B.D. Carlson, L.M. Goodman, J. Austin, M.W. Ganz, and L.O. Upton

An Ultralow-Sidelobe Adaptive Array Antenna

One of the most important functions of a radar antenna is to provide spatial filtering that maximizes the radar's sensitivity in the desired surveillance direction while suppressing interference signals that enter the radar from other directions. This article describes a UHF radar antenna system with exceptional interference-rejection capabilities that are achieved through a combination of precision passive beamforming in the azimuth plane, and active, digital adaptive nulling in elevation. The antenna is composed of 14 stacked rows; each row contains a stripline, low-sidelobe, passive corporate-feed network. For signal reception, a separate receiver and analog-to-digital converter are used at each row output. The system adaptively combines the digitized row signals to form an elevation pattern with nulls at the elevation angles of interference sources. The antenna system is part of an advanced air surveillance radar that Lincoln Laboratory is developing for the Navy. This article describes the design of the antenna system and presents performance analyses and measured results.

Modern radars use a variety of electronic counter-countermeasures to suppress both inadvertent interference and intentional jamming. When the desired target and the interference source are not at the same location, a radar can take advantage of the separation to suppress the interference without suppressing the target return. This selective suppression can be achieved by a radar antenna with a radiation pattern that provides maximum gain in the direction of the target and a very low response in the interference direction.

There are several techniques for controlling the shape of an antenna radiation pattern. Passive, fixed tapering of the aperture illumination can produce a pattern with uniformly low sidelobes outside the direction of the main beam. A pattern of this type is useful for suppressing both interference and clutter. A key parameter in low-sidelobe design is the radiation-pattern sidelobe level expressed in relation to isotropic gain, which is the gain of a lossless antenna that radiates uniformly in all directions. Without undue difficulty, electrically large, high-gain antennas can achieve radiation patterns exhibiting sidelobes that are low with respect to the peak gain, but not with respect to the isotropic level. On the other hand, when the peak gain is limited by physical aperture size constraints, as is commonly the case for radar antennas, the desired sidelobe level is often well below isotropic. Low isotropic sidelobe levels can be difficult to attain because they require very accurate aperture illumination (described in the next section, "Azimuth Beamforming"). Aperture-illumination errors become even more critical when they are correlated within the aperture illumination.

Another interference-rejection technique is adaptive nulling. An adaptive antenna generates nulls in the pattern only at the locations of interferers. Adaptive nulling is an especially attractive option for cases in which the required long-term illumination tolerances for fixed networks become impractical, or when the beamwidth increase and loss associated with severe aperture tapers cannot be accommodated.

When carefully implemented, adaptive nulling can achieve extremely deep nulls at the locations of interfering signals. In phased-array applications, the nulls produced with adaptive nulling can be much deeper than the sidelobe level that can be achieved in practice with an

amplitude taper alone. The price of the superior performance is the increased complexity of the antenna system. Adaptive-nulling systems typically require a separate receiver behind each adaptive antenna element. Also required is a complex electronic control system to set the element weights adaptively. That is, when an adaptive-nulling radar changes its frequency of operation, a new set of weights must be computed and applied to the elements of the phased array to null any interfering signals that are present at the new frequency (described in the section "Elevation Pattern Control" on p. 299). In this way, the antenna system automatically moves the nulled positions as the positions of the interfering sources move with respect to the antenna. In addition to adjusting for such movement, adaptive nulling also automatically compensates for component aging, temperature effects, and mutual coupling between the individual antenna elements.

This article describes an antenna system for an experimental, UHF air surveillance radar that has been designed to provide long-range detection of high-altitude targets in the presence of severe interference and clutter. The antenna uses a combination of passive (i.e., fixed amplitude taper) and active (i.e., adaptive nulling) beamforming techniques to achieve excellent spatial rejection of interference in a system of reasonable complexity.

The antenna is a rotating planar array that consists of 14 stacked rows with 24 radiating elements per row (Fig. 1). A precision corporate

Fig. 1—UHF ultralow-sidelobe array antenna installed at Lincoln Laboratory.

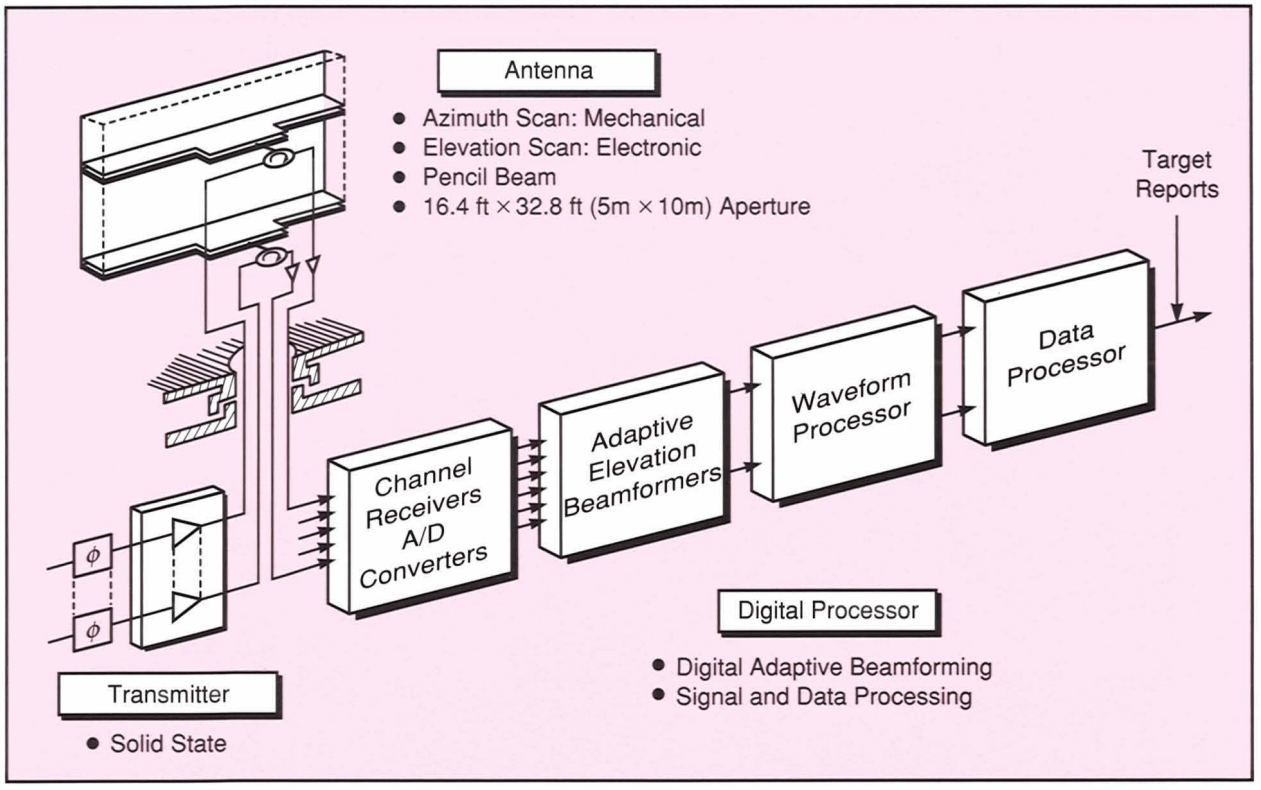

Fig. 2—Radar configuration.

feed is used to apply a fixed amplitude taper to the 24 radiators in each of the rows. Thus each row produces an antenna pattern that is a fan beam in elevation and has very low sidelobes in azimuth directions other than that of the main beam (which is broadside to the row). The fixed, precision design of the azimuth corporate-feed network produces azimuth sidelobes low enough to suppress even strong interferers. Fourteen matched receivers are used to down-convert and amplify the row output signals prior to A/D conversion. A special-purpose high-speed digital processor combines the digitized signals from the rows to form a pencil beam in elevation. Interference that is separated from the desired target in elevation is adaptively nulled.

Although the overall antenna system described here is currently under development, construction of the ultralow-sidelobe passive array has been completed, as shown in Fig. 1. Under subcontract to Lincoln Laboratory, Westinghouse Electric Corp. in Baltimore developed and tested the 14-row passive antenna, which will become a central part of an experimental radar system. Lincoln Laboratory is developing the system for the Navy to demonstrate advanced radar technology for shipborne applications. The adaptive elevation beamforming technology has been developed within Lincoln Laboratory.

Figure 2 shows the configuration of the surveillance radar. Each row of the antenna is associated with a separate transmitter channel. A multichannel rotary coupler feeds the coherent transmit signals, which are phase-shifted for elevation scanning, to the 14-row array. Solid state amplifiers in the transmitter facilitate subdivision of the transmitted power between the different rows. Subdivision of the transmitted power permits low-power phase shifting to precede the power amplification.

The antenna system performs the beamforming on reception in two stages. Azimuth beamforming occurs in the analog domain within each row of the 14-row, passive antenna. The 14-row outputs are preamplified and cabled, via the multichannel coupler, to 14 re-

ceivers and A/D converters. Elevation beamforming, with electronic scanning capability, is performed adaptively in the digital domain. Special-purpose digital hardware processes the digital signals at the adaptive beamformer output to perform the conventional radar signal- and target-processing functions.

