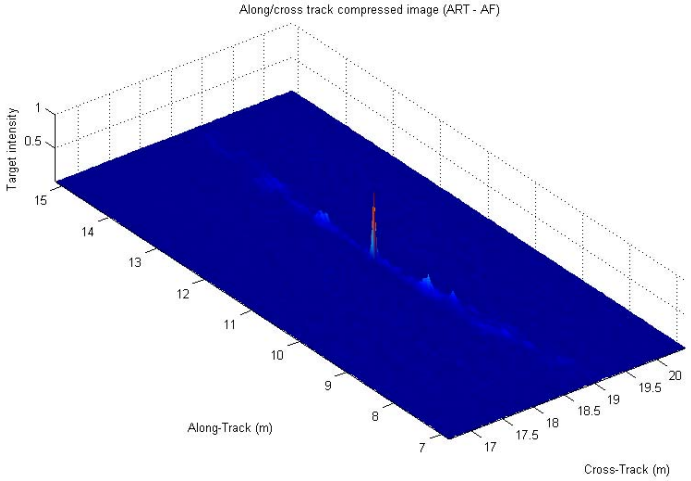

Fig. 20. Sonar image of the artificial target placed in the river bottom.

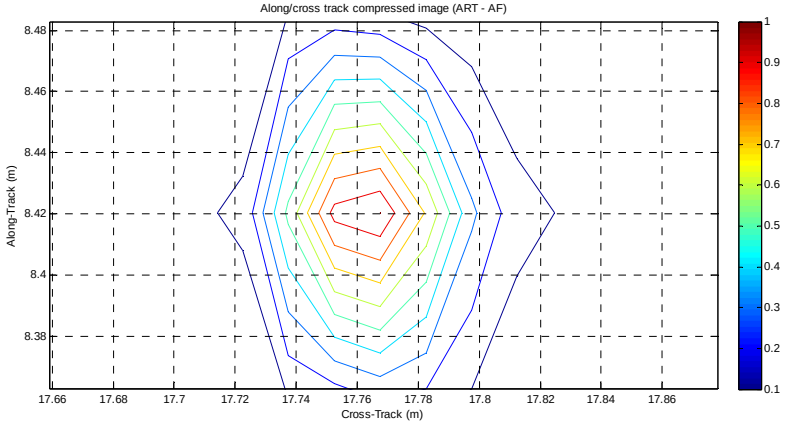

Fig. 21. Synthetic aperture sonar resolution.

As can be seen if Fig. 21, after auto-focus the image obtained from the artificial target presents sharp point like response, achieving the theoretical maximum resolution of the sonar system: 2.5x2.5 cm.

Fig. 22 shows an image obtained near the river shore before synthetic aperture processing and Fig. 23 show the same image processed using the described back-projection algorithms. It is possible to see several hyperbolic like target responses from rocks in the river bed that, after synthetic aperture image processing, assume the correct point like form.

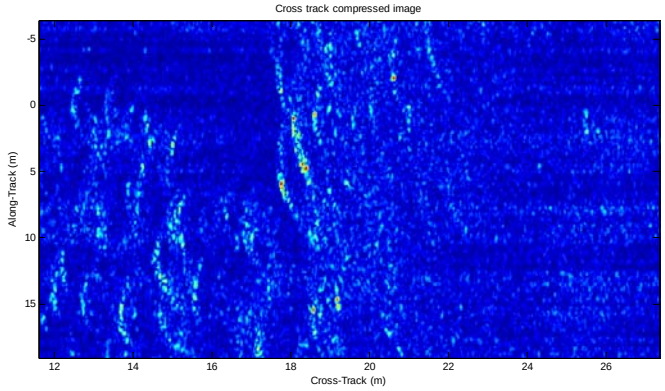

Fig. 22. Cross-track compressed reflectivity map an area near the river shore.

