SSPI
STATISTICAL SIGNAL PROCESSING, INCORPORATED

SIGNAL INTERCEPTION BY EXPLOITATION OF CYCLOSTATIONARITY: AN OVERVIEW AND COMPILATION OF RESEARCH

by

William A. Gardner*

prepared for

The Naval Postgraduate School
Monterey, CA

Technical Report SSPI-92-1
January 1992

*The author is with Statistical Signal Processing, Inc. and the Department of Electrical and Computer Engineering, University of California, Davis.

ACKNOWLEDGEMENT

The author expresses his gratitude to Professor Herschel H. Loomis, Jr. of the Naval Postgraduate School for his substantial role in helping to secure support for the research reported herein, and for his service as coadvisor with the author of the Ph.D. thesis research, at the University of California, Davis, of Drs. William A. Brown and Randy S. Roberts, which contributed to the research reported herein. The author also gratefully acknowledges Drs. C. K. Chen, S. V. Schell, and C. M. Spooner for the contributions of their Ph.D. thesis research at the University of California, Davis, and the contributions of their work as consultants to Statistical Signal Processing, Inc. The specific contributions of these colleagues are reflected in the coauthorships of and acknowledgements in each of the individual reports making up this compilation. Finally, the author acknowledges the many M.S. thesis students at the Naval Postgraduate School, supervised by Professor Loomis, whose thesis work provided independent verification of the performance characteristics of a number of the algorithms invented during the course of this research.

ABSTRACT

This report consists of an extensive overview of the research reported in twenty-five individual research reports and publications and an appendix containing these twenty-five reports and publications. The objectives of this research were to develop theory and method for signal interception, with focus on the problems of signal detection, recognition, and geolocation, particularly for direct-sequence spread-spectrum phase-shift-keyed signals but also for frequency-hopped and hybrid frequency-hopped direct-sequence spread-spectrum signals. The approach taken was to develop a theory for exploiting the cyclostationarity properties exhibited by the signals of interest to obtain a generally applicable methodology that yields programmable algorithms that can meet evolving requirements. The period of research covered is primarily 1989-1991, with some background research from the period 1986-1988 included for reference.

INTRODUCTION

This report is a compilation of annual reports for research conducted during the three-year period 1989-1991, with some background research conducted during the preceding three-year period 1986-1988 included for reference. The objectives of this research were to develop theory and method for signal interception, with focus on the problems of signal detection, recognition, and geolocation, particularly for direct-sequence spread-spectrum phase-shift-keyed signals but also for frequency-hopped and hybrid frequency-hopped direct-sequence spread-spectrum signals. The approach taken was to develop a theory for exploiting the cyclostationarity properties exhibited by the signals of interest to obtain a generally applicable methodology that yields programmable algorithms that can meet evolving requirements. This approach to the research grew out of the relatively new theory of cyclostationary random signals developed by the author and reported in his two books published in 1985 and 1987.[1]

Much of the research described in the annual reports preceding this compilation has since been published in professional journals. Some of this published research was jointly supported by several government and industrial sponsors. This compilation incorporates the published versions in order to provide the broadest treatment. Funding sources are acknowledged in these published versions.

This report consists of an extensive overview of the research reported in twenty-five individual research reports and publications and an appendix containing these twenty-five reports and publications. The overview is organized into an introductory section and six major sections corresponding to six categories of research specified by the following section titles:

---

1 W. A. Gardner, *Introduction to Random Processes with Applications to Signals and Systems*, Macmillan, 1985 (second ed., McGraw-Hill, 1989); W. A. Gardner, *Statistical Spectral Analysis: A Nonprobabilistic Theory*, Prentice-Hall, 1987.

The contents of the Appendix are listed in the References section at the end of this overview.

1. OVERVIEW OF CYCLOSTATIONARITY AND ITS EXPLOITATION

Although most of the individual papers and reports summarized in the following six sections incorporate concise reviews of the relevant basic concepts and definitions in the theory of cyclostationary random signals, which is central to all the research described in this overview and compilation, these reviews do not provide adequate background for those who are not already familiar with this relatively new theory. Therefore, in order to make this compilation self-contained, the first paper [1] provides a tutorial introduction and survey of this theory, together with an overview of some of the various ways to exploit cyclostationarity for signal processing tasks, most of which are encountered in signal interception. A brief introduction to and summary of this paper follows.

Many conventional statistical signal processing methods treat random signals as if they were statistically stationary, in which case the parameters of the underlying physical mechanism that generates the signal would not vary with time. But for most manmade signals encountered in communication, telemetry, radar, and sonar systems, some parameters do vary periodically with time. In some cases even multiple incommensurate (not harmonically related) periodicities are involved. Examples include sinusoidal carriers in amplitude, phase, and frequency modulation systems, periodic keying of the amplitude, phase, or frequency in digital modulation systems, and periodic scanning in television, facsimile, and some radar systems. Although in some cases these periodicities can be ignored by signal processors, such as receivers which must detect the presence of signals of interest, estimate their parameters, and/or extract their messages, in many cases there is much to gain in terms of improvements in performance of these signal processors by recognizing and exploiting underlying periodicity. This typically requires that the random signal be modeled as *cyclostationary*, in which case the statistical parameters vary in time with single or multiple periodicities.

The paper [1] explains that the cyclostationarity attribute, as it is reflected in the periodicities of (second-order) moments of the signal, can be interpreted in terms of the property that enables generation of spectral lines from the signal by putting it through a (quadratic) nonlinear transformation. It also explains the fundamental link between the spectral-line generation property and the statistical property called *spectral correlation*, which corresponds to the correlation that exists between the random fluctuations of components of the signal residing in distinct spectral bands. The paper goes on to explain the effects on the spectral-correlation characteristics of some basic signal processing operations, such as filtering, product modulation, and time sampling. It is shown how to use these results to derive the spectral-correlation characteristics for various types of manmade signals.

Some examples of signals that can be appropriately modeled as cyclostationary can be interpreted as the response of a linear or nonlinear system with some periodically varying parameters to stationary random excitation. Specific examples include stationary random modulation of the amplitude, phase, or frequency of a sinewave; stationary random modulation of the amplitudes, widths, or positions of pulses in an otherwise periodic pulse train; periodically varying Doppler effect on a stationary random wave; and periodic sampling, multiplexing, or coding of stationary random data. In addition to these examples of manmade signals, some natural signals also exhibit cyclostationarity due, for example, to seasonal effects in time-series data sets obtained in meteorology, climatology, atmospheric science, oceanography, and hydrology, as well as astronomy. Numerous examples are given in the references in [1].

Finally, and most importantly, the paper describes some ways of exploiting the inherent spectral redundancy associated with spectral correlation to perform various signal processing tasks. These include detecting the presence of signals buried in noise and/or severely masked by interference; recognizing such corrupted signals according to modulation type; estimating parameters such as time-difference-of-arrival at two reception platforms and direction of arrival at a reception array on a single platform; blind-adaptive spatial filtering of signals impinging on a reception array; reduction of signal corruption due to cochannel interference and/or channel fading for single-receiver systems; linear periodically time-variant prediction; and identification of linear and nonlinear systems from input and output measurements. The descriptions include brief explanations of how and why the signal processors that exploit spectral redundancy can outperform their more conventional counterparts that ignore spectral redundancy or, equivalently, ignore cyclostationarity.

2. OVERVIEW OF RESEARCH ON SIGNAL DETECTION

a. Introduction

Interception of communications is attempted for a variety of reasons including reconnaissance, surveillance, and other intelligence gathering activities, as well as position fixing, identification, and communications jamming. For example, an aircraft might attempt to intercept the communications between a submarine or ship and a satellite, or a satellite might attempt to intercept ground-to-ground communications. Typically, the interceptor has knowledge of no more than the communicator's frequency band, modulation format, and modulation characteristics such as bandwidth and hop rate or chip rate. In the past, it was commonly held that the most appropriate approaches to the detection task for signal interception must be based on radiometry, that is, measurement of received energy in selected time and frequency intervals. However, it is commonly recognized that such radiometric methods can be highly susceptible to unknown and changing noise levels and interference activity. There have been many proposals for methods of countering such complications, including various approaches to adjusting or adapting threshold levels, and adaptive filtering, cancelling, and directional nulling of interfering signals. But these problems remain as the most serious impediment to signal detection and other signal interception tasks.

The radiometric approach to detection is based on the use of stationary random processes as models for the signals to be intercepted. However, for the purposes of signal interception, the signal of interest is more appropriately modeled as a *cyclostationary* random process, that is, a random process whose probabilistic or statistical parameters vary periodically with time. The message contained in the modulated signal is unknown, and is usually modeled as a stationary random process (discrete time or continuous time). This stationarity coupled with the periodicity of sine wave carriers, pulse trains, repeating spreading codes, etc., results in a cyclostationary model for the signal. However, these cyclostationary signals typically do not exhibit spectral lines because the spectral lines of the unmodulated carriers and/or pulse trains are spread out over relatively broad bands by the stationary random modulation.

