THE APPLICATION OF ADAPTIVE NULLING, SPECTRAL CANCELLATION AND CROSS POLARIZATION TO TERRESTRIAL NETWORKS AT VLF, LF AND HF

Peter Sielman

Eaton Corporation AIL Division
Commack Road
Deer Park, New York 11729

July 1984

84-26

EATON Advanced Electronics

Eaton Corporation
AIL Division
Deer Park, N.Y. 11729

THE APPLICATION OF ADAPTIVE SPATIAL NULLING, SPECTRAL CANCELLATION AND CROSS POLARIZATION TO TERRESTRIAL NETWORKS AT VLF, LF AND HF

Peter Sielman

Eaton Corporation AIL Division
Commack Road
Deer Park, New York 11729

ABSTRACT

The paper addresses problems, ideal solutions as a function of frequency, limits to achievable performance, flight and field test results and system architectural considerations associated with nulling, cancelling and cross polarization at VLF, LF and HF.

PROBLEM

Over-the-horizon communications frequencies VLF (3-30 kHz), LF (30-300 kHz) and HF (3-30 MHz) are widely used by the military because of the range they provide and, in the case of VLF, the ability to penetrate sea water to modest depths.

The strength of these frequencies is also a weakness in that interference (both intentional and nonintentional) can be effective at great distances. At VLF particularly, lightning, which can generate very large impulses with energy peaks near 10 kHz, is often the limitation to communications throughput.

At LF and HF there are multipath problems of interference between groundwaves (direct) and skywaves (multipath). If the delay differential is right this can lead to both RF cancellation and intersymbol interference.

Besides external interference the over-the-horizon frequencies (particularly, VLF) are susceptible to onboard generated EMI. Both machinery and 400-Hz power supplies are potential sources of interference found particularly on airplanes and submarines. These problems are pictorially summarized in Figure 1.

IDEAL SOLUTIONS

Figure 2 depicts three types of system solutions that can be applied to the problems identified above for over-the-horizon communications.

- Nulling
- Cancelling
- Cross Polarization

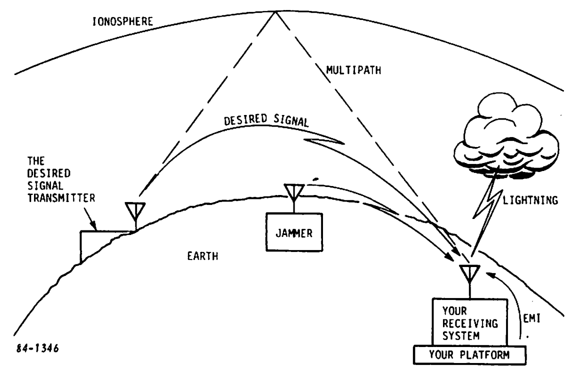

FIGURE 1. OVER-THE-HORIZON COMMUNICATIONS BANDS PATTERN

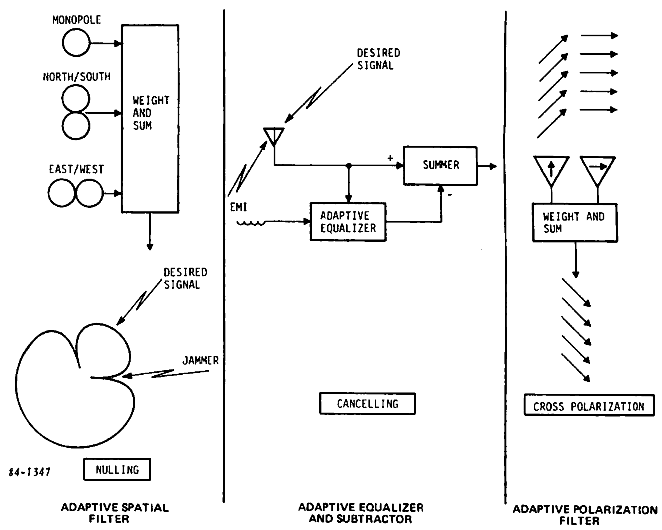

FIGURE 2. SOLUTIONS - NULLING, CANCELLING AND CROSS POLARIZATION

Nulling, or spatial filtering, is the generation of an antenna pattern which varies with time in such a manner as to provide (at all times) nulls on interfering signals while simultaneously maintaining maximum gain on the desired signal. A 3-element orthogonal array composed of H-field crossed loops and an E-field monopole all responsive to TM waves is shown in the figure. By suitable weighting and summing of the antenna inputs it is possible to form a cardioid of arbitrary orientation and null separation such that one or two jammers are placed in nulls and the desired signal is provided with gain.