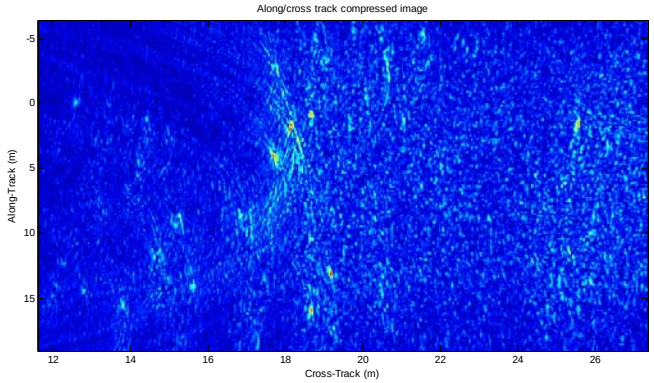

Fig. 23. Along/Cross-track compressed reflectivity map an area near the river shore.

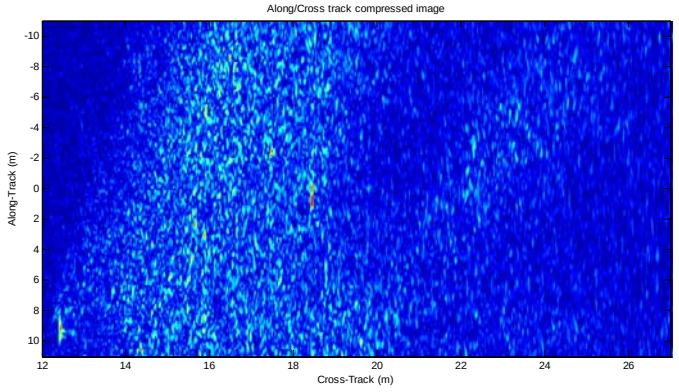

Fig. 24. Reflectivity map of harbour entrance.

The hyperboles are wavy due to the uncontrolled platform motion but, because the platform motion is known, the image nevertheless correctly synthesized.
Fig. 24 shows another sonar image obtained with the described system, this time it's a sand bank near a small harbour entrance in the Douro river.

11. Conclusion

As demonstrated, synthetic aperture sonar is a technique that enables attainment of high quality, high resolution underwater images.
Autonomous surface vehicles provides several advantages for synthetic aperture imagery. Not only it is possible to control the boat motion in this way, it is also possible to obtain navigation measurements with precisions in the order of the wavelength used in high resolution sonar systems. Furthermore unsupervised surveillance applications that combine the high quality sonar images with the effectiveness of an autonomous craft are possible.
Sonar images obtained in this way can be easily integrated in geographical information systems.
Using back-projection algorithms one is no longer restricted to linear paths, and deviations from this path are not treated as errors, but simply as sampling positions.
Phase errors due to navigation uncertainties and medium fluctuations cause blurring of the image. Nevertheless, the results can be further enhanced through auto-focus procedures that iterate the solutions until convergence to a predefined image quality parameter is achieved.
The use of high frequency signals imposes demanding restrictions in motion estimation and medium stability due to the sensibility of the image formation process to phase errors. A clever combination of the received signals enables the creation of a new one with an equivalent frequency equal to difference of the centre frequencies of the previous ones. This longer wavelength signal effectively masks phase uncertainties and enables efficient auto-focus of the synthetic aperture sonar image.
The synthetic aperture sonar thus enables enhanced imagery of underwater realms that combined with suitable platforms, such as an autonomous boat, can be obtained at low cost and high availability.

12. Future work

Synthetic aperture sonar imagery is a powerful technique for underwater imaging that is becoming widespread as the problems and difficulties inherent to it are solved.
The use of multiple receiver systems will enable a higher area coverage ratio and ease navigation system precision requirements through the use of micronavigarion techniques. With the maturing of image formation algorithms, real-time image formation will further extend the application possibilities of synthetic aperture sonar systems.
Bottom height mapping is possible through the use of a double array of transducers and also by exploring the possibility of dual-pass interferometry. In this case the combination of images of the same scene obtained from different positions of the platform will allow the construction of three dimensional maps of the analyzed surfaces.

13. References

Bellettini, A.; Pinto, M. A., (2002) . "Theoretical accuracy of the synthetic aperture sonar micronavigation using a displaced phase centre antenna", *IEEE Journal of Oceanic Engineering*, 27(4), pp. 780-789,, October 2002.