Communication signals have traditionally been modeled as stationary random processes. Although communication signals typically involve one or more periodicities underlying their random fluctuations, due to sine-wave carriers and repetitive pulsing or keying, by introducing random phase variables uniformly distributed over one period of each periodicity, a stationary model can be obtained. On the surface this seems appropriate if the receiver has no knowledge of carrier phase or clock timing and is, therefore, unsynchronized to the periodicities. However, by looking beneath the surface we find that carrier and clock periodicity can indeed by exploited by receivers that have no knowledge of the associated phases, and that make no attempt to estimate those phases. Superior detection performance can be obtained by exploiting the single frequency of some harmonic of a periodicity, such as a doubled carrier frequency, or a keying rate. This is partially analogous to the common procedure of detecting the presence of an additive sine wave in noise by measuring the magnitude of the correlation of the noisy data with a complex sine wave of arbitrary phase. But, for detection of communication signals, the problem is more subtle since such signals rarely contain additive sine-wave components (which result in spectral lines). The spectral lines must be regenerated by using a quadratic (or in some cases higher-order) nonlinear transformation.

Motivated by the potential utility of exploiting cyclostationarity for signal detection, we have developed a general theory of weak random signal detection within the unifying framework of the spectral correlation theory of cyclostationary signals. This theory clearly reveals the relationships among the variety of detectors that have been proposed, or are in the development stage, or are in use. The performance studies we have conducted illustrate the superiority of *cyclic-feature* detectors over energy detectors for accommodating the problems associated with unknown and changing noise levels and interference activity. Cyclic features result from the characteristic property of cyclostationarity called *regenerative periodicity*, which means that spectral lines can be regenerated from the signal with the use of appropriate nonlinear transformations.

b. Research Findings

In this research, we have used the theory of cyclostationary signals to derive optimum detectors using several criteria for optimality. We have developed practical implementations, analytically studied their output signal-to-noise-ratio performance, and conducted simulations to evaluate their probability-of-detection performance in terms of receiver operating characteristics.

i) Unifying Theoretical Framework for Cyclic-Feature Detection

In the signal detection techniques we have studied, the underlying periodicities of communication signals are exploited by quadratically processing the received data. Certain simple quadratic processors are widely used to regenerate sinewaves; for instance, a squarer is often used to regenerate a sinewave with frequency equal to the doubled carrier, and a delay-and-multiply device with delay equal to half of the chip interval can be used to regenerate a spectral line at the chip rate. By using the mathematical framework of cyclostationarity, we have characterized all spectral lines that can be regenerated and solved for the quadratic transformation of the data that yields the strongest possible spectral line at a given frequency. Thus, we have shown that the theory of cyclostationarity provides a rigorous framework for understanding ad hoc spectral-line generators, as well as allowing for the specification of optimal spectral-line generators, which are more general than squarers and delay and multipliers. More generally, we have developed in [2] a unifying approach to the design and analysis of (single-sensor) quadratic detectors for signal interception. The methods of detection that are incorporated in this unification include radiometry, delay-and-multiply chip-rate feature detection, filter-and-square carrier feature detection, ambiguity-plane feature detection, Wigner-Ville time-frequency plane feature detection, spectral-correlation plane feature detection, likelihood-ratio detection for weak signals, maximum-deflection, and optimum spectral-line regeneration detection. The unification reveals the fundamental role that spectral correlation and spectral-line regeneration play in the signal interception problem, and shows how to modify various ad hoc techniques to make them optimum. It also suggests the spectral-correlation plane approach as a general approach to interception that offers great flexibility as well as inherent tolerance to one of the most challenging problems in interception, namely, accommodating unknown and changing noise level and interference activity.

ii) Performance Evaluation of Cyclic-Feature Detectors

The purpose of the research reported in [3] is to quantify the gains in detection performance that are attainable by exploiting the underlying periodicity that is present in communication signals. More specifically, we consider phase-shift-keyed signals, which are used for some spread-spectrum communications, and analytically evaluate output signal-to-noise ratio and determine by simulation receiver operating characteristics for operation in a relatively strong white-Gaussian-noise (WGN) background with random fluctuation in the noise power-level. In particular, we consider a nominal in-band SNR of 0 dB (total SNR of –9 dB), corresponding to the expected noise level, and we consider a coefficient of variation (variance normalized by squared mean) of 10% for the random fluctuation about the expected noise level. In a receiving system that monitors the noise level and adjusts the detection threshold accordingly, this 10% variation represents the error in noise-level tracking. This variation can also represent nonstationary activity of broadband interfering signals. In addition, a highly variable narrowband interference background with SIR of –10 dB and fixed-level WGN with SNR = 10 dB is also considered. The detectors whose performances are compared are the standard optimum radiometer, which is the weak-signal likelihood-ratio detector for the stationary model of the signal, a modified radiometer that provides the weak-signal joint maximum-likelihood estimate of the noise power-level and detection of the signal (modeled as stationary), and two maximum-SNR cyclic-feature detectors (or simply cycle detectors) one of which exploits the doubled carrier frequency and the other of which exploits the keying rate. It is shown by simulation and analysis that cycle detectors (for PSK signals) can outperform the radiometer and the joint signal detection and noise-level estimation radiometric detector when the background noise or interference is variable. Also, differences between our exact results on output SNR and previously published approximations for PSK signals are pointed out.

iii) Simplifications of the Maximum-Likelihood Detector for Weak Cyclostationary Signals

In this part of our research, we generalize the weak-signal maximum-likelihood detector described in [2] from single-antenna reception to dual-antenna reception, for which signal detection and joint time-difference-of-arrival estimation are performed [4]. We also carry out extensive simulations of the maximum-likelihood processors and numerous suboptimum but more practical partial implementations of the relatively complex maximum-likelihood processors, including multi-cycle as well as single-cycle detection of PSK signals. Both constant and variable noise levels are simulated.

iv) Robust Cyclic-Feature Detection

The purpose of this part of our research is to study the robustness of the quadratic detector consisting of a filter followed by a delay-and-multiply device that multiplies the filtered signal by a delayed and conjugated replica of itself. This delay-and-multiply (DM) detector is commonly used to regenerate a spectral line at a frequency equal to the pulse rate of a PCM signal for purposes of detection of the signal's presence or synchronization to the signal's pulse timing phase. Optimizing the DM detector with respect to the filter bandwidth and the delay requires knowledge of the pulse rate and carrier frequency. When the pulse rate or carrier frequency assumed for optimization is in error, the bandwidth and/or delay that maximizes detection performance for the erroneous pulse rate or carrier frequency is suboptimum for the actual pulse rate and carrier frequency. This suggests maximizing the robustness of the detector, as measured by the degree of tolerance to error in pulse rate and carrier frequency. In this research, we have maximized the robustness numerically, and we have used both output SNR and receiver operating characteristics as measures of performance for studying the robustness of the DM detector for baseband PCM signals, modeled as real PAM, and passband PCM signals (digital QAM, BPSK, QPSK), modeled as complex PAM, in white Gaussian noise (WGN). Since the primary application motivating this work is detection of spread-spectrum signals buried in noise, only low-SNR conditions are considered. To put the DM detector in perspective, it is pointed out that the weak-signal likelihood-ratio detector for PAM in WGN is a matched-filter followed by a magnitude squarer.

This detector also maximizes the output SNR for the regenerated spectral line at the pulse-rate frequency. This detector is the special case of the DM detector for which the prefilter is the matched filter and the delay is zero. This detector, however, is shown to be not as robust to errors in pulse rate and carrier frequency as that which is optimized for robustness.

v) Detection Using the Cyclic Spectrum

As explained in [2] and expanded on in [8] – [12] (and as discussed in the following Section 3), the cyclic spectrum (or spectral correlation function) provides a powerful and general means for signal detection and other interception tasks such as parameter estimation and modulation recognition. Since the cyclic spectrum is, ideally, the Fourier transform of the cyclic autocorrelation, which is in practice equivalent to the radar ambiguity function, the objective of this part of our research was to compare and contrast the cyclic features that appear in the cyclic spectrum and radar ambiguity surfaces. Some discussion of the results is given in the following Section 3.

vi) Multi-Hop Chip-Rate Detection for Hybrid FH/DS Signals

The most straightforward approach to detection of a hybrid frequency-hopped direct-sequence (FH/DS) spread-spectrum signal is to simply use a wideband delay-and-multiply chip-rate detector. However, the large number of inactive channels present in the wide passband results in severe noise-cross-noise product terms in the output, which reduce the output SNR by the factor $(1/M)^2$ for M channels. We have devised an attractive alternative algorithm that collects over the multiple channels after quadratic processing rather than before. This algorithm results in an SNR-degradation factor (relative to the SNR of a non-hopped DS signal) of only 1/M, which is the square root of that which results from the straightforward wideband delay-and-multiply detector. This new multi-hop chip-rate detector has been implemented in software and initial simulations have been conducted. A brief report of the SNR analysis and some receiver operating characteristics obtained from initial simulations are given in [7].

c. *Conclusions*

In this research, we have developed a unifying theoretical framework for second-order cyclic feature detection by characterizing a variety of ad hoc and optimum detectors explicitly in terms of the spectral correlation function (or cyclic spectrum) or its inverse Fourier transform, the cyclic autocorrelation function. We also have analytically evaluated output SNR and obtained receiver operating characteristics by simulation for various detectors and various types of DS PSK signals, in various noise and interference environments. We conclude that single- and multi-cycle detectors show great promise for robust detection of specific DS signals, and that the spectral correlation analyzer or, equivalently, cyclic spectrum analyzer (cf. Section 3), shows great promise for signal search and recognition. Future work should extend the catalog (cf. references in [1]) of spectral correlation functions to include signal types of more recent interest, such as hybrid FH/DS types, and should use these analytical characterizations to evaluate detectability of these newer signals. Continuing work is also needed in the development of algorithms for automatic modulation recognition using the cyclic spectrum and for general-purpose cyclic spectrum (spectral correlation) analysis, and in the development of digital implementations of cyclic spectrum analyzers (see Section 3). Finally, continuing work on developing the theory for and design of detectors that exploit higher-than-second-order cyclic features is needed (cf. Section 7).