Cancelling provides direct subtraction of locally generated EMI that has found its way into the receiving antenna. Since EMI may be broadband and may enter the antenna through a variety of paths, it is necessary to provide an adaptive equalizer prior to effective subtraction.

Cross polarization takes advantage of polarization differences between desired and interfering signals by cross polarizing the weighted sum of orthogonally polarized antennas so as to cancel the interfering signal.

Clearly combinations of these techniques can be applied to provide optimized performance. The remainder of this paper is devoted to defining some limits to performance, providing the field and flight test results and setting forth some system architectural considerations in applying these techniques to the over-the-horizon communications bands.

LIMITS TO ACHIEVABLE PERFORMANCE

Cancelling effectiveness is dependent upon the degree to which the adaptive equalizer matches the two channels prior to subtraction.

Cross polarization effectiveness is dependent upon the degree to which the desired signal and the interference differ in polarization and the ability of the system to provide resolution in the polarization of the output.

In contrast, nulling performance is limited by a number of factors, which are distinct from the limits of other forms of AJ protection and interference immunity.

Figure 3 provides the relationship between nulling and angular separation. At low frequencies, the antenna aperture is small and the resultant antenna beams are flat, so that often nulling an interfering signal also impacts the gain towards the desired signal. The left-most part of

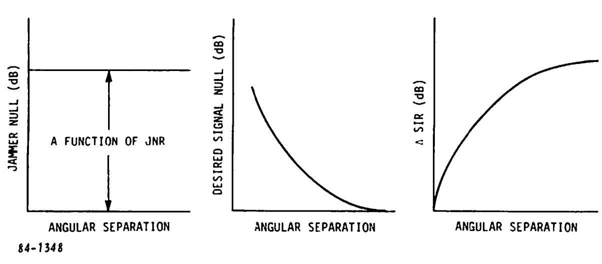

FIGURE 3. PERFORMANCE vs ANGULAR SEPARATION

Figure 3 indicates that null depth on the jammer (or interferer) is independent of angular separation between the desired signal and the jammer. It is a function of jammer-to-noise ratio - or how well can the jammer be measured. The middle part of Figure 3 shows that the loss of gain (from peak) on the desired signal is very much a function of the angular separation from the jammer - with the angle at which there is no loss of gain being a function of the antenna aperture. The right-hand portion of Figure 3 combines these concepts giving the net improvement (delta signal-to-interference ratio) versus angular separation.

Given the ability to generate nulls does not guarantee that the appropriate null will be created at the appropriate time. The logical process is called the adaptive array algorithm. It computes the appropriate combining antenna weights based on some form of measurement - often a correlation. The literature abounds with descriptions of various algorithms - none perfect - each effective in different situations. In order for the algorithm to be effective, it is necessary that there be a means of discriminating between the desired signal and the interfering signal, so that the resultant Δ SIR can be maximized. Figure 4 lists several algorithms and discriminants that have been used singly or in combination by AIL.