Chatillon, J.; Adams, A. E.; Lawlor, M. A.; Zakharia, M. E. (1999). "SAMI: A Low-Frequency Prototype for Mapping and Imaging of the Seabed by Means of Synthetic Aperture", *IEEE Journal of Oceanic Engineering*, vol. 24, no. 1, pp. 4-15, January 1999.

Cruz, N.; Matos, A.; Cunha, S.; Silva, S. (2007). "ZARCO - An Autonomous Craft for Underwater Surveys", *Proceedings of the 7th Geomatic Week*, Barcelona, Spain, Feb 2007.

Cutrona, L. J. (1975), "Comparison of sonar system performance achievable using synthetic-aperture techniques with the performance achievable by more conventional means", *The Journal of the Acoustical Society of America*, Volume 58, Issue 2, August 1975, pp. 336-348.

Douglas, B. L.; Lee, H. (1993) "Synthetic-aperture Sonar Imaging with Multiple-element Receiver Array", *IEEE International Conference on Acoustics*, Speech, and Signal Processing, vol. 5, pp. 445-448, April 1993.

Foo, K.Y.; Atkins, P. R.; Collins, T. (2003) "Robust Underwater Imaging With Fast Broadband Incoherent Synthetic Aperture Sonar", *Proceedings of IEEE International Conference on Acoustics*, Speech, and Signal Processing 2003, Volume 5, 2003 pp. V - 17-20 vol.5.

Fornaro, G. (1999). "Trajectory deviations in airborne SAR: analysis and compensation", *IEEE Transactions on Aerospace and Electronic Systems*, vol.35, no.3, pp.997-1009, Jul 1999.

Fortune, S. A.; Hayes, M. P.; Gough, P. T. (2001). "Statistical Autofocus of Synthetic Aperture Sonar Images using Image Contrast Optimization", *OCEANS '01 Conference Proceedings*.

Gough, P. T.; Hayes, M. P. (1989), "Measurements of acoustic phase stability in Loch Linnhe, Scotland", *The Journal of the Acoustical Society of America*, Volume 86, Issue 2, August 1989, pp.837-839.

Gough, P. T. (1998). "Unified Framework for Modern Synthetic Aperture Imaging Algorithms". *The International Journal of Imaging Systems and Technology*, Vol. 8, pp. 343-358, 1998.

Hansen, R.E.; Saebo, T.O.; Callow, H.J.; Hagen, P.E.; Hammerstad, E. (2005) "Synthetic aperture sonar processing for the HUGIN AUV", *Oceans 2005 - Europe*, vol.2, no., pp. 1090-1094 Vol. 2, 20-23 June 2005.

Hawkins, D. W.; Gough, P. T. (2004). "Temporal Doppler Effects in SAS", *Sonar Signal Processing*, Vol. 26, p. 5, 2004.

Hunter, A. J. ; Hayes, M. P. ; Gough, P. T. (2003). "A Comparison of Fast Factorised Back-Projection and Wavenumber Algorithms For SAS Image Reconstruction", *Proceedings of the World Congress on Ultrasonics*, Paris, France, September 2003.

Kundur, D.; Hatzinakos, D. (1996). "Blind image deconvolution", *Signal Processing Magazine*, IEEE, vol.13, no.3, pp.43-64, May 1996.

Liu, W.; Pokharel, P. P.; Principe, J. C.. (2006). "Correntropy: A Localized Similarity Measure", *International Joint Conference on Neural Networks*, 2006, pp. 4919 – 4924.

McHugh, R.; Shaw, S.; Taylor, N. (1998). "Spatial sampling ambiguities in synthetic aperture sonar for broadside and squint look operations", *OCEANS '98 Conference Proceedings*, vol.2, no., pp.960-964 vol.2, 28 Sep-1 Oct 1998.

Nelson, M. A. (1998). "DARPA Synthetic Aperture Sonar", *Proceedings of the Adaptive Sensor Array Processing (ASAP) Workshop*, vol. 1, pp. 141-155, 15 May 1998.

