

A Maximum-Likelihood Beamspace Processor for Improved Search and Track

Richard M. Davis, *Senior Member, IEEE*, and Ronald L. Fante, *Fellow, IEEE*

Abstract—The authors describe a two-stage digital beamforming architecture for use in radars. The first stage uses digitized subarray voltages to form a number of highly overlapped adaptive beams. The beams are themselves adaptively weighted during the second stage. Maximum-likelihood processing is used in the second stage to estimate target amplitude and angle of arrival. It is shown that the new architecture supports improved search and track performance compared to traditional systems that use a single receive beam and monopulse processing to measure angle.

Index Terms—Adaptive array, beamformer, maximum-likelihood, radar.

NOMENCLATURE

ADC	Analog-to-digital converter.
AOA	Angle of arrival.
DBF	Digital beamforming.
INR	Interference-to-noise ratio.
ML	Maximum-likelihood.
MLBP	Maximum-likelihood beamspace processor.
PC	Pulse compression.
PRF	Pulse repetition frequency.
SIR	Signal-to-interference ratio.

I. INTRODUCTION

MODERN radars utilize phased-array antennas and analog beamformers. The search function is performed using a single beam to interrogate a grid of spatial positions. If a target is detected, its angle of arrival (AOA) is measured using two orthogonal difference beams. Simultaneous ratioing of the difference-beam voltages to the sum-beam voltage provides an instantaneous estimate of target location. The process is referred to as monopulse.

The technology currently exists to do digital beamforming (DBF). Direct RF digitization at the element level is possible in low-frequency systems. At higher frequencies analog beamforming can be used to form subarrays whose outputs can be downconverted and digitized. The authors believe that with DBF comes the possibility of using multiple highly overlapped receive beams together with maximum-likelihood (ML) processing to improve both search and track performance. We have developed a maximum-likelihood beamspace processor (MLBP) that implements the new approach and is the subject of this paper. The processor and its advantages are described

in Section II. The processor is analyzed in Section III and numerical results are presented in Section IV. The work is summarized in Section V.

II. MAXIMUM-LIKELIHOOD BEAMSPACE PROCESSOR (MLBP)

The MLBP architecture (Fig. 1) entails dividing the antenna aperture, which may consist of a large array containing thousands of elements, into a small number of subarrays (M), and digitizing their outputs. A typical value of M would be 16. The digitized signals are then passed through two stages of processing. In the first stage, the subarray voltages are multiplied by N sets of complex weights and summed to form N highly overlapped beams. A typical value of N would be four. The centers of the beams are all located within the 3 dB contour of the transmit beam. Their exact location is determined *a priori*. Interference is removed from each of the beams by choosing the subarray weights to minimize the output mean square power subject to a beam-pointing constraint, or maximize the output signal-to-interference ratio (SIR). Our procedure is a variation of an approach discussed by Baranowski and Ward [1]. In particular, our first stage beams are formed adaptively and take the form of highly overlapped, fixed sum beams. Our choice of architecture was driven by our desire to both remove interference and eliminate beam shape loss. The first level of processing is equivalent to what has historically been referred to as a fully-adaptive array of subarrays. Applebaum was one of the first to demonstrate the virtue of the adaptive array architecture [2]. Many techniques exist for choosing the adaptive weights used to form the beams. The weights are chosen in such a way as to constrain the gain in the pointing direction of each of the N beams. The constraint preserves the gain in the pointing direction and reduces beam shape distortion. Pulse compression (PC) is done after the first stage processing to prevent the first stage adaptive arrays from attempting to cancel target returns. Target (often referred to as signal) cancellation can be prevented if the signal to noise ratio in each of the beams is much less than unity prior to adaptive nulling (see Appendix A for additional discussion). Performing pulse compression at the beam level also reduces computational complexity (only four channels must be compressed versus 16 had we done it at the subarray level).

During the second stage of processing, the N beams are themselves adaptively weighted and combined to form a single output beam. The beam weights in the second stage are chosen to maximize the probability of detecting a target return. The processing used to choose the second set of weights is referred to as ML. Kelly, Reed, and Root were among the first to demonstrate the virtues of ML. The ML process entails performing a search over all possible AOA directions of the

Manuscript received November 17, 1999; revised August 28, 2000.

The authors are with The MITRE Corporation, Bedford, MA 01730-1420 USA.

Publisher Item Identifier S 0018-926X(01)03171-4.

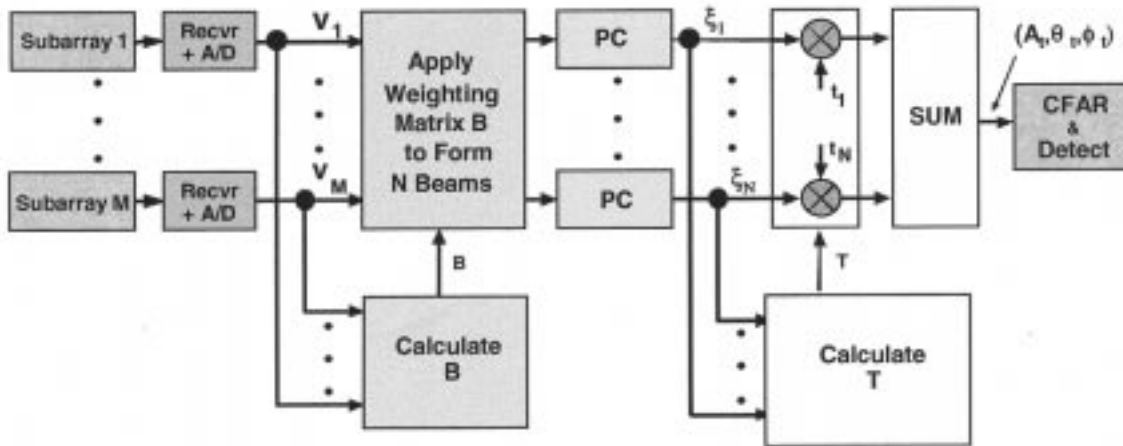


Fig. 1. MLBP.

target return and picking the direction that yields a set of beam weights that produces the highest probability of declaring a target to be present. The search can be constrained to a region slightly larger than the 3 dB contour of the transmit beam. An independent search must be performed in every range cell within the range sweep. Thus, there will be a separate set of optimum beam weights for each range cell. We assume, herein, that the target geometry and range resolution of the radar is such that there will only be one target per cell.

The authors will show that under the assumption of Gaussian interference, the second stage processing is equivalent to operating a fully adaptive array of beams where the weights are chosen to maximize SIR in the composite output beam. The observation emphasizes the commonality of the ML and adaptive array processing paradigms.

