A ROBUST ALGORITHM FOR REVERBERATION SUPPRESSION IN DOPPLER SENSITIVE TRANSMISSIONS

NH Parsons Thales Underwater Systems Ltd., Cheadle Heath, Stockport, UK

1 INTRODUCTION

As is well known, the Doppler shift of the signal received on a moving active sonar, from the sea bed, is dependent on the angle of arrival. Using CW, for example, to a good approximation the Doppler frequency of the return varies as

$$f = f_0(1+2(v/c)\cos\theta)$$

Equation 1

where f_0 is the transmit frequency, v is the platform speed, c is the speed of sound and θ is the angle of arrival relative to forward endfire. This relationship defines the so called reverberation ridge. Inverting this formula we see that the bottom reverberation has a direction of arrival which depends on the Doppler frequency ($f - f_0$) as

$$\cos\theta = (f - f_0)c/(2vf_0)$$

Equation 2

Of course, for more general Doppler sensitive pulses, there will be a similar relationship between angle of arrival of sea bed reverberation and Doppler bin number for Doppler bins corresponding to relative speeds between $\pm v$. Therefore, in a particular Doppler bin, the performance of the sonar may be severely reverberation limited in beams pointing in directions close to the direction given by Equation 2 (the so-called A-zone, in commonly used parlance). Even for beams which do not point directly at the reverberation ridge, the reverberation may seriously limit the detection performance in Doppler bins corresponding to relative speeds between $-v$ and $+v$ (the so-called B-zone) as there may be significant reverberation power entering the beam sidelobes from the direction of the reverberation ridge.

Figure 1 shows the reverberation ridge for some typical CW data. The A-zone ridge clearly shows the dependence between Doppler and angle of arrival implied by Equation 1 & Equation 2. One can also clearly see the vertical striations in the B-zone caused by the leakage through beam sidelobes. For Doppler bins corresponding to speeds outside the range $-v$ to $+v$ (the so called C-zone) the noise level corresponds to the ambient noise.

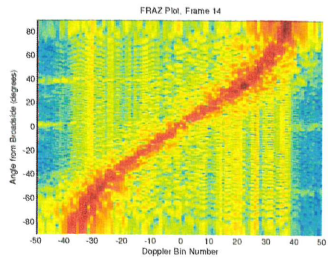

Figure 1: The Reverberation Ridge

In principle, the 'B-zone' reverberation can be significantly reduced by forming appropriate nulls in the beam patterns over the range of angles affected by the reverberation ridge. The formation of such nulls is quite straightforward, in theory, but in practice their quality (e.g. depth and width) may be very susceptible to unknown mismatch errors between different hydrophones. The algorithm described here uses beam-space adaptive beamforming techniques in order to use incoming data to essentially calibrate the array so that

good quality nulls can be formed in beam sidelobes without having closely matched hydrophones. It is a development of an algorithm described in an earlier publication[1]. A short calibration phase is found to be adequate: typically two pings.

2 ALGORITHM DESCRIPTION

The algorithm is based on the minimum variance distortionless response[2] (MVDR) adaptive algorithm applied in beam-space. This algorithm uses the calibration data to estimate the covariance matrix between different beams. The covariance matrix is used to calculate a set of complex weights which recombine the beams in order to produce a new set of beams which minimise the total energy of the calibration data, subject to a constraint of unit response in the look direction. The original beam in the look direction is regarded as the 'primary' beam, whilst the rest are regarded as 'reference' beams. The weights are calculated as[2]

$$\mathbf{b} = \frac{\mathbf{R}^{-1}.\mathbf{a}}{\bar{\mathbf{a}}.\mathbf{R}^{-1}.\mathbf{a}}$$

Equation 3

where **R** is the covariance matrix of the set of beams selected for recombination and **a** is the constraint vector, given by the complex response of each beam in the look direction. Because this algorithm minimises the energy, if the calibration data is from an acoustic field with very strong reverberation in some narrow sector, the beams resulting from the recombination with this set of weights will have a reduced response in this narrow sector. That is, the new beams will possess a null in the direction of the reverberation. The only place in which hydrophone amplitude and phase errors have an effect is in the estimation of the constraint vector **a**. This can affect the accuracy of the response in the look direction, but is unlikely to affect the ability to form a good quality null. The effect of the mismatch errors is also likely to be diminished due to the fact that the elements of **a** are beam responses, which are subject to averaging over hydrophones.