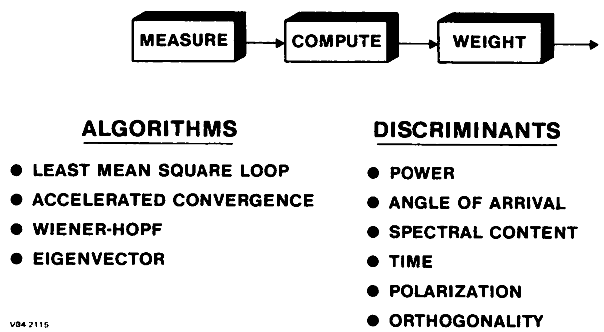

FIGURE 4. ALGORITHMS AND DISCRIMINANTS

If correlation between array output and element input or among input elements (the covariance matrix) is used as the measurement input to the algorithm - then theoretically, the achievable null depth is twice the input jammer-to-noise ratio. (See Figure 5). The simple heuristic argument that underlies this concept is an input x dB above noise when correlated with a null x dB below the noise results in a signal with 0 dB SNR and hence, not capable of providing inputs for further nulling. Limitations to this theoretical null depth include: (1) the fact that all signals other than the jammer to be nulled are part of the noise; (2) the discriminant that prevents nulling of the desired signal may limit the null depth; (3) the finite bandwidth or correlation time of the measurement may introduce an error; and (4) the actual implementation of both the correlation measurement and the weighting process may lack resolution or dynamic range sufficient to realize the full theoretical null depth.

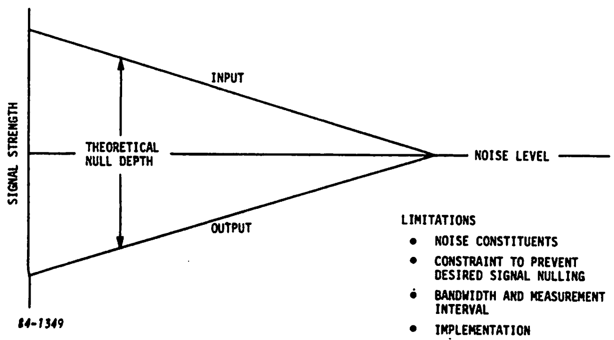

FIGURE 5. NULL DEPTH vs SIGNAL STRENGTH

Figure 6 shows a generic array of elements, each weighted and then summed to provide an adapted array output. The weighting can be real, complex or tapped delay line, depending upon the aperture and bandwidth of the problem. A general rule of thumb is that n-1 independent nulls can be formed by an n-element array. That is, the more signals that are to be simultaneously nulled, the more elements (degrees of freedom) are required - as well as more electronics and more complex processing. An adaptive array to be effective need not provide high gain towards the desired signal and very low side lobes in all other directions. Rather, it can be very effective if there is adequate gain towards the desired signal, sufficient null depth towards the jammer(s) and what happens in other directions is immaterial.

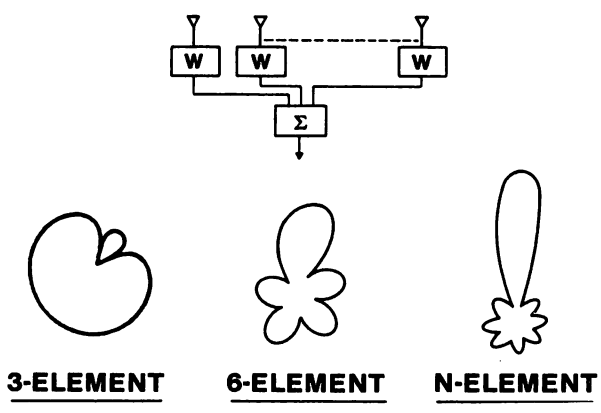

FIGURE 6. ELEMENTS, PATTERNS AND DEGREES OF FREEDOM

FLIGHT AND FIELD TEST RESULTS

The root stock system of our nulling, cancelling and cross polarization systems is the Adaptive Antenna Receiving System (ADARS) a VLF system developed for RADC. It is depicted in Figure 7. The top half of the figure shows the two (H-field) crossed loop antennas - which we label North/South and East/West - and the (E-field) monopole antenna. Through real weighting of these antenna inputs in ADARS a cardioid pattern is generated that provides gain towards the desired signal (160°) and nulls the jammer (30°). The lower half of the figure shows the spectrum analyzer pictures of the three inputs and the adapted output. For purposes of exposition on a spectrum analyzer the desired signal (left) and the jammer (right) have been separated by 1 kHz in frequency; the system bandpass, however, encompasses both signals. The actual signal processing unit hardware is shown in the middle. It is packaged into a standard Air Force VLF receiver chassis with electronics on the left and power supplies on the right. Operation is performed through a 16-key pad and an alphanumeric display.