Two impediments to performing the maximum-likelihood search in the past have been the computational burden and the possibility of getting stuck in local minima. Advances in high speed computing, together with the idea of limiting the maximum-likelihood search to only a few beams, make the computation problem tractable. Performing the search in beam space with highly overlapped beams also appears to reduce the probability of getting stuck in local minima. The MLBP architecture offers many advantages, notably:

- 1) The ML search eliminates the need to form azimuth and elevation difference beams to measure target AOA, thereby, reducing the required RF beamforming, receiver, and analog-to-digital conversion hardware by up to a factor of three.
- 2) Using multiple highly overlapped beams on receive, instead of a single beam in conjunction with ML, nearly eliminates beam shape loss on receive and increases the volume of space, which can be searched for a given number of transmissions. The MLBP can search a given volume of space using approximately 40% fewer pulses compared to a traditional single-beam system. This provides the radar with additional time and energy to perform other functions.
- 3) The MLBP architecture supports improved angle accuracy compared to monopulse. The improvement varies

from a few percent when the target is at the peak of the transmit beam to approximately 30% when the target is at the 3-dB (half power) point on the transmit beam for a system using four highly overlapped beams.

- 4) The MLBP supports angle estimation over a larger angular volume than monopulse. Monopulse systems can usually only be relied upon when the target is within the 3-dB contour. The MLBP, however, degrades gracefully as the target moves outside the 3-dB contour and down the skirt.
- 5) The fully adaptive array of subarrays followed by the fully adaptive array of beams supports mainlobe and sidelobe interference suppression while minimizing beam shape distortion and signal cancellation.

A few comments regarding the first advantage are in order. The antenna designer is faced with a difficult decision if he wants to do digital beamforming and monopulse processing. The problem is how to simultaneously apply aperture weighting (tapering) to sum and difference channels to achieve low sidelobes. One option is to form a large number of small subarrays, digitize their output voltages, and then split each one three ways and apply the appropriate weighting (Taylor or Bayliss) digitally. The large number of small subarrays is required to appropriately sample the tapering function. The first option requires a large number of receivers and analog-to-digital converters (ADCs). The second option is to split the output of each individual element three ways, apply tapering near the array face, and duplicate the subarray beamforming three times. The second option supports digitization of a smaller number of larger subarrays, but trades the smaller number of receivers and ADCs for additional RF beamforming stages (to form the larger subarrays), weighting circuitry, and possibly analog time delay units to support wideband modes. A small number of large subarrays can be used since there is no sampling problem associated with applying the taper weights. The MLBP supports the best of both options. Since we don't form difference beams, we can apply taper weights to the elements within each subarray and use a small number of receivers and ADCs by digitizing a small number of large subarrays. Although we do not consider wideband operation in this paper, we expect that

fewer time delay steering units may be required with the MLBP compared to the second option because the multiple overlapped beams cover a larger volume of space than a single beam.

III. ANALYSIS

The MLBP depicted in Fig. 1 uses two stages of sequential processing on receive. In the first stage, the subarray voltages are adaptively weighted and summed to form N highly overlapped beams. The weights are chosen to maximize SIR in each beam separately. The second stage entails adaptively weighting and summing the beam voltages to form a single composite output beam. A search process is used to find the adaptive beam weights that maximize the probability of a target being present in the output beam. Interference that is not cancelled during the first stage of processing, but is correlated beam to beam, can be nulled in the second stage. In the analysis to follow the narrow-band approximation (signal bandwidth \ll carrier frequency) is implicit.

A. The First Stage

Suppose we are given the voltages out of M subarrays and we wish to form N interference-free beams centered on the directions $(\theta_1, \phi_1), (\theta_2, \phi_2), \dots, (\theta_N, \phi_N)$, in spherical coordinate system, where θ is the polar angle and ϕ is the azimuthal angle. If $v(k) = [v_1(k) \ v_2(k) \ \dots \ v_M(k)]^T$ is a vector of the voltages at the outputs of the subarrays at discrete time k , we wish to multiply this set of voltages by an $N \times M$ beamforming matrix

$$B = \begin{bmatrix} w_1^H \\ w_2^H \\ \vdots \\ w_N^H \end{bmatrix} \quad (1)$$

that will remove interference while simultaneously preserving the gain in the N beam pointing directions. It is well known [2], [6] that the weight vector $w_n = [w_{n1} \ \dots \ w_{nM}]^T$ given by

$$w_n = \frac{R^{-1} S_n}{S_n^H R^{-1} S_n} \quad (2)$$

maximizes SIR in the beam pointing direction (θ_n, ϕ_n) , thus, producing the desired result. In (2), R is the $M \times M$ covariance matrix of subarray voltages defined as $R = \langle v v^H \rangle$, where $\langle \cdot \rangle$ denotes expectation. R can be approximated as

$$R = \frac{1}{K} \sum_{k=1}^K v(k) v^H(k) \quad (3)$$

where $K \geq 4M$ and the superscript H denotes conjugate transpose [7]. Also, $S_n = [s_{n1}, s_{n2}, \dots, s_{nM}]^T$ is a $M \times 1$ vector that defines the steering direction (θ_n, ϕ_n) of the n th beam. The numerator of the weight equation is often referred to as a Wiener whitening filter, while the denominator implements a point constraint by scaling the adaptive pattern to provide the desired gain in the beam steering direction. A typical component of S_n is

$$s_{nm} = g_m(\theta_n, \phi_n) \exp[jk(x_m \sin \theta_n \cos \phi_n + y_m \sin \theta_n \sin \phi_n + z_m \cos \theta_n)] \quad (4)$$

where $g_m(\theta_n, \phi_n)$ is the voltage gain of the m th subarray in the direction (θ_n, ϕ_n) , k equals $2\pi/\lambda$, λ is the wavelength, and (x_m, y_m, z_m) is the location of the center of the m th subarray in a Cartesian coordinate system. The voltages in each of the output beams (see Fig. 1) from the first stage of processing are denoted by $\xi_1, \xi_2, \dots, \xi_N$. If we define $\xi = [\xi_1 \ \xi_2 \ \dots \ \xi_N]^T$, then the first stage beamformer output voltage vector is

$$\xi = Bv. \quad (5)$$

B. The Second Stage

Once the beams have been formed and most of the interference removed, PC is performed on the beam voltages. Since there will be fewer beams than subarrays, performing PC at the beam level is computationally efficient. As previously noted, performing PC at the beam level, rather than the subarray level, also ensures that the first stage adaptive array will not attempt to cancel target returns. Prior to pulse compression, target returns in the subarrays will be well below thermal noise and the adaptive array will, to first order, ignore them (see Appendix A). There may be some applications involving high pulse repetition frequency (PRF) radars where the presence of clutter will force PC to be performed at the subarray level. If a clutter-free region does not exist, then PC followed by Doppler processing may have to be performed at the subarray level in order to generate clutter-free Doppler cells to sample the interference.