Of course, as discussed in the last section, the direction of arrival of the reverberation is dependent on the Doppler bin. Thus it is necessary to form nulls in different directions. This implies that several covariance matrices are required, each one corresponding to a different direction of arrival of the reverberation, and, therefore, producing a null in a different direction. Thus a set of N_C covariance matrices, of size $N_{beams} \times N_{beams}$, are defined, with each one to be used for reverberation arriving from a different range of angles in order that a null can be formed to reduce the beam responses over the corresponding sector.

The algorithm begins with conventional beamforming, to produce N_{beams} beams, and matched filtering using Doppler shifted replicas into separate Doppler channels for each range bin. The beams with highest power in each Doppler bin, in the current range bin (we call these 'Level 1' reference beams), are used to estimate the instantaneous slope of the reverberation ridge. That is the choice of which beams have highest power, as a function of Doppler bin, is used in a straight line fit of Doppler versus $\cos\theta$ in order to estimate the tow speed, v. Knowing the instantaneous tow speed allows us to assign, or map, each Doppler bin to a 'closest' covariance matrix defined as the one corresponding to a direction of arrival closest to the value of θ given by Equation 2. Of course as tow speed varies, the mapping will vary correspondingly. This mapping is recalculated at regular intervals (e.g. in CW processing at every range slice or FFT frame) both during the calibration phase and afterwards in the actual operational phase.

2.1 Calibration Phase

During the calibration phase, each covariance matrix estimate is integrated by summing the contributions of cross products of beam data with complex conjugated beam data, from all Doppler bins with the appropriate mapping. Thus, if Doppler bins $n_{c,min}$ to $n_{c,max}$ are mapped to covariance matrix R_c, ($1 \le c \le N_C$) the estimate is updated by

$$R_{c,kl}^{new} = R_{c,kl}^{old} + \sum_{n=n_{c,\min}}^{n_{c,\max}} x_{n,k} x_{n,l}^{*}$$

Equation 4

where $x_{n,k}$ is the signal in Doppler bin n, beam k in the current range slice and the * operation denotes complex conjugation. A 'popularity score' is also kept, and updated, of the number of times each beam is selected as a 'Level 1' reference beam over all Doppler bins mapping to R_c (see Figure 2). Thus N_C popularity scores are calculated.

2.2 Calculation of Weights

After the calibration phase the weights for each covariance matrix are calculated. However, instead of using the full covariance matrix a sub-matrix is selected consisting of a primary beam and a set of reference beams. These are selected by using the 'popularity score' to select the most popular beams (Figure 2). This defines a set of so-called 'Level 2' reference beams. The differences between 'Level 1' and 'Level 2' beams are listed in Table 1. A sub-matrix of the full covariance matrix, based on the rows and columns associated with the primary beam and the 'Level 2' reference beams is selected. This is inverted and used in Equation 3 to calculate the set of weights for the primary beam and the 'Level 2' reference beams for this covariance matrix. The same calculation is performed for each of the beams and each of the N_C covariance matrices as primary beam in order to produce a null in the sidelobe pattern of every beam. The advantage of the selection of such a sub-matrix is that the rank of the covariance matrix is reduced. This results in a more stable set of weights and requires less time samples of data to estimate the covariance matrix to sufficient accuracy. The sets of weights (and, therefore, the matrix inversions) only need to be calculated once after calibration as the N_C sets of 'Level 2' reference beams are fixed.

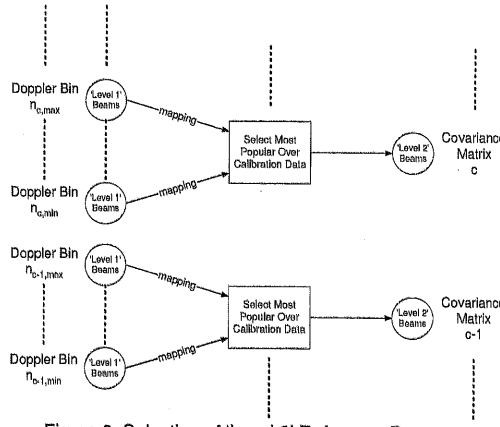

Figure 2: Selection of 'Level 2' Reference Beams.

Table 1: Differences between "Level 1" and "Level 2" reference beams.