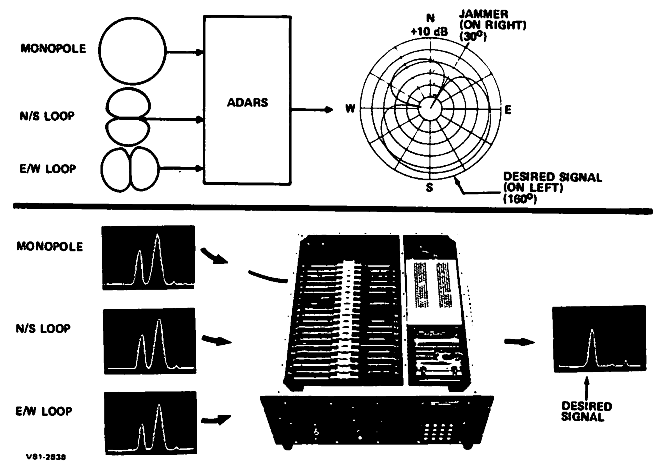

FIGURE 7. ADAPTIVE ANTENNA RECEIVING SYSTEM

Figure 8 shows a picture of the ADARS Signal Acquisition Unit with the radome removed. The antenna is 37 inches on a side and 12 inches high. Aperture was selected to ensure that ADARS was atmospheric (rather than thermal) noise limited even under winter arctic night (very low atmospheric noise) conditions.

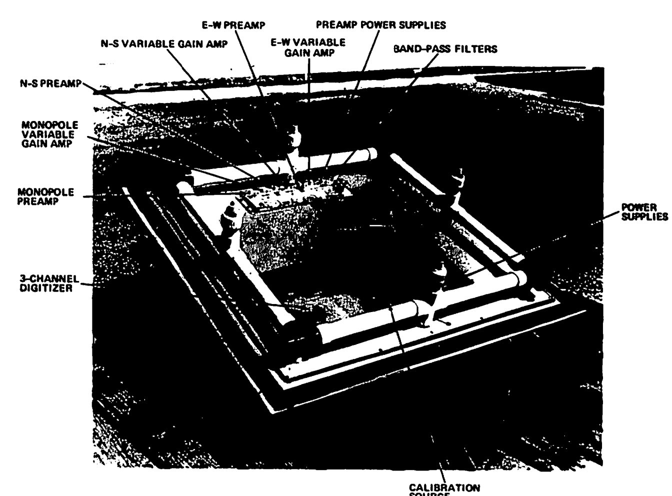

FIGURE 8. ADARS SIGNAL ACQUISITION UNIT

RF amplification and digitization are performed at the antenna. The digital streams are multiplexed and sent via a fiber optic cable to the signal processing unit. Calibration using a local source is used to periodically equalize the channels to ensure broadband nulls.

Testing at very low frequencies presents some unique challenges. Because of the very long wavelengths, generating a controlled far-field signal requires more than a signal generator attached to a whip. As a consequence, most VLF testing is performed using existing fixed VLF transmitters, which owing to the long range of VLF communications, can be used anywhere in the country. However, for controlled tests where power, modulation, bandwidth and angle of arrival need to be systematically varied, it was necessary to develop far-field simulators. Two of these are shown in Figure 9 in the melting snows of the Air Force's Verona test site. The ADARS antenna (with radome) is seen immediately to the right of the right-most U-shaped far-field simulator.

FIGURE 9. TESTING WITH FAR-FIELD SIMULATORS

Many tests have been run in San Diego, Long Island, Hawaii and Point Mugu. These tests have been conducted with the ADARS on the ground and when flown on a TACAMO aircraft.

One test is shown in Figure 10. The ADARS was situated on Eaton Corporation's Melville facility. A TACAMO transmitter was off the coast of the Carolinas transmitting at 23.0 kHz. The Hawaii transmitter was operating at 23.4 kHz and the Annapolis transmitter at 21.4 kHz. For the purposes of this test TACAMO was considered the interferer and Hawaii the desired signal.

