PROC. INST. ACOUSTICS VOL 6 PT 3 1984

PASSIVE SONAR ARRAYS: SOME ASPECTS OF THEIR DESIGN AND CALIBRATION

M J Earwicker

Admiralty Research Establishment, Portland, Dorset.

By its very nature such a paper as this has to be rather selective in the topics concerned. However it is hoped that even so it does give some indication of the range of factors that need to be considered in the design of passive arrays and their calibration. Particular emphasis has been given to those topics that have been of interest to the author and where possible references are given. It should be remembered that although the final design of any array is always a compromise the success of that compromise relies greatly on a proper understanding of all the factors involved.

The design of passive sonar arrays requires an understanding of not only the acoustics of the array itself but also of the signal and noise environment in which they are required to work and the nature of the signal processing used on the array outputs. In addition the mechanical aspects such as the ability to withstand shock and vibration, hydrostatic pressure and problems such as corrosion must not be forgotten. The array designer must be aware of these wider considerations and think not only as an acoustician but also as a systems engineer.

The function of a passive sonar is in principle very simple. The acoustic field in the sea is generated by many factors such as ambient noise due to the sea itself, locally generated hydrodynamic flow noise due to the motion of the ship carrying the sonar etc and that generated by the targets ie the ships which the sonar is designed to detect. The passive sonar is required to resolve these different components in both frequency and direction, and hence expose the "targets" against the background of "noise". To do this an array of sensors is used as a phased array to give increased resolution by forming a fan of directional beams.

The sensors in the array are spaced half a wavelength apart to remove aliasing effects. The output of each sensor is then weighted by a complex weight, usually in the form of an amplitude shading which is designed to control the beampattern and a time delay to provide beam steering. These outputs are then summed for each steer direction, squared and short time averaged, the outputs being displayed as an intensity modulated time-bearing plot.* Now the output power may be written in terms of the input signals in the following way,

$$R = \sum_{i}^{N} \sum_{j}^{N} R_{ij}(\tau_{ij}) \quad (1)$$

where $R_{ij}(\tau_{ij})$ is the cross correlation function coefficient of the acoustic field with relative time delay τ_{ij} as measured between a pair of sensors i,j.

---

*There are of course many different types of processing. However many of the variations can be considered as special cases of the simple power detector.

The directionality of the array arises because of the spatial correlation present in the signal and noise fields. The signal being taken as that component of the total field originating from a target to be detected. The whole system therefore behaves just like a diffraction grating, that is it produces a spatial Fourier transform of the field at the array. The array in a passive system has two principal functions: to provide an increase in the signal to noise ratio by way of its array gain and to provide bearing resolution.

The resolution of such arrays is fundamentally determined by the size of the array in wavelengths. Shading the array can provide some increase in resolution particularly in the form of adaptive array processing. However, one simple way of increasing the resolution of the array is to use sparse array techniques. The array output, relation (1), consists only of estimates of the correlation coefficient between pairs of sensors within the array, and if the signal and noise fields are spatially stationary then these coefficients depend only on the separation between sensors and not on their exact positions. Thus the output of the array could be synthesised by measuring each of these correlation coefficients once for each unique separation. The output of the array would then consist of these coefficients weighted by the number of times they would have appeared in the double summation (1) ie

$$R_{sparse} \doteq NR_o + \sum_{n=1}^{N-1} 1(N-n)R_n \qquad (2)$$

where the single subscript is the difference in the sensor indices, (i-j).

The advantage of such a sparse array arises from the fact that these coefficients can be measured from an array consisting of many fewer sensors than the fully filled array. Indeed given k sensors an array of approximately $k^2/4 + k$ sensors can be synthesised. Clearly if the limiting factor is the number of sensors available and not the aperture over which the array can extend then this technique can be very valuable.

However the question arises as to what is lost to pay for the increased angular resolution? The answer is simple: signal to noise gain. It should be remembered that the signal to noise gain of an array processor stems from two parts. The directionality of the array that gives it is array gain, and the post detector averaging which confers further gain dependent on the integration time-bandwidth product and the input signal to noise ratio.