3. OVERVIEW OF RESEARCH ON CYCLIC SPECTRAL ANALYSIS (for Detection, Recognition, and Parameter Estimation)

a. Introduction

As explained in [1], the cyclic spectral density function for a random signal process (abbreviated to *cyclic spectrum*) is one and the same as the spectral correlation density function (abbreviated to *spectral correlation function*). The spectral correlation concept has associated with it four fundamental properties of processes that are of significant practical value: 1) different types of modulated signals (such as BPSK, QPSK, and SQPSK) that have identical power spectral density functions can have highly distinct spectral correlation functions; 2) stationary noise and interference exhibit no spectral correlation (the spectral correlation function is identically zero); 3) the spectral correlation function contains phase and frequency information related to timing parameters in modulated signals; and 4) the existence of spectral correlation in a signal means that some possibly corrupted spectral components can be estimated using other correlated but less corrupted spectral components of the signal. Furthermore, all optimum signal processors (such as estimators and detectors) that are specified in terms of the cyclostationary autocorrelation function can be interpreted and implemented in terms of the spectral correlation function. For example, the low-SNR likelihood ratio for detection and the low-SNR likelihood function for parameter estimation for random signals in white Gaussian noise are, in general, completely specified by the signal autocorrelation function. Similarly, the optimum periodic or multiply-periodic time-variant filter for a cyclostationary signal in stationary noise is completely specified by the signal and noise autocorrelation functions. Also, the optimum quadratic transformation for generation of maximum-SNR spectral lines from a cyclostationary process for purposes of detection and synchronization is completely specified by the spectral correlation function for the signal and the power spectral density for the noise. Consequently, properties 1) - 4) can be exploited for detection, classification, parameter estimation, and extraction of signals buried in noise and further masked by interference. Although the spectral correlation function is a second-order (quadratic) statistic or probabilistic parameter like the power spectral density function, these four properties enable it to be used to accomplish tasks that are impossible to accomplish with the power spectral density function. This includes synchronization and noise and interference rejection for detection, recognition, parameter estimation, and signal extraction. For example, modulated signals that are severely masked by other interfering signals as well as noise can, in some applications, be more effectively detected by detection of spectral correlation rather than detection of energy. This is so, for instance, when the energy level of the background noise or interference fluctuates unpredictably (as explained in Section 1). Also, there are situations in which cyclostationarity is problematic, rather than being a property to be exploited. For example, nonlinearities in transmission systems (e.g., traveling-wave tube amplifiers and noise limiters) can inadvertently generate spectral lines from cyclostationary signals, and these spectral lines can cause severe interference effects. As explained in [1], the characteristics of spectral lines (their frequencies, phases, and strengths) that can be generated with quadratic nonlinearities are determined by the spectral correlation function. Consequently, the spectral correlation function or cyclic spectrum plays a central role in various signal detection, analysis, and extraction problems.

The need for computationally efficient algorithms for cyclic spectral analysis is becoming increasingly evident as cyclic spectral analysis grows in importance as a signal analysis tool. For many signal analysis problems the computational complexity of cyclic spectral analysis far exceeds that of conventional spectral analysis. The reason for the computational complexity of cyclic spectral analysis lies in the nature of the estimation problem. Essentially, cyclic spectral analysis algorithms estimate the correlation between spectral components of signals. It is the potentially large number of correlation computations, rather than the computation of the spectral components, that makes cyclic spectral analysis computationally complex.

Several computationally efficient cyclic spectral analysis algorithms have evolved from the original methods introduced in [8]. Cyclic spectral analysis algorithms generally fall into two classes: those that average in frequency (frequency smoothing) and those that average in time (time smoothing). Although both classes of algorithms produce similar approximations to the cyclic spectrum, time smoothing algorithms are more computationally efficient for general cyclic spectral analysis. Frequency smoothing algorithms can be computationally superior to time smoothing algorithms in certain restricted cases, e.g., for estimation of the cyclic spectrum for a relatively small number of values of cycle frequency or estimating the cyclic spectrum for relatively small values of the time-frequency resolution product, or when one-bit multipliers are used for the correlation operation [11], [12]. Combinations of time- and frequency-smoothing can also be used to meet special requirements.

b. Research Findings

In this research, we have developed a variety of methods for doing cyclic spectral analysis or, equivalently, spectral correlation analysis. The first work [8] lays the conceptual and mathematical foundation for all of the basic methods of spectral correlation measurement. The next contribution [9] presents the major algorithms we have developed for discrete-time (digital) implementation. Subsequent work [10] - [12] presents in-depth performance analyses and more detailed discussions of implementation issues.

i) Spectral Correlation Measurement

Various methods for measurement/computation of spectral correlation functions are presented in [8] in a unifying theoretical framework, and a general analysis of the resolution and reliability properties of these methods is carried out. Some of these methods are amenable to digital hardware or software implementations, others are amenable to analog electrical or optical implementations, and other implementation types used for conventional spectral analysis also are possible. It is explained that because of novel problems of computational complexity, cycle phasing, cycle leakage and aliasing, and cycle resolution, the measurement/computation of spectral correlation functions is more challenging than is the measurement/computation of the conventional spectral density function, and also the conventional cross-spectral density function for stationary data. Nevertheless, it is concluded that empirical spectral correlation analysis is indeed practically feasible in many applications. It is also explained that with sufficient averaging the spectral correlation function for a signal can be accurately measured, even when the signal is buried in noise or masked by interference, and that this contrasts with the fact that noise and interference have an unremovable (in general) masking effect on the measured power spectral density of a signal.

ii) Computationally Efficient Algorithms for Cyclic Spectral Analysis—An Overview

With computational efficiency for general cyclic spectral analysis as the primary motivation, the signal processing, computational, and structural properties of time-smoothing algorithms are described in [9]. Several representations of the basic time smoothing algorithm are developed there, and the transformed kernel representation, which is perhaps the most important characterization of cyclic spectral analysis algorithms due to its generality and conceptual utility, is described. Using this representation, the signal processing attributes of an algorithm, such as frequency and cycle frequency resolution, and phenomena such as cycle leakage and aliasing can be studied. Additionally, the computational complexity of cyclic spectral analysis algorithms is readily determined using this representation. Starting from the basic time smoothing algorithms, several improved algorithms are developed in [9]. Modifications to reduce the computational complexity of the basic algorithm include decimating the complex demodulates prior to forming product sequences and time averaging the frequency-shifted product sequences with FFTs. With the addition of an FFT-based input channelizer, the computationally efficient FFT Accumulation Method is obtained. Another computationally efficient algorithm, the Strip Spectral Correlation Algorithm, is also developed as an alternative to FAM. Both of these algorithms have highly parallel structures and are readily implemented on general purpose computers or, if execution time is critical, specialized multiprocessor signal analyzers.

iii) Digital Implementations of Spectral Correlation Analyzers—An In-Depth Analysis

In [10], the general quadratic system representation for spectral correlation analyzers based on conventional cross spectral analysis techniques is studied. The analyzer tuning and resolution parameters and cycle frequency leakage performance are shown to be simply related to the system kernel transform. An expression for the output signal-to-noise ratio for Gaussian input waveforms is given in terms of the kernel transform and the spectral correlation function of the input waveform. The kernel transforms associated with alternative configurations that are conducive to efficient digital implementation are studied, and a detailed hardware-complexity analysis for the frequency-smoothing method and two fast time-smoothing methods is given. The most straightforward implementation, which is obtained from the frequency-smoothing method that uses an FFT of size N, is found to have a complexity on $o(N^2)$ for N time-samples of data, and would require a parallelism factor of $\log_2(N)$ complex butterflies and N/4 complex correlators in order to be computed in real time. Or, viewed as a function of $\Delta t \Delta f$ (where Δt is the collect time and Δf is the spectral resolution) for fixed Δf, the complexity is on $o((\Delta t \Delta f)^2)$. The two computationally-efficient structures based on time smoothing, which use an input FFT of size N/M and an output FFT of size M/4, have complexity on $o((N^2/M)\log_2(N))$. Furthermore, when the complexity is viewed as a function of $\Delta t \Delta f$, for fixed Δf, the complexity is on $o((\Delta t \Delta f)\log_2(\Delta t \Delta f))$. A specific example involving detecting the presence of a cyclostationary signal buried in noise is discussed.