This article describes the basic theory and design techniques associated with the achievement of low-sidelobe patterns and deep nulls. It was not possible, however, to present all of the achieved results in an open-literature article. For example, the sidelobe specification for the antenna is not stated here, but it is a value more than 50 dB below the main-beam peak, or 21 dB below the isotropic level (i.e., a value less than –21 dBi). Because we wanted to illustrate some degree of measured performance, the azimuth sidelobe patterns shown have been intentionally truncated at –50 dB (–21 dBi). Furthermore, this article does not show adapted 14-channel elevation patterns; however, we present the measured results from a non-real-time, four-channel, adaptive-nulling test bed built for demonstrations of adaptive nulling.

This article contains four sections: an introduction; the section "Azimuth Beamforming," which describes the theory, hardware design, and experimental results for the passive, fixed, azimuth-row beamformers; the section "Elevation Pattern Control," which focuses on the theory, hardware limitations, and experimental results of digital adaptive beamforming; and a summary.

Azimuth Beamforming

Description of Array Hardware

To achieve reasonable beamwidths, UHF operation requires an antenna with a large aperture; the antenna described in this article is a rotating planar array that is 16.4 ft × 32.8 ft (5 m × 10 m) and less than 3 ft deep. At UHF, a tapered aperture of this size provides about 29 dB of gain above the isotropic level. The required sidelobe performance is a value more than 21 dB below the isotropic level. Careful design and manufacture of the row corporate-feed beamformers was necessary to achieve the very low aperture-illumination tolerances that were required.

The 14-row array design has an open-ground-plane structure: approximately 40% of the aperture is open to minimize weight and to allow wind, but not electrical energy, to pass through. Radomes on each of the rows protect the radiating elements and deflect the wind. The effect of the radomes on the antenna pattern is insignificant at UHF. Except for the fiberglass radomes, the antenna is constructed of aluminum. The entire antenna, excluding the strong back and the pedestal, weighs 4700 lbs.

A basic trade-off in the design of the passive-antenna rows was that of sidelobe level versus bandwidth. A wide operating bandwidth provides a radar with the ability to frequency-hop over a wide range, which thereby dilutes the energy of barrage (i.e., wideband) interference. However, since the added circuit complexity associated with very wideband design is most likely to raise the sidelobe level at a rate faster than that at which the benefits of wider bandwidth accrue, the attainment of unusually wideband performance was not emphasized. Nevertheless, we anticipate operation over a bandwidth greater than 50 MHz.

The rotating antenna is composed of 14 air-dielectric stripline distribution networks that divide the energy from the single feed point for each row to the row's 24 radiating elements. The stripline consists of parallel ground planes of aluminum with a center conductor running between them. Dielectric spacers support the center conductor, which is separated from the ground planes by air gaps above and below. Air-dielectric stripline was used for the distribution network to minimize losses and to avoid the illumination errors that result from nonhomogeneity of solid dielectrics.

Azimuth Pattern of the Array

The azimuth pattern of the full array is the complex superposition of the individual-row azimuth patterns, weighted by the elevation taper. Great care, therefore, must be taken to

ensure that the azimuth pattern of each row has low sidelobes in order to ensure that the complete array also has low sidelobes. If each row of the array is absolutely identical, the array will have exactly the same azimuth pattern as each of the individual rows. If, however, each of the rows has independent random errors that cause different error sidelobes to occur in each row pattern, the superposition of the rows will result in an array pattern in which the error sidelobes add up incoherently or with random phase. In contrast, the mainlobes of the rows will all be phase-aligned and will thus combine coherently. Therefore, the net effect is for the relative sidelobes of the array to decrease as more rows are added to the array. The isotropic sidelobe levels actually remain unchanged as rows are added, but the main-beam gain increases. This effect can be quantified by the well-known equation of J. Ruze [1]. The equation predicts the expected value of the error sidelobes (relative to the main-beam peak) in an array in terms of the independent radiator errors:

$$\overline{SLL} = \frac{\sigma_\phi^2 + \sigma_A^2}{N\eta},$$

where $\overline{SLL}$ = the average sidelobe level,
σ_ϕ^2 = the rms phase error in radians,
σ_A^2 = the rms amplitude error expressed as a fraction of the nominal amplitudes,
η = the taper efficiency, and
N = the number of elements with independent errors.

The effects of errors that are correlated from row to row are much more severe than the effects of random errors, since the correlated errors do not decrease when the rows are combined into an array.

In the experimental antenna array, close attention was paid to reducing the excitation errors (especially the correlated errors) as much as possible in an effort to determine to what extent the sidelobes could be reduced. The theoretical-design azimuth pattern for each row was chosen to be a Chebyshev pattern of –65 dB in order to keep the mathematical sidelobes below the expected sidelobes. Careful budgeting of the amplitude, phase, and positional error tolerances was necessary to ensure that the total illumination error was low enough to produce the desired sidelobe level.

The feed network in each row is a corporate design. To ensure that all radiators will be in phase across the bandwidth, the electrical distance to each of the 24 radiators from the input is designed to be equal. Resistors in the Wilkinson power dividers used in the air-dielectric stripline network dissipate much of the unwanted internally reflected energy. The dissipation prevents the energy from introducing illumination errors at the radiators. Modified Wilkinson couplers are used in the first few high-power divisions in the network because the modified couplers allow a higher power dissipation than that permitted by the standard Wilkinson design. The mode of operation of the modified couplers is the same as that of the standard design except that the modified couplers direct any reflected energy to resistors that are mounted remotely where heat can be dissipated externally.

The air-dielectric stripline was built as a two-layer package in order to minimize the physical space required for each row. The first few high-power divisions are done in the upper package, which has a 3/4-in ground-plane spacing. The outputs from this package are fed through at six points to the lower package where the remaining power splits are done with a ground-plane spacing of 1/2 in. This configuration is compact enough that the whole stripline package is less than 2 ft deep and 2 in thick.

Numerically controlled aluminum stamping and milling machines manufactured the stripline couplers. The center conductor for each of the couplers was produced and individually tested to fine-tune each power divider's electrical characteristics. Once satisfactory coupler designs had been achieved, a design for a complete row composed of the tested couplers was produced. The numerically controlled machines then constructed a single row, which was tested in the laboratory with a network analyzer to measure the amplitude and phase errors. On the basis of the test measurements, errors could be traced back through the divide network to the

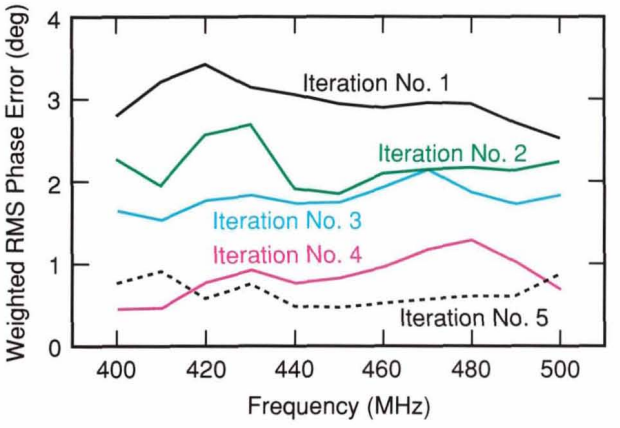

Fig. 3—The rms excitation phase errors of the antenna row as a function of frequency. The graph plots results from the five iterations of the antenna-row design.

couplers that caused the errors. Corrections to the row design were then made, and a new row with improved performance was produced. This entire design process was repeated through five iterations until the remaining errors appeared to be random. Figure 3 shows the root-mean-square (rms) level of the excitation phase errors as a function of frequency for the five iterations. The results of the five iterations allowed the choosing of a final row design for the array. The first three couplers of each of the individual rows were designed to be slightly adjustable by leaving access covers in the ground planes where teflon inserts could be placed. The inserts, which allow the adjusting of the amplitude and phase at the coupler outputs, can be used to trim random manufacturing errors. Figure 4 shows the 1/2-in stripline package during the assembly process.

Careful attention was also given to the design of the radiating elements. Trough radiators were chosen because of their very low cross-polariza-

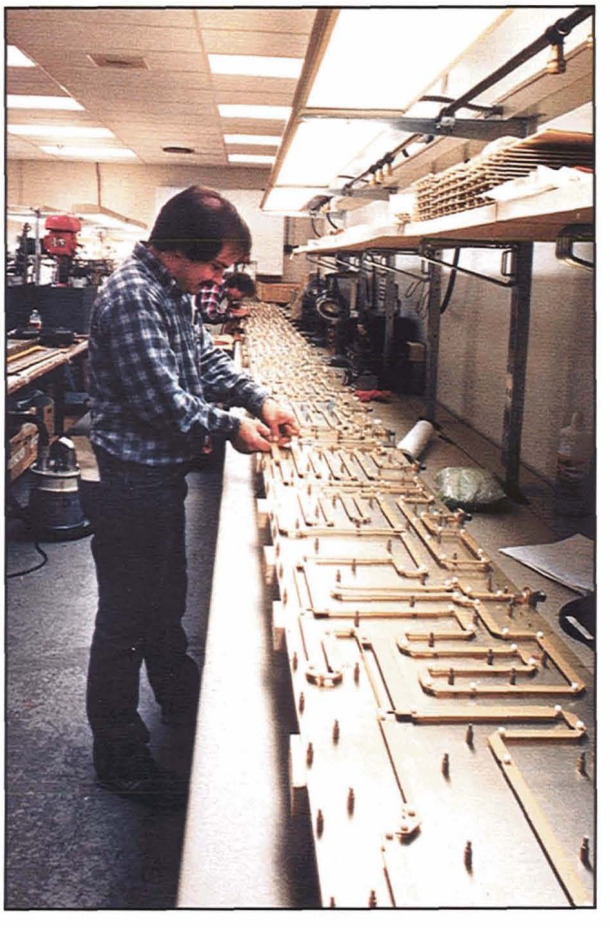

Fig. 4—Construction of a single row of the 1/2-in air-dielectric stripline package. (Photo courtesy of Westinghouse Electric.)