Now the array gain gives a measure of the increase in the signal to noise, provided by the array over a single sensor. From (1) this can be seen to be given by,

$$\text{array gain} = \frac{\sum_i^N \sum_j^N R^s_{ij}/R^s_{ii}}{\sum_i^N \sum_j^N R^n_{ij}/R^n_{ii}}$$

where the superscripts s and n refer to the signal and noise components. If a simple case is chosen where the signal is perfectly correlated across the array and the noise perfectly uncorrelated then it can be seen that the signal to noise at the input to the detector has increased by a factor N (the number of sensors in the array)

In fact for a simple sum square array processor the output signal to noise ratio is given by,

$$(S/N)_{out} = BT\ (\bar{S}/\bar{N})^2_{in} \times N(N-1) \quad (3)$$

where the output signal to noise is defined by,

$$(S/N)_{out} \equiv (\bar{R}_{S+N} - \bar{R}_N)^2/\sigma^2, \quad (4)$$

where $\bar{R}$ is the mean output.

The factor N(N-1) is due to the array gain where the noise background has been assumed to be spatially uncorrelated, and $\bar{S}$ and $\bar{N}$ are respectively the signal and noise powers at the sensor. The output signal to noise is a measure of the uncertainty in the estimates of the correlation coefficients in (1). The variance of the output may be determined directly from

$$\sigma^2 = E\left[\hat{R}^2\right] - E^2\left[\hat{R}\right], \quad (5)$$

and substituting directly from (1) into (5). Evaluating the expected values for a band limited white noise spectrum gives the well known result (3).

Now for the sparse array the same analysis may be performed by substituting (2) into (5) and the variance becomes,

$$\sigma^2_{sparse} = N^2\ \mathrm{Cov}\left[\hat{R}_o\ \hat{R}_o\right] + \sum_{n=1}^{N-1}\sum_{n=1}^{N-1} 4(N-n)(N-m)\ \mathrm{Cov}\left[\hat{R}_m\ \hat{R}_n\right]$$

$$+ \sum_{n=1}^{N-1} 4N(N-n)\ \mathrm{Cov}\left[\hat{R}_n\ \hat{R}_o\right],$$

where Cov $(\hat{R}_n, \hat{R}_m)$ is the covariance between the fluctuations in the estimates of the correlation function coefficients, $\hat{R}$. The covariance will depend on whether there are any common sensors used to form the original correlation function coefficients. As there may be many alternative array geometries for any given sized sparse array no unique form for the variance can be defined except at low signal to noise ratios, where it is given as,

$$\sigma^2_{sparse} \approx \frac{\bar{N}}{BT}\left[\frac{2N^3 + N}{3}\right] \qquad \text{for } \bar{S} \ll \bar{N}$$

where N is the number of sensors in the filled array of the same aperture. However upper and lower bound can be given for all signal to noise ratios and for large N and are given in the references. The resulting output signal to noise ratio of a sparse array using N sensors with unrestricted aperture can be written as,

$$(S/N)_{out\ sparse} \approx \frac{3N^2}{8} BT (\bar{S}/\bar{N})^2_{in}$$

and should be compared with (3) for a fully filled array using the same number of sensors. The comparison shows that the sparse array has less gain by a factor 3/8 or about 4 dB, yet has increased the resolution, due to the effective increase in aperture.

Regardless of the detailed form of the array processor the array gain will depend critically on the spatial correlation of the "signal" and "noise", and for any given field will vary with frequency, bandwidth, steer angle and array geometry.

The signal itself is rarely well correlated having been generated by a finite source, been decorrelated during the propagation through the ocean and distorted by any structure covering the array, such as hydrodynamic windows or domes. Indeed such factors, particularly propagation effects, do in practice limit the spatial coherence length of the signal and hence the array gain of the array and as a consequence limit the useful array size to the order of the spatial coherence length.