Although, the amplitude and direction of a target return is unknown, we can postulate a unit-amplitude target in range bin n in direction (θ, ϕ) . This produces an $M \times 1$ voltage vector $h = [h_1 \ h_2 \ \dots \ h_M]^T$ at the output of the subarrays, where

$$h_m = g_m(\theta, \phi) \exp[jk(x_m \sin \theta \cos \phi + y_m \sin \theta \sin \phi + z_m \cos \theta)]. \quad (6)$$

This hypothesis also produces a voltage vector $G = Bh$ in beamspace. If the target is nonfluctuating¹ and has unknown amplitude and phase denoted by the complex quantity A , the voltage vector in beamspace produced by the postulated target is

$$\xi_t = ABh = AG. \quad (7)$$

If we let the vector ρ denote the component of ξ containing uncancelled interference and thermal noise in the second stage beams, then the first stage output vector ξ can be written as

$$\xi = \rho + \xi_t = \rho + AG. \quad (8)$$

Assuming that ρ is a circular, normal random vector, it can be shown [5], [9] that the probability density function for ρ is

$$p(\rho) = \pi^{-N} [\det(C)]^{-1} \exp(-\rho^H C^{-1} \rho) \quad (9)$$

where C is the $N \times N$ covariance matrix of beam voltages defined by

$$C = \langle \rho \rho^H \rangle = BRB^H. \quad (10)$$

¹For a fluctuating target we must estimate A_1, A_2, \dots, A_K , in addition, to (θ, ϕ) , where A_k is the amplitude at time sample k . The probability function to be maximized becomes $p(\xi_1, \dots, \xi_K | A_1, \dots, A_K, \theta, \phi)$. The analysis is more difficult, but straightforward provided the time samples are sufficiently spaced to make the noise samples independent.

Solving (8) for ρ and substituting into (9) gives the conditional probability density

$$p(\xi|A, \theta_t, \phi_t) = \pi^{-N} [\det(C)]^{-1} \cdot \exp[-(\xi - AG)^H C^{-1} (\xi - AG)]. \quad (11)$$

The maximum-likelihood target parameter estimator $(\hat{A}, \hat{\theta}_t, \hat{\phi}_t)$ occurs when (11) is a maximum or when

$$\Gamma = (\xi - AG)^H C^{-1} (\xi - AG) \quad (12)$$

is a minimum. We can solve for the estimator \hat{A} , by differentiating Γ with respect to A and setting the result to zero. Performing the differentiation we find

$$\hat{A} = T^H \xi \quad (13)$$

where the beam weighting vector T is given by

$$T = \frac{C^{-1}G}{G^H C^{-1}G}. \quad (14)$$

We can show that the estimator given by (13) is unbiased by substituting (8) into (13) and taking the expected value. Since $\langle \rho \rangle = 0$, then $\langle \hat{A} \rangle = A$.

At this point in the analysis we make an important observation. T is functionally identical to the weight vector that maximizes SIR in the output beam of a fully adaptive array. Let us compare T in (14) with w_n in (2). In (14) the covariance matrix C of beam voltages replaces the covariance matrix R of subarray voltages in (2) and G , the steering vector that coheres the beam voltages to form a composite beam in the direction of the target AOA, replaces S_n , the steering vector that coheres the subarray voltages in the predetermined direction of the n th beam. The numerator of T is again a Weiner whitening filter where $C^{-1/2}\xi$ is the whitened data and $C^{-1/2}G$ is the whitened steering vector in beamspace. The denominator is a scalar that acts to constrain the gain of the composite beam in the target direction. Based upon this observation, we could stop the analysis at this point and simply search for the target AOA that maximizes SIR in the output beam. Instead, we will continue the analysis to more firmly establish the equivalency between adaptive array and maximum-likelihood theory.

Substituting (13) back into (12) gives

$$\Gamma = \xi^H C^{-1} \xi - (G^H C^{-1} G) |T^H \xi|^2. \quad (15)$$

Because $\xi^H C^{-1} \xi$, $G^H C^{-1} G$ and $|T^H \xi|^2$ are positive definite quantities, it is evident that Γ is minimized when $(G^H C^{-1} G) |T^H \xi|^2$ is maximized. That is, the estimate of target location is that direction (θ, ϕ) for which $(G^H C^{-1} G) |T^H \xi|^2$ is a maximum. Let us now discuss the physical meaning of this quantity.

Referring to Fig. 1 we see that the voltage in the composite output beam after weighting is $T^H \xi$. In a range bin where a target is present $T^H \xi$ yields an estimate of the target amplitude \hat{A} [from (13)]. However, in range bins where there is no target the mean output interference power is

$$I = \langle |T^H \rho|^2 \rangle = T C T^H = \frac{1}{G^H C^{-1} G}. \quad (16)$$

Taking the ratio of the power in the range bin with the target present to the power in a bin where there is no target yields an estimate of SIR in the output beam. This is

$$\text{SIR}_{\text{est}} = (G^H C^{-1} G) |T^H \xi|^2. \quad (17)$$

It should be noted that (17) is actually the ratio of signal-plus-interference power to interference power, but we identify it as SIR because we assume that the tandem adaptive arrays effectively null the interference, so that signal is much larger than the residual uncanceled interference in the output beam. We can now see that maximizing $(G^H C^{-1} G) |T^H \xi|^2$ is equivalent to searching over (θ, ϕ) to find that value $(\hat{\theta}, \hat{\phi})$ that maximizes SIR in the output beam. The value of SIR at $(\hat{\theta}, \hat{\phi})$ is readily calculated to be $\xi^H C^{-1} \xi$. After estimating the target AOA, $(\hat{\theta}, \hat{\phi})$, the target's amplitude and phase can be estimated using (13)

$$\hat{A} = T^H (\hat{\theta}, \hat{\phi}) \xi. \quad (18)$$

Although (17) looks mathematically formidable, we show in Appendix B that it reduces to a familiar result in the absence of any interferers.

C. Target Detection

Target detection is achieved by comparing the value of \hat{A} in each range cell with a threshold τ determined using conventional constant false alarm rate (CFAR) processing. If $\hat{A} \geq \tau$, a "target" is declared and if $\hat{A} < \tau$, "no target" is declared.

IV. NUMERICAL EXAMPLES

A MATLAB computer simulation was developed to verify the analysis and predict performance of the MLBP. A 1024 element square array divided into 16 square subarrays was used in the simulation. The elements were spaced $\lambda/2$ apart and the array was uniformly weighted on both transmit and receive. The phase shifters within each subarray were set to steer the subarrays to the pointing direction of the transmit beam. Target returns, interferers, and thermal noise were modeled as tones (zero bandwidth). Thermal noise voltages were complex with their real and imaginary components chosen from a unit variance, zero mean, Gaussian distribution. Target returns were assumed to be non-fluctuating. Angle errors were normalized to the 3 dB boresight beamwidth (θ_3) equal to 0.886 times λ/D , where the diameter of the array D was 16λ .