iv) One-Bit Spectral Correlation Algorithms

Direct implementations of cyclic spectrum analyzers can be computationally intensive, resulting in substantial hardware complexity. In [11], in an attempt to simplify the hardware, a one-bit (the sign bit) multiplier is introduced into the correlation operation in each of the two basic algorithms for cyclic spectral analysis (or spectral correlation analysis). For the algorithm that produces frequency-averaged estimates, this results in substantial hardware simplifications while producing very accurate estimates of the cyclic spectrum. For the algorithm that produces time-averaged estimates, the hardware simplifications are more moderate and some accuracy in the estimates is sacrificed. It is concluded that the one-bit spectral-correlation algorithm for producing frequency-averaged estimates shows great promise for the realization of practical spectral-correlation analyzers. In [12], partitioning schemes for the one-bit spectral correlation algorithm based on frequency smoothing are presented and used to develop a system architecture that uses parallel computational structures.

v) Comparison of the Spectral Correlation and Ambiguity Surfaces

The usefulness of the radar ambiguity function (AF) for determining the accuracy of both detection of radar signals and estimation of their parameters has long been established. Ambiguity is a function of two target-related variables: the range rate (or Doppler) which is related to the target velocity, and the range which represents the distance to the target. The AF bears a close relationship to the *cyclic autocorrelation function* [1] which arises naturally in more general signal processing contexts. The cyclic autocorrelation is also a function of two variables: the *cycle frequency* which is somewhat analogous to the Doppler, and the *lag* which is analogous to the radar range. The Fourier transform of the cyclic autocorrelation can be shown to be a kind of *spectral correlation function* (SCF), and is also called the *cyclic spectrum* by analogy with the nomenclature of the Wiener relation for stationary random processes, for which the spectrum is the Fourier transform of the autocorrelation [1]. SCF measurements have been shown to be useful for the signal processing problems of signal parameter estimation (e.g., time-difference-of-arrival, and carrier and pulse frequency and phase), and signal detection and recognition [1].

Since the AF and the SCF are nearly a Fourier transform pair, and much work has focused on exploiting various properties of the AF, it is well to ask which function better differentiates between signal and noise under various conditions. In the report [13], we endeavor to answer this question quantitatively using both auto (one sensor) and cross (two sensor) versions of the AF and SCF by comparing the auto AF to the auto SCF and the cross AF to the cross SCF. The environment considered consists of signals that exhibit spectral correlation and broadband background noise that does not exhibit spectral correlation (stationary noise).

We investigate the performance of the AF and SCF surfaces in terms of their peak SNRs. The general formulas derived are valid for any weak cyclostationary signal, and as such are not especially useful for insight into performance comparisons. However, for the case of PSK signals, these formulas reveal that the relative performance of the two surfaces is strongly dependent on the baud rate and carrier frequency of the signal as well as the spectral resolution, reception bandwidth, and collect time of the measurement. The dependence of the relative feature-strengths on these four parameters can be summarized by the following three dimensionless quantities: the band-position ratio (BPR), which is the ratio of the receiver bandwidth to the carrier frequency; the fractional bandwidth (FB), which is the ratio of the signal bandwidth to the receiver bandwidth; and the fractional resolution (FR), which is the ratio of the smoothing window width to the signal bandwidth.

The spectral correlation surface performance is superior to the ambiguity surface performance if the FB is sufficiently small (corresponding to exploitation of the baud-rate cycle frequency) or if the BPR is sufficiently small (corresponding to exploitation of the doubled-carrier cycle frequency). In general, the larger the FR, the easier it will be to meet these conditions. For applications in which it is desired to intercept signals with unknown bandwidth and/or carrier frequency, or if an entire AF or SCF surface must be produced over large intervals of cycle frequency, spectral frequency, or lag, it is likely that the SCF and cross SCF will be the surfaces of choice. For applications that need only a single *line* of the surface (for a single cycle frequency), the various SNRs must be evaluated using the appropriate signal and measurement parameters in order to make the proper choice.

c. Conclusions

In this research, we have developed a unifying theoretical framework for cyclic spectral analysis or, equivalently, spectral correlation analysis, and we have analytically characterized the resolution and reliability properties of the basic methods. We have developed computationally efficient algorithms for digital implementation, and we have analyzed the performance of these algorithms. SSPI has developed a commercially available software package that rapidly computes and displays spectral correlation surfaces for moderate numbers of cycle frequencies by making efficient use of computational, memory, and disk-storage resources. Continuing work is directed at the development of hardware implementations and of algorithms for processing the spectral correlation surface for automatic modulation recognition.

4. OVERVIEW OF RESEARCH ON SOURCE LOCATION BY TIME-DIFFERENCE-OF-ARRIVAL ESTIMATION

a. Introduction

The problem of passively locating the source of a propagating radio signal received at multiple platforms has a number of applications including direction finding for navigation by land, air, or sea, tracking of moving emitters for surveillance (e.g., for research, law enforcement, or national security), monitoring and locating illegal and/or hostile communicators (e.g., contraband, violators of communications regulations, or enemies in a battlefield), military reconnaissance and intelligence, and so on. Unlike active source-location systems, including some Radar, Sonar, and Lidar, which transmit signals and then process the received reflections from the objects to be located, passive systems simply process whatever signals are emitted from the object to be located.

In many of these passive source-location applications, the receivers in the source-location system are subject to a variety of types of interference and noise, including natural and manmade signals other than the signal of interest (SOI). These signals not of interest (SNOI) can severely degrade the performance of conventional source-location systems when they are present at the same time and also occupy the same spectral band as the SOI. This can be particularly problematic when the SOI is weak relative to the SNOI and/or noise. The problem is further exacerbated when the locations of the sources of the SNOI are unknown and/or close to that of the SOI.

One of the objectives of this research was to study a new class of multiple-platform passive source-location methods that exploit inherent signal properties, collectively called *cyclostationarity*, that are characteristic of radio signals used for communications and telemetry to obtain substantial tolerance to all types of interference and noise (except possibly for some interfering signals that are intentionally designed to be of the same type, e.g., in communication networks). Like most conventional methods that require some degree of tolerance to noise, the new methods are based on crosscorrelation of time-shifted measurements of data from multiple receivers. However, unlike conventional methods, the new methods crosscorrelate frequency-shifted as well as time-shifted versions of the received data in order to exploit the unique cyclostationarity properties of the SOI.

(Although cross-ambiguity methods, which compensate for frequency-difference-of-arrival due to Doppler effects, do crosscorrelate frequency-shifted data, the frequency shifts are relatively small since they correspond to Doppler shifts, whereas the frequency shifts used in the new methods are much larger than typical Doppler shifts and accomplish an entirely different task.)

Within the general class of source-location methods that are based on crosscorrelation measurements, there are two distinct subclasses: There are those methods that are designed for an array of closely spaced antenna elements on a single platform (e.g., a ship, aircraft, satellite, ground vehicle, or fixed ground station), for which the element spacing is typically less than half a wavelength for all signals received and phase-alignment methods for beam/null steering are employed; and there are those methods that are designed for two or three widely spaced antenna elements, each often (but not always) on a separate platform, where time-difference measurements are used to obtain location information. Whereas the former class of methods (often called *direction finding methods*) exploits phase differences (less than π radians in order to avoid ambiguities), from element to element, of relatively narrowband signals (or wideband signals decomposed into narrowband components) to estimate angle of arrival, the latter class of methods (often called *time-difference location methods*) exploits relative time differences (the larger the better since root-mean-squared-error of location is approximately inversely proportional to the separation between platforms) of preferably wideband signals from platform to platform to estimate location (both angle of arrival and range). In practice, however, both methods can be applied to all bandwidths used in typical communications and telemetry systems.

The requirements on accuracy and spatial resolution capabilities of array-based methods become more stringent as the distance between each source to be located and the reception platform increases, since this decreases differences between angles of arrival at the array. In contrast, the requirements on accuracy and temporal resolution capabilities of time-difference-of-arrival (TDOA)-based methods become less stringent as the separation between reception platforms increases, since this increases differences between times of arrival.

The need for high resolution arises primarily when (relatively) closely spaced multiple sources give rise to multiple received signals that cannot be separated by preprocessing (prior to processing for location). For instance, TDOAs of multiple signals that are not separated by more than the widths of their crosscorrelation peaks (whose locations on the time-delay axis correspond to the TDOAs) usually cannot be resolved by conventional TDOA-based methods. To minimize this problem, the distance between platforms is typically made as large as is practically feasible so that the magnitudes of the TDOAs (which are proportional to the distance between platforms) will be as large as possible, thereby minimizing the overlap of adjacent peaks (whose widths depend only on the signal bandwidth—being inversely proportional—not on the distance between platforms).