However the noise fields are in general more of a problem. The fundamental noise for low ship speeds will be that due to the ambient noise in the ocean. However as the speed of the ship carrying the sonar increases other noise sources will become important and may well dominate. For example the noise generated by the propulsion machinery of the ship will give rise to an acoustic field directly due to vibration of the ships hull locally to the machine. However some of this vibrational energy will travel along the ship's structure to the array site and re-radiate from the impedance discontinuities locally to the array. Clearly where possible this energy should be reduced by damping etc before it reaches the array site. However there is a limit to how much can be done in this way and the array designer still has a problem in reducing the residual energy that actually reaches the sensors.

One way of doing this is to mount the sensors on an acoustically opaque baffle. This serves to screen the sensors from the radiation from behind the array. The baffle often needs to be vibration isolated so that it does not itself vibrate and re-radiate. Even when the baffle is perfectly opaque there will still be diffraction around the baffle edges as well as sources of radiation from in front of the baffle. Additionally the baffle will itself have an effect on the sensor response which may not always be beneficial!

The design of such baffles plays an important role in passive array design and will be discussed further below: before doing so however it is useful to consider the flow noise component of the acoustic field. If the array were mounted directly in the flow there would in general be considerable turbulence generated which would cause flow noise. To combat this the array is normally mounted behind a window whose primary purpose is to present a smooth surface to the flow and reduce the turbulence generated. In practice the flow over the window may still be turbulent due to other factors and hence the array will still see a flow noise field.

Now this field contains most of its energy in high wavenumber components well above the range of radiating wavenumbers. These subsonic components do not radiate to the farfield but do generate an exponentially decaying field which the array can detect. The peak of the power spectrum occurs at a wave number corresponding to the mean flow velocity, typically much greater than the limiting radiating wavenumber. Although the array will see these components their effect can be reduced by separating the sensors of the array from the window thus allowing the natural exponential decay to reduce their effect.

For wavenumber components near to the radiating region the stand-off required becomes large, and in practice impracticably large. The reason why these components are seen at all by the array is due to the aliasing caused by having separate sensors. The choice of half wavelength separation of the sensors does stop aliasing within the range of real angles but still gives rise to aliases at wavenumbers $2\pi/d$ apart (d is the separation of the sensors) some of which will occur in regions of the wavenumber spectrum with significant intensity. The use of closer spaced and finite sized sensors can help. For example reducing the separation of the sensors will cause the aliases to move apart in wavenumber, and if the separation is sufficiently small will not cause any within the range of significant wavenumbers. The use of larger sensors will also help as the sensor pattern will reduce the effect of any aliases that do remain.

The window itself also has another effect in addition to the standoff. This arises because the dome being a simple plate acts as a wavenumber filter, and has a passband of wavenumbers determined by the velocity of the waves present. It is often the case that the first antisymmetric or flexural wave causes the window to appear as a low pass wavenumber filter with a cut off determined by this flexural wave velocity. In spite of these beneficial effects the dome can of course also radiate energy when excited as it always is either directly from the flow or from its mountings.

The design of the array itself may seem rather trivial by comparison with trying to reduce the effect of the noise mechanisms discussed above. However it is not quite as simple as it may at first appear. The array basically must provide an undistorted view of the total field of interest whilst rejecting unwanted components. A simple array of omnidirectional sensors in an unbounded medium would be simple enough, although the design of the sensors themselves to cover more than two decades of frequency (typical of the bandwidth required) is not always trivial. Such an array would be unable to distinguish between the field originating from the front and that from behind, unless a three dimensional array is used. However the real problems arise when these sensors are mounted near structure.

The presence of any reflecting surface near to the array will cause the sensor to see not only the direct acoustic field but also that scattered by the local structure. The resulting interference effects can be severely limiting. The combination of the desire to remove these ambiguities and the need to remove the effects of locally generated or scattered fields leads to the use of baffles. As already mentioned the baffles can be used to screen the sensors from some of the unwanted fields. They can also, by having a well defined boundary, allow the array response to be better controlled by only allowing scattering from the baffle.

