The Institute of Electrical and Electronics Engineers, Inc.

Fifth

ASSP Workshop on Spectrum Estimation and Modeling

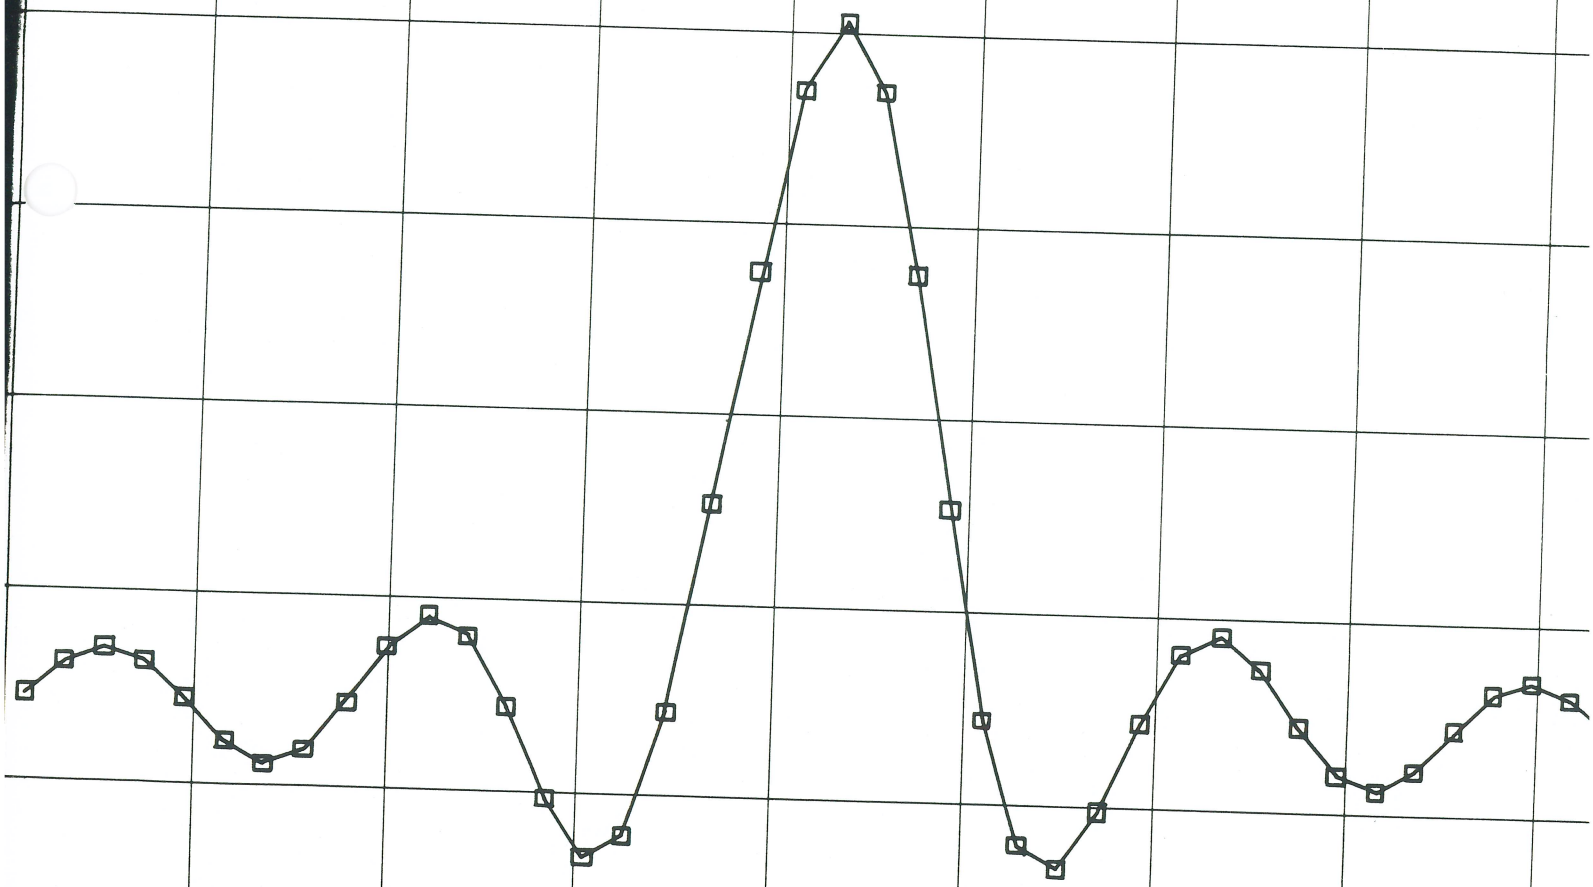

October 10-12, 1990

Woodcliff Conference Center
Rochester, New York

90TH0331-9

PROGRESS ON SIGNAL–SELECTIVE DIRECTION FINDING*

Stephan V. Schell and William A. Gardner
Department of Electrical Engineering and Computer Science
University of California, Davis, CA 95616, USA

Abstract

The recently discovered Cyclic MUSIC algorithm for narrowband signal–selective direction finding using antenna arrays circumvents many drawbacks of conventional techniques by exploiting known spectral correlation properties (namely, known cycle frequencies such as the baud rate or carrier frequency) of the desired signals to reject undesired signals, interference, and noise. In this paper, two recent advances in the capabilities of Cyclic MUSIC are described. The first enables Cyclic MUSIC to simultaneously estimate the directions of arrival of signals having different cycle frequencies instead of having to sequentially process each separate frequency in a list of cycle frequencies (either known a priori or measured). The second advance reduces the sensitivity of Cyclic MUSIC to error in the knowledge of the cycle frequency of interest by estimating the frequency of a quadratically–regenerated sine wave and then using that estimate as the cycle frequency parameter in the computation of the cyclic autocorrelation matrix, which is then processed to estimate the directions of arrival.

1 Introduction

The need to estimate the directions of arrival (DOAs) of propagating waves arises in many areas, including intelligence, surveillance, commercial communications monitoring, sonar, radar, astronomy, and geophysics [6]. Many techniques that use an array of sensors have been proposed for solving this problem, and, recently, emphasis has been placed on high–resolution and super–resolution