The best-performing array-based methods attempt to circumvent this resolution problem in locating multiple signal sources by *simultaneously* estimating multiple angles of arrival rather than individually estimating the angle of arrival of each signal (as commonly done by conventional beamformers and TDOA-based methods) [22]. Although the array-based methods offer the advantage of high spatial resolution (with respect to the spatial extent of the source-location system—the array size), and the ability to simultaneously locate a number of signals up to one less than the number of elements in the array, their complexity is typically much higher than that of the time-difference-of-arrival-based methods because of the need for measurement, storage, and usage of large amounts of array calibration data (the recently proposed ESPRIT method is an exception because of its special array design), and because of the use of computationally intensive algorithms: the best-performing algorithms require singular value decomposition of crosscorrelation matrices of possibly high dimension, and/or require the solution of a multidimensional optimization problem.

The new cyclostationarity-exploiting TDOA-based methods we have introduced and studied alter this tradeoff between highly sophisticated high-resolution array-based methods and the relatively simple TDOA-based methods that require widely separated multiple platforms by eliminating the resolution-limitation problem for TDOA in many types of environments. That is, because the spectral correlations used make the new methods signal selective (in spite of temporal and spectral overlap between the SOI and SNOI), they often do not need to resolve multiple crosscorrelation peaks along the time-delay axis. As a result of this novel signal selectivity, the requirements on separation of platforms can be eased considerably. Also, the problem of deciding which peak among multiple crosscorrelation peaks is associated with the SOI, and which with the SNOI, is eliminated. The only requirement of the new methods is that they must know a carrier frequency, or keying rate, or possibly some other *cycle frequency* associated with the cyclostationarity of the random SOI, although such cycle frequencies can be estimated using the same data to be used for TDOA estimation. The cycle frequencies are the frequency shifts used to obtain signal selectivity.

Unfortunately, when multiple SOIs sharing the same cycle frequencies arrive at each receiver, the TDOA-resolution problem remains. In this most difficult situation, although the high-resolution array based methods can be used, an attractive alternative that avoids the need for calibrating the array and can provide more accurate estimates is to use an array on each of two (or more) platforms to separate the multiple SOIs by blind-adaptive spatial filtering using our new spectral-correlation-restoring (SCORE) algorithms (which require no calibration) [26], and then apply TDOA-based methods to the separated signals. (This is discussed further in Section 5.)

Motivated by the potential utility of the cyclostationarity-exploiting TDOA-based approach to emitter location, we have developed various algorithms to perform signal-selective TDOA estimation and we have evaluated their performance. The results obtained indicate that very substantial improvements in the performance of source-location systems can be obtained by exploiting cyclostationarity.

b. Research Findings

In this research, our primary effort has been directed toward applying the least-squares approach to the problem of exploiting second-order cyclostationarity for obtaining signal-selective methods of TDOA estimation. In the first part [14] of a two-part paper, we have derived a number of methods and corresponding algorithms by performing least-squares fitting on measurements of cyclic correlations or cyclic spectra (spectral correlations). In [16], we have shown that some of these algorithms are partial implementations of the weak-signal maximum-likelihood TDOA estimator for signals in white Gaussian noise. In the second part [15] of a two-part paper, we have evaluated the RMSE performance of a number of new algorithms in various signal and noise/interference environments by simulation, generalized several of the algorithms for multiple platform-pairs, studied the problem of implementing these algorithms using FFTs, and investigated the effects of some departures from the idealized models used to derive the algorithms. In other reports described in the following subsections, we also have compared and contrasted the performance characteristics of algorithms that use only cyclic autocorrelation measurements with those that also use cyclic crosscorrelation measurements, initiated the study of higher-order cyclostationarity for TDOA estimation involving signals with weak second-order cyclostationarity, developed algorithms for signal-selective joint TDOA-FDOA estimation, and compared and contrasted the performance of DF methods and TDOA methods.

i) The Least Squares Approach to Exploiting Cyclostationarity for TDOA Estimation

Using idealized models for the interference/noise-corrupted signals received at two platforms, we calculated the cyclic correlations and cyclic spectra to reveal their dependence on the TDOA. We then used these results to set up various ways of minimizing sums of squared errors between measurements of the cyclic correlations (or cyclic spectra). This least-squares approach led to the SPECCORR, SPECCOA, SPECCON, CLP and CPD algorithms. Derivations and performance evaluations are reported in the two-part paper [14] - [15].

ii) The Maximum-Likelihood Approach to TDOA Estimation

By studying the maximum-likelihood joint signal-detector and TDOA-estimator for a weak signal in white Gaussian noise, we found that the SPECCOA algorithm is a partial but much simpler implementation of the relatively complicated weak-signal maximum-likelihood TDOA estimator. Similarly, the relatively simple single-receiver spectral-correlation-exploiting detector was shown in [2] to be a partial implementation of the single-receiver weak-signal maximum-likelihood detector. Generalizing on this, we have shown in [4] that the weak-signal maximum-likelihood joint detector and TDOA estimator is made up of a multiplicity of components that can be partitioned into a variety of detectors and estimators, each of which is optimal in either a maximum-likelihood sense or a maximum-SNR sense. The results of extensive simulations that compare the performances of some of these alternative detectors and estimators were then obtained and it was concluded that the performance of the relatively complex maximum-likelihood joint detector/estimator, which requires a two- or three-dimensional search over phase and timing parameters (for PSK signals, these include the carrier phase [for BPSK only], chip timing, and time-difference-of-arrival), can be closely approximated in many cases (for data record lengths that are normally used in weak-signal detection and source location) with much simpler suboptimum detectors/estimators that avoid all parameter searches for detection, and reduce the search to one dimension for TDOA estimation.

iii) Performance of Signal-Selective TDOA Estimators

Because of the considerable difficulty in obtaining (non-asymptotic) analytical results on performance (RMSE) of TDOA estimators for the highly corruptive environments of interest in this study, we chose to use simulations to evaluate RMSE for a variety of signal and noise/interference environments. In [15], we graphed statistical samples of the TDOA estimation functions whose peak-locations provide the TDOA estimates. This clearly revealed the noise and interference rejection capabilities of our signal-selective algorithms, and provided a qualitative assessment of their performance. We then averaged over many Monte Carlo trials to obtain graphs of RMSE versus averaging time in order to quantify performance. The results presented in [14] show that new signal-selective algorithms that use the keying-rate cycle frequency exhibit excellent robustness for BPSK signals in a wide range of interference and noise environments and operating conditions. (Even better performance results from using the doubled-carrier cycle frequency.) Results not reported show the same level of performance for QPSK signals when the keying-rate frequency is used, and experimentation with other signal types such as FSK suggests that comparable performance is attainable in those cases where the signals of interest exhibit substantial spectral correlation. These new algorithms are tolerant to both interfering signals and noise, and they can outperform conventional algorithms that achieve the Cramer-Rao lower bound on variance for stationary signals because the signals considered here are nonstationary (cyclostationary) and the algorithms exploit the nonstationarity to discriminate against noise and interference. The most important issue regarding signal type and corresponding performance that has been discovered so far is the strength of the signal's cyclic feature (or spectral correlation feature) to be exploited. For example, bandwidth-efficient digital signals can have relatively weak keying-rate features in the sense that the spectral coherence function is close to unity over only a relatively small band. This limits the band over which the linear phase-vs.-frequency characteristic of the spectral correlation functions can be used, thereby limiting reliability of the estimate of the slope of this line. Although this can in principle always be compensated for by increasing the collection time, there are of course practical limits to this.

In connection with this, we have also studied in [17] the performance degradation that results from narrowband interference excision and from using only part of the signal band in order to reduce the required bandwidth of the crosslink between platforms. We also have studied in [17] the need for and performance of time-interpolation within TDOA algorithms for close platforms (or a single platform containing both sensors), the appropriate width for the spectral smoothing window used for the spectral correlation estimates employed in the TDOA algorithms, and the relative performances of a variety of TDOA algorithms described in [14] (not all of which are evaluated in [15]).

iv) Comparison of Auto- and Cross-Correlation Methods for Signal-Selective TDOA and FDOA Estimation

The CPD method of TDOA estimation, which uses the cyclic autocorrelation (or cyclic auto-spectrum) and the SPECCOA method, which uses the cyclic cross correlation (or cyclic cross-spectrum) are compared in [18]. The CPD has the advantage of requiring only a very low-capacity cross-link, whereas the SPECCOA method requires a cross-link with bandwidth equal to that of the signal. However, the CPD algorithm produces ambiguous TDOA estimates because these estimates are given modulo $1/\alpha$, where α is the cycle frequency (e.g., chip rate). The CPD method also produces considerably larger RMSE (e.g., 10 dB, that is, 20 dB in MSE, larger for collects ranging from 1024 to 8192 for a BPSK signal in noise with 0 dB SNR). Nevertheless, this ambiguity and reduced accuracy might be acceptable when the capacity of the cross-link must be small. It is also shown in [18] that for both methods studied, RMSE of TDOA is independent of the distance between sensors.