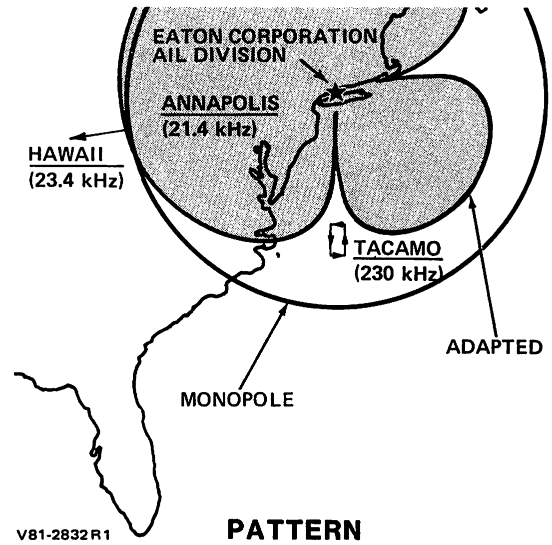

FIGURE 10. ADARS TEST SENARIO

Figure 11 shows the results of that test. The upper trace represents the monopole response (circular pattern around Eaton Corporation's AIL Division in Figure 10) and the lower trace represents the adapted response corresponding to the cardioid patterns of Figure 10. It is seen that the adapted response has nulled the TACAMO interferer by at least 25 dB, while maintaining constant gain on the desired Hawaii signal.

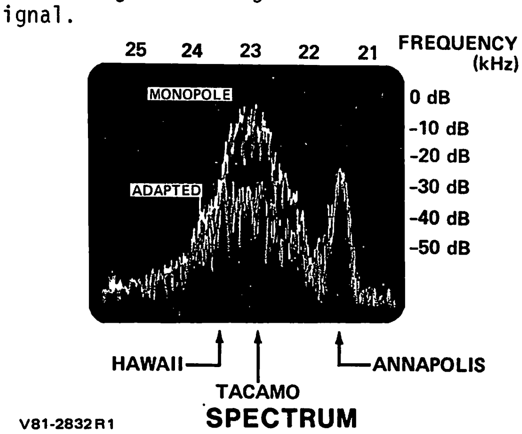

FIGURE 11. SPECTRUM ANALYZER OUTPUTS

A derivative of ADARS is the Automatic Nulling Network (ANN) which uses the same high speed digital processing modules developed for ADARS. Two successive cancellations are performed by subtracting an adaptively equalized locally generated signal (in this case the aircraft 400 Hz main and essential power busses) from the antenna channel (HF-3) which contains both the desired signal and the undesired EMI (400 Hz harmonics). The process is depicted in Figure 12. The processor takes the appropriately filtered inputs, generates the correlations necessary to derive the weights of the adaptive equalizers which are implemented as digital finite impulse response filters. This process is the same as that used in the calibration of ADARS - except that it is accomplished on line, in real time.

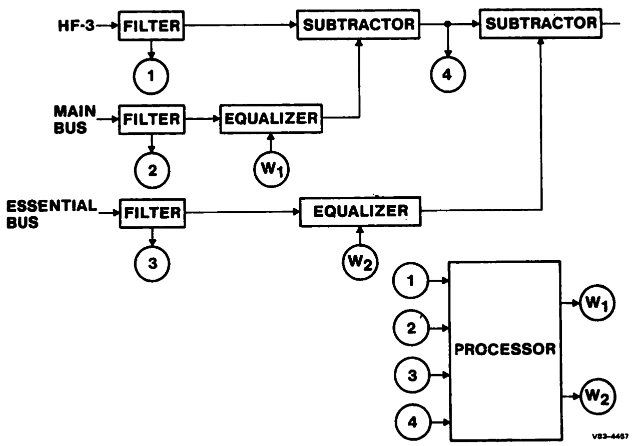

FIGURE 12. TWO-CHANNEL TANDEM CANCELLER

The results of one of the ANN flight tests are shown in Figure 13. The upper trace is the uncancelled antenna line and shows 400 Hz harmonic spikes. In the lower trace the desired signal (the Cutler, Maine transmitter operating at 17.8 kHz) is unaffected but the 400 Hz harmonics have been cancelled.