DOA estimation techniques which exhibit resolution much finer than that attainable by conventional beamforming. One popular example of the super–resolution techniques, the Multiple Signal Classification (MUSIC) method [7], exhibits very desirable properties, including asymptotic (as the number of data samples goes to infinity) unbiasedness of the DOA estimates and mean–square–error (MSE) comparable to the Cramer–Rao Lower Bound (CRLB) in many signal environments [4].

*This material is based upon work supported in part by the National Science Foundation under Grant No. MIP–88–12902 and by the United States Army Research Office under Contract DAAL03–89–C–035 sponsored by the Army Communications Electronics Command Center for Signals Warfare.

However, the excellent performance of the MUSIC method depends on several conditions being satisfied, and this can be difficult or impossible in practice. First, the number of signals impinging on the array must be less than the number of sensors. Second, the spatial characteristics of the noise must be known or exhibit a transformational invariance. Third, since the DOAs of *all* signals impinging on the array must be estimated, some post–processing must typically be applied to determine which DOA estimates, if any, correspond to the signals of interest. Thus, computational effort is often wasted in obtaining and processing the DOAs of undesired signals. An important consequence of the third condition is that an undesired signal can arrive from a direction that is very nearly the same as the direction from which a desired signal arrives, and this can require that very long data sets be collected to resolve the signals. Yet another consequence is that the presence of undesired signals that are fully correlated among themselves can prevent the use of the MUSIC method even if the desired signals are not fully correlated among themselves or with the undesired signals.