When the available bandwidth for communication between two platforms is severely limited, TDOA and/or FDOA algorithms that require cross-correlation or cross-spectrum (cyclic or conventional) measurements cannot be used at all. In this case cyclic autocorrelation (or cyclic autospectrum) measurements can be used to estimate cycle phases and frequencies (e.g., chip-rate clock phases and carrier frequencies), and these can be used to obtain estimates of not only TDOA, as with the CPD algorithm, but also FDOA in the case of moving platforms or moving sources. The FDOA can be used to help resolve the ambiguity in TDOA by intersecting the family of iso-TDOA loci with iso-FDOA loci. In [19], several cyclic-autocorrelation algorithms for estimating TDOA and FDOA are derived, and relative strengths and weaknesses are briefly discussed.

v) Implementation and Sources of Degradation for TDOA Estimation

We have developed algorithms for implementing all methods of TDOA estimation that we have introduced. These algorithms, reported in [15], utilize FFTs to obtain computational efficiency and they use the spectral-smoothing method of obtaining reliable estimates of spectral correlation functions. Sources of degradation that were studied and reported on in [15] include inappropriate spectral-smoothing window-width (cf. [17]), insufficient averaging time, error in knowledge of cycle frequency, and propagation-channel mismatch. It was found that the amount of data often needed to obtain acceptable TDOA estimates is more than adequate to obtain sufficiently accurate estimates of cycle frequencies with no prior knowledge.

vi) Comparison of Performance of Array-Based AOA Estimation and TDOA Estimation

Our simulations with a 5-sensor array in [20] showed that using multiple pairs of sensors in the M-SPECCOA algorithm provides relatively little improvement in performance over that attainable using the single pair of sensors with the greatest separation in the SPECCOA algorithm.

Furthermore, for close sensors, e.g., a sensor array on a single platform, the array-based AOA estimation algorithm, Cyclic MUSIC, provides much more accurate AOA estimates than the multiple-sensor-pair TDOA-based algorithm M-SPECCOA; e.g., for a collect of 4096 time samples (512 keying intervals of the BPSK signal), SPECCOA produces RMSEs of 15° and 2° for SNRs of 0 dB and 10 dB, respectively, whereas Cyclic MUSIC produces RMSEs of 0.7° and 0.2° in the same situations. It is estimated that SPECCOA would produce an RMSE as small as 0.2° for an SNR of 10 dB if either the collect or the sensor separation were increased by a factor of 10. For widely separated sensors, e.g., on separate platforms, it is clear that SPECCOA can outperform Cyclic MUSIC. Also, although CPD has an RMSE that is about a factor of 10 greater than that produced by SPECCOA, it too should outperform Cyclic MUSIC if it is used with widely separated platforms. However, the ambiguity in the CPD estimate of AOA must be resolved by some other means. On the other hand, when used in conjunction with SPECCOA or CPD, Cyclic MUSIC can be useful (e.g., to help resolve ambiguity) if each of two (or more) platforms contains a sensor array.

vii) Exploitation of Higher-Order Cyclostationarity

For bandwidth-efficient signals, like some PSK, CPFSK, and partial-response encoded signals with small (or zero) excess bandwidth, the degree of second-order cyclostationarity can be relatively weak (or zero). In [21], we have shown with preliminary simulations of CPFSK obtained by dehopping a frequency-hopped (FH) signal that signal-selective TDOA estimation can be performed by exploiting higher-order cyclostationarity (4-th order in the simulations). Motivated by this, we have initiated the development of a general theory of the time-domain and frequency-domain characteristics of higher-order cyclostationarity [28] - [29]. The approach considered in [21] involves the use of an FFT channelizer for dehopping the FH signal to produce a CPFSK signal, and also the use of an FM demodulator to convert the CPFSK signal to an ASK signal. The ability of a simple sorting algorithm to separate out multiple FH signals is demonstrated by simulation for SINCGARS FH signals, and the sorted output is then studied. The cyclic features of CPFSK signals are derived analytically, and the cyclic features of the

CPFSK output of the FH sorter and its squared version are measured from simulations, as are the cyclic features of its FM demodulated version and its square. It is seen that both squaring and FM demodulation can enhance baud-rate features. Two-platform reception is then considered, and it is demonstrated that TDOA estimation can be performed on these four versions of sorted dehopped FH signals. Two TDOA estimation algorithms are demonstrated for an environment consisting of several FH signals in white noise. These are a conventional generalized crosscorrelation algorithm and the SPECCOA algorithm. Both algorithms perform satisfactorily.

c. Conclusions

In this research, we have developed and simulated methods that exploit cyclostationarity, which is exhibited by most communication and telemetry signals, to estimate the locations of emitters by estimating time-difference-of-arrival at two or more reception platforms. Exploitation of cyclostationarity renders the new methods signal-selective, which enables them to perform well even in the presence of temporally and spectrally overlapping interference and noise, regardless of the proximity of the interfering emitters to the emitters of interest. These methods can substantially outperform conventional methods and the Cramér-Rao lower bound based on a stationary model of the signal, even if the cycle frequencies are unknown and must be detected and estimated. Future work should attempt to analytically characterize the performance of these new algorithms and pursue even better-performing algorithms.

5. OVERVIEW OF RESEARCH ON DIRECTION FINDING AND SPATIAL FILTERING WITH ANTENNA ARRAYS

a. Introduction

Broadly, one of the questions we have attempted to answer in this research is "How can data from multiple receiving antennas in a sensor array be used to describe the spatial and temporal characteristics of the signals in a radio-signal environment?" In particular, the signals of interest are man-made communication signals which typically exhibit a statistical property called *cyclostationarity* or, equivalently, *spectral correlation* [1]. Unlike stationary signals, cyclostationary signals are correlated with frequency-shifted as well as time-shifted versions of themselves, a property that enables some signal processing algorithms to estimate parameters of such signals even when they are severely masked by interference and noise.

To this end, we have developed and analyzed signal processing algorithms to estimate the directions of arrival (DOAs) of the signals and to estimate the signal waveforms themselves by spatial filtering, with emphasis on algorithms that, by exploiting cyclostationarity, can operate properly when (1) little or no prior knowledge of the signal environment is available, (2) sufficient prior knowledge is available to classify signals as being signals of interest (SOIs), and the SOIs constitute only a subset of the set of all received signals, (3) physical or economic constraints force the number of deployed antennas in the array to be small, perhaps smaller than the number of signals arriving at the array, and (4) in some cases the array to be used for spatial filtering is uncalibrated. It is explained in this report that, in some situations, the sufficient prior knowledge required in condition (2) can be very modest (e.g., a keying rate or carrier frequency), and can even be replaced with estimates obtained from the data. Even though these three conditions can arise in surveillance, intelligence, and reconnaissance applications or when the array size is limited, they are not addressed by conventional direction-finding (DF) methods or spatial filtering methods. Together, the estimated DOAs, the various intermediate parameters found in the process of estimating the DOAs, and the extracted signal waveforms can form a reasonably complete description of the spatial locations and temporal characteristics of the signals in the received environment.

b. Research Findings

Within the broad scope of DF and spatial filtering, our primary effort has been directed toward estimating the directions of arrival of cyclostationary SOIs having sufficiently high carrier frequencies and narrow bandwidths that the narrowband approximation is valid (i.e., that the response of the sensor array over the frequency band of interest is approximately constant).

In the survey article [22], we unify a wide range of algorithms for narrowband DF by interpreting them in terms of spatial filtering instead of the less physically motivated abstraction of subspace fitting. Therein we also provide an overview of the DF methods we have developed and explain how they are able to perform properly in many applications in which conventional methods perform poorly or fail completely. These explanations are simplified by postponing discussion of the related problems of estimating the number of SOIs present and, for the cyclostationarity-exploiting methods, estimating one or more cycle frequencies of the SOIs. However, in practice these quantities must be estimated prior to processing the data with one of the DF algorithms; subsequent discussion considers the problems of estimating these quantities and of accommodating other departures from ideality, such as the need to adaptively adjust array calibration data to correct for effects of perturbation in the sensor position and age-induced and temperature-induced drift in component characteristics. This order of presentation parallels the order in which we addressed the corresponding problems in our research and proceeds from the more easily understood to the more difficult.

i) DF Methods for Cyclostationary Signals

In this part of our research, we have developed and analyzed the Cyclic MUSIC methods and the Cyclic Least Squares (CLS) methods for estimating the DOAs of cyclostationary signals. We derive the methods, explain their capabilities, and analyze their performance (both analytically and using computer simulations) in [23] - [25]. Briefly, the forms of the Cyclic MUSIC methods resemble the form of the conventional MUSIC methods, whereas the forms of the CLS methods resemble the form of the conventional maximum likelihood method for unknown signals in white Gaussian noise. Both types of methods estimate the DOAs of the signals having a specified cycle frequency: after estimating the corresponding cyclic autocorrelation matrix from the received data, each method finds the DOAs that "best" describe the spatial characteristics represented by the matrix, where the quality of the solution is determined differently by each method.