"Level 1" reference beams	"Level 2" reference beams
Associated with individual Doppler bins.	Associated with covariance matrices
Sets change regularly (e.g. every FFT frame in CW).	Sets constant after completion of covariance matrix integration.
Used to determine tow speed and, during integration, tallied to determine "Level 2" reference beams.	They are the beams which are combined with 'primary' beams to suppress reverberation (by creating a null).

2.3 Operational Phase

At each time step the variable mapping between Doppler bins and the N_C covariance matrices is used to select the set of beam recombination weights, **b**, for each Doppler bin, according to tow speed. These weights are then used to recombine the complex data in the 'primary' and 'Level 2'

reference beams. In this phase the reverberation suppression algorithm requires very little processing over and above the conventional tasks of beamforming and matched filtering.

2.4 Stability of Sidelobe Nulls with Tow Speed Variations

The formation of a null in a beam sidelobe pattern, by combining with a set of reference beams whose mainlobes point in the direction of the required null, using a set of complex weights, depends on some delicate cancellations and, therefore, on the exact amplitude and phase relationships between the mainlobes and sidelobes. At different tow speeds, a fixed set of reference beams will, in general, receive reverberation with different Doppler frequencies. In order to use the same set of adaptive weights (or, equivalently, the same covariance matrix), to form a null in some particular direction, the change in the phase and amplitude relationships between different beams, as Doppler frequency changes, needs to be sufficiently small.

Let us first consider operation in a CW mode. Assuming that the amplitude and phase mismatches between hydrophones do not vary too much, in a time delay or broadband beamformer, as the frequency increases, the amplitudes and phases of the sidelobe features remain the same, but the sidelobe features themselves move closer to the beam maximum response axis (MRA) direction, i.e. the beam pattern becomes more compressed in bearing; however, the beam MRA direction remains fixed. Therefore, with a broadband beamformer, the sidelobes of one beam at the MRA of another will change with Doppler frequency (i.e. beam sidelobes move relative to other beam MRAs). This is illustrated in Figure 3.

In a phase shift or narrowband beamformer, however, the beam sidelobes and MRAs all move together towards broadside as the frequency increases. Therefore, the beam patterns remain invariant relative to each other as Doppler frequency changes. This is illustrated in Figure 3. This result implies that a covariance matrix or set of complex weights used to form a null using a set of reference beams at one tow speed can be used to form an almost identical null, using the same reference beams at any tow speed. The only difference is a slight bearing change in the null position. Thus, if we use a narrowband beamformer, we should be able to use a single set of covariance matrices, or complex beam combination weights, *independent of tow speed.*

Now let us return to the case of the broadband beamformer. If the algorithm is used as a 'bolt-on' to an existing sonar system, it may be necessary to cope with an existing broadband beamformer. In this case we may need to integrate and store several sets of covariance matrices or complex beam combination weights for different tow speeds. The amplitude and phase relationship between two beams changes in a progressive way with frequency, and, therefore, there must be a range of frequencies for which the amplitude and phase are sufficiently similar for an adequate null to be maintained. If we can estimate this range of frequencies, we can also estimate the number of sets of covariance matrices needed for a particular range of tow speeds.

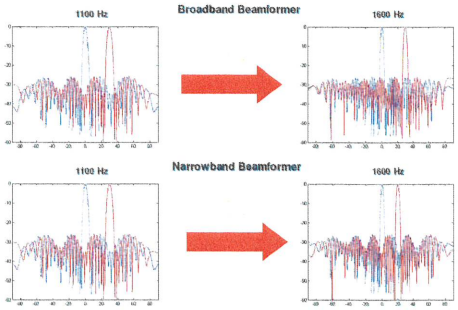

Figure 3: Illustration of Difference in Frequency Dependence Between Broadband and Narrowband Beamformers.

The restriction that needs to be imposed, in order to form an adequate null, is that a sidelobe feature in one beam should not change by more than a certain fraction, α, of a beamwidth (of another beam pointing in the same direction as this sidelobe feature) over the range of frequencies for which the amplitude and phase are sufficiently similar for an adequate null to be maintained. This fraction is likely to be of the order of $\alpha = 1/10$.

A short analysis reveals that the restriction that $\cos\theta$ change by less than $\alpha \times$ beamwidth implies that

$$\delta f \leq \frac{\alpha c}{2L}$$

Equation 5

where L is the aperture length. Note that the ratio L/c is just the time of transit for a sound wave to travel the length of the array. This, then defines the number of covariance matrix sets required for a particular range of tow speeds with a broadband beamformer.