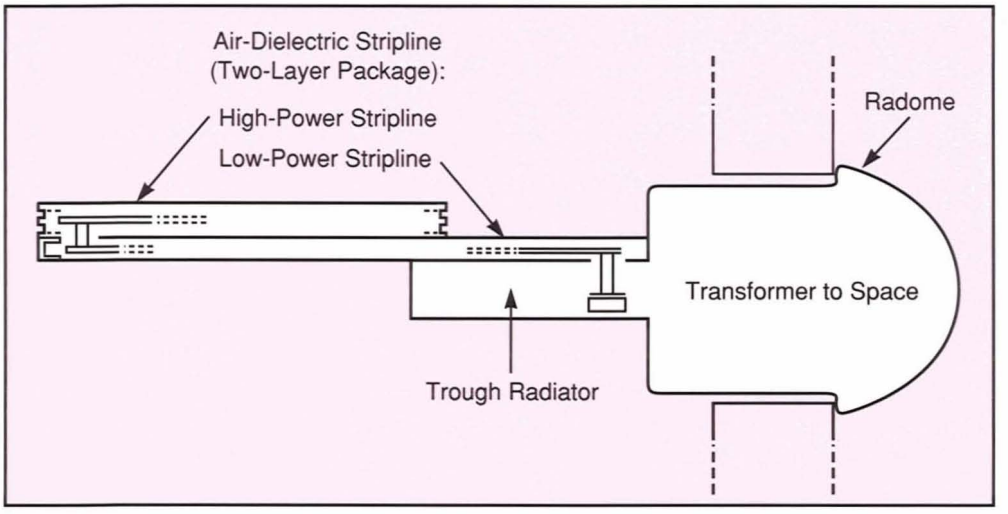

Fig. 5—Cross section of the end of a row.

tion characteristic. When dipoles are used in a low-sidelobe array, currents can be excited on the posts of neighboring dipoles, which can then reradiate and interfere with the desired low-sidelobe response of the array in both polarization directions. The use of trough elements avoids this problem.

Figure 5 shows a cross section of a row. The diagram includes the row radome and the trough, which runs the entire length of the row and has radiating posts that are fed from the stripline above it. The energy launched in the trough section is matched to free space by the quarter-wave transformer section. The resulting transmission has linear, vertical polarization.

Figure 6 is a photograph of a preliminary four-row array during the assembly process. The radome can be seen on the bottom row, and the troughs and feed points on the upper rows.

To match the impedance and measure the mutual coupling of an array of trough radiators, a waveguide simulator was constructed. The simulator consisted of two reflecting plates mounted orthogonally to the ground plane of a row of seven elements. The reflecting plates allowed the elements to interact with their reflections as if the reflections were actually other elements in an infinite array. This technique helped to achieve a good impedance match for the radiating elements, and it helped to facilitate measurement of the mutual coupling of the array. The measured data on mutual coupling were applied to the theoretical aperture distribution to find the actual excitation for the 14-row array and to predict the perturbation of the ideal far-field pattern. Because of the severe taper on the aperture, the calculated effect of mutual coupling on the azimuth pattern was found to be negligibly small. In the elevation pattern, mutual-coupling effects are relatively unimportant because the adaptive system, in effect, measures the errors and compensates for them when the system forms the adapted elevation beam.

Experimental Results

Once the total row design had been completed, a subarray of four rows was built to conduct an initial performance evaluation. During the evaluation, the row design was high-power-tested with a UHF solid state transmitter to a peak input-power level of 120 kW to a single row. After the design passed this test, pattern measurements were made on Westinghouse's far-field antenna test range. Following the successful four-row measurement, the full 14-row array was assembled and pattern-tested in a similar manner.

The Westinghouse far-field antenna range is a 1700-ft-long ground reflection range. Reflections from hillsides and trees at the sides of the range complicate antenna testing. These multi-

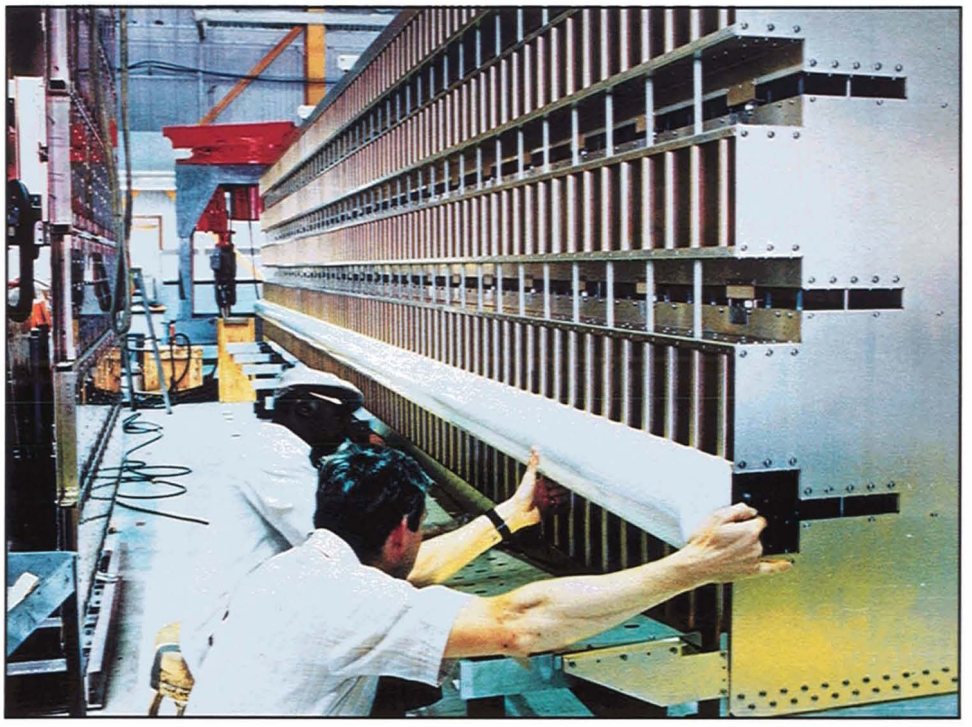

Fig. 6—Construction of the experimental four-row array. (Photo courtesy of Westinghouse Electric.)

path reflections can enter the main beam and effectively cause an apparent degradation of the measured azimuth sidelobe pattern. Careful test planning and the removal of some trees were required to ensure meaningful measurements of sidelobe levels.

To obtain accurate antenna measurements, we must be able to distinguish between the true antenna pattern and the range reflections. This differentiation can be accomplished by measuring the antenna pattern and then turning the antenna upside down and repeating the measurement. Figure 7 shows the azimuth principal-plane pattern of the complete 14-row array when measured at 440 MHz, which is close to the center frequency of 450 MHz. For this mea-

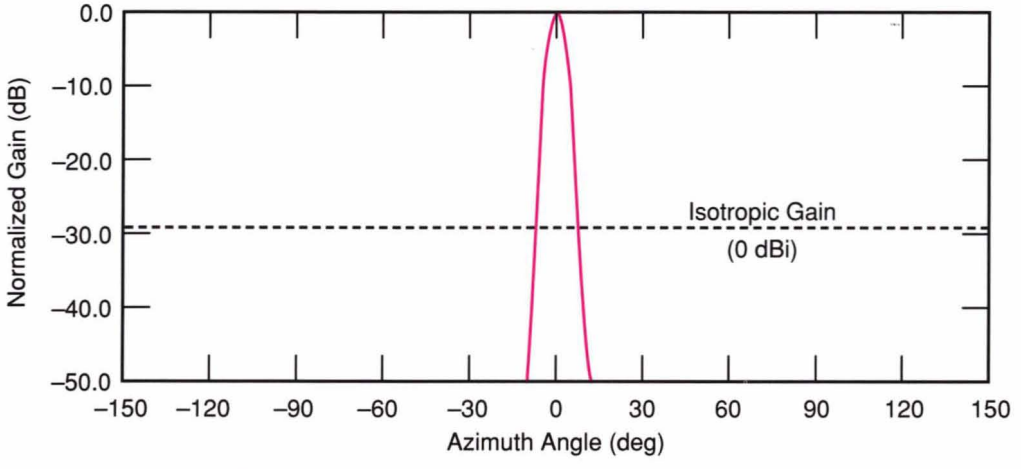

Fig. 7—Azimuth pattern for 14-row array at 440 MHz.