Neudorfer, M.; Luk, T.; Garrood, D.; Lehtomaki, N.; Rognstad, M. (1996) "Acoustic Detection and Classification of Burried Unexploded Ordinance (UXO) Using Synthetic Aperture Sonar", *PACON' 96*, 1996.

Putney, A.; Anderson, R.H. (2005) "Reconstruction of undersampled SAS data using the WIPE algorithm", *Proceedings of the OCEANS 2005 MTS/IEEE Conference*, vol., no., pp. 111-118 Vol. 1, 2005.

Saebo, T.O.; Callow, H.J.; Hansen, R.E.; Langli, B.; Hammerstad, E.O. (2007) "Bathymetric Capabilities of the HISAS Interferometric Synthetic Aperture Sonar", *Proceedings of the OCEANS 2007 MTS/IEEE Conference*, vol., no., pp.1-10, Sept. 29 2007-Oct. 4 2007.

Sammelmann, G. S.; Fernandez, J. E.; Christoff, J. T. (1997) "High Frequency / Low Frequency Synthetic Aperture Sonar", *Proocedings of SPIE – International Society for Optical Engineering*, vol. 2079, pp-160-171, 1997.

Shippey, G.; Banks, S.; Pihl, J. (2005). "SAS image reconstruction using Fast Polar Back Projection: comparisons with Fast Factored Back Projection and Fourier-domain imaging", *Proceedings of the OCEANS 2005 MTS/IEEE-Europe Conference*, vol.1, no., pp. 96-101 Vol. 1, 20-23 June 2005.

Silkaitis, J.M.; Douglas, B.L.; Hua Lee. (1995). "Synthetic-aperture sonar imaging: system analysis, image formation, and motion compensation", *Conference Record of the Twenty-Ninth Asilomar Conference on Signals, Systems and Computers 1995*, vol.1, no., pp.423-427 vol.1, 30 Oct-1 Nov 1995.

Silva, S.; Cunha, S.; Matos, A.; Cruz, N. (2007 a). "An In-SAS System For Shallow Water Surveying", Proceedings *of the 7th Geomatic Week, Barcelona*, Spain, Feb 2007.

Silva, S.; Cunha, S.; Matos, A.; Cruz, N. (2007 b). "An Autonomous Boat Based Synthetic Aperture Sonar," *Proceedings of the OCEANS 2007 MTS/IEEE Conference*, vol., no., pp.1-7, Sept. 29 2007-Oct. 4 2007.

Silva, S.; Cunha, S.; Matos, A.; Cruz, N. (2008 a). "An Algebraic Approach to Synthetic Aperture Sonar Image Reconstruction", *Proceedings of the OCEANS 2008 MTS/IEEE Conference*, Sept. 15 2008-Sept. 18 2008.

Silva, S.; Cunha, S.; Matos, A.; Cruz, N. (2008 b). "Sub-Band Processing of Synthetic Aperture Sonar Data", *Proceedings of the OCEANS 2008 MTS/IEEE Conference*, Sept. 15 2008-Sept. 18 2008.

Silva, S.; Cunha, S.; Matos, A.; Cruz, N. (2008 c). "Shallow Water Height Mapping With Interferometric Synthetic Aperture Sonar", *Proceedings of the OCEANS 2008 MTS/IEEE Conference*, Sept. 15 2008-Sept. 18 2008.

Tomiyasu, K. (1978). "Tutorial review of synthetic-aperture radar (SAR) with applications to imaging of the ocean surface", *Proceedings of the IEEE*, vol.66, no.5, pp. 563-583, May 1978.

Wang, L.; Belletini, A.; Fioravanti, S.; Chapman, S. ; Bugler, D. R.; Perrot, Y.; Hétet, A. (2001). "InSAS'00: interferometri SAS and INS aided SAS imaging", *Proceeding of the IEEE Oceans 2001 Conference*, volume 1, pp. 179-187, MTS/IEEE, 2001.

Yang, C.; Duraiswami, R.; Gumerov, N. A.; Davis, L. (2003). "Improved Fast Gauss Transform and Efficient Kernel Density Estimation", *Proceedings of the Ninth IEEE International Conference on Computer Vision (ICCV'03)*, 2003, pp. 664 - 671 vol.1.