One means of softening or eliminating the aforementioned constraints is to select only a subset of signals to be processed. In particular, the cyclostationarity properties exhibited by most communication and telemetry signals [2, 1, 3] can be exploited to select a desired subset of signals and to discriminate against undesired signals, interference, and noise. The Cyclic MUSIC direction–finding (DF) method [5] can attain much better performance in some environments than the MUSIC method even when MUSIC is operating properly, and can obtain the desired DOA estimates in some environments in which MUSIC fails. For example, in order to resolve two signals arriving from nearly the same direction, Cyclic MUSIC can require much less data than conventional MUSIC does if the two signals have different cycle frequencies. As another example, MUSIC fails if more signals are received than there are sensors, but Cyclic MUSIC can still estimate the DOAs of the signals having known cycle frequencies provided that the number of signals having any given cycle frequency is less than the number of sensors. Furthermore, substantial savings in computation can be achieved by using the Cyclic MUSIC method because the DOAs of only the desired signals are estimated.

One drawback of Cyclic MUSIC is, in fact, a result

TH0331-9/90/0000-0144$01.00©1990 IEEE

of its signal–selectivity. If the signals of interest do not share a common cycle frequency, then Cyclic MUSIC must be applied separately for each cycle frequency of interest, which requires the estimation of the corresponding cyclic autocorrelation matrix, computation of its singular values and vectors, and location of the minima in the null spectrum. To minimize the extra computation, an extension of Cyclic MUSIC which is called *Multi–Cyclic MUSIC* can be applied instead. This extension is introduced in this paper.

Another drawback in some applications is the need to know a cycle frequency of the signal(s) of interest. However, the true cycle frequency can differ from the assumed cycle frequency due to symbol clock drift in the transmitter, unknown and/or time–varying Doppler shift, drift in the local oscillator, and other causes of imprecise estimates of the true cycle frequency. In some applications, knowledge of the true cycle frequency can be completely unavailable, although the benefits of signal selectivity are still desired. To address these issues, a new method for adapting to unknown cycle frequencies is introduced in this paper.

2 Cyclic MUSIC

The Cyclic MUSIC method exploits both the spatial and spectral correlation properties of the received data, which is assumed here to be accurately described by the narrowband model

$$\mathbf{x}(n) = \sum_{l=1}^{L} \mathbf{a}(\theta_l)\, s_l(n) + \mathbf{i}(n), \tag{1}$$

where $\mathbf{x}(n)$ is the Mx1 vector of sampled complex envelopes at the sensor outputs, $\mathbf{a}(\theta)$ is the array response vector for a signal arriving from angle θ, $s_1(n), \ldots, s_L(n)$ are the L sampled complex envelopes of the signals arriving from angles $\theta_1, \ldots, \theta_L$, respectively, and $\mathbf{i}(n)$ is interference and noise that is uncorrelated with the impinging signals. It is also assumed that L_α of the signals have cycle frequency α, where $L_\alpha \leq L$. The Cyclic MUSIC method computes the $M - L_\alpha$ left singular vectors corresponding to zero–valued singular values of the cyclic autocorrelation matrix $\mathbf{R}^{\alpha}_{\mathbf{xx}}$ and searches for the angles for which the array response vectors are orthogonal to those singular vectors:

1. Measure $\mathbf{R}^{\alpha}_{\mathbf{xx}}(\tau) = \langle \mathbf{x}(n)\,\mathbf{x}^H(n-\tau)\, e^{-j2\pi\alpha n}\rangle_N$, where $\langle\cdot\rangle_N$ denotes the time average over N samples,

2. Compute the SVD $\mathbf{R}^{\alpha}_{\mathbf{xx}}(\tau) = \mathbf{U}\Sigma\mathbf{V}^H$, where $\mathbf{U}$ and $\mathbf{V}$ are unitary matrices whose columns are the left and right singular vectors, respectively, and Σ is the diagonal matrix of real, non–negative singular values,