Two main problems arise in designing such baffles. The first is just to design a baffle with a large transmission loss over a very wide range of frequencies, hydrostatics pressures and physical environments. The second is to design this baffle so that it admits to a well understood response from the sensors. There are basically two choices open to the designer. The first is to provide the transmission loss by having a large impedance mismatch between the baffle and the water and the second to provide a high absorption loss. The latter is very difficult in practice particularly when considering the large range of frequencies over which the array is required to work.

The provision of a large impedance mismatch can be achieved two ways. The simplest of these is to use a material like steel with a relatively high characteristic impedance compared to water, and to place the sensors flush with the baffle surface. The sensors will then see the direct field and that reflected from the baffle surface in phase which will give rise to a 6 dB increase in the output. Such baffle-sensor combinations do not give rise to much distortion of the sensor response either in terms of frequency or direction. However the transmission loss afforded by such baffles is very poor. The total field seen by the baffle rises from the free field value 0 dB at low frequencies to a maximum of 6 dB. The fact that the response is not always 6 dB above the free field is just because the baffle becomes transparent at low frequencies. The transmission loss of the baffle will be zero at low frequencies, for the same reason, and even at higher frequencies only becomes minimal.

One way around this lack of transmission loss could be to increase the baffle thickness. However although this does give at first an improved transmission loss it can make things worse due to increased significance of the higher order waves that the plate can now sustain. As the plate thickness in wavelengths increases additional symmetric and antisymmetric waves appear which cause the plate to become transparent. These higher order waves at first appear with high velocities which with increasing frequency thickness product tend towards the Rayleigh velocity. The effect of this is to give a high transmission loss at angles beyond the Rayleigh angle but for angles near normal to the baffle the transmission loss is not improved. Although techniques can be found for modifying the velocities of the waves present the rather poor transmission loss and large weight of such baffles makes their use limited.

An alternative is to make use of a low impedance baffle with correspondingly low wave velocities. Such materials generally make use of air (contained in some matrix material) to provide the low impedance. For this type of baffle the reflected field is phase shifted by π radians and so the sensor must be mounted a quarter wavelength in front of the baffle so that the reflected and direct fields add coherently. Even assuming that the baffle is an ideal low impedance surface such a combination has its problems. Clearly whenever the frequency is such that the sensor-baffle separation is half a wavelength or multiple thereof the field at the sensor will be zero. As the frequency falls below the quarter wavelength separation the total field at the sensor falls until the baffle becomes transparent when it returns to the free field value. In addition the directional response above the quarter wavelength separation frequency becomes distorted with more than one maxima. As such this type of baffle even in its ideal form does not look very promising, especially as such low impedance materials tend to be physically soft and therefore deform under hydrostatic pressure! However the transmission loss obtainable from such materials can be very high, and it is this factor remembering the importance of reducing unwanted noise in passive sonars that makes them interesting in spite of the above difficulties.

Some control of the behaviour can be obtained by varying the thickness, impedance and internal damping of the material. An additional complication arises however when the finite extent of the baffles is considered. Edge diffraction effects can dominate the transmission loss at low frequencies and internal resonances of say cylindrical or spherical baffles can cause unwanted disturbances in the array response.

Having designed and built the array in whatever form was appropriate to the requirement in hand it is still necessary to calibrate it to prove its performance. As many of these arrays are very large in wavelengths conventional farfield techniques can be very limiting, especially where the array is required to be calibrated over a large range of depths or in confined spaces. As a result nearfield calibration techniques have become increasingly useful.

One such technique relies on generating an acoustic field from points on a surface surrounding, and in the nearfield of, the array and measuring the corresponding array output as the source is moved from point to point in the nearfield. The data thus collected is a mapping of the field on the nearfield surface to the array output. This may then be used as data for a modified Helmholtz integral transformation to predict the response of the array for a farfield source.