These arguments can all be generalised quite easily to broadband Doppler sensitive pulses, such as pulse train FM (PTFM) or Cox Comb. A narrowband beamformer may be impractical for broadband pulses because of the variation of the beam MRA with frequency. However, the spectrum of these pulses consists of a set of tonals. Thus a "semi-narrowband" beamformer could be implemented by splitting the incoming signal into bands centred on the tonals of the transmitted waveform. In each band a fixed set of phases, rather than delays, would be applied to the hydrophones so that within each band a narrowband beamformer would be implemented. The arguments above can now be applied to argue that we should be able to use a single set of covariance matrices or complex beam combination weights, *independent of tow speed*, provided the hydrophone mismatches are sufficiently constant with variations in Doppler.

If, on the other hand, the reverberation suppression algorithm is being 'bolted-on' to an existing system, it is most likely that a broadband beamformer will already exist. In which case additional sets of covariance matrices, each covering a restricted range of tow speeds, may be required. The estimation of the number of sets required will be carried out by the analysis referred to above.

3 RESULTS

This algorithm has been applied to CW data recorded on a towed array. Data from three sea-trials have been used to test the algorithm. Data from the third trial was recorded at various tow speeds varying between 4.1 and 4.9 m/s, thus providing some test of the stability of the beam patterns with tow speed. In all cases, the performance of the algorithm was tested using a synthetic target echo injected into the actual recorded data at hydrophone level. The data was beamformed using Dolph-Chebyshev shading with –30dB sidelobes. The SNR of the synthetic target was measured before and after reverberation suppression, using a two dimensional double boxcar algorithm to measure the noise background and a measurement of the power at the expected position of the target to assess the signal level. The target signal level was chosen to be difficult or impossible to detect without reverberation suppression.

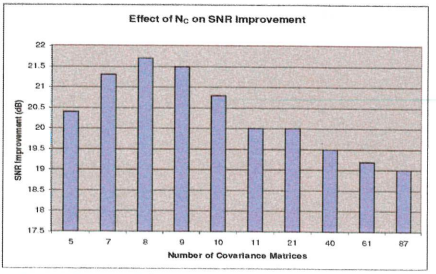

Figure 4: SNR Improvement versus N_C.

Figure 4 shows the variation of the mean improvement in SNR with the number of covariance matrices, N_C, for the data from the first sea trial. It is seen that the best performance, an improvement of about 21.7dB is achieved with 8 different covariance matrices. Theoretical analysis has revealed that the initial rapid improvement, between $N_C = 5$ and $N_C = 8$ is due to the fact that, for $N_C = 5$, the width required for the null is significantly greater than the width of the reverberation ridge. Thus, during the integration of the covariance matrices, the height of the ridge is underestimated as the power is effectively smeared over a wider sector. As the number of

covariance matrices is increased beyond N_C = 8, the required width of the null is well matched to the reverberation ridge (as N_C increases further, the beam pattern nulls from consecutive covariance matrices will overlap) but the performance degrades due to two effects. Firstly, for a fixed amount of calibration data (e.g. two pings) the number of data samples integrated per covariance matrix will decrease as the number of covariance matrices is increased. Thus the estimation error will increase. Secondly, the number of 'Level 2' reference beams will decrease, as the 'popularity score' histogram used to select them will become peakier. This will lead to a less complete sampling, by the 'Level 2' reference beams, of the full extent of the reverberation ridge.

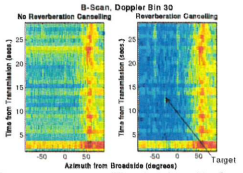

Figure 5: Typical B-Scan before (left) and after (right) reverberation suppression.

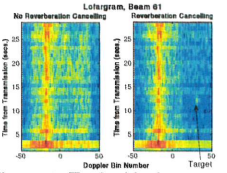

Figure 6: Typical Lofargram before (left) and after (right) reverberation suppression.

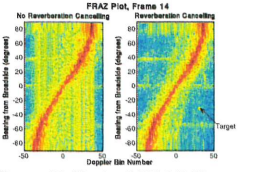

Figure 7: Typical FRAZ Plot before (left) and after (right) reverberation suppression.

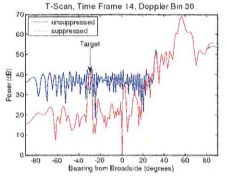

Figure 8: Typical Bearing Plot before (blue) and after (red) reverberation suppression.