The number of beams in the second stage of the processor was varied from two to eight. The beams were offset from the subarray steering direction (u_0, t_0) by some fraction of θ_3 . The direction cosines with the x - and y - axis in the plane of the array are defined in Fig. 2. Fig. 2 depicts a typical beam geometry for the baseline case of four-beams. MATLAB function "FMINS" with an error bound of 10^{-4} was used to perform the search over all possible target AOA (θ, ϕ) . FMINS uses the Nelder-Mead simplex direct search method. Nelder-Mead is an excellent engineering tool, but is inefficient. The mean and variance of the final error were found to be relatively insensitive to both the starting point of the search and the error used to determine the stopping point. Since Nelder-Mead is not a gradient-based technique, it may be less likely to get stuck in a local minima than other approaches. We did not notice a problem with local

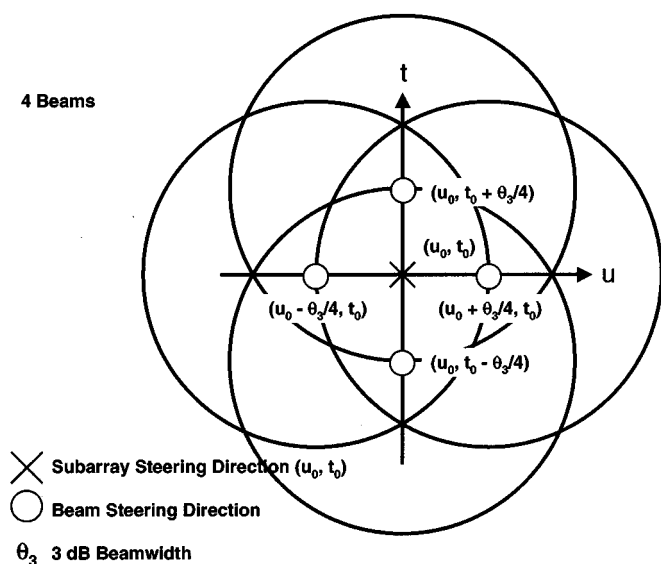


Fig. 2. Beam geometry in direction cosine space.

minima in any of the studied cases. Studied cases included one and two mainlobe interferers. We conjecture that the use of highly overlapped beams, each having a constraint in the pointing direction, limits beamshape distortion and the presence of multiple deep nulls.

The number of floating point operations per second (flops) required to implement the MLBP depends upon the algorithm used to search for the AOA. Even though Nelder–Mead is inefficient, we calculated the throughput using it to obtain a benchmark. The number of nonrecurring floating point operations per beam is due mainly to formation and inversion of the sample covariance matrix of subarray voltages and to beamforming (multiplying subarray voltages by the beam weights for each range cell). The recurring computation per beam is equal to the number of operations needed to calculate SIR [see (17)] for each hypothesized AOA, times the number of hypothesized AOAs per range cell, times the number of range cells. Assuming 500 range cells, 16 subarrays, and 4 beams, it was found that 37 hypothesized AOAs at 1.7 K flops per iteration were typically needed. The total nonrecurring flop count was roughly 200 K flops per beam compared to 31.5 megaflops per beam for the recurring. If the radar puts out 100 beams per second the total throughput (100 times the sum of nonrecurring plus recurring flops) would be 3.1 gigaflops. Note that the flop count is dominated by the recurring count and other search algorithms exist that are more efficient than Nelder–Mead (see, for example, [10]).

A. Improved Clear Environment Search Performance

Fig. 3 shows computer generated comparisons of the two-way (transmit plus receive) loss in signal-to-noise ratio (SNR) between a conventional architecture that uses a single receive beam and MLBP architectures, which use three (right insert) and eight (left insert) receive beams. The figure was generated by moving a single target return out from the peak of

the transmit beam (also equal to the subarray steering direction (u_0, t_0)) along the t -axis ($ut = utarget = u_0$). No interference except receiver noise was present. The curves represent an average of 200 Monte Carlo runs using different receiver noise realizations on the subarray voltages. The figure shows that if search beams are laid down in a line and overlapped at the 5-dB (two-way) loss point, the coverage in one- dimension can be increased by about 29%–35% as the number of receive beams is increased from three to eight using the MLBP architecture compared to the conventional approach. The improvement is due to the elimination of receive beamshape loss. Fig. 4 shows beam geometry for the three and eight beam cases.

B. Improved Clear Environment Angle Estimation

Fig. 5 presents a comparison of the angle-estimation capability of a conventional-monopulse processor with that of the MLBP for the three (left insert) and eight (right insert) beam architectures. The figure was generated by moving a single target return out from the peak of the transmit beam (also equal to the subarray steering direction (u_0, t_0)) along the t -axis ($ut = utarget = u_0$). The transmit loss was taken into account in the processing and the curves represent an average of 200 Monte Carlo realizations with different receiver noises. The results are compared to a normalized monopulse error (σ) in a thermal noise environment. The one-sigma error was assumed to equal the 3-dB beamwidth divided by 1.6 times the square root of twice the rms SNR [11].

Fig. 5 shows that the performance of both the conventional (monopulse) processor and the MLBP approach the monopulse error for high SNR (target near subarray pointing direction). At low SNR (as the target moves off the peak and down the skirt of the transmit beam), the MLBP produces significantly smaller one-sigma angle errors than the monopulse processor. The improvement when the target is located at the 3-dB point, for example, varied from about 16% for the three-beam architecture to 33% for the eight-beam architecture. Using more than eight beams did not materially improve performance over the eight-beam system. In fact, a four-beam system did almost as well as the eight-beam system in the clear environment. A four-beam MLBP system was interfaced to a six-state Kalman filter tracking simulation and the resulting track errors were compared with a monopulse system that used the same Kalman-filter tracker [12]. The performance of the two systems (results not shown) was nearly the same for on-axis tracking, but the track errors of the MLBP system were nominally 50% lower than those of the monopulse system for off-axis tracking. In an operational radar, the beam will normally be pointed at one of the largest objects² in the beam and most targets will be tracked off-axis. It is also noteworthy that while monopulse systems can fail catastrophically as the target return approaches the sum beam null, the MLBP should degrade gracefully and angle measurements made outside the 3-dB transmit contour should be reliable. In summary, the improvement in tracking performance compared to monopulse is due mainly to the fact

²For example, when tracking a reentering missile complex, the radar might track the booster (tank) and the separating reentry vehicles would be tracked off axis.

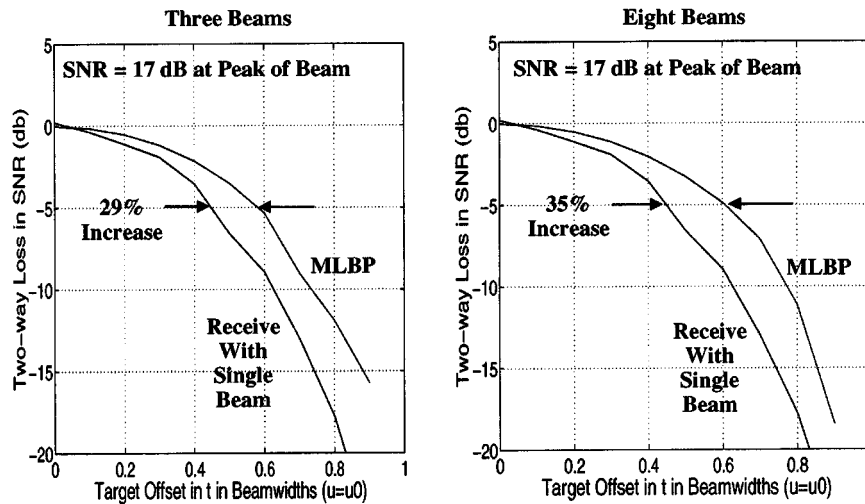


Fig. 3. Comparison of two-way beam shape loss between MLBP and conventional receiver.




Fig. 4. Three and eight beam geometry in direction cosine space.

that the ML search for the target AOA results in the composite output beam being pointed directly at each tracked target.

C. Performance Against a Single Mainlobe Interferer

Fig. 6 shows typical computer simulated performance of a four-beam MLBP operating against a single mainlobe interferer. A target was located a distance of $\theta_3/8$ from the subarray steering direction along a cut that made a 30° angle with the t -axis. A single mainlobe interferer was stepped from the subarray-steering direction along the cut, through the target's location, and out to a distance of one beamwidth. The target SNR and the INR (interferer-to-noise ratio) were set equal to 20 and 15 dB, respectively, when received at the peak of a beam pointed in the subarray steering direction. A nonfluctuating jammer waveform (constant amplitude and phase) was used when generating the figure. Although, the figure represents only a single realization of receiver noise, nearly identical SIR performance (not shown) was obtained using a complex stochastic waveform and averaging the signal and interference over 100 Monte Carlo runs. The left insert in Fig. 6 plots the signal, interference, and SIR in the output beam, while the right insert shows SIR in each of the four first stage beams. Note that the SIR in the composite beam in Fig. 6 is usually larger than

any of the SIRs in the four first stage beams. The observation emphasizes the fact that the second-stage processing does not merely pick the first-stage beam having the highest SIR ratio. The first-stage beams are adaptively weighted and summed to form a composite adapted output beam.

A significant improvement in SIR was observed in all computer runs involving a single interferer except when it was very close to the target return—and even then the degradation in SIR was always less than 2 dB of the value in the absence of any adaptivity (which was 5 dB in Fig. 6). The degradation observed for small target-interferer separations was able to be eliminated by adding more beams. One might have expected more degradation than observed when the interferer approached the target. We hypothesized that the gain constraints placed in the pointing direction of each of the first-stage beams were limiting their ability to null the interferer. This prompted us to examine more closely why the fully adaptive arrays of subarrays operating on the first-stage beams did not result in more interferer nulling and more nulling of the target return along with it.

Fig. 7 shows adapted antenna patterns for the first stage beam that is pointed to $(u0 + \theta_3/4, t0)$ as a single interferer located at u_i is moved toward the peak of the beam along the u -axis. The location of the gain constraint is denoted by u_c . The interference-to-noise ratio (INR) was set to 15 dB at the peak of

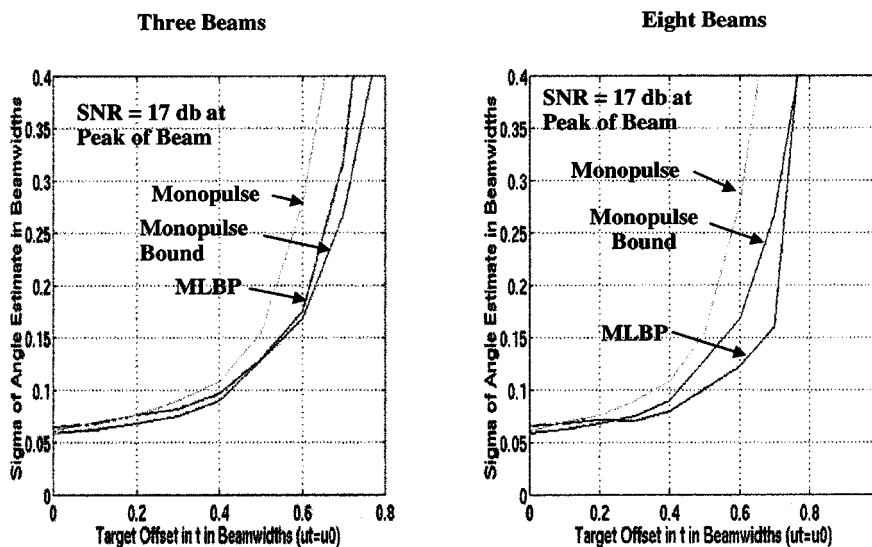


Fig. 5. Comparison of angle accuracy between MLBP and conventional monopulse receiver in clear environment.

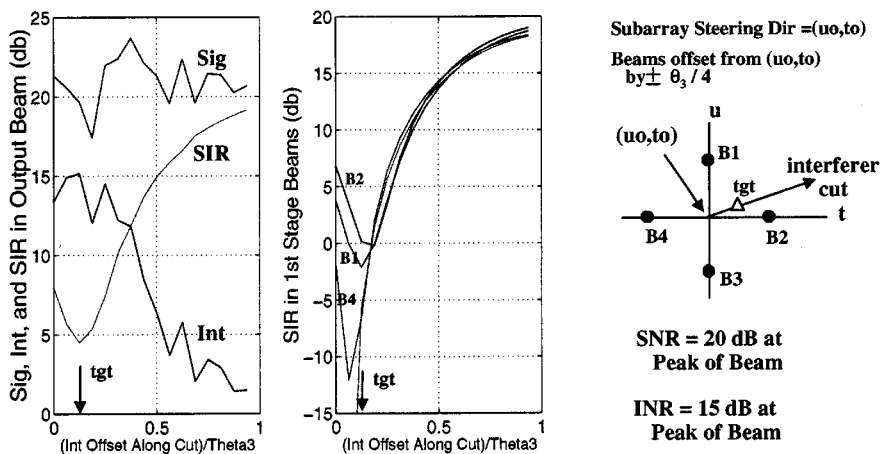


Fig. 6. Four-beam MLBP operation against a single-mainlobe interferer.

the beam and the fluctuation noise was removed from the interferer. Each insert represents a single realization (no Monte Carloing). The inserts demonstrate that when the interferer is brought closer than about $\theta_3/4$ to the constraint point, the array does not have enough resolution to maintain the constraint and place the null on the interferer. In fact, as the separation between the interferer and the constraint point is changed from $\theta_3/8$ to zero, the null is observed to move back out to where it was in the absence of the interferer. Thus, we find as anticipated, that the point constraints placed at the peak of the first-stage beams prevent significant interferer nulling over a relatively large area.

Although the patterns in Fig. 7 can be used to estimate the exact amount of interferer and target nulling, they cannot be used to determine the overall interference after nulling. The total interference is due to the sum of the uncanceled interferer and the thermal noise in the beam. As the interferer approaches the constraint point, the adaptive subarray weights used to form the first-stage beam become large. Multiplication of the weights by the thermal noise causes an increase in the noise figure in the beam. Fig. 8 plots the norm of the subarray weights (the thermal

noise) and the total interference in the first-stage beam versus separation of the interferer from the constraint point. Note that the weights increase until the separation reaches about $\theta_3/8$ and then decrease as the separation goes from about $\theta_3/8$ to zero. Increasing INR will allow the null to move slightly closer to the constraint point before reversing direction. The reversal point appears to be the point at which the norm of the weights (and the output noise figure) is a maximum.

V. SUMMARY

The authors have described a maximum-likelihood beamspace processor for use with an antenna architecture that performs digitization at the subarray level. The MLBP uses two stages of processing. The first stage entails operating N fully adaptive arrays of subarrays in parallel to form N highly overlapped beams. The second stage entails operating a single fully adaptive array of beams to form a single composite output beam pointed in the direction of the target return. The pointing direction of the output beam is determined by performing a

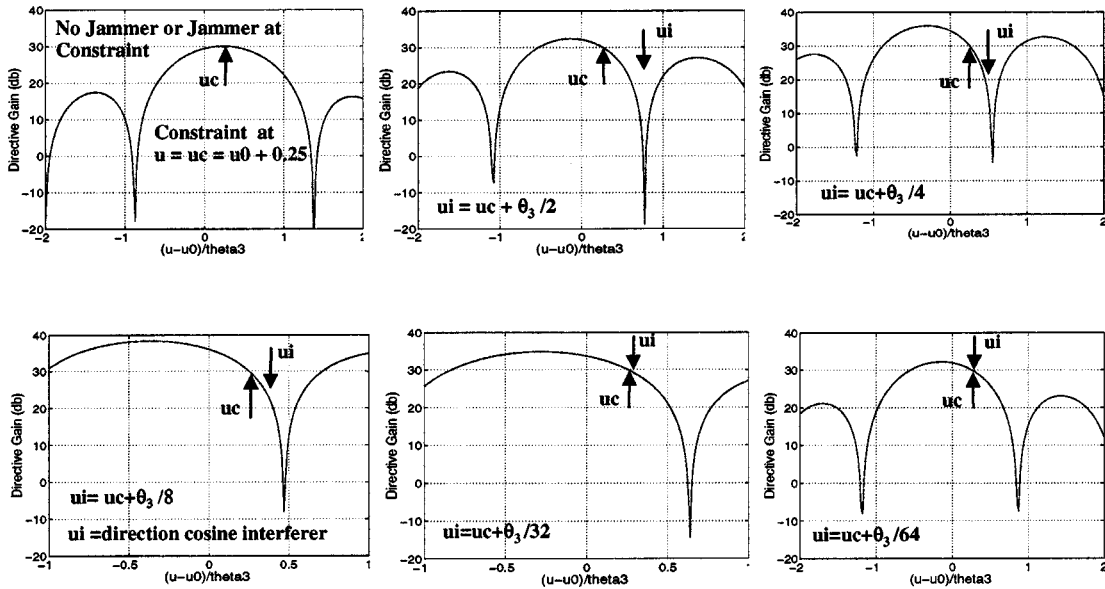


Fig. 7. Adapted antenna patterns versus separation of interferer from constraint point: INR = 15 dB at peak.

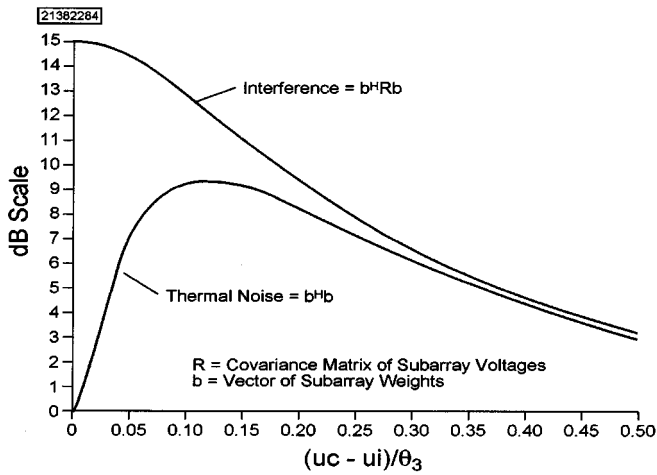


Fig. 8. Total interference and thermal noise in an adapted beam versus separation of interferer from constraint point.

search over a constrained region of possible target AOAs and declaring the angular location that produces the highest SIR in the output beam to be the target location. The processing must be repeated for each range cell. A typical system might use 16 subarrays and four-output beams. The architecture supports improved radar performance by eliminating beamshape loss on receive and enabling on-axis tracking of all targets.

Simulation results predict that a four-beam MLBP will be able to search roughly 70% more volume for the same number of transmissions compared to a conventional approach that uses a single receive beam. The improvement is due to the elimination of receive-beamshape loss. The improvement in angle accuracy compared to monopulse varied from a few percent for a target located at the peak of the transmit beam, to roughly 30% for a target located at the 3-dB point for a four-beam MLBP system. The improvement in angle accuracy translated into roughly a 50% reduction in Kalman filter tracking errors for a target whose track averaged about a third of a beamwidth off-axis.

The new architecture supports nulling of both mainlobe and sidelobe interference. Most nulling occurs in the first stage processing, but residual—interference that is correlated beam to beam—can be further nulled in the second stage. Significant improvement in SIR was observed using a four-beam MLBP operating against a single-mainlobe interferer except for very small target-interferer separations. The small degradation observed for small separations, compared to using no adaptivity, was able to be eliminated by adding more beams. The fact that more degradation was not observed for small separations was shown to be due to the fact that the gain constraints placed in the first-stage beams preclude interferer nulling over a relatively large region.

We believe that future radars will abandon monopulse in favor of maximum-likelihood target location and that a beamspace architecture using highly overlapped beams supports practical and efficient implementation.

APPENDIX A

SIGNAL CANCELLATION IN FIRST STAGE BEAMFORMER

In this appendix, we demonstrate that if beam formation is done before pulse compression, there will generally be no signal cancellation as long as SNR in the beam before nulling is less than unity (target return less than receiver noise). In order to demonstrate the point, we have modeled a linear array of M elements separated by one-half wavelength (i.e., each subarray consists of only one element). Suppose we adaptively weight the M elements so as to form a beam steered to θ_0 in the presence of a nonfluctuating signal incident at $\theta_s = \theta_0 + \alpha\theta_3$, where θ_3 is the 3 dB beamwidth. In the absence of interference, the voltage on element n at discrete time i is

$$v_n(i) = s_n + x_n(i) \quad (\text{A-1})$$

where the signal voltage on the n th element is given by

$$s_n = A \exp[j\pi n(\sin \theta_s - \sin \theta_0)] \quad (\text{A-2})$$

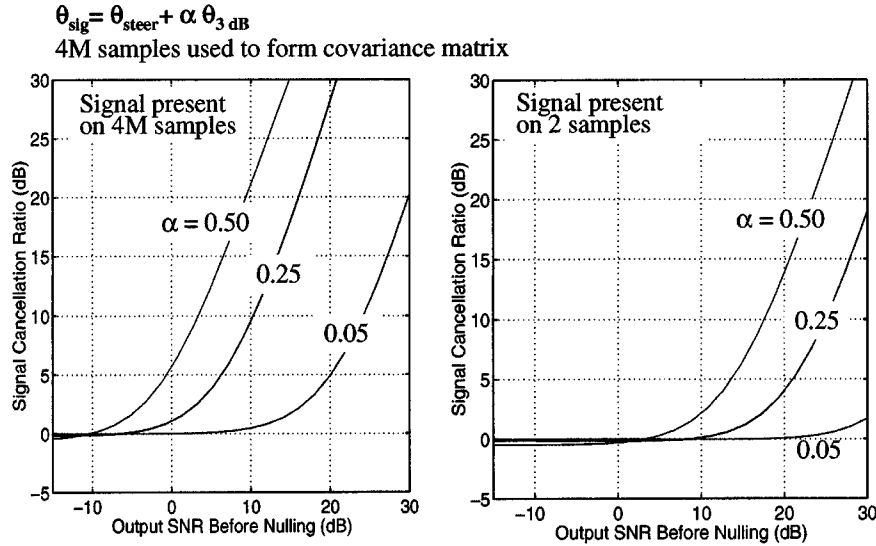


Fig. 9. Signal cancellation prior to pulse compression, with no diagonal loading.

and x_n is the complex, zero mean, unit variance noise voltage on element n . The covariance matrix of the M elements is then given by

$$R_I = I^{-1} \sum_{i=1}^I V_i^* V_i^T \tag{A-3}$$

where

$$V_i^T = [v_i(i) v_2(i) \dots v_M(i)].$$

In order to form a beam steered to θ_0 , we apply the weight vector

$$W = \frac{M^{1/2} R_I^{-1} P}{P^H R_I^{-1} P} \tag{A-4}$$

to the received voltage V_i , where P is an $M \times 1$ steering vector of ones. The numerator of (A-4) is the whitening filter and the denominator is a scale factor that constrains the gain at the peak of the beam. The signal power in the beam (in the absence of any interference) before nulling is

$$\xi_o = \frac{|S^T P|^2}{M} \tag{A-5}$$

where

$$S^T = [s_1 s_2 \dots s_M] \tag{A-6}$$

The signal cancellation ratio is then given by

$$CR_s = \frac{|W^T S|^2}{\xi_o}. \tag{A-7}$$

A MATLAB program was written for a 16 element linear array to study the signal cancellation problem. The beam was steered 20° off normal and the angle of arrival and SNR of the incident signal were varied. The left insert in Fig. 9 plots signal cancellation ratio versus output SNR before nulling, for a case where 4M samples ($I=64$) were used to generate the sample covariance matrix and the target return was present in each sample.

The right insert shows the same cases, but the target was present in only two of the 64 training samples.

The left insert shows that if the signal is present in all samples used to form the covariance matrix, no signal cancellation occurs until SNR in the beam before nulling (adaptive weights set to unity) is within 10 dB of the receiver noise level, and even at that signal level, nulling only begins if the signal is near the 3-dB point on the beam. As the signal approaches the peak of the beam (the constraint point), SNR before nulling must get larger before any signal cancellation occurs. The right insert demonstrate that if the signal is not present in all samples, the SNR at which cancellation begins increases, as the number of samples containing the signal decreases.

The issue at hand is whether to perform pulse compression at the subarray level or the beam level. Let us first assume that compression is done at the subarray level prior to beam formation and adaptive weighting. After compression, the signal will typically straddle two adjacent range cells. The right insert in Fig. 9 shows that significant signal cancellation will occur a large percent of the time since SNR will be in the 10–20 dB region, in two cells, at the beam level, after compression but prior to nulling. One solution is to recalculate both the first and second stage covariance matrices and their inverses for every range cell in which we attempt a detection; and, omit the detection cell and its nearest neighbors from the training region. The latter requirement would substantially increase the computational complexity. A more reasonable approach appears to be to do compression at the beam level—at least in radars having high pulse compression gain.

If compression is done at the beam level, then in radars having high pulse compression gain, SNR will usually be much less than unity in each of the first-stage beams prior to compression and adaptive nulling, and no signal cancellation will occur. If the signal is so large that it is above noise at the beam level prior to compression and the compression gain is large, the radar could insert attenuation to lower the signal or could just tolerate a little cancellation. If, however, the waveform has low compression gain, even if we compress the signal at the beam level, it may

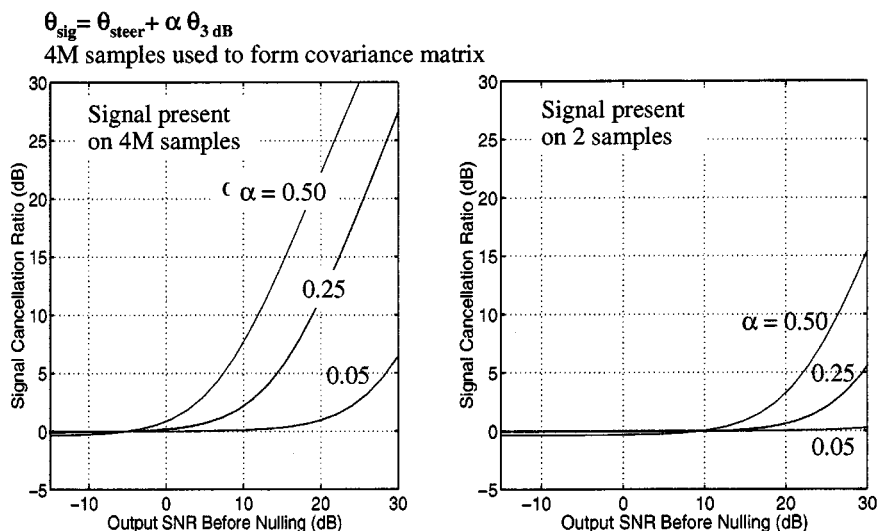


Fig. 10. Signal cancellation prior to pulse compression, with 8 dB of diagonal loading.

still be above noise prior to compression and we could be in trouble. In the latter case, we have two options: 1) recalculate the covariance matrix for each range cell in which a detection is attempted, omit the detection cell and its nearest neighbors from the training region, and live with the increased computation, or 2) apply diagonal loading to the covariance matrix of subarray voltages [13]. Fig. 10 shows that adding 8 dB of diagonal loading solves the problem. The diagonal loading was accomplished by adding a term $d|z_i|^2$ to the i th diagonal component of the covariance matrix, where d is a scalar and z_i is a complex, zero-mean, unit variance Gaussian variate. Computer simulation predicted that in the presence of sidelobe interference, d could be as large as one-tenth the trace of the covariance matrix prior to loading before seriously affecting the ability of the canceller to null the interference.

APPENDIX B INTERFERER-FREE LIMIT OF (17)

Suppose that there are no interferers present and that each subarray consists of only a single isotropic element, so that $g_m(\theta, \phi) = 1$. Then, if the noise in each subarray is independent and a target of amplitude A is located at the center of a beam, it is well known that the output signal to noise ratio is $M|A|^2/\sigma^2$ where M is the number of elements and σ^2 is the noise power on each element. Let us show that (17) correctly reduces to this limit.

In the absence of interferers, $R = \sigma^2 I$, where I is the $M \times M$ identity matrix. Therefore, (10) becomes

$$C = \sigma^2 BB^H. \quad (\text{B-1})$$

If (B-1) and (8) are used, we can calculate

$$T^H \xi = A + \frac{G^H C^{-1} \rho}{G^H C^{-1} G}. \quad (\text{B-2})$$

Because ρ includes only noise, then if the SNR after pulse compression is large, we can ignore the second term on the

right-hand side of (B-2). Consequently, if we use (B-1) and (B-2), along with the defined quantity $G = Bh$ in (17), we get

$$\text{SIR} = \frac{|A|^2}{\sigma^2} h^H B^H (BB^H) Bh. \quad (\text{B-3})$$

Next, suppose there is only a single constraint (beam) at (θ_1, ϕ_1) and the target is located at the peak of the beam. Then, because $S_1^H S_1 = M$, we see from (1) and (2) that $w_1 = S_1/M$ and $B = S_1^H/M$. Also, upon comparing (4) and (6), it is evident that when $\theta = \theta_1, \phi = \phi_1$, we have $h = S_1$. Therefore, if these results are used in (B-3), we find

$$\begin{aligned} \text{SIR} &= \frac{|A|^2}{\sigma^2} \left(\frac{S_1^H S_1}{M} \right) \left(\frac{S_1^H S_1}{M^2} \right)^{-1} \left(\frac{S_1^H S_1}{M} \right) \\ &= M \frac{|A|^2}{\sigma^2} \end{aligned} \quad (\text{B-4})$$

which is the correct limit.

REFERENCES

- [1] E. Baranoski and J. Ward, "Source localization using adaptive subspace beamformer outputs," in *1997 IEEE Int. Conf. Speech Signal Processing (ICASSP 97)*, vol. 5, pp. 3773–3776.
- [2] S. Applebaum, "Adaptive arrays," *IEEE Trans. Antennas Propagat.*, vol. AP-24, pp. 585–598, Sept. 1976.
- [3] E. Kelly, I. Reed, and W. Root, "The detection of radar echoes in noise II," *J. Soc. Ind. Appl. Math.*, vol. 8, pp. 481–510, Sept. 1960.
- [4] H. Urick, "The accuracy of maximum likelihood angle estimates in radar and sonar," *IEEE Trans. Mil. Electron.*, vol. 8, pp. 39–45, Jan. 1964.
- [5] R. Davis, L. Brennan, and I. Reed, "Angle estimation with adaptive arrays in external noise fields," *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-12, pp. 179–186, Mar. 1976.
- [6] A. Whalen, *Detection of Signals in Noise*. New York: Academic, 1971.
- [7] L. Brennan and I. Reed, "Theory of adaptive radar," *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-9, pp. 237–252, 1973.
- [8] I. Reed, J. Mallet, and L. Brennan, "Rapid convergence rate in adaptive antennas," *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-10, Nov. 1974.
- [9] B. Picinbono, *Random Signals & Systems*. Englewood Cliffs, NJ: Prentice Hall, 1993.
- [10] G. F. Hatke, "Superresolution source location with planar arrays," *Lincoln Lab. J.*, vol. 10, no. 2, 1997.
- [11] D. K. Barton and H. R. Ward, *Handbook of Radar Measurement*. Englewood Cliffs, NJ: Prentice-Hall, 1969, p. 24.

- [12] R. M. Davis, R. L. Fante, W. J. Crosby, and R. J. Balla, "A maximum likelihood beamspace processor," in *Rec. IEEE Int. Radar Conf.*, May 2000.
- [13] B. D. Carlson, "Covariance matrix estimation errors and diagonal loading in adaptive arrays," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 24, July 1988.



Richard M. Davis (S'62–M'72–SM'85) received the B.S.E.E. degree from the University of Rochester, Rochester, NY, in 1964, and the B.A. and M.S.E.E. degrees from Syracuse University, Syracuse, NY.

He was with the Syracuse Research Corporation in Syracuse, New York for 24 years. After leaving the research corporation, he was Technical Director of the Analytical Studies Center, which specialized in electronic counter countermeasures. He is currently a Senior Principal Engineer with the MITRE Corporation, Bedford, MA. He has published 29 papers in

the areas of signal and adaptive processing.

Ronald L. Fante (F'79) received the B.S. degree from the University of Pennsylvania, the M.S. degree from MIT, and the Ph.D. degree from Princeton University.

He is currently a MITRE Fellow, with the MITRE Corporation, Bedford, MA. At MITRE he is leading tasks on applying adaptive processing to ECCM for GPS, to space-time adaptive processing for clutter cancellation and adaptive techniques for SAR image formation in the presence of interference. He is the author of *Signal Analysis and Estimation* (New York:Wiley, 1988). He has written 140 journal papers and holds 2 patents in the fields of adaptive signal processing, electromagnetics, optics, and wave propagation.

Dr. Fante is a Fellow of the Optical Society of America and the Institute of Physics. He recently received the Third Millennium Medal from the IEEE. He has been an IEEE Distinguished Lecturer, Chairman of the Awards and Fellows Committee, and Editor-in-Chief of the IEEE TRANSACTIONS ON ANTENNAS AND PROPAGATION.