We show in the cited papers that, unlike conventional methods, these methods can operate properly when (1) the number of all signals is greater than the number of sensors, provided that the number of SOIs having the cycle frequency of interest is less than the number of sensors, (2) SNOIs are arbitrarily closely spaced to SOIs, and (3) the spatial characteristics of the noise are arbitrary and unknown. Furthermore, by estimating the DOAs of only the SOIs having a particular cycle frequency, these DF methods reduce or eliminate the need for post-processing steps that would otherwise be needed to classify the DOAs according to signal type.

Although the Cyclic MUSIC and CLS methods are treated as if the cycle frequencies and numbers of signals having each of these cycle frequencies were known, the methods described in the next two sections can estimate these parameters and thus allow the DF methods to be applied if these parameters are unknown. In this sense, the benefits of signal-selective DF that accrue from exploiting the cyclostationarity of the received signals can be enjoyed even when essentially nothing is known about the environment beforehand.

ii) Cycle Frequency Estimation

Cyclostationary signals are correlated with frequency-shifted versions of themselves for certain values of the frequency-shift parameter. Equivalently, the lag-product waveforms of these signals contain finite-strength additive sine waves having frequencies equal to the values of the frequency-shift parameter, which are called *cycle frequencies*. Thus, one means of estimating the cycle frequencies is to estimate the frequencies of the sine waves present in lag-product waveforms computed from finite records of received data. We demonstrate this simple technique in [24], where the resulting cycle frequency estimates are then used by Cyclic MUSIC and CLS to estimate the corresponding DOAs. We show that these estimated DOAs have approximately the same mean-squared error (MSE) as those obtained using the exact values of the cycle frequencies.

iii) Estimating the Number of Signals Present

Conventional DF methods estimate the DOAs of all signals present and thus require an estimate of their number, which is typically obtained by analyzing the eigenvalues of the autocorrelation matrix of the data. In contrast, the cyclostationarity-exploiting DF methods operate only on a subset of the signals and thus require an estimate of the number of signals in each subset. Unlike the statistics of the singular values of the autocorrelation matrix, those of the cyclic autocorrelation matrix are not well understood. Also, the results in the literature on canonical correlation analysis in general (also called common factor analysis) are meager in comparison with those on principal component analysis (also called factor analysis); the former are applicable to the signal-selective DF problem, whereas the latter are applicable to the conventional DF problem. Consequently, the Cyclic Correlation Significance Test (CCST) that we have proposed is only an initial attempt at solving this difficult problem by using the singular values of the cyclic autocorrelation matrix. Among its noteworthy features is the fact that it uses a penalty function (similar in form to that used by the conventional AIC, MDL, and EDC methods) instead of a subjectively chosen threshold. Alternatively, we show in [23, 24] that the CLS objective function, after maximization over the DOAs for each possible number of cyclostationary signals, can be used to estimate the number of cyclostationary signals, even in the presence of perfectly correlated multipath. However, this method requires a subjectively chosen threshold and has not yet been extensively tested.

iv) Performance Bounds

In addition to developing and analyzing specific DF algorithms, we have also investigated the Cramer-Rao Lower Bound (CRLB) on the covariance of unbiased parameter estimates of cyclostationary Gaussian signals. Apart from being of theoretical interest (almost all relevant work has dealt with stationary Gaussian signals), the CRLB for this problem can be used to gauge the efficiency of the cyclostationarity-exploiting methods and to indicate the potential for such methods to exhibit far better performance than that of methods for stationary signals. We show in [25] that the CRLB for two closely spaced cyclostationary signals having different cycle frequencies can be several orders of magnitude less than that for two closely spaced stationary signals. Computer simulations therein show that the Cyclic MUSIC and CLS methods yield DOA estimates in some environments with MSE that is comparable to the CRLB for cyclostationary signals and much less than that for stationary signals, which indicates that the CRLB for cyclostationary signals is not overly optimistic in its prediction of potentially huge performance gains due to exploitation of cyclostationarity.

v) Spatial Filtering

In addition to their primary function of estimating the DOAs of the SOIs, the CLS method and one of the Cyclic MUSIC methods can also compute the coefficients of spatial filters that attenuate noise and SNOIs and enhance SOIs. In [23, 24], we show that the signal-to-interference-and-noise ratio (SINR) of the signal waveforms estimated by CLS converge to the maximum attainable SINR if the signals are uncorrelated. However, the multidimensional search and array calibration data needed by CLS can be prohibitively expensive. In contrast, the Phase-SCORE Cyclic MUSIC method [23] is based on the Phase-SCORE method for blind adaptive spatial filtering [26]. Phase-SCORE requires neither a multidimensional search nor any array calibration data, albeit at the expense of lower output SINR than that of CLS in some cases. When the spatial resolution provided by phase-SCORE is adequate, this algorithm can be used on each of two or more platforms to separate multiple signals from closely spaced emitters whose signals have the same cycle frequency. The TDOA method of source location can then be used on the separated signals. This could yield better source-location performance than that obtainable using the TDOA method alone, which would have to be able to resolve multiple closely-spaced TDOAs.

c. Conclusions

In this research, we have developed and analyzed methods that exploit cyclostationarity, which is exhibited by most communication and telemetry signals, to estimate the directions of arrival of signals arriving at a sensor array without requiring many of the prohibitive assumptions or potentially costly prior knowledge required by conventional direction-finding methods. The inherently signal-selective methods reduce or eliminate the need to perform post-processing to classify the direction estimates as corresponding to desired or undesired signals. Furthermore, when signals are closely spaced and have different cycle frequencies, the methods we have developed can substantially outperform conventional methods and the Cramer-Rao lower bound for stationary signals, even if the cycle frequencies are unknown and must be detected and estimated. In addition to performing DF, some of the signal-selective methods can also extract high-quality estimates of the signal waveforms from the received signal environment. Thus, the methods that we have developed and studied in this project can provide a reasonably complete description of the spatial locations and temporal characteristics of received signals in a much wider variety of received environments, especially those for which little or no prior knowledge is available, than can be accommodated by previously developed methods. Further work is needed on the analytical characterization of performance of the new algorithms and in the pursuit of even better-performing algorithms.

6. OVERVIEW OF RESEARCH ON COCHANNEL SIGNAL SEPARATION (By Frequency-Shift Filtering)

Conventional time and space filtering of stationary random signals, which amounts to forming linear combinations of time translates and space translates, exploits the temporal and spatial coherence of the signals. By including frequency translates as well, the spectral coherence that is characteristic of cyclostationary signals can also be exploited. In [27], we develop some of the theoretical concepts underlying this generalized type of filtering, called *FREquency-SHift (FRESH) filtering*, summarize the theory of optimum FRESH filtering, which is a generalization of Wiener filtering (optimum time-invariant filtering) to *cyclic Wiener filtering* (optimum multiply-periodic filtering) and illustrate the theory with specific examples of separating temporally and spectrally overlapping communications signals, including AM, BPSK, and QPSK. The structures and mean-squared-error performances of optimum FRESH filters are presented, and adaptive adjustment of the weights in these structures is discussed. Also, specific results on the number of digital QAM signals that can be separated, as a function of excess bandwidth, are obtained. The results show that the spectral redundancy inherent in excess bandwidth can be used effectively to improve system performance, and this suggests that excess bandwidth is a richer system-design parameter than previously recognized. In other words, high bandwidth-efficiency can be more costly in system performance when cochannel interference is present than one might have thought.

7. OVERVIEW OF RESEARCH ON HIGHER-ORDER CYCLOSTATIONARITY

a. *Introduction*

Regeneration of a sine wave from received data is a familiar idea in communication and interception systems. Sine waves are regenerated for a variety of purposes, including bit and carrier synchronization for demodulation, parameter estimation, weak-signal detection, and identification of modulation type. Since a sine wave contains no information in the sense of a message signal, it is desirable to avoid transmitting sine-wave components along with the modulated signal to reduce transmitted power. Therefore, in power-efficient systems it is necessary to generate needed sine waves from the received data using nonlinear transformations, and to design the transmitted signal so that this can be accomplished. Also, in applications that require covert communication, it is desirable to design the signal so that an unintended receiver cannot easily regenerate sine waves from the transmitted data.

The theory behind generating sine waves by using second-order nonlinearities is briefly surveyed in [1] and comprehensive treatments are given in the references cited in [1]. Signals from which sine waves can be generated by using quadratic devices are called *second-order cyclostationary* or just *cyclostationary*. The central parameters of the theory of cyclostationary signals are the cyclic autocorrelation and the cyclic spectral density. Both of these parameters are generalizations of conventional statistical parameters, namely the autocorrelation and the power spectral density (PSD). There are many signal processing applications wherein proper exploitation of second-order cyclostationarity (SOCS) can provide better performance relative to conventional algorithms that ignore SOCS, such as algorithms that utilize only the information in the PSD. However, there are signals from which no sine waves can be generated by using quadratic nonlinearities, but from which sine waves can be generated by using higher-order nonlinearities. Signals from which sine waves can be generated using an nth-order nonlinearity are called *nth-order cyclostationary* and, for $n > 2$, are said to exhibit higher-order cyclostationarity (HOCS).