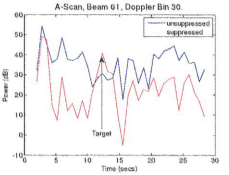

Figure 9: Typical A-Scan before (blue) and after (red) reverberation suppression.

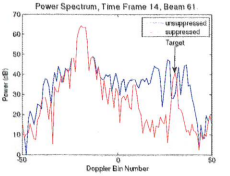

Figure 10: Typical Power Spectrum before (blue) and after (red) reverberation suppression.

Figure 5 to Figure 12 show various cuts through the data recorded from one of the pings of the first of the three sea trials before and after reverberation suppression. The cuts are all appropriately arranged to intersect with the synthetic target in order to provide an indication of the SNR improvement. We only show one out of the three possible 3D cuts (Figure 11 & Figure 12), which display the bearing-azimuth cells in the immediate vicinity of the target, because the other two cuts look very similar. The pseudo-colour plots, Figure 5, Figure 6 & Figure 7 give a qualitative feel for the greatly improved visibility of the target after reverberation suppression, while the line plots, Figure 8, Figure 9 & Figure 10 clearly show an improvement in SNR of 20-30 dB. A similar qualitative impression is given by the 3D plots, Figure 11 & Figure 12, in which, before reverberation suppression a jumble of noise peaks is observed, while after reverberation suppression a single target peak sticks up above an almost flat plane. Finally, Figure 13 & Figure 14 attempt to represent all the data, respectively before and after reverberation suppression, by essentially combining Doppler and range along the y-axis and bearing along the x-axis. The two plots are scaled to the same pseudo-colour axis and, although the target is too small to be seen on these plots, they give a feeling for the scope of these three co-ordinates over which the algorithm is effective. In Figure 7, Figure 13 & Figure 14 it is apparent that, over about 50% of the B-zone, the reverberation has been reduced to about the same level as the ambient noise (which is visible in these figures in the C-zone, at high negative and positive Doppler, before and after reverberation suppression). It is also apparent that there is a region of the B-zone around the A-zone reverberation ridge, extending to zero Doppler at forward and rear endfire directions, where the reverberation suppression has been somewhat ineffective. We have good reason to believe that the reverberation in this region is not predominantly entering the beams through their sidelobes. Indeed this contribution to the reverberation is thought to be mainly caused by double scattering of the transmitted pulse off

features on the sea bottom, and enters the beams mainly through their main lobes. Thus forming sidelobe nulls cannot be expected to suppress this contribution.

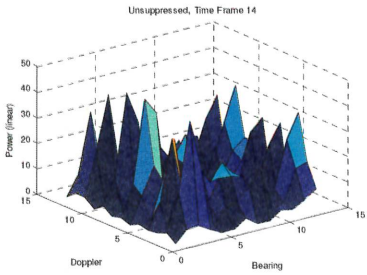

Figure 11: Typical Doppler-Azimuth 3D Plot before reverberation suppression.

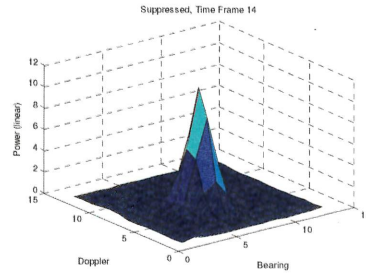

Figure 12: Typical Doppler-Azimuth 3D Plot after reverberation suppression.

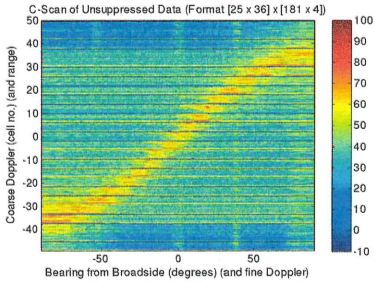

Figure 13: Combined Range-Doppler-Azimuth Plot before reverberation suppression.

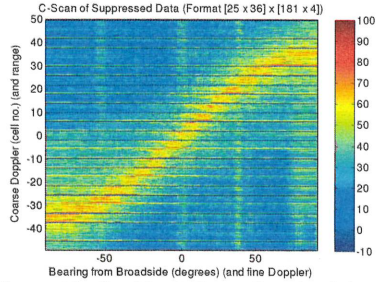

Figure 14: Combined Range-Doppler-Azimuth Plot after reverberation suppression.

Table 2 gives the mean measured improvement in SNR due to reverberation suppression for an approaching target, at a range of 7.5km, 30° aft of broadside, in the middle of the B-zone (but avoiding the double scattering region). Of course these numbers depend on environmental conditions, such as reverberation level and ambient and flow noise levels, as well as characteristics such the actual sidelobe level achieved in the beamformer. In Sea Trial 1, the estimated tow speed in ping number 13 deviates by up to 13% from the tow speed estimated from the data used for calibration. Yet, the measured SNR improvement in ping 13 is 21.7dB, coinciding exactly with the average improvement over all the pings. Thus the tow speed (and hence frequency) variation has had no apparent effect on the performance.

Table 2: Mean Improvement in SNR of a Synthetic Target in Each Ping.

Dataset	Sea Trial 1	Sea Trial 2	Sea Trial 3
Number of Valid Pings	14	43	13
Mean Improvement	21.7dB	12.2dB	11.6dB

In Sea Trial 2, the towing vessel was very noisy and the B-zone noise has a different appearance on the false colour plots. We believe that the noise in the B-zone and C-zone is dominated by tow ship noise. The reverberation suppression algorithm successfully suppresses the tow ship noise in both B-zone and C-zone, as well as the B-zone reverberation. Unfortunately however, because of the high tow ship noise, we do not consider this data set to be a good test of reverberation suppression, but it does demonstrate the versatility of the algorithm.

In Sea Trial 3 the tow speed was estimated at about 4.1m/s over the first six pings, but varied between 4.5 and 4.9m/s over the remaining 7 pings. Initially, the first 2 pings were used for calibration. It was found that the mean SNR improvement over pings 3-6 was 13.2dB, while over pings 7-13 it was only 10.9dB, suggesting that the change in tow speed had degraded the performance. However, when we used pings 7 and 8 for calibration, the mean improvement over pings 7-13 actually decreased to 10.8dB. Thus the lower performance over these pings appears to be unrelated to the tow speed used for calibration. (It may be caused by a higher noise background due to an increase in flow noise.) There were an additional 5 pings in this data set which we rejected as invalid, due to the fact that the data displayed evidence of a very non-straight tow (i.e. an own ship manoeuvre). Even so, and despite the fact that no array shape correction was implemented in our MATLAB code, we decided to measure the performance of the reverberation suppression algorithm on these invalid pings. The algorithm gave a mean improvement of 8.6dB in SNR.

4 CONCLUSIONS

In this paper we have described an algorithm which is able to suppress reverberation from Doppler sensitive pulses by the formation of nulls, which prevent leakage through the beam sidelobes. The algorithm is a form of array calibration and, during the calibration phase, is able to adapt to reasonable amplitude and phase mismatches between hydrophones. After calibration the adaptive freedom is limited to a mapping between Doppler and null direction which is constantly varied according to the tow speed estimated from the data. This estimate could be augmented by own ship's log data if necessary. This makes the algorithm very robust, as no precise matching of hydrophones is required and it compensates automatically for any variations in tow speed. Furthermore, it does not suffer from the fragility of many adaptive algorithms as the 'adaptive' weights are fixed unless the operator decides to recalibrate the array. The only automatic adaptivity is the tow speed estimate.

The new algorithm is highly effective at suppressing B-zone reverberation, being only limited by the single (A-zone) and so called double scattering reverberation entering the beam main lobes. Because the sets of 'adaptive' weights are only calculated once, the processing requirements are low. Also, as only a handful of 'adaptive' weights are required (N_C is small) the memory requirements are modest.

As the algorithm uses beamformed match filtered data, it can be easily 'bolted-on' to existing systems. In the case of time delay beamforming this may require a slight increase in the number of covariance matrices required, as the beam patterns change relative to each other with Doppler, but it is still quite simple to implement.

Further research is required to consider ways of reducing the effect of reverberant double scattering, which is the main limitation away from the main reverberation ridge.

For further work, we are keen to test the algorithm on recorded data from high bandwidth Doppler sensitive pulses, such as pulse train FM and Cox Comb.

5 REFERENCES

1. N.H. Parsons, Adaptive Reverberation Suppression for Towed Arrays, Proc UDT Conference, La Spezia, Italy, (June 2002).
2. D.H. Johnson, The Application of Spectral Estimation Methods to Bearing Estimation Problems, Proc. IEEE, Vol. 70, 1018-1028 (1982).