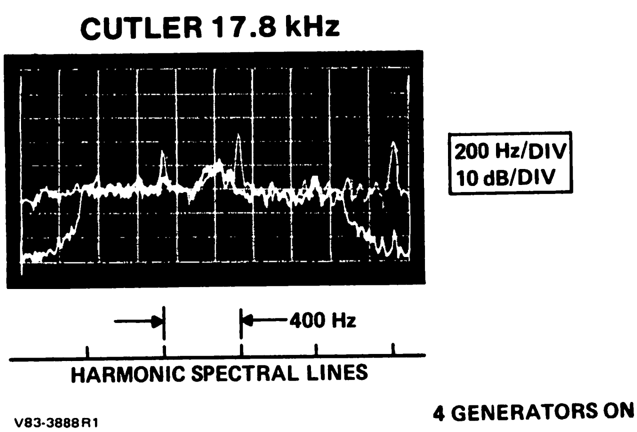

FIGURE 13. IN-FLIGHT PERFORMANCE OF ANN

Tests of cross polarization were also conducted using a transverse electric (H-field) loop in place of the monopole in ADARS. Even with the airplane in a steep bank the ADARS system was able to isolate signals of different polarizations.

SYSTEM ARCHITECTURAL CONSIDERATIONS

Three types of digital adaptive array (and canceller and cross polarizer) architecture are shown in Figure 14. The closest to the conventional analog loop is shown at the top of the figure. It is called a hybrid system because it employs conventional analog weights and performs the measurements in analog correlators. The correlations are digitized and the algorithmic processing is performed digitally in a microprocessor. The input is RF. The output is RF. Hybrid systems are appropriate for wideband signals such as GPS.

The middle portion of Figure 14 depicts the all-digital system which is applicable to narrow-band signals at low frequencies such as VLF and LF. In this case, the RF from each element is digitized and used to generate the covariance matrix which, in turn, is manipulated by the algorithm to generate the weights. The summed (adapted) output is a digital representation of the adapted pattern output which then can be (digitally) downconverted and/or demodulated.

A way of realizing the advantages of the all-digital system at frequencies where it is impractical to digitize the RF is shown as the converted all-digital system in the bottom of Figure 14. Here an analog downconverter (with common LO for all channels) is used to convert the band of interest down to where digitization of the IF is practical.

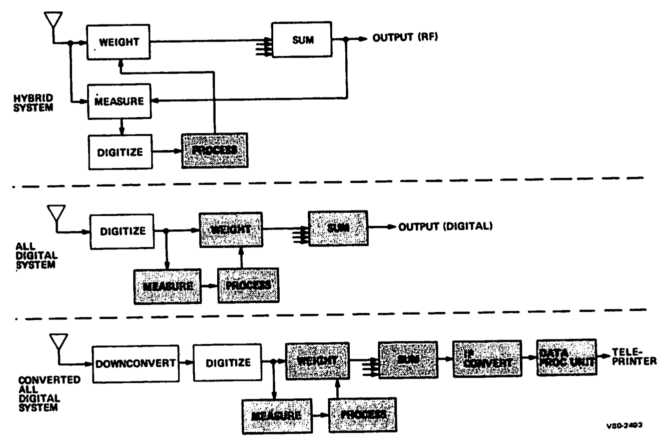

FIGURE 14. ADAPTIVE ARRAY SYSTEM ARCHITECTURE

An example of this type of architecture is HF ADARS which covers the range up to 30 MHz. Figure 15 shows the performance of HF ADARS using the Eigenvector algorithm. This algorithm provides as many outputs as there are elements with the characteristic that the outputs are made maximally orthogonal. In the picture, Radio Paris and a rock music station are shown as mutual interferers in the input bandwidth. The three outputs (in the same bandwidth) have separated the signals, without reliance on any a priori information, so that depending upon the spigot one listens to one gets either Radio Paris (EV = 3), or rock music (EV = 2) or neither (EV = 1).

A byproduct of digital processing is that the transmission can be stored while processing is ongoing so that the results of processing (in this case applying the appropriate weights to each antenna input prior to summing) can be applied to the data from which the measurements are made. The value of this capability at VLF is that it permits the ADARS to null highly impulsive lightning strokes. This is demonstrated in Figure 16 where lightning strokes were counted (on an unused frequency). Both the monopole and North-South antennas received over 200 lightning strokes in successive 4 minute periods, whereas the adapted output (sum of three antennas) had only seven lightning strokes in a comparable 4 minute period. The lower right hand portion of the figure pictures a superposition of resultant antenna patterns showing the nulls to be southeast of Long Island, corresponding to the expected location of storms during the month of September.

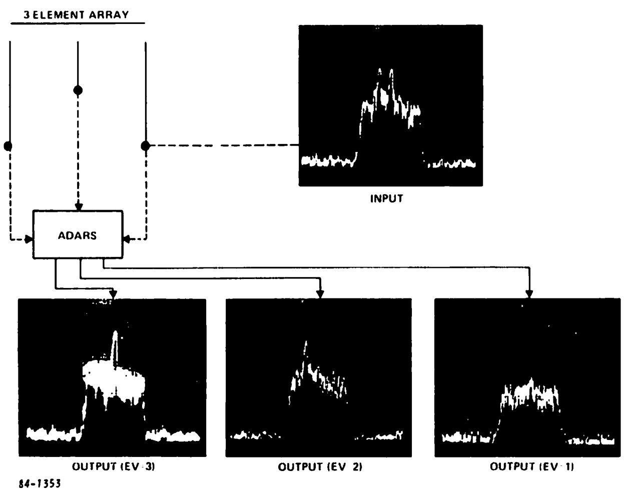

FIGURE 15. HF ADARS PERFORMANCE (17.94 MHz)

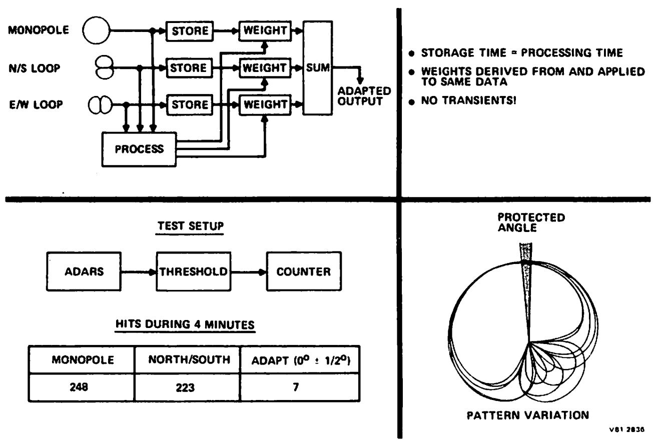

FIGURE 16. INSTANCY

Digitizing RF at 100 to 400 kilosamples per second with 12 bit A to D conversion for each channel of a multielement array results in a lot of bits to process. Tens of megops per second are required - well beyond the current microprocessor state-of-the-art.

Our solution to this problem is the Distributed Processor. Its relation to the interface is depicted in the upper portion of Figure 17. The make-up of the Distributed Processor is shown in the lower half of the figure.

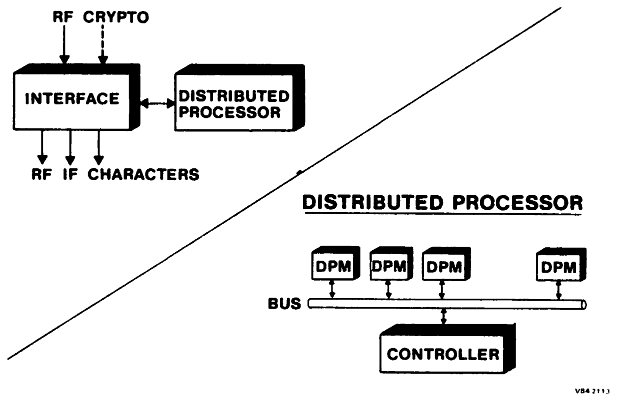

FIGURE 17. THE DISTRIBUTED PROCESSOR

The Distributed Processor allocates functions into two categories:

- Fast repetitive processing requiring no decisions
- Slow processing requiring decisions

The former is accomplished with Digital Processing Modules (DPM's) that are capable of 5 million multiplications and accumulations per second. They perform the functions of filtering, frequency conversion, correlation, weighting, summing and AGC.

The latter is accomplished with 16-bit micro-processors that apply the algorithm, perform demodulations and provide the operator interface.

The DPM's are RAM controlled with program downloaded from the processor. All DPM's are tied to a common bus so that the systems can be electronically reconfigured by the processor. The resultant system is very flexible and capable of being reconfigured to provide automatic fault tolerant operation.