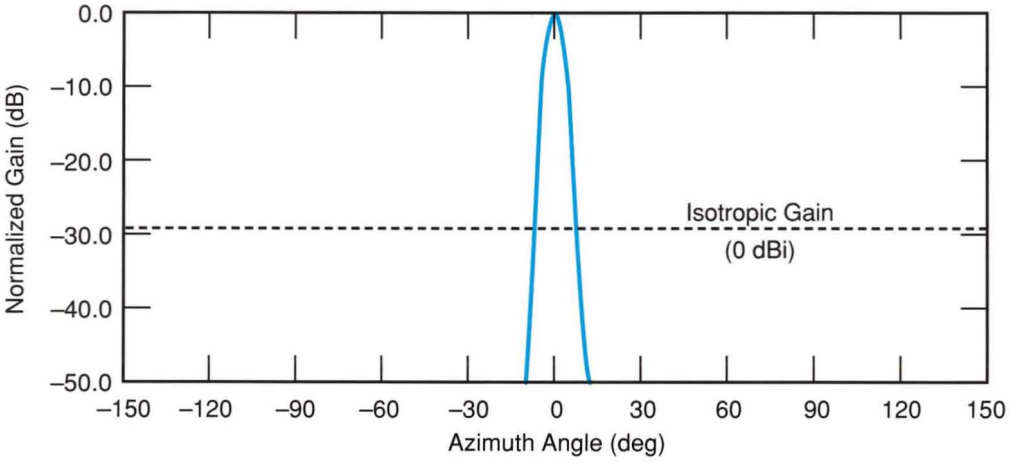

Fig. 8—Azimuth pattern for 14-row array averaged over 410 to 460 MHz.

surement, the outputs of the 14 rows were combined in elevation with an analog combiner that provided uniform illumination except for a –6-dB weighting on the top and bottom rows. Figure 7 shows that all azimuth sidelobes are below the pattern truncation level of –50 dB (–21 dBi). Outstanding performance was also obtained over the 410-to-460-MHz band (Fig. 8). Note that the curve of Fig. 8 is virtually identical to that of Fig. 7. In fact, after truncation to the level of –50 dB, single-frequency patterns between 410 and 460 MHz are essentially the same as that shown in Fig. 7.

We found that the measured sidelobes occurring at wide azimuth angles were related to illumination errors that were random from row to row. We know that they were random because the measured level of the errors in the full array was lower than the level in the partial array of four rows. (Figures 7 and 8 do not show the sidelobes because they are below the truncation level of –50 dB.) In contrast, the sidelobes that occurred at smaller azimuth angles were slightly higher and were related, in part, to errors at least partially correlated between the rows. We know that they were correlated because, in going from the four-row to the full 14-row array, we saw a smaller reduction in the small-azimuth-angle sidelobes than would have occurred if the errors had not been correlated.

The antenna also exhibited excellent cross-polarization performance. The measured cross-polarization principal-plane pattern was 40 dB below the primary polarization within the main beam, and approximately 15 dB below the primary polarization in the azimuth sidelobes.

Diffraction fences mounted on the upper and lower edges of the array minimized the edge diffraction and helped to control the pattern response directly behind the array. As a result, the measured backlobe was more than 50 dB below the beam peak.

In summary, the 14-row array successfully met the various design goals for controlling aperture illumination errors. The array achieved low azimuth sidelobe levels with circuits designed for high-power handling and with production-orientated manufacturing techniques.

Elevation Pattern Control

The previous section described the design of the passive 14-row ultralow-sidelobe array that suppresses those interference signals which enter the radar at the azimuth sidelobe locations. It is also necessary to discriminate against jammers that are at the same azimuth as the main beam but separated from the desired scan direction in the elevation plane. This case arises, for example, when a radar searches for a high-altitude target in the presence of a jammer located near the horizon. This section describes how the array antenna adaptively combines row signals to form elevation beams with nulls at the elevations of jammers.

The antenna system performs the elevation beamforming digitally on reception and samples the digital signals from each of the 14 rows in

order to obtain an estimate of the interference environment. The Sample Matrix Inversion (SMI) algorithm computes weights for each of the rows. When applied to the row signals, the SMI weights produce an antenna pattern with nulls at the elevation angles of interfering sources, while simultaneously maintaining the maximum gain in the desired direction. The antenna system must frequently recalculate the adaptive weights to keep up with changes in the interference environment, in the operating frequency of the radar, and in the antenna position.

The following subsections describe the theory of adaptive nulling and discusses the hardware implementation and practical limitations on the depths of nulls that are achievable.

Theory of Sample-Matrix-Inversion Adaptive Nulling

This subsection briefly describes the mathematical basis of adaptive nulling. Consider, for example, a four-element linear array of antenna elements that are spaced $\lambda/2$ apart, where λ is the wavelength. (In this case, the four elements can be thought of as four rows.) Assume that a desired signal is incident upon the array from some elevation angle θ_d, as measured in a direction broadside to the array. Also assume that thermal noise signals with power σ_n^2 exist on each element and that the signals are mutually uncorrelated. In addition to the thermal noise signals, there may also be interfering signals incident upon the array.

A signal vector, whose components are the complex signals on the rows, is denoted by $\mathbf{X} = [x_1(t), x_2(t), x_3(t), x_4(t)]^T$, in which T denotes the transpose operation. $\mathbf{X}$ is multiplied by a complex weight vector $\mathbf{W} = [w_1, w_2, w_3, w_4]^T$ to produce the array output. The array output $s(t)$ is thus given by

$$s(t) = \mathbf{W}^T\mathbf{X}.$$

The signal-to-interference-plus-noise ratio (SINR) is maximized at the array output if the weight vector is set according to Ref. 2:

$$\mathbf{W} = \mathbf{R}^{-1}\mathbf{T} \qquad (1)$$

where $\mathbf{R}$ is the covariance matrix given by

$$\mathbf{R} = E[\mathbf{X}\mathbf{X}^\dagger]$$

and $\mathbf{T}$ is a *steering vector* given by

$$\mathbf{T} = \begin{pmatrix} 1 \\ e^{j\phi_d} \\ e^{j2\phi_d} \\ e^{j3\phi_d} \end{pmatrix}. \qquad (2)$$

In the above equations, E[. . .] denotes the expected value of a quantity, † denotes the conjugate-transpose operation, and ϕ_d is the element-to-element phase shift for a signal that arrives from an angle θ_d; i.e.,

$$\phi_d = \pi \sin(\theta_d)$$

for $\lambda/2$ element spacing.

We can gain insight into the behavior of the weights by examining the various terms in Eq. 1. First, in the absence of interference, the covariance matrix is given by

$$\mathbf{R}_Q = \sigma_n^2\mathbf{I}, \qquad (3)$$

where $\mathbf{I}$ is the identity matrix and the subscript Q is added to $\mathbf{R}$ to indicate that $\mathbf{R}_Q$ is the quiescent, or noise-only, covariance matrix. Therefore, in the noise-only case we can combine Eqs. 1 and 3 to produce

$$\mathbf{W} = \frac{1}{\sigma_n^2}\mathbf{T}.$$

Thus, in the case in which no interference is present, the array weight vector $\mathbf{W}$ is simply equal to a scalar constant multiplied by $\mathbf{T}$. Note that $\mathbf{T}$ is the weight vector that would be applied to a conventional (i.e., nonadaptive) phased array to maximize the array gain on a signal that arrives from the θ_d direction.

We can modify the quiescent pattern by changing the steering vector $\mathbf{T}$ that is used in the calculation of the weight vector. The quiescent pattern specified by the vector $\mathbf{T}$ in Eq. 2 is that of a uniform aperture illumination. This pattern has a peak sidelobe level that is 13 dB lower than the main-beam peak. By applying an amplitude taper to $\mathbf{T}$, we can obtain other quiescent patterns. Chebyshev, Taylor, and Hamming, as well as other weight tapers, can be used to obtain low

sidelobes with only a small corresponding loss in main-beam gain.

Before we can use Eq. 1 to calculate the optimum (i.e., the maximum SINR) weight vector **W**, we must obtain an estimate for the covariance matrix **R**. Because of the stationary and ergodic conditions of the input signals, the maximum likelihood estimate of **R** (denoted by $\hat{\mathbf{R}}$) is given by Ref. 3:

$$\hat{\mathbf{R}} = \frac{1}{K}\sum_{i=1}^{K} \mathbf{X}_i \mathbf{X}_i^{\dagger} \qquad (4)$$

where $\mathbf{X}_i$ is an instantaneous sample, or snapshot, of the signal vector **X**, and K is the number of samples taken.

The process for computing adaptive weights proceeds as follows. First, the signals on the elements are sampled and $\hat{\mathbf{R}}$ is formed with Eq. 4. Next, $\hat{\mathbf{R}}$ is numerically inverted and postmultiplied by a steering vector **T** that has been chosen to give the desired quiescent pattern.

A central issue that must be addressed during the implementation of the SMI algorithm is the determination of how large K must be (i.e., how many samples must be taken) to ensure that the estimate $\hat{\mathbf{R}}$ of **R** is good enough. A K value that is too small results in a suboptimal weight vector typically producing reduced main-beam gain and increased sidelobes [4].

Figure 9 illustrates a simulation of this problem for a 14-element array. For the case shown in the figure, we set the steering vector **T** to produce a quiescent pattern at a beam-steering direction of 25°. We then applied an amplitude taper to the steering vector to produce a Chebyshev pattern shape with sidelobes of –40 dB (the magenta curve in Fig. 9). The blue curve in Fig. 9 shows the pattern when $\hat{\mathbf{R}}$ is used to calculate the weights with $K = 50$ and a strong interfering signal at 0°. Note that the finite sampling causes a distortion of the main beam and an increase in the peak sidelobe level to about –15 dB. Thus it appears that more than 50 samples are required to obtain a good adapted pattern. Note, however, that the adapted pattern does have a null at 0°, the direction of the interference source.

There is a way to achieve a much better adapted pattern for the case described above without an increase in the required number of samples. The improved performance can be achieved by diagonal loading, a technique in which the diagonal elements of the sample matrix are augmented prior to inversion. That is, instead of inverting $\hat{\mathbf{R}}$, we invert $\hat{\mathbf{R}}'$ where

$$\hat{\mathbf{R}}' = \hat{\mathbf{R}} + \beta \mathbf{I}. \qquad (5)$$

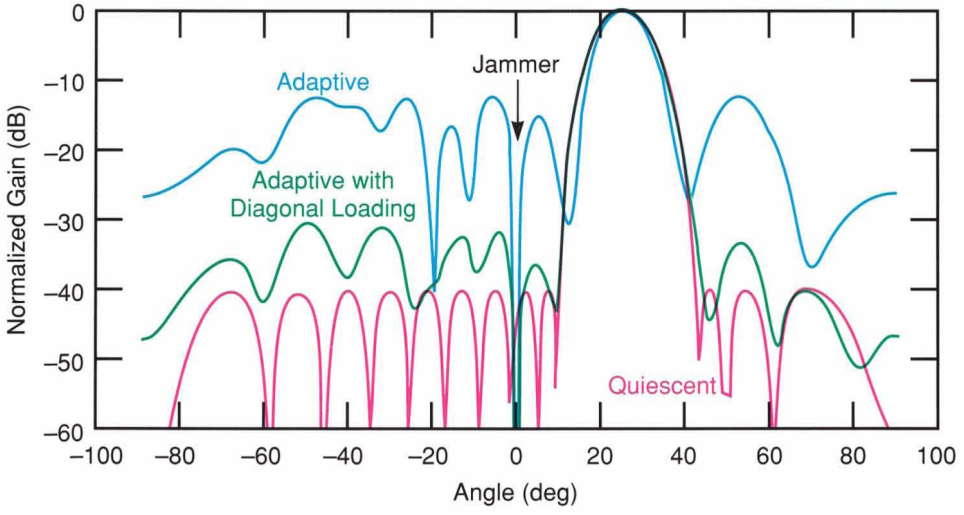

Fig. 9—Simulation of 14-element array with beam steered to 25° and a strong interfering signal at 0°. The magenta line represents the quiescent pattern, the blue line the adaptive pattern, and the green line the adaptive pattern with 10 dB of diagonal loading, a technique in which the sidelobe levels are lowered while nulls in the direction of interference signals are maintained.

Typically, we chose β to be 3 to 10 dB greater than the thermal noise power on each element. The green curve in Fig. 9 shows the performance of the array with K = 50 samples and $\beta = 10\sigma_n^2$. The curve shows that the diagonally loaded covariance matrix produces weights that result in an antenna pattern much closer to the quiescent pattern. Note that the diagonally loaded pattern also retains the null at 0°, the direction of the interference source.

An examination of Eqs. 3, 4, and 5 leads to an understanding of the diagonal-loading operation. From the equations, we see that diagonal loading replaces the covariance matrix by a weighted sum of the covariance matrix and the quiescent covariance matrix. Therefore, the adapted pattern when $\hat{\mathbf{R}}'$ is used retains more of the features of the quiescent pattern than when $\hat{\mathbf{R}}$ is used. One undesired effect of this averaging of the quiescent and adapted patterns is a reduction in the depth of the null when $\hat{\mathbf{R}}'$ is used. Even though the null depth is reduced, however, the null is still deep enough so that the interference has a negligible effect. B.D. Carlson presents a detailed discussion of the effects of diagonal loading in Ref. 4.

Null-Depth Limitations

The level of interference rejection and the ability to maintain a reasonably shaped beam are the primary measures of performance of an adaptive array. In practice, there are several effects that can limit the depth of adaptively generated nulls. One major effect is element-to-element signal decorrelation. In discussing the effects of element-to-element signal decorrelation, it is important to remember that the adaptive array cancels an interfering signal by applying a complex weight to the signal that is incident on each element, and the array then sums the weighted signals. Thus, if the signals on the various elements are not perfectly correlated (i.e., if the signals differ from one another in some manner other than a complex multiplicative constant), the array will not be able to null the signals perfectly.

There are two mechanisms by which the signals on the various elements can become decorrelated. The first of these mechanisms is signal dispersion, which occurs when a signal with nonzero bandwidth arrives upon the array from some angle other than the broadside direction. The element-to-element propagation time delay between each pair of elements produces a phase shift proportional to the signal frequency. Therefore, the complex weights required to null a signal at one frequency will be slightly different from the weights required to null a signal at another nearby frequency. If the interfering signal has a significant bandwidth and if the signal arrives from some angle other than broadside, the array will require several closely spaced nulls (i.e., the use of several degrees of freedom) to null all frequency components simultaneously.

A similar dispersion effect occurs when signals with nonzero bandwidth undergo undesired reflections in the RF hardware. The reflections, which can result from cable and device interconnections, are time delayed by the round-trip propagation delays. Mutual coupling can allow signals that are reflected back from one channel to enter other channel. Therefore, there might be many very low-level versions of a signal present at the element inputs and the versions can have a wide variety of time delays. If the signals exist at a high level and if they have a nonzero bandwidth, the array may not be able to null a single jammer effectively, even if the array uses several of its available degrees of freedom.

The second mechanism that can cause signal decorrelation from element to element is a mismatch in the frequency responses of the receivers on each channel. It is important to note that when the frequency responses do not match exactly, it becomes impossible to choose a set of complex weights that will perfectly cancel any interference. In the experimental adaptive array, the problem is overcome with the insertion of 32-tap transversal digital equalizers at each receiver output. The initial injection of a calibration signal into all channels adjusts the equalizers to match each channel's frequency responses. A least-squares algorithm is used to perform the matching.

Another major limit on null-depth perfor-

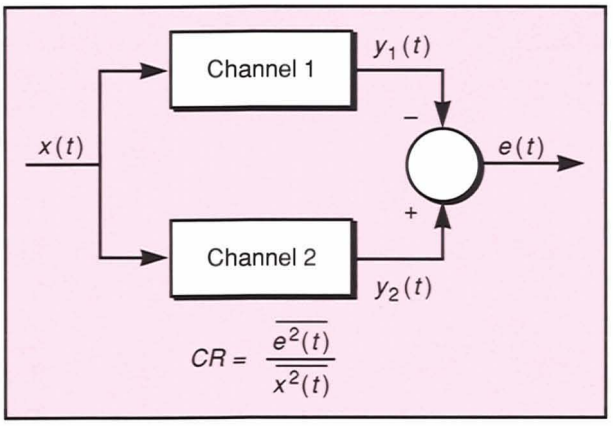

*Fig. 10—Definition of two-channel cancellation ratio (*CR*). The overline notation indicates the mean value of a quantity.*

mance is the numerical accuracy of the system. For the experimental adaptive array, the primary contributor to numerical inaccuracy is the word length of the A/D converter.

Cancellation Ratio

To a large extent, the degree to which the frequency responses of the array channels can be matched determines the level of interference rejection. The two-channel cancellation ratio (*CR*) is a measure of the level of interchannel matching. Figure 10 depicts the definition of *CR*, which is given mathematically by

$$CR = \overline{e^2(t)} \Big/ \overline{x^2(t)}$$

where $\overline{e^2(t)}$ = the mean square value of $e(t)$, the difference between the outputs, and

$\overline{x^2(t)}$ = the mean square value of $x(t)$, the input to each channel.

In general, the error between the output of the two channels is caused by either independent internal noise in each of the channels or by differences in the channels caused by hardware mismatches.

Assume each channel has independent thermal noise with power σ_n^2 at the input. Then, for the case of equal input interference signals with power J on each channel and perfect channel matching, it can be shown that

$$CR = 2\left(\frac{\sigma_n^2}{J}\right). \qquad (6)$$

Hence, even in the case of perfect hardware matching, the front-end system noise along with the maximum expected interference power always provides a bound on the cancellation ratio *CR*.

The contributions of each critical component in the channel must be measured, or estimated from measurements or specifications. The sum of the contributions of the individual components represents the hardware limit to the cancellation ratio. The ultimate goal of the hardware design is to make the cancellation ratio that is due to the sum of all hardware contributions much less than the bound imposed by Eq. 6.

For an n-element array, it can be shown that the relationship between null depth *ND* and *CR* is given by

$$ND = \frac{CR}{2n}. \qquad (7)$$

Hence, in many applications the two-channel *CR* is sufficient to estimate the nulling capability of a full n-element array.

Four-Channel Test Bed

In the final 14-row adaptive antenna system, the output of each row is preamplified and fed through a multichannel rotary coupler to the 14 receivers. To minimize reflections over long cables, circulators are inserted in the receive paths. The final 1-MHz IF outputs of the receivers are filtered to a 200-kHz bandwidth and then sampled by 14-channel high-speed (4 MHz) 14-bit A/D converters. The conversion to baseband inphase and quadrature (I/Q) signals is performed digitally to avoid the imbalance of analog I/Q circuits [5]. The SMI weights are computed with a Givens transformation [6] in a special-purpose linear systolic array.

Figure 11 shows a simplified block diagram of the four-channel adaptive-nulling test bed that was designed and built by Lincoln Laboratory to measure the performance of the adaptive-array hardware configuration described above. The test bed also supported the development of mathematical relationships for predicting null-

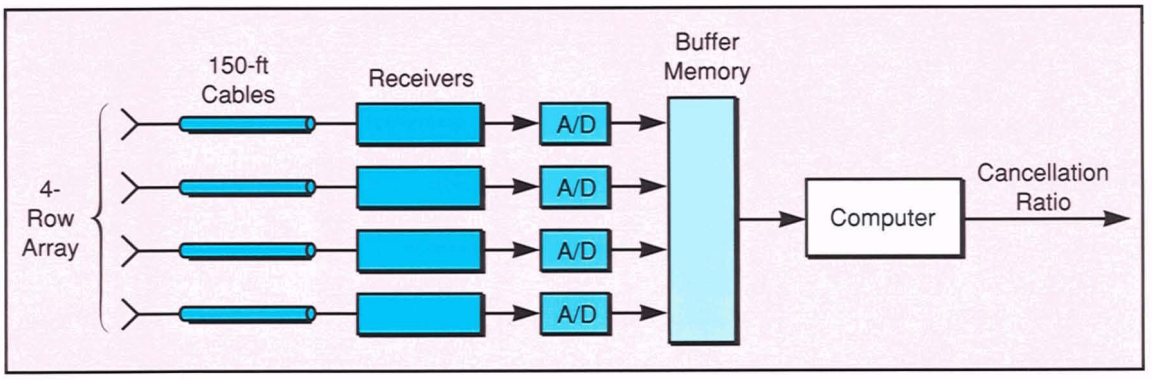

Fig. 11—Simplified block diagram of the four-channel test bed.

depth limitations. The following subsection presents those relationships.

The test bed was implemented with the initial four rows of the ultralow-sidelobe antenna, receivers composed of off-the-shelf components, 12-bit instrumentation A/D converters (14-bit converters were used in the final system), and a commercially available minicomputer for non-real-time processing. The receivers provided three-stage down-conversion from UHF to the final 1-MHz IF. Each receiver included a 30-MHz bandpass filter in the UHF front end, a 3-MHz Chebyshev filter in the first IF, and a 200-kHz Gaussian filter in the second IF. Test signals could be injected into the system at various calibration points. Table 1 lists some of the key parameters of the test bed hardware.

Bounds on Cancellation Ratio Due to Hardware Limitations

The following subsections review the mathematical relationships that describe the limitations imposed by the adaptive-array hardware.

Table 1. Four-Channel Test-Bed Parameters

Component	*Specification*
Antenna	Full row (17-dB gain)
RF input	435 MHz (30-MHz bandwidth)
Final IF	1-MHz frequency (200-kHz bandwidth)
A/D converter	12 bits (4-MHz sample rate)

Digital subsystem limitations. The fundamental limitation on system performance is the transfer-function mismatch between the multiple receive channels. Because digital equalization techniques can match the transfer functions to any desired level, the ultimate limit on performance is established by the accuracy of the digital system. This accuracy is, in practice, set by the performance of the A/D converters. In the case of perfect A/D converters with very large word sizes and high sampling rates, it is possible, in principle, to match the channels perfectly, given ample digital processing precision and an equalization filter of sufficiently high order. We can calculate the limitation imposed by a practical A/D converter from the effective number of bits (n_e) of the device. The value of n_e, which is slightly less than the actual number of bits, is a performance measure that combines the effects of quantization error as well as errors due to nonlinearities, timing jitter, and noise. For the case of digital I/Q conversion [5], the two-channel equalized bound on the cancellation ratio for a full-scale sine wave due to the A/D converter is

$$CR = \frac{2^{-2(n_e-1)}}{12}. \qquad (8)$$

(Equation 8 assumes a digital-filter bandwidth that is matched to the I/Q sampling rate.)

Figure 12 shows a typical match between two analog channels. As can be seen from the figure, the worst-case analog-channel matching of the IF band is typically on the order of –30 dB, with typical maximum-to-minimum spacing of 50 kHz. Digital-channel equalization can

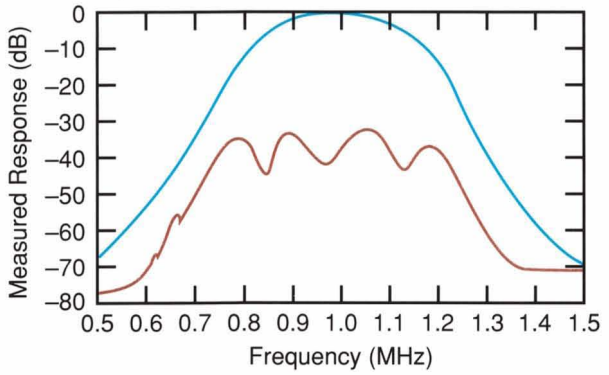

Fig. 12—Typical match between two channels. The upper curve shows a typical channel response dominated by the receiver's 200-kHz filter; the lower curve shows the difference between two typical analog channels.

correct this type of mismatch between analog channels.

Figure 13 shows the experimental measurements of *CR* as a function of the number of taps in the digital equalizer. As shown in the figure, the A/D converter sets a *CR* floor of –55 dB that can be achieved with the four-channel test bed. The test signal that was canceled in this experiment was a Gaussian-noise signal 13 dB below full-scale (dBFS) power. (This level was only 10 dB below a full-scale sine wave since the power of a full-scale sine wave is one-half that of a full-scale DC input.) Using experimental data similar to that shown in Fig. 13, we arrived at an empirical relationship between average cancellation ratio out (CR_o) and average cancellation ratio in (CR_i):

$$10 \log_{10} \frac{CR_i}{CR_o} = 30 N_t \left(\frac{f_p}{f_s} \right)$$

where f_s = I/Q sampling frequency (equal to 1 MHz),
f_p = minimum spacing between minimum and maximum on analog error curve (MHz), and
N_t = equalization filter order.

From the results shown in Fig. 13 and the configuration of off-the-shelf transversal-filter ICs, we decided that 32 taps should be used in the equalizers for the experimental radar. Consequently, all the tests that followed used 32 taps. For this number of taps, a rough empirical estimate for the *CR* limit due to imperfect equalization is –76 dB. The limit is based on extrapolation of the Fig. 13 data beyond the A/D converter floor.

The value of *CR* is also limited by the total number of bits in the A/D converter. For Gaussian-noise jammers, unless the input average power is substantially below full scale, severe degradation in *CR* performance can occur from the clipping of the signal peaks. The maximum CR_o is bounded by

$$\frac{CR_o}{CR_i} = \frac{1}{\sigma} \sqrt{\frac{2}{\pi}} \exp\left(-\frac{1}{2\sigma^2} \right) \qquad (9)$$

where σ is defined by $20 \log_{10}(1/\sigma)$ = the power level below full scale (dBFS) at which the average input power is set, and the value CR_i is the analog cancellation ratio shown in Fig. 10.

Figure 14 shows the effect of signal level on *CR*. For constant-envelope signals, the *CR* performance essentially follows the A/D-converter limit. For Gaussian noise, however, clipping degrades performance unless the noise is less than 13 dBFS. Figure 14 also shows the dynamic-range limit, which is given by Eq. 9, for $CR_i = 0.01$ (i.e., –20 dB). In practice, the thermal noise at the A/D-converter input is typically set approximately 2 dB above $q^2/12$, where q is the size of the quantization step. With this rule of thumb and the 13-dBFS requirement, the bound on *CR* due to the total number of bits in the A/D converter (n_b) is

$$CR = 2^{-2(n_b - 4)}.$$

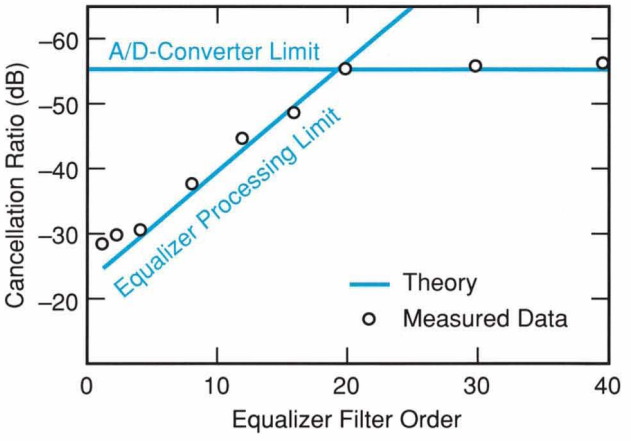

Fig. 13—CR as a function of the number of taps in the digital equalizer.

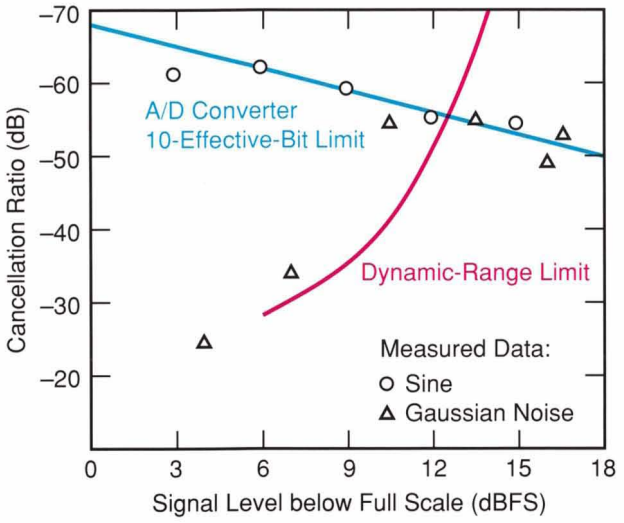

Fig. 14—Effect of signal level on cancellation ratio for a constant-envelope signal (a sine wave) and for Gaussian noise.

Antenna and RF subsystem limitations. In theory, the nulling process compensates perfectly for a fixed amplitude and phase difference between two channels but does not compensate for linear (or higher order) amplitude and phase variations across the IF bandwidth (W). Such variations can be caused by differential delay between the channels, RF-filter mismatch, or poor impedance matching, which can combine with high mutual coupling at the antenna end of the cables. Because it is desirable to transmit and receive sequentially at a number of RF frequencies without having to reequalize the receiver channels, front-end mismatches of these types must be minimized.

Figure 15 shows the typical response (top curve) of a 30-MHz front-end RF filter. The bottom curve shows the analog CR measurement result for a typical pair of RF filters, i.e., the cancellation obtained when the analog outputs of the two filters are subtracted after the same signal is applied to each filter input. This worst-case cancellation ratio (CR_{ip}) is on the order of –20 dB, and the spacing from valley to peak is typically 20 MHz (i.e., f_p = 20 MHz). The beamforming calibration procedure and/or the adaptive nulling operation must compensate for this frequency-dependent variation in analog-RF-filter matching.

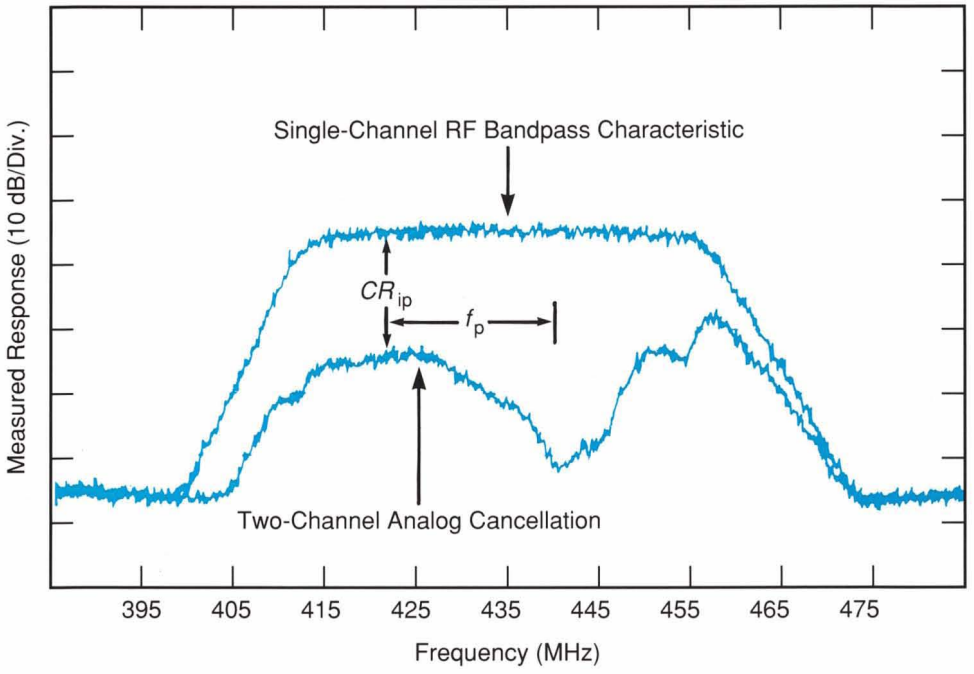

Fig. 15—Typical response (top curve) of a 30-MHz front-end RF filter and the two-channel analog cancellation (bottom curve) for a typical pair of RF filters. The cancellation curve represents the difference between the analog outputs of two filters with the same signal applied to each filter input. Note that CR_{ip}*, the worst-case cancellation, is about –20 dB, and* f_p*, the typical spacing from valley to peak, is about 20 MHz.*

For the case of RF-filter mismatch, it is straightforward to show that for Gaussian IF filters of bandwidth W, the value CR_o can be given by the following equation:

$$CR_o = \frac{CR_{ip}}{4 \ln 2}\left(\frac{W}{f_p}\right)^2. \qquad (10)$$

For the experimental radar filters, $CR_{ip} = 0.01$ and $f_p = 20$ MHz. Substituting these values into Eq. 10 yields a two-channel CR bound of –64 dB.

As previously described, impedance mismatches at either (or both) ends of a delay element can also introduce subtle channel mismatches. The cables that connect the antenna elements to the receivers are an example of such a delay element. For a system with a single impedance-mismatched cable with reflections between components separated by a delay, the two-channel CR can be shown to be

$$CR = |L|^2|\Gamma_A|^2|\Gamma_R|^2 \frac{\pi^2}{2 \ln 2}[W(2\tau)]^2 \qquad (11)$$

where L is the cable one-way loss, Γ_A is the voltage reflection coefficient at the antenna end, Γ_R is the voltage reflection coefficient at the receiver end, τ is the one-way cable delay, and W is the IF bandwidth.

The cables behind each antenna element in the experimental system are 150 ft long. To minimize reflections, circulators have been provided at each end of the cables. The measured circulator mismatch corresponds to a randomly phased return loss of –25 dB. The 150-ft cables have a one-way loss of 2 dB and a one-way delay of 0.17 μsec. For two randomly mismatched channels, Eq. 11 predicts a CR limit due to the mismatch reflections of –66 dB.

Similar subtle mismatch effects may be introduced by other reflections from impedance mismatches over paths that include the mutual coupling between antenna rows. An example of such a path is a signal that is reflected at a corporate-feed output and introduced into other channels via mutual coupling. For the effect of mutual coupling through adjacent antenna elements, CR can be shown to be

$$CR = |L|^2|M|^2|\Gamma_R|^2 \frac{\pi^2}{2 \ln 2}[W(2\tau)]^2. \qquad (12)$$

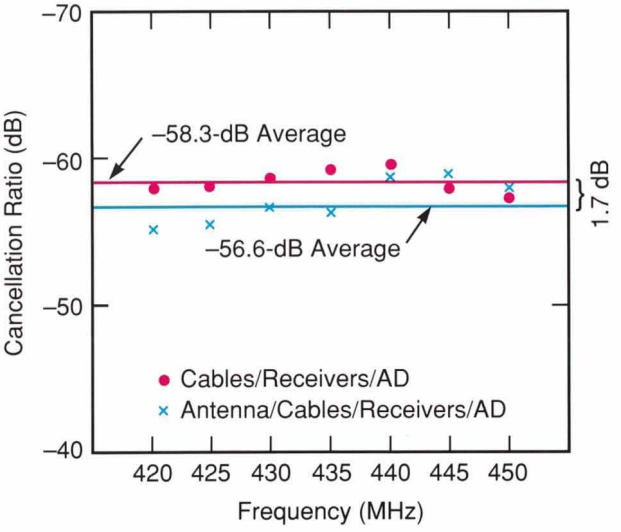

Fig. 16—Cancellation ratio for two complete channels as a function of frequency. The data include the effects of all hardware from the antenna elements to the output of the digital signal processor. For reference, the results for two channels without the effects of the antenna elements are shown.

Equation 12 is identical to Eq. 11, except that M, the interelement mutual-coupling coefficient, has replaced Γ_A.

For the experimental antenna, the mutual-coupling coefficient between adjacent elements was measured to be –12 dB, and the reflection coefficient was measured to be –25 dB. From these values, Eq. 12 predicts a CR limit of –68 dB due to the combined effects of element mutual coupling along with an impedance mismatch. (The value of –68 dB assumes a 32-ft corporate-feed length with no loss). The overall CR limit due to mutual coupling with impedance mismatches is actually approximately 6 dB higher. Of the 6-dB figure, 3 dB results from the fact that the expression given in Eq. 12 assumes that only one of the two channels in the CR test is corrupted. The remaining 3 dB must be added to account for mutual coupling from other near neighbors if the two elements under test are part of a larger array. Therefore, the overall CR limit due to mutual coupling with impedance mismatches is approximately –62 dB.

Far-Field Digital Adaptive Nulling

Figure 16 shows measured CR data for two complete channels. The data include the ef-

fects of all hardware, from the antenna elements to the output of the digital signal processor. For reference, the figure also shows the results for two channels without the antenna elements. Twelve-bit A/D converters were used to obtain the experimental measurements. To avoid A/D clipping of interference noise peaks, the rms level of a Gaussian input signal must be set well below the full-scale A/D-converter input level. In order to test the experimental system, high-input interference signals from a non-Gaussian envelope-limited noise source were used. The experimental interference signal was a 400-kHz chip-rate maximal-length pseudo-noise (P/N) code that modulated the RF carrier in a biphase manner. This envelope-limited signal allowed the A/D converters to run very close to full scale, which thereby extended the A/D-induced limit on *CR* to –61 dB. The limit imposed by the RF-filter mismatch for this experiment was –64 dB, the cable impedance mismatch limit was –66 dB, and the antenna impedance with mutual-coupling limit was –62 dB. If we assume that the error introduced by each of the error sources adds noncoherently, then the combined effect of each of the sources would produce an overall limit of –57 dB, which is close to the measured value. The difference between the two measured curves in Fig. 16 shows the effect of antenna mutual coupling.

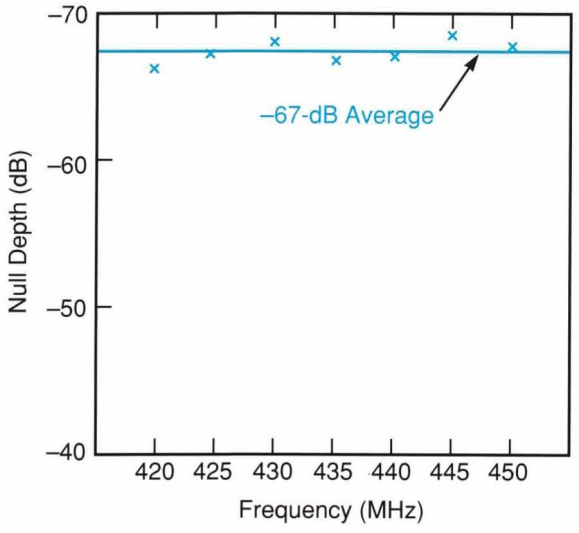

Fig. 17—Experimentally measured null depth as a function of frequency for a four-row array.

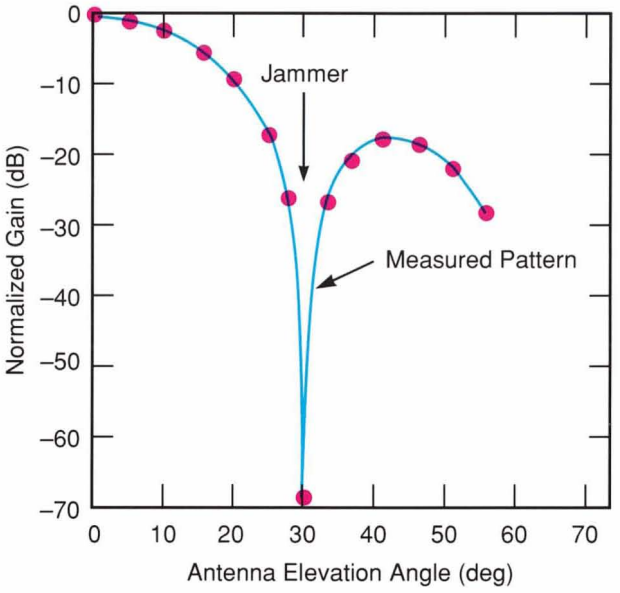

Fig. 18—Experimentally measured elevation pattern for the four-row array antenna at 445 MHz. Note that a null depth of –68 dB was obtained.

The measurements of two-channel antenna mutual coupling described earlier showed a two-channel *CR* of almost –57 dB. From this *CR* value, Eq. 7 predicts a four-channel null depth of –66 dB. Figure 17 shows the results from an experiment that measured the null depth for the four-row array as a function of RF frequency. During the experiment, the channels were equalized at 445 MHz, and the RF frequency was then varied. Across the UHF band, the data show an average null depth of –67 dB, which is very close to the value predicted by Eq. 7.

Figure 18 shows a pattern that was taken at the Westinghouse far-field test range at a frequency of 445 MHz with 400-kHz P/N. The pattern was obtained by first measuring the received power from the far-field noise jammer through channels that were previously digitally equalized. Next, digital antenna-element weights were computed to form both a beam perpendicular to the array and a null in the direction of the jammer. The antenna was then mechanically rotated in elevation and the received elevation power was measured with the previously computed weights. Note that a null

depth of close to –70 dB was obtained.

Summary

This article has described a 14-row array-antenna system that provides a high level of spatial discrimination against interference. In the azimuth dimension, rejection is accomplished with array rows that use precision, passive, fixed feed networks to produce ultralow-sidelobe patterns. The sidelobe requirements are low with respect to both the isotropic level and the peak level of the pattern. To our knowledge, the antenna has achieved sidelobe levels unmatched for a fixed passive array of this type. Excellent performance is achieved in both the copolarized and the cross-polarized patterns. The array has been designed to include high-power handling, weight and wind-load reduction, and the use of manufacturing techniques suitable for production. In the elevation dimension, interference and jammer rejection is achieved with digital adaptive nulling techniques. Significant improvements in the understanding of digital adaptive nulling have been described, including measurements of the hardware limitations on nulling performance. Measurements of the azimuth and four-channel adaptive elevation beams show that very high levels of rejection have been achieved. The antenna system's ability to reject interference dramatically reduces the potential impact of radar sidelobe interference.

Acknowledgments

The support provided by James Hall of the Office of Naval Technology is gratefully acknowledged.

This article summarizes the results of efforts undertaken by a number of people at both Lincoln Laboratory and Westinghouse Electric Corp. It would be impossible to list all of those who participated; however, we wish to name a few people whose help was instrumental to the success of our efforts. The technical efforts at Westinghouse were led by Peter Hrycak and Tim Waterman. The experimental testing of the antenna and 4-channel nulling test bed at Lincoln Laboratory would not have been possible without the help of Robert Baker, Gene Baszkiewicz, James Bertram, Richard Dec, Thomas Emberly, Hugh Maguire, David McQueen, Frank Pichler, Ken Teitelbaum, and Tim Wallace.

References

1. J. Ruze, "The Effect of Aperture Errors on the Antenna Radiation Pattern," *Supplemento al Nuovo Cimento* **9,** 364 (1952).
2. R.T. Compton Jr., *Adaptive Antennas, Concepts and Performance* (Prentice Hall, Englewood Cliffs, NJ, 1988).
3. I.S. Reed, J.D. Mallett, and L.E. Brennan, "Rapid Convergence Rate in Adaptive Arrays," *IEEE Trans. on Aerospace and Electronic Systems* **AES-10,** 853 (1974).
4. B.D. Carlson, "Covariance Matrix Estimation Errors and Diagonal Loading in Adaptive Arrays," *IEEE Trans. on Aerospace and Electronic Systems* **AES-24,** 397 (1988).
5. C.M. Rader, "A Simple Method for Sampling In-Phase and Quadrature Components," *IEEE Trans. on Aerospace and Electronic Systems* **AES-20,** 821 (1984).
6. G.H. Colub, C.F. Van Loan, *Matrix Computations*, 2nd ed. (Johns Hopkins University Press, Baltimore, 1989).

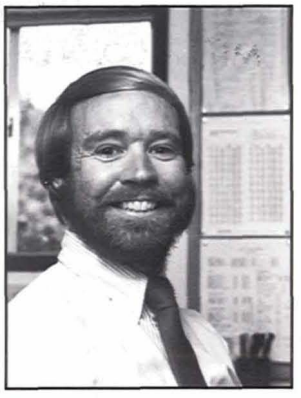

BLAIR D. CARLSON is a staff member in the Radar Systems Group. He received a B.S. in engineering science from the Pennsylvania State University in 1976 and an M.E. in electrical engineering from the University of Virginia in 1978. Before joining Lincoln Laboratory in 1982 he worked in the area of controls and dynamic systems. His current research is in the area of radar and adaptive antenna systems.

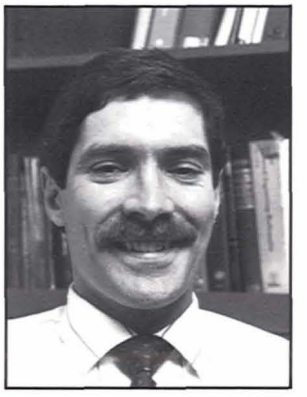

MATTHEW W. GANZ is an assistant group leader in the Radar Systems Group. He received a B.S., an M.S., and a Ph.D. in electrical engineering from Ohio State University in 1981, 1982, and 1986, respectively. Prior to joining Lincoln Laboratory in 1986 he worked on satellite communication and tracking systems at the Johns Hopkins University Applied Physics Laboratory. His current research interests are in radar applications of adaptive arrays.

LEONARD M. GOODMAN is a senior staff member in the Radar Systems Group. He received a B.S., an M.S., and an E.E. in electrical engineering from MIT. Leonard is an IEEE member and a past chairman of the Boston section.

LEE O. UPTON received a B.S.E.E. from Tufts University in 1965 and an M.S.E.E. from the University of Pennsylvania in 1967. Before joining Lincoln Laboratory in 1978, he worked for RCA in Moorestown, N.J. At Lincoln Laboratory, Lee was the Section Leader of the Millimeter Wave Radar for the Kiernan Reentry Measurement System (KREMS) at Kwajalein Atoll in the Marshall Islands, and Group Leader of the Radar Systems Group. On leave of absence from Lincoln Laboratory, he is currently a program manager at the Defense Advanced Research Projects Agency.

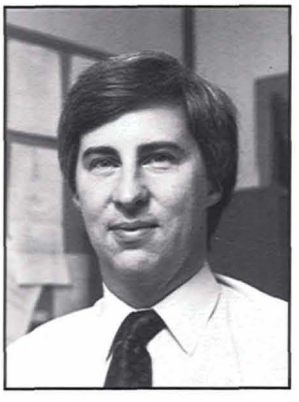

JOHN AUSTIN received a B.Sc. from the University College of North Wales in 1973, and a Ph.D. from University College London in 1978. He worked at Plessey Radar Ltd. and University College London before joining Lincoln Laboratory in 1979. John is currently Group Leader of the Radar Systems Group.