Some examples follow. M-ary PSK (phase-shift-keyed) signals with $M \geq 4$ and various digital QAM (quadrature amplitude modulated) signals typically exhibit SOCS associated with the keying rate, but not the frequency of the sine-wave carrier. However, they do exhibit fourth- or higher-order cyclostationarity (depending on the particular amplitude/phase constellation) associated with the frequency of the carrier. Real and complex digital PAM (pulse-amplitude modulated) signals (which include the complex envelopes of idealized PSK signals and digital QAM signals) exhibit SOCS associated with the keying rate only if the bandwidth exceeds half the keying rate, and can exhibit only weak SOCS if the bandwidth is only slightly larger than half the keying rate. However, these signals always exhibit HOCS, of some order, associated with the keying rate. In addition, some FM (frequency-modulated) and CPFSK (continuous-phase, frequency-shift-keyed) signals exhibit weak SOCS, but strong HOCS.

The study of HOCS is the study of the sine waves contained in nth-order nonlinearly transformed time-series and as such it encompasses the study of the nth-order transformations of a purely stationary time-series, from which only sine waves with zero frequency can be generated (for any order nonlinearity). An important statistical parameter of this latter theory is the *polyspectrum*, which is a limiting form of a joint spectral cumulant function and is equivalent to the multidimensional Fourier transform of the joint temporal cumulant function. Loosely speaking, a cumulant is that portion of a moment that is independent of lower-order moments. For example, think of the second centralized moment—the variance—of a random variable: it does not depend on the mean, whereas the second moment does. The theory of HOCS is not only a generalization of the theory of SOCS, but also of the temporal and spectral cumulant theory of purely stationary time-series. It introduces the *cyclic temporal cumulant function* and the *cyclic polyspectrum*, which include the conventional temporal cumulant function and polyspectrum as special cases.

b. Research Findings

We are in the process of developing the theory of higher-order cyclostationarity. This theory provides the foundation and framework for understanding how to use higher than second-order (quadratic) nonlinearities (e.g., quartic) to intercept signals with weak or no second-order cyclic features (such as QPSK with substantial clock jitter). In [28], the foundation for the theory of higher-order cyclostationarity is laid in a concise overview. (A more detailed presentation is currently in preparation.) In [29], we present the preliminary findings of our first investigation into methods for measuring the cyclic polyspectrum.

i) The Temporal and Spectral Cumulant Theory of Cyclostationary Time-Series

The relatively recent theory of second-order cyclostationary (SOCS) time-series has been applied to a variety of signal processing problems, such as direction-finding, weak-signal detection, and system identification. The most basic and most unifying concept in the theory of SOCS is that of sine-wave generation: a signal that has no additive sine-wave components exhibits SOCS if a sinewave can be generated by using a quadratic transformation. Another important concept is spectral correlation: certain distinct spectral components of a SOCS time-series exhibit temporal correlation. In [28], the concepts and formalism associated with SOCS are generalized to higher order to form the theory of higher-order cyclostationary (HOCS) time-series. Specifically, the two limit statistical parameters, the cyclic autocorrelation which quantifies regenerated sine-wave strength, and the cyclic spectrum which quantifies spectral correlation, are directly generalized to the cyclic temporal moment function (CTMF) and the spectral moment function (SMF), respectively. It is shown in [28] that these higher-order parameters are not well-behaved mathematically and suffer from the problem of inherent mixing of sine waves from lower orders with those from the order of interest. It is also shown that the introduction of cumulants in both the time and frequency domains solves these problems. The claim is made that cumulants are the proper functions upon which to base the theory of HOCS, and a brief tutorial treatment of cumulants is included (since they are not as familiar as moments, although they contain the same information in the sense that the first n moments determine exactly the first n cumulants and vice versa). Thus, the cyclic temporal cumulant function (CTCF) is the appropriate generalization of the cyclic autocorrelation and the spectral cumulant function (SCF), as well as its reduced-dimension equivalent—the cyclic polyspectrum (CP)—is the appropriate generalization of the cyclic spectrum. The CTCF and CP are calculated for the particularly interesting and important complex pulse-amplitude-modulated time-series model.

In continuing work, not included in [28], the effects on the CTMF, CTCF, SMF, and CP of various transformations on the time-series have been derived; these include convolution, multiplication, addition, and periodic time-sampling. A qualitative discussion of cycloergodicity has been used to bridge the gap between our time-average parameters and the ensemble-average parameters obtainable in a stochastic process framework. Finally, approaches to signal detection and TDOA estimation that are based on the CTCF or CP have been proposed.

ii) Estimating Cyclic Polyspectra

In [29], we consider the problem of estimating cyclic polyspectra. The cyclic polyspectrum can be viewed as the Fourier transform of a higher-order cyclic temporal cumulant, or as a cumulant of spectral components of the time-series. In the latter case, the cyclic polyspectrum is seen to be a linear combination of products of spectral moment functions which include Dirac deltas, all of which cancel out in the cyclic polyspectrum. Because of the impracticality of this cancellation using finite precision arithmetic, it is concluded that the cyclic polyspectrum is best measured by first measuring the cyclic temporal cumulant, tapering it, and then transforming it. The noise tolerance of the cyclic polyspectrum, which results from the tolerance to Gaussian noise exhibited by any cumulant (of order greater than 2), combined with the tolerance to any stationary noise exhibited by any cyclic statistic, is illustrated qualitatively by considering estimation for a noisy pulse-amplitude-modulated time-series. Also, the tolerance to cyclostationary interference with cycle frequencies unequal to those of the signal of interest is illustrated by considering estimation for a time-series consisting of the sum of two PAM signals.

8. REFERENCES

[1] W. A. Gardner, "Exploitation of spectral redundancy in cyclostationary signals," *IEEE Signal Processing Magazine*, Vol. 8, No. 2, pp. 14-36, 1991 (Invited Paper).

[2] W. A. Gardner, "Signal interception: A unifying theoretical framework for feature detection," *IEEE Transactions on Communications*, Vol. COM-36, No. 8, pp. 897-906, 1988. (Awarded S. O. Rice Prize Paper in Field of Communication Theory, from IEEE Communications Society.)

[3] W. A. Gardner and C. M. Spooner, "Signal interception: Performance advantages of cyclic-feature detectors," *IEEE Transactions on Communications*, Vol. 40, 1992 (in press).

[4] W. A. Gardner and C. M. Spooner, "Weak-signal detection and source location: Simplifications of the maximum-likelihood receiver," in press (same as item 16).

[5] C. M. Spooner and W. A. Gardner, "Robust feature detection for signal interception," *IEEE Transactions on Communications* (in press).

[6] C. M. Spooner and W. A. Gardner, "A performance comparison of two signal analysis surfaces: Ambiguity and spectral correlation" *SSPI Technical Report*, February 1989. (same as item 13).

[7] W. A. Gardner and C. M. Spooner, "Performance of multihop detectors for hybrid FH/DS signals," *SSPI Technical Report*, Sept. 1991.

[8] W. A. Gardner, "Measurement of spectral correlation," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, Vol. ASSP-34, No. 5, pp. 1111-1123, 1986.

[9] R. S. Roberts. W. A. Brown, and H. H. Loomis, Jr., "Computationally efficient algorithms for cyclic spectral analysis," *IEEE Signal Processing Magazine*, Vol. 8, No. 2, pp. 38-49, 1991 (Invited paper).

[10] W. A. Brown and H. H. Loomis, Jr., "Digital implementations of spectral correlation analyzers," *IEEE Transactions on Signal Processing*, Vol. 40, 1992 (in press).

[11] W. A. Gardner and R. S. Roberts, "One-bit spectral correlation algorithms," *IEEE Transactions on Signal Processing* (in press).

[12] R. S. Roberts and H. H. Loomis, Jr., "Parallel computational structures for a class of cyclic spectral analysis algorithms," in preparation (not included herewith).

[13] C. M. Spooner and W. A. Gardner, "A performance comparison of two signal analysis surfaces: Ambiguity and spectral correlation" *SSPI Technical Report*, February 1989 (same as item 6).

[14] W. A. Gardner and C. K. Chen, "Signal-selective time-difference-of-arrival estimation for passive location of manmade signal sources in highly corruptive environments. Part I: Theory and method, *IEEE Transactions on Signal Processing*, Vol. 40, 1992 (in press).

[15] C. K. Chen and W. A. Gardner, "Signal-selective time-difference-of-arrival estimation for passive location of manmade signal sources in highly corruptive environments. Part II: Algorithms and performance, *IEEE Transactions on Signