Proceedings of The Institute of Acoustics

SOME RESULTS USING A SPLIT-WINDOW CFAR PROCESSOR TO NORMALISE SONAR RETURNS

David Morris

AUWE(N) Portland, Dorset

1. Introduction

A sonar signal processor which was designed to normalise the post-detector reverberation background, thereby producing a constant false-alarm rate (CFAR), and also enhancing the signal to interference (S/I) ratio of target echoes is briefly described in this paper. Some typical results using this processor are presented in the form of the probability of false-alarm (PFA) and the S/I improvement obtained when sonar returns from an experimental multimode sonar were processed by the system.

2. The Split-Window Normaliser/Averager

The signal processor to be described is shown in diagrammatic form in figure 1. It was designed around a digital delay line which stored a running time segment of the current sonar return from a particular beam. This delay line was tapped at a number of appropriate points to form the time windows W1, W2, G and S as shown on the diagram. The CFAR output from the processor was fed directly to a refreshed CRT sonar display and to an automatic detection and classification system for further processing.

In its operation the processor achieved both signal enhancement and CFAR by the proper utilisation of the data stored in the delay line and hence the windows W1, 2, G and S. Data stored in the equal time windows W1, 2 was used to form estimates of the local reverberation conditions prevailing around the echo. For example, a mean reverberation level $\bar{r}$ and power level σ_R^2 were estimated. The central window S was used to average the signal prior to normalisation, and the signal gap window G was adjusted to admit only signals of time duration less than or equal to the length of the window G.

The output from the processor can be expressed in terms of the signal and estimates of the background as follows. If $\bar{s}(t)$ is the output of the averager window S, and $y(t)$ the output of the processor then,

$$y(t) = \left(\frac{\bar{s}(t) - \bar{r}}{\sigma_R}\right) \quad . \tag{1}$$

The output $y(t)$ was therefore a running measurement of a detection index applied to the time segment stored in the delay line.

An alternative and simpler form to equation (1) can be obtained if the assumption is made that the quadrature outputs, $x_1(t)$ and $x_2(t)$, from the correlator are normally distributed with zero mean and variance σ_x^2. The envelope detector output is then Rayleigh distributed [1] and,

$$\bar{r} = \sigma_x\sqrt{\frac{\pi}{2}}, \text{ and } \sigma_R^2 = 2\sigma_x^2(1 - \pi/4). \tag{2}$$

Substituting (2) in (1), $y(t)$ becomes

$$y(t) \propto \left(\frac{\bar{s}(t)}{\sigma_R} - 1.9\right) \tag{3}$$

If alternatively the output is formed as

$$y(t) = \frac{\bar{s}(t)}{\bar{r}} = \frac{\bar{s}(t)}{1.9\sigma_R}\text{, then} \tag{4}$$

aside from an offset and scale change, (1) and (4) both depend upon the ratio $\frac{\bar{s}(t)}{\sigma_R}$. However, the implementation of (4) is simpler and gives a significant speed advantage over (1) for digital implementation with a computer, for example. Equation (4) will be recognised as an unknown-level CFAR processor [2], in which it is assumed that the reverberation in the i^{th} cell is the product of a constant, $A^{\frac{1}{2}}$, and a random variable z. That is, all cells have equal interference power and,

$$r_i = A^{\frac{1}{2}} z_i,\ i = 0, \pm 1, \pm 2, \ldots \pm M\ . \tag{5}$$

The voltage $\bar{r}$ is given by,

$$\bar{r} = \frac{1}{2M} \sum_{i=\pm p}^{\pm(M+p)} r_i\text{ , where M and p refer to} \tag{6}$$

the size and start of the windows W1, 2 respectively.

The set $\{z_i\}$ are members of a unit variance, identically distributed random variable, which in this case was assumed Rayleigh distributed. The magnitude and statistical parameters of the scale factor A are assumed unknown. If $\bar{s}(t)$ is a linear combination of the r_i, substitution of (5) and (6) into (4) shows that y(t) depends only on the statistics of the set $\{z_i\}$.

By definition, y(t) will fluctuate around a mean level, $\overline{y(t)} = 1.0$, when no signal echo is present in the delay line. The fluctuations are due to random variations in both the output of the averager S and the normalising voltage $\bar{r}$. Assuming initially that only one sample is stored in S, the fluctuations in S will be the same as at the input to the processor. Because $\bar{r}$ is formed from a large number of independent samples (> 100 here), to a first approximation it will be normally distributed with a mean $\bar{r}$ and variance $\sigma_R^2(2M)^{-1}$. The number of samples used to form $\bar{r}$, ie 2M, is at least an order of magnitude greater than the number of samples usually stored in S, so the voltage $\bar{r}$ is sufficiently slowly varying to allow the assumption that the fluctuations in y(t) are Rayleigh distributed, as per the set $\{z_i\}$.

If more than one sample is averaged in S, and provided 2M>>N, where N = number of samples averaged in S, the variance of y(t) will be reduced and

$$\text{var}\ [y(t)] \simeq \sigma_y^2/N \tag{7}$$

The response of the system to rapid changes in the reverberation level depends on the value of 2M. The upper cut-off frequency of the normalising voltage $\bar{r}$ being approximately,

$f_c \simeq \frac{1}{2M\Delta t}$ in Hertz, where Δt is the delay line tap spacing in seconds. When deciding upon the size of W1, 2 therefore, the following two factors should be considered,

a. The rate of change (or frequency response) of the overall reverberation envelope during the ping cycle, and

b. the average duration of discrete reverberation returns of a target-like nature.

The windows W1, 2 must be small enough to follow the reverberation envelope, but sufficiently large such that $\bar{r}$ is not dominated by discrete reverberation returns which might cause suppression of small signals in the immediate vicinity of such reverberation. Fortunately this latter effect is much less severe than in lagging-window agc's because of the better estimation of the reverberation slope afforded by the split-window. In this system $\bar{r}$ was formed from approximately 600 data samples, and Nitzberg [2] shows that if 2M = 64 samples, the probability of detection for a Swerling case I target is less than 0.5 dB different from M = ∞ which is the theoretically optimum case.

When a target echo is present in the averager window S, the processor output is,

$$y(t) = \left(\frac{\bar{s}(t)}{1.9\sigma_R}\right), \text{ where } \bar{s}(t) = \frac{1}{N}\sum_{i=1}^{N} s_i(t)$$

The output S/I ratio is therefore

$$S/I_{out} = \left(\frac{y(t) - \bar{r}}{\sigma_y/\sqrt{N}}\right) = \sqrt{N}\ S/I_{in}\ . \qquad (8)$$

S/I improvement is therefore achieved exactly in the manner of a simple boxcar averager, and is a function of the number of samples integrated.

3. Application of this processor to sonar returns

A multibeam version of this processor was used to normalise and enhance sonar returns obtained with a system capable of operating in a bottom-bounce mode. It was necessary to normalise the very large dynamic range of the reverberation envelope, and reduce the high data rate from the correlators before presentation on a refreshed CRT display.

A feature of echoes from submerged targets received via the bottom reflection path is the triplet nature of the echo which is due to the four distinct multi-paths from the source to the target and back. A typical triplet echo is shown in figure 2 and illustrates certain important characteristics.

a. The total return can have a down-range extent of between 50 and 80 basic system resolution cells.
b. Each individual echo is time smeared over 10 to 15 resolution cells.
c. The structure of each individual echo shows large peak-to-trough ratios, ie a large variance.
d. The spacing between each echo is a function of the depth of the echoing object.

These characteristics influenced the way the processor was set up. In particular, (a) influenced the size of the central gap window G, which was adjusted to admit echoes of 60-80 cells; this meant that G was about 100 resolution cells in practice. The size of each individual echo, and the desirability of maintaining the triplet structure tended constrained S to about 10-15 cells, although the results show the effect of longer averaging times as well. The fluctuations in power level within each echo changed the effectiveness of the integration and this is discussed further below. The overall reverberation envelope contained frequencies not greater than approximately 2 Hertz and the windows W1, 2 were adjusted accordingly.

If the target echo is modelled as three rectangular pulses, each of duration T, and separation t, the enhancement in S/I ratio due to averaging over a range of from one to eighty resolution cells is easily calculated. A calculated plot for a triplet echo with parameters T = 10 cells, and t = 35 cells is shown in figure 3 as curve (1). A 10 dB enhancement occurs with only a ten cell averager. As the averager window is increased no significant increase in this figure is obtained, but the data rate and of course the false alarm rate will be reduced.

In practice, these calculated gains were a long way from reality. A more realistic representation of the individual echoes in the triplet assumes that only a fraction, $\sqrt{F}$, of the T cells actually contain target energy, and these are randomly distributed throughout the down-range extent. In this particular context F can be thought of as a power fluctuation parameter, and is sometimes expressed in terms of an echo splitting loss. The inclusion of a fluctuation loss F moves the curves of gain vs averager size down the y-axis a corresponding amount, so that, using this modified model of the echo triplet, the improvement in S/I ratio is given by,

$$I(N) = 20 \log\left(\sqrt{\frac{F}{N}} \cdot R(N)\right) \quad \text{in dB, } N = 1, 2, 3 \ldots \quad N \neq 0 \tag{9}$$

where I(N) is the improvement in dB, N is the size of the averager window in resolution cells, and the function R(N) is the area under the triplet echo,

$$S(N) = \text{rect}\left(\frac{N - T/2}{T}\right) + \text{rect}\left(\frac{N - t}{T}\right) + \text{rect}\left(\frac{N - 2t}{T}\right). \tag{10}$$

The results obtained from an ensemble average of the measured gain achieved with a 120 degree aspect target can be compared with the calculated gain assuming F = -8 dB (ie an 8 dB fluctuation within the echo), $T \simeq 15$ cells and $t \simeq 37$ cells. Agreement to within 0.5 dB was obtained except at $N >$ about 60 cells, and the discrepancy here was attributed to the fact that long averaging times tended to also enhance the discrete reverberation components and therefore reduce the measured S/I ratio. If the fluctuation parameter is increased to -9 dB (ie F = 0.12) and T reduced to about 12 cells reasonable agreement with typical ensemble average beam aspect results was obtained as shown on the figure, except at $N >$ about 60 cells for the same reason.

4. <u>Measured probability of false alarm</u>

In order to determine suitable data threshold values, measurements of the PFA were made at the output of the processor when operating with bottom reflected sonar returns. Some representative results are shown in figure 4. The solid curves show the measured results for three S/I ratios of 5, 10 and 15 dB, and at each threshold two curves showing the effect of wind speed have been plotted.

At the 5 dB threshold and with less than about 12 knots of wind speed, the PFA was about 10 per cent independent of the size of the averager window. When the returns were surface reverberation dominated at more than 18 knots, a small upward trend is apparent in the PFA as the averager window is increased. At a 10 dB threshold a marked difference occurs in the PFA, according to whether the reverberation was surface dominated or not. An almost constant PFA of about 1.5 per cent was measured at less than 12 knots of wind, which increased to 2 per cent to 3 per cent with more than 18 knots, depending on the amount of averaging.

At a 15 dB threshold the upward trend in PFA as the averager window increased was from 0.12 to 0.5 per cent at less than 12 knots and from 0.18 to 0.7 per cent at more than 18 knots. The implication in this instance was that if a low PFA was required, say less than 1 per cent, integrating the returns to reduce the false-alarm rate would actually increase the probability of a false alarm in any resolution cell. This effect was present in the lower thresholds but to a very much smaller degree.

In order to put these curves into perspective, the dotted curves show calculated PFA's for the same thresholds and for the nearest threshold which gave approximately the same PFA as that measured. The raw data points were calculated assuming a Rayleigh distributed random variable, and the values after integrating ten cells or more by assuming a normally distributed random variable. It can be seen that for the same threshold a much higher PFA was actually measured, and a threshold of 2 dB was needed to bring the calculated PFA in line with that measured at 5 dB for instance.

5. An example of sonar returns before and after processing

An example of a typical sequence of twenty successive pings from the same beam, both before and after processing, is shown in the remaining figures 5 and 6. Figure 5 is a sequence of stacked A-scans, offset for clarity, and showing the reverberation environment around the echo track. The time window encompasses about a 20 dB dynamic range of amplitudes, and for intelligibility purposes a ten cell averager has been applied to the data prior to plotting. Figure 6 shows the same twenty pings after passing the raw data through the processor described above, and with the S window set to average ten cells. This figure, plotted with a 5 dB threshold, clearly shows the target track together with two spurious reverberation tracks and random false-alarms.

6. Conclusions

The system described briefly above gave very good results when used in conjunction with a refreshed CRT display with 15 grey levels. Using a number of minicomputers to implement the processor, real-time multibeam operation was possible under a wide range of reverberation environments. The self-adaptive nature of the system meant that no a priori knowledge regarding the reverberation envelope was necessary (except its approximate frequency cut-off) and the non-critical nature of parameters such as the size of W1, 2 meant that a sensible tolerance was available in setting these up. Lagging-window agc systems cannot properly estimate local reverberation conditions under unknown slope or impulsive reverberation conditions, and this can result in the suppression of small targets to below the display threshold. The split-window system does a better interpolation of unknown slope reverberation and is less sensitive to impulsive reverberation dominating the normalising voltage and suppressing weak echoes.

The system described here was neither designed to, nor able to deal with target-like reverberation, by which is meant reverberation echoes which have the same down-range extent as the true target triplet. Averaging enhances these false echoes and a different technique must be used to remove them. The task of dealing with false tracks is best left to a higher order process, involving spatial as well as temporal data, such as might be incorporated into a computer aided detection system (CAD). In this respect, the figures here are threshold crossing rates and the final false-alarm rates will be much lower if CAD follows this processor.

More complex algorithms were possible with this processor because all the necessary data was stored in the delay line. Experiments using an energy normalising algorithm [3] showed that the important performance indicators of false alarm and detection rates did not change significantly. This augurs well that equation (5) was an adequate description of the reverberation environment, and a computationally straightforward algorithm therefore gave optimum results.

The adage of resolve and recombine (incoherently) was found to be suspect from the point of view of S/I gain at least. The author adduces the results of figure 3 in this respect. In this context other results [3] have indicated that the fluctuations in echo level from ping to ping are not necessarily reduced by the resolve and recombine rule.

Higher PFA's were measured compared with the assumption that a Rayleigh distributed random variable was a good description of the set $[z_i]$. The higher PFA's were due to an excess of high strength target-like false-alarms which could not be normalised out, and which were considered to be due to surface wind generated reverberation and discrete bottom features. The effect of wind speed was most apparent at the higher thresholds, ie lower PFA, and in general, wind generated false-alarms tended to occur in a band of S/I ratios from about 10 to 15 dB; this may be a usable statistic from the target classification viewpoint.

In conclusion, the system is a good candidate for implementation using the ubiquitous microprocessor based hardware now available. In the author's opinion, the improved performance compared with lagging-window agc's with respect to CFAR operation, the low thresholds possible, and the ability to cope with varying reverberation envelopes, makes this processor an almost mandatory part of any sonar system.

<u>References</u>

1. Radar signal analysis, W S Burdic, Prentice-Hall 1968
2. Constant-false-alarm-rate signal processors for several types of interference, Ramon Nitzberg, IEEE transactions on aerospace and electronic systems, Vol AES-8, No 1, Jan 1972
3. Internal report in preparation, D Morris, AUWE, April 1980

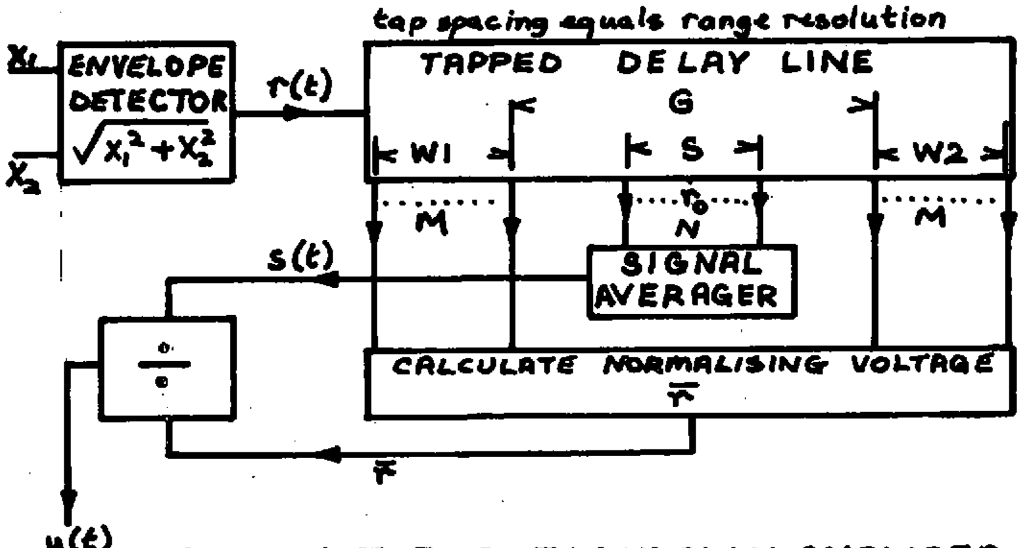

FIGURE 1. THE SPLIT-WINDOW NORMALISER

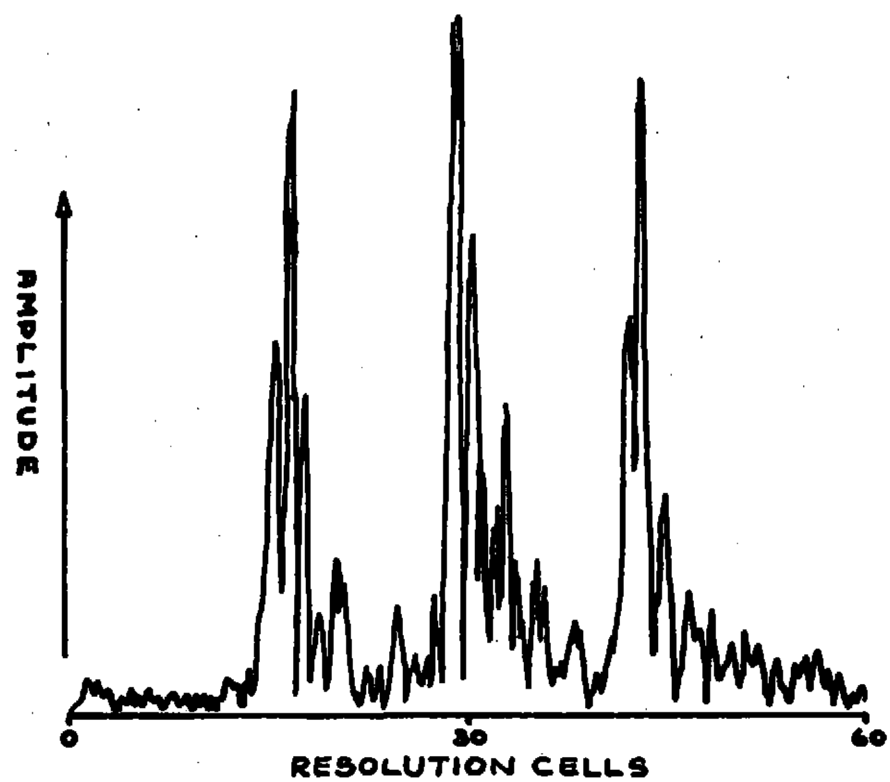

FIGURE 2. BEAM ASPECT TRIPLET ECHO

SOME RESULTS USING A SPLIT-WINDOW CFAR
PROCESSOR TO NORMALISE SONAR RETURNS

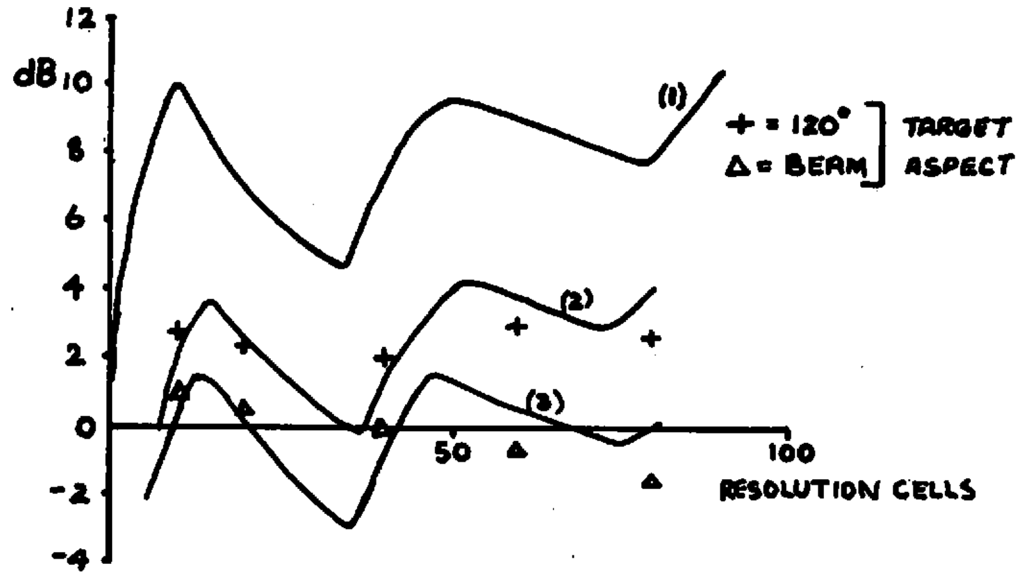

FIGURE 3. SIGNAL TO INTERFERENCE GAIN

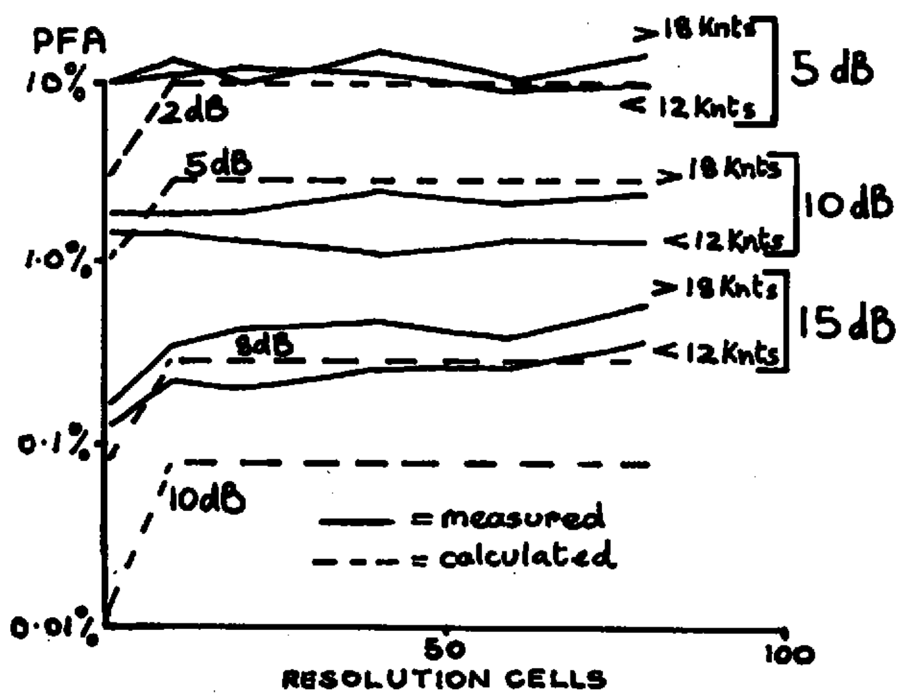

FIGURE 4. PROBABILITY OF FALSE ALARM

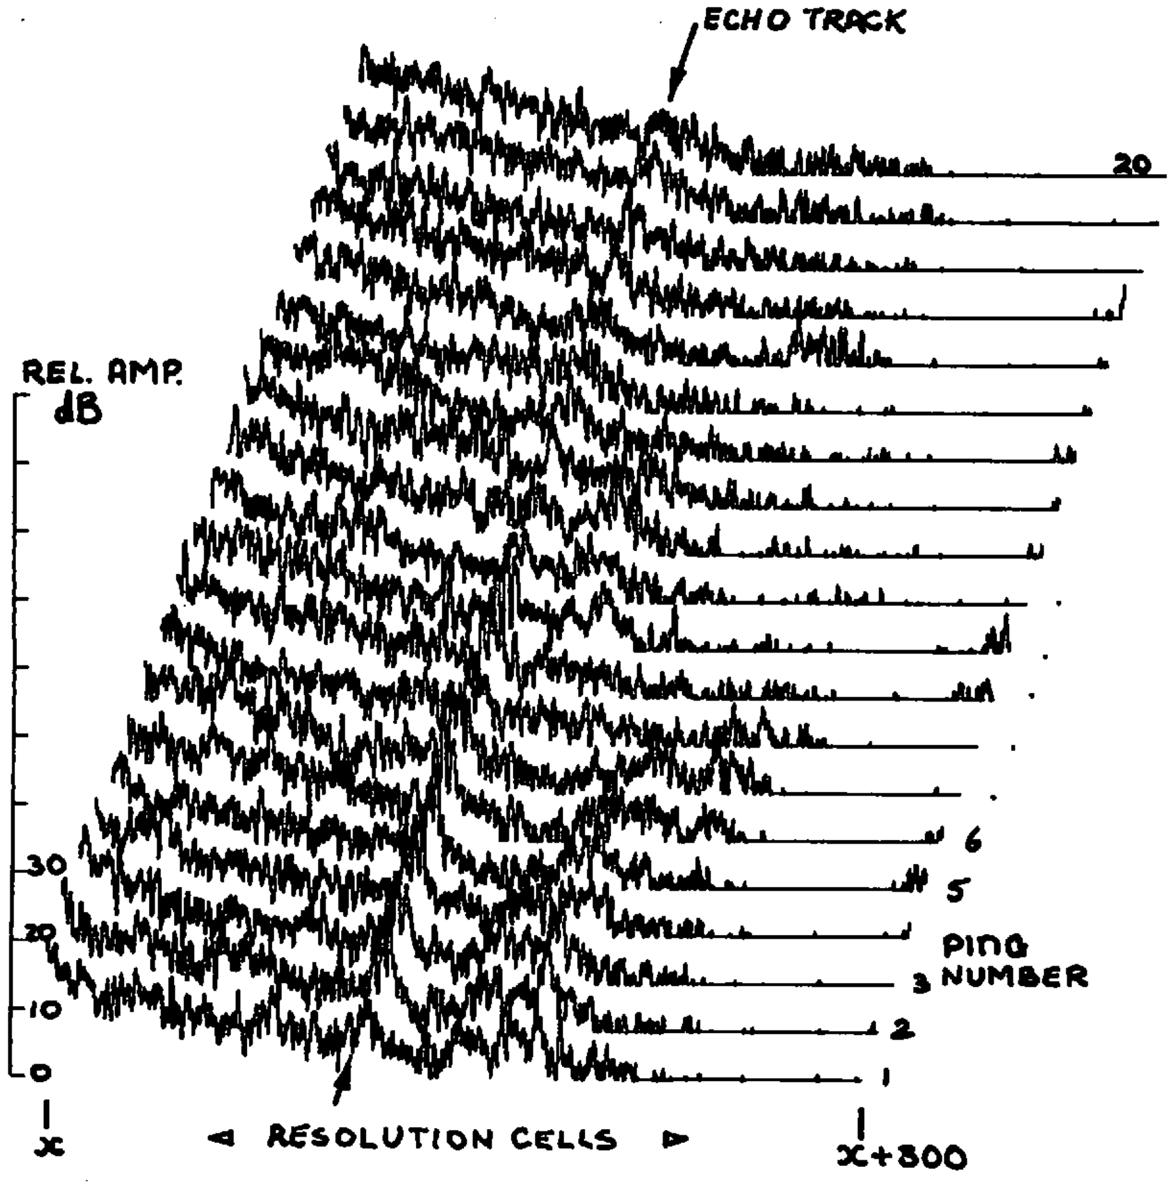

FIGURE 5. 20 PING SEQUENCE BEFORE PROCESSING BY NORMALISER

SOME RESULTS USING A SPLIT-WINDOW CFAR
PROCESSOR TO NORMALISE SONAR RETURNS

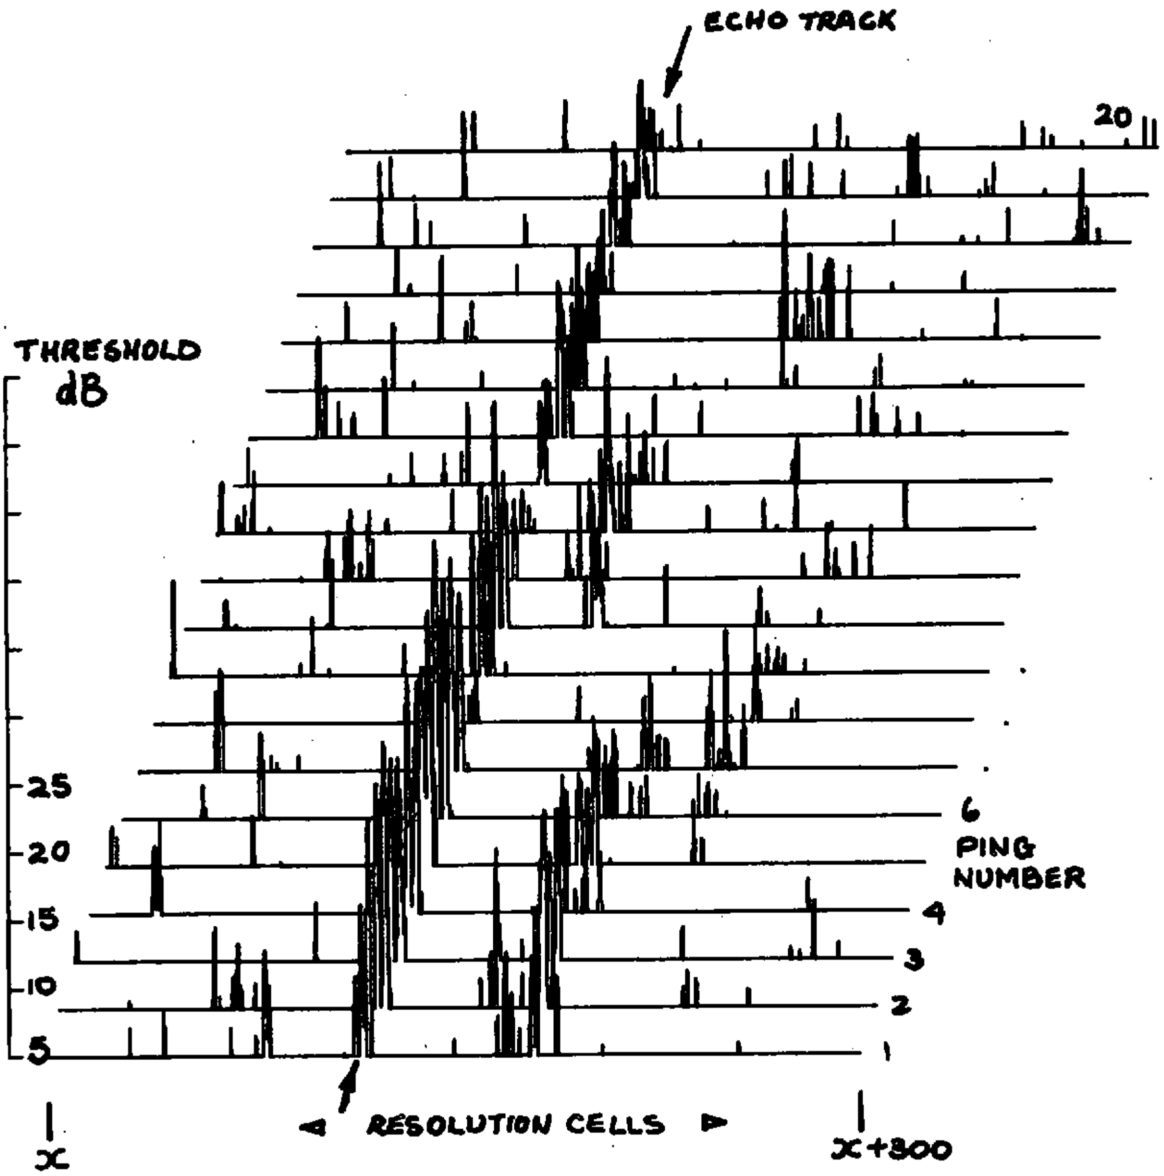

FIGURE 6. 20 PING SEQUENCE AFTER PROCESSING BY NORMALISER