The transformation from nearfield data to farfield response requires the calculation of the Green's function for the appropriate surface geometry with the boundary condition that the field is zero on the measurement surface. The use of an exact Green's function removes the need to measure the pressure gradient response of the array. Although the method does require some computation it does allow measurements to be made on large arrays without the use of large measurement distances. The exact shape of the nearfield surface is not of itself critically important and tends to be chosen on the basis of ease of mathematical computation and geometry of the array under test. The use of a cylindrical measurement surface only slightly larger than the array in question often proves convenient.

As with most calibration systems the success of the technique relies on a good understanding of the errors involved. In the case of nearfield systems these can be grouped into three types. Firstly those concerned with the geometry of the measurement surface ie what shape surface, how far away and what extent; then those associated with measurement errors, both amplitude and phase; and finally those due to the sampling of the surface which are related to the familiar Nyquist sampling criteria.

The shape of the measurement surface is of little importance except that it should be one for which it is possible to evaluate the modified Helmholtz Integral and which encloses the array under test. In practice the surface need only cover that region where the array has a finite response. The distance of the surface from the array will depend on how much space is available (for example one could do the measurements in the farfield and then use them to predict the farfield although it would not be very sensible!) More important is the question of how near to the array the surface can be. This depends on the nature of the array under test and the sampling spacing used on the nearfield surface. The response of the array will in general contain wavenumber components which are above the radiating range and which would not play a part in a conventional farfield calibration. However because the nearfield surface is sampled and will cause aliasing these high wavenumber responses will therefore be included as spurious components that will corrupt the data used in the farfield prediction, in exactly the same way as the problem of overlapping side bands in time sampled signals and their spectra or the response of an array to high wavenumber components of flow noise. These responses will decay exponentially with distance from the array and will be reduced if closer samples are used. These two factors can of course be adjusted independently according to the constraints of the system.

In the case of cylindrical measurement surfaces the Nyquist sampling criteria which are sufficient for planar measurement surfaces, have to be modified as the circumferential field does not exhibit a non radiating region. All circumferential wavenumbers can radiate, although as the circumferential wavenumber increases the efficiency of the radiation decreases. The sampling criteria then for the circumferential field is in fact more restrictive than for planar surfaces. That for planar surfaces being

$$\frac{2\pi}{\bar{x}} > k + k_{x\,max} ,$$

where $k_{x\,max}$ is the highest wavenumber present in the nearfield and k the free space wavenumber, and that for cylindrical surfaces being given by,

$$\frac{2}{\bar{\phi}} > \omega_{\phi\,max} + 1.4\ ka \sin\gamma + 16. \qquad \text{for } 10 < ka < 50$$

where a is the radius of the measurement surface, γ the angle of the farfield point for prediction from the axis of the cylinder and $\omega_{\phi\,max}$ the highest circumferential wavenumber component of the nearfield response. It can be seen that these criteria for the cylinder depend not only on the circumferential wavenumbers but also the size of the cylinder in wavelengths. The axial sample spacing can be decided as for a planar surface.

Assuming that one has chosen a satisfactory surface geometry and sample spacing there still remains the problems of errors in the measurement of the amplitude and phase response of the array. The effect of these errors can be analysed by considering an ensemble of nearfield data sets each with some specified variance on the data. The modified Helmholtz Integral can then be cast in a form that relates these variances in the nearfield data to the variances in the predicted farfield, thus defining the probability distribution for each farfield point in the prediction. Integration of these probability density functions over appropriate regions will then give the probability of the prediction being within specified limits of amplitude, phase or both, as a function of the variances of the nearfield data amplitude and phase.

Many such nearfield measurement systems have been designed and used at ARE and are now regarded as a well understood calibration technique. Indeed one such system has for some time been fitted as an integral part of a RN sonar system thereby allowing calibration of the sonar array whilst the ship is at sea.

It is hoped that this paper has given an idea of the range of topics considered in the design of passive sonar arrays, and that it may encourage some of the readers to take an interest in sonar design!

References

1. "Sampling errors in nearfield measurements on on planar surfaces." Journal of Sound and Vibration (1979) 66(2). M J Earwicker.

2. "Signal to noise ratio gain of sparse array processors." Journal of the Acoustical Society of America 68(4) Oct 1980. M J Earwicker.