3. Locate minima of $\left\| [\mathbf{u}_{M-L_\alpha} \ldots \mathbf{u}_M]^H\, \mathbf{a}(\theta) \right\|^2$.

3 Multi–Cyclic MUSIC

Clearly, the time average in step 1 estimates the complex coefficient of a single sine wave of frequency α in the lag–product waveform $\mathbf{x}(n)\,\mathbf{x}^H(n-\tau)$. Thus, replacing the single sine wave of frequency α in the time average with a sum of sine waves at K different frequencies $A = \{\alpha_1, \ldots \alpha_K\}$ estimates the sum of the complex coefficients of the corresponding sine waves in the lag product, and the resulting *multi–cyclic autocorrelation matrix* is the sum of the cyclic autocorrelation matrices corresponding to those cycle frequencies:

$$\mathbf{R}^{A}_{\mathbf{xx}}(\tau) = \left\langle \mathbf{x}(n)\,\mathbf{x}^H(n-\tau) \left(\sum_{k=1}^{K} e^{-j2\pi\alpha_k n} \right) \right\rangle_N \tag{2}$$

$$= \sum_{k=1}^{K} \mathbf{R}^{\alpha_k}_{\mathbf{xx}}(\tau). \tag{3}$$

If the frequencies coincide with the cycle frequencies of the signals of interest, then $\mathbf{R}^{A}_{\mathbf{xx}}(\tau)$ describes the spatial characteristics of those signals, because each constituent term $\mathbf{R}^{\alpha_k}_{\mathbf{xx}}$ describes the spatial characteristics of the signals having cycle frequency α_k. Provided that the number of signals exhibiting spectral correlation at the K cycle frequencies of interest is less than the number of sensors, steps 2 and 3 of the conventional Cyclic MUSIC algorithm can be performed on $\mathbf{R}^{A}_{\mathbf{xx}}(\tau)$ to estimate the DOAs of those signals simultaneously.

Multi–Cyclic MUSIC exhibits all of the same benefits of signal–selectivity as Cyclic MUSIC, such as the ability to operate properly in the presence of noise and interference having unknown spatial characteristics, in the presence of more signals than sensors (provided that the number of signal having the desired cycle frequency or frequencies is less than the number of sensors), and in the presence of undesired signals that are spaced too closely to the desired signals for conventional MUSIC to resolve them using a practical number of data samples.

Furthermore, substantial computational and hardware savings can result from using Multi–Cyclic MUSIC instead of Cyclic MUSIC because only one cross–correlation matrix is estimated, only one SVD is computed, and only one search is performed.

However, some degradation in the performance of Multi–Cyclic MUSIC as compared to Cyclic MUSIC can occur when two closely–spaced desired signals having different cycle frequencies are received. Since Cyclic MUSIC estimates the DOAs of those two signals separately, having resolved them first on the basis of differing cycle frequencies, it can do better than Multi–Cyclic MUSIC, which can resolve the desired signals only in the spatial domain.

3.1 Multi–Cyclic MUSIC Simulations

The behavior of Multi–Cyclic MUSIC is illustrated here by Monte Carlo simulations. It is shown that for environments in which conventional MUSIC performs poorly

or fails, Multi–Cyclic MUSIC can still obtain good estimates of the DOAs of the desired signals, although Cyclic MUSIC can do better.

In all of the simulations presented here, a four–element circular array having diameter equal to one half of the carrier wavelength receives signals from sources in the far–field. Two binary PAM signals, each having a raised–cosine pulse–transform of 100% excess bandwidth and an SNR of 10dB, arrive from 10 degrees and -15 degrees, respectively. The bit rates vary from environment to environment, and are specified relative to the sampling rate of unity. Spatially and temporally white complex Gaussian noise is also present. Cyclic and Multi–Cyclic MUSIC are given exact knowledge of the cycle frequencies of the desired signals, and MUSIC is given exact knowledge of the ideal spatial autocorrelation matrix of the noise. All algorithms are given exact knowledge of the appropriate number of signals. For each environment of interest, 1000 independent trials are conducted for each of 64, 128, 256, ..., 8192 complex data samples.

In the first environment, the signal arriving from 10 degrees has bit rate 0.25, and the signal arriving from -15 degrees has bit rate 0.1875. As shown in Figure 1, Multi–Cyclic MUSIC obtains estimates having relatively high RMSE when estimating both DOAs simultaneously, whereas Cyclic MUSIC performs better because it first resolves the signals on the basis of their cycle frequencies. Finally, conventional MUSIC performs better here than both cyclic methods, indicating that it is the best choice of the three in this simple environment, primarily because only the desired signals and known noise are present, and the quadratically regenerated spectral line at zero frequency (see Figure 4 for an example) that MUSIC exploits is much stronger than that at a non–zero cycle frequency which is exploited by Cyclic MUSIC.

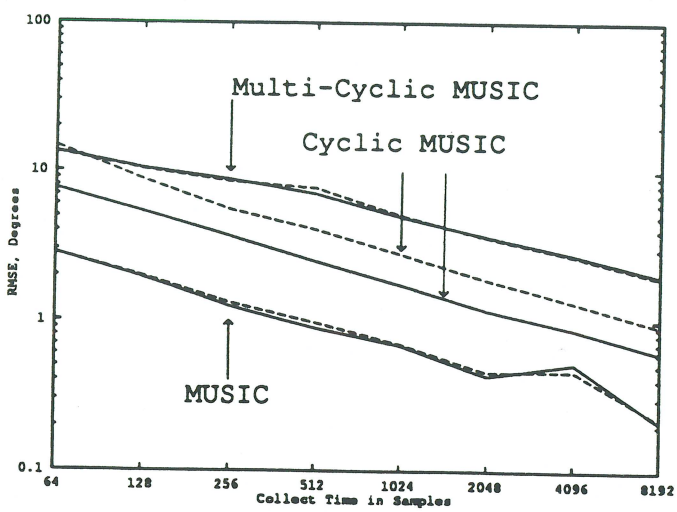

Figure 1: RMSE of DOA estimates for two signals from 10 degrees and -15 degrees having different baud rates.

In contrast, the cyclic methods do much better when an additional, undesired signal is present. In addition to the two desired signals in the first environment, we now consider a stationary SSB–AM signal having 10 dB SNR and arriving from 0 degrees. As shown in Figure 2, both Multi–Cyclic MUSIC and Cyclic MUSIC outperform conventional MUSIC, for which the probability of success is less than 1% even when using 8192 samples. It should be noted that results obtained here for the four–element circular array should not be compared with those that could be obtained for a uniform linear array having four (or more) elements due to the differing physical apertures. Thus, although MUSIC is unable to resolve the 3 signals arriving within a 25–degree sector using the circular array, it might well resolve them using a linear array. However, these results do indicate that Cyclic MUSIC can accommodate arrays with smaller apertures, because the signal–selectivity exhibited by both cyclic DF methods allows them to discriminate against the undesired signal *before* computing the SVD. In particular, it should be noted that the RMSEs of the estimates obtained by the cyclic methods, shown in Figure 3, are very close to those obtained in the previous environment *without the interference*, shown in Figure 1.

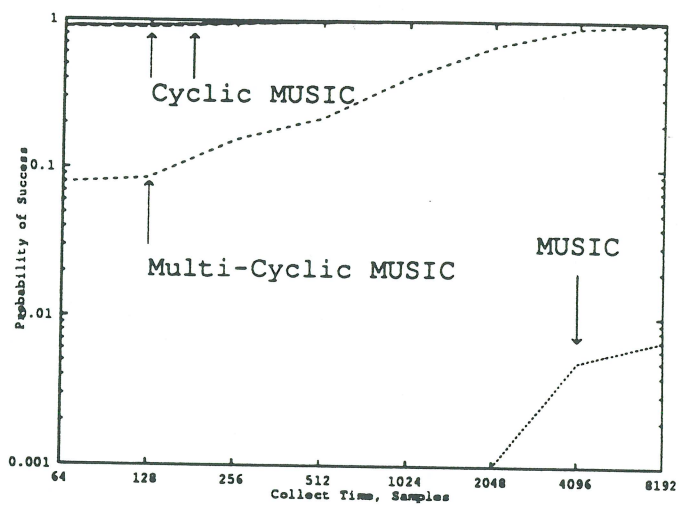

Figure 2: Probabilities of success for Cyclic MUSIC, Multi–Cyclic MUSIC, and conventional MUSIC for two signals arriving from 10 degrees and -15 degrees, having different baud rates, and a stationary undesired signal arriving from 0 degrees.

4 Adaptive–α Cyclic MUSIC

As mentioned in the Introduction, exact knowledge of the cycle frequency of interest can be unavailable, in which case the performance of Cyclic MUSIC falls off as the averaging time exceeds (roughly) the reciprocal of the error in the assumed cycle frequency. This phenomenon can be illustrated by looking at the frequency response $H(f)$ of the rectangular averaging window of length N

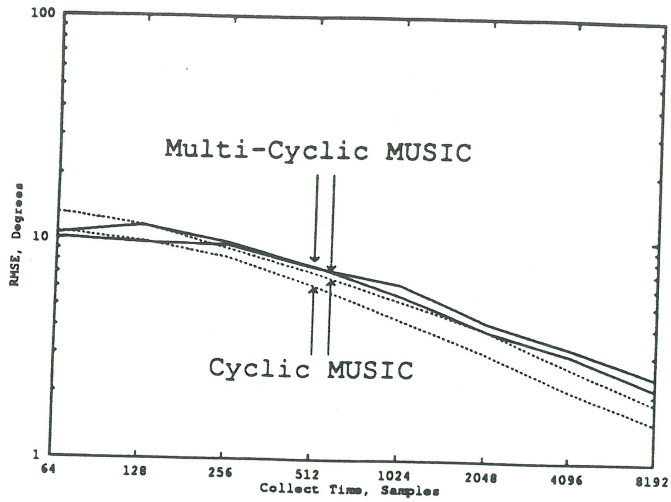

Figure 3: RMSE of DOA estimates obtained by Cyclic MUSIC and Multi–Cyclic MUSIC for two signals from 10 degrees and -15 degrees having different baud rates. A stationary undesired signal arrives from 0 degrees.

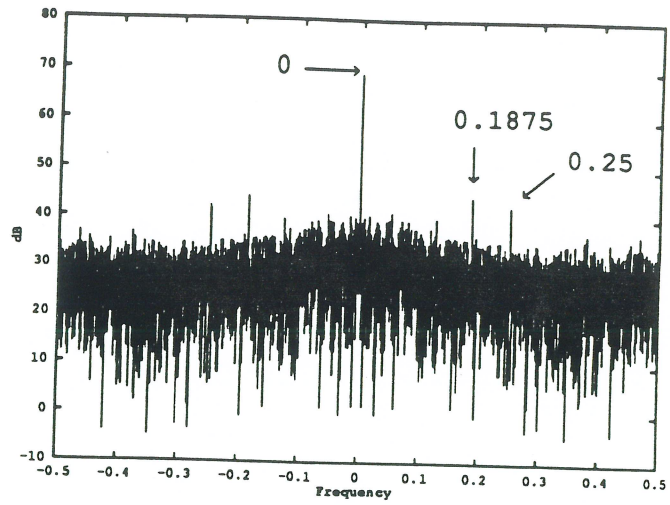

Figure 4: Magnitude of FFT of $x_1(n)\,x_1^*(n)$ for signal environment used in Figures 2 and 3. Regenerated spectral lines are clearly visible at frequencies ± 0.1875 and ± 0.25, corresponding to the bit rates of the two cyclostationary signals.

used in $\langle\cdot\rangle_N$,

$$H(f) = \frac{\sin(\pi f N)}{\pi f N}\frac{\pi f}{\sin(\pi f)}e^{-i\pi f(N-1)}. \qquad (4)$$

If the sine wave of interest in the frequency–shifted lag product $\mathbf{x}(n)\,\mathbf{x}^H(n-\tau)\,e^{-j2\pi\alpha n}$ is not at zero but is instead at $\Delta\alpha = \alpha - \alpha_{true}$, then $H(\Delta\alpha)$ equals zero when $N = 1/\Delta\alpha$, at which point the sine wave of interest is rejected. In fact, performance begins to degrade noticeably for $N > 1/2\Delta\alpha$ as the main lobe of $H(f)$ shrinks in width and the sine wave of interest is attenuated to a greater degree.

One possible means of reducing this sensitivity to error in α is to use Multi–Cyclic MUSIC with the multiple cycle frequencies $\alpha_1, \ldots, \alpha_K$ being spread uniformly across a specified band in which the true cycle frequency is assumed to be. This scheme would seemingly perform better than using only one α, because the error in α for the constituent term $\mathbf{R}_{\mathbf{XX}}^{\alpha_k}$, for which α_k is closest to the true value α_{true}, is as little as $1/(K-1)$ times that incurred when using only one α. However, simulations in which the α_k are spread across a band of width 0.1% of the true cycle frequency, and also across a band of width 1%, show that the DF performance is not dominated by the term for which the assumed cycle frequency is closest to the true cycle frequency. In particular, simulations show performance degrades for $K = 2$ due to cancellation between terms, and that the performance improvement due to smaller error in one term for $K > 2$ (values of K up to 16 were tested) is offset by extra noise and interference that is contributed by the other terms.

However, another approach can perform much better, albeit at the expense of increased computation, by merely estimating the frequency of the sine wave in the lag product and then using it in Cyclic MUSIC. For example, the magnitude of the FFT of the lag product of the first element of $\mathbf{x}(n)$ from the second environment in the Multi–Cyclic MUSIC simulations (which Figures 2 and 3 correspond to) is shown in Figure 4. In addition to the line at $\alpha = 0$, the sine waves at the bit rates of the signals (and at their negatives) are clearly visible. Thus, with essentially no prior knowledge of the cycle frequencies present in the received data, either of these frequencies can be estimated and then used to compute the corresponding cyclic autocorrelation matrix for use by Cyclic MUSIC.

Using the classical periodogram method, the frequency estimates are taken to be the frequencies at which local maxima occur in the magnitude of the FFT of the lag–product waveform.

Depending upon the constraints on the amount of available data and memory for storing that data, the data sets used to estimate the cycle frequency and the directions of arrival can be the same, or the data set used to estimate the cycle frequency can precede the data set used to estimate the directions of arrival. The former implementation is used in the following simulations.

4.1 Adaptive–α Simulations

The same simulated environment as used in the second Multi–Cyclic MUSIC simulation in Section 3.1 is used here. However, the cycle frequency 0.1875 is known only to lie in the band $[0.15, 0.22]$, and the cycle frequency 0.25 is known only to lie in the band $[0.22, 0.5]$. Effectively, this implies that the error in the knowledge of the cycle frequencies is 15% and 100%, respectively. When the collect time is small, the estimates of the cycle frequencies are corrupted by leakage from each other

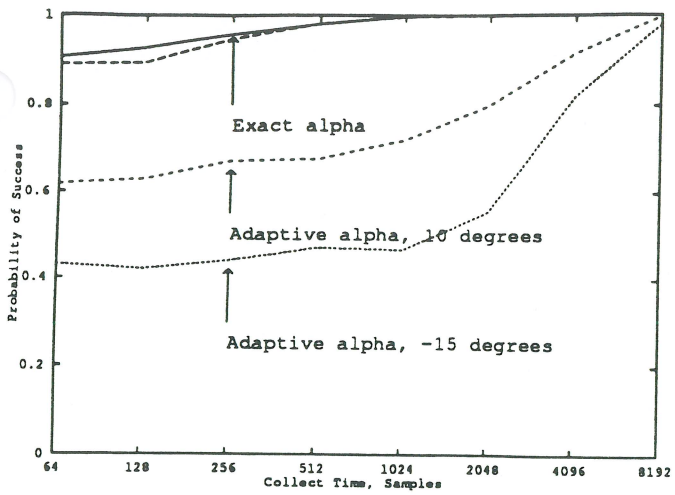

Figure 5: Probability of success of adaptive–α Cyclic MUSIC compared with Cyclic MUSIC using exact knowledge of α.

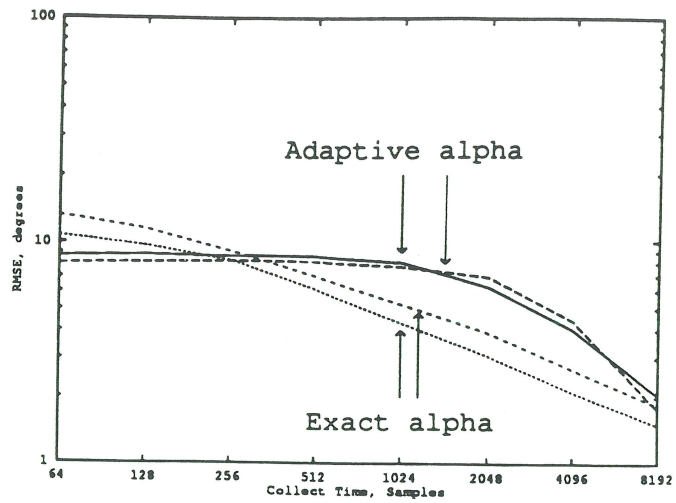

Figure 6: RMSE of θ estimates computed by Cyclic MUSIC using estimated values of α compared with Cyclic MUSIC using exact knowledge of α.

and from the feature at cycle frequency 0, in addition to random effects. The resulting poor estimates of the cycle frequencies yield correspondingly low probabilities of success for the Cyclic MUSIC algorithm that uses them, as shown in Figure 5. However, as the number of samples becomes large, the probability of success rises, becoming comparable to that obtained with exact knowledge of the ycle frequencies. This behavior is also reflected in the RMSE of the DOA estimates from the successful trials, since the high–accuracy estimates of the cycle frequencies obtained from using a large number of samples yield correspondingly good estimates of the DOAs, as shown in Figure 6. Thus, even though the cycle frequencies are essentially unknown, Cyclic MUSIC can perform well in this environment, whereas conventional MUSIC fails in greater than 99% of the trials.

Simulations (the results of which are not shown here) show that much better performance is obtained in the absence of the stationary interference. In fact, when using more than 1024 samples, the algorithm obtains estimates of the cycle frequencies that are exact. Thus, the RMSE of the corresponding DOA estimates is exactly the same had the cycle frequencies been known *a priori*.

5 Conclusion

Two extensions to the Cyclic MUSIC direction–finding method are presented and shown to yield acceptable performance while alleviating certain drawbacks of Cyclic MUSIC, thus extending applicability. First, the Multi–Cyclic MUSIC method is shown to perform DF simul- neously on multiple signals having different cycle frequencies , which reduces computational load at the expense of higher RMSE. Second, a technique for greatly reducing the sensitivity of Cyclic MUSIC to error in the knowledge of the cycle frequency of interest is shown to yield excellent performance for a sufficiently large number of samples, even when the cycle frequencies are essentially unknown at the outset.

References

[1] W.A. Gardner. Spectral correlation of modulated signals: Part I — analog modulation. *IEEE Trans. Comm.*, COM–35(6):584–594, June 1987.

[2] W.A. Gardner. *Statistical Spectral Analysis: A Nonprobabilistic Theory.* Prentice-Hall, Englewood Cliffs, NJ, 1987.

[3] W.A. Gardner, W.A. Brown, III, and C.-K. Chen. Spectral correlation of modulated signals: Part II — digital modulation. *IEEE Trans. Comm.*, COM–35(6):595–601, June 1987.

[4] B. Porat and B. Friedlander. Analysis of the asymptotic relative efficiency of the MUSIC algorithm. *IEEE Trans. Acoust., Speech, Signal Processing*, ASSP-36(4):532–544, April 1988.

[5] S.V. Schell, R.A. Calabretta, W.A. Gardner, and B.G. Agee. Cyclic MUSIC algorithms for signal-selective DOA estimation. In *Proc. IEEE Int. Conf. Acoust., Speech, Signal Processing*, pages 2278–2281, Glasgow, Scotland, May 1989.

[6] S.V. Schell and W.A. Gardner. High–resolution direction finding. *ASSP Magazine*, 1990. Submitted.

[7] R.O. Schmidt. *A signal subspace approach to multiple source location and spectral estimation.* Ph.D. dissertation, Stanford Univ., Stanford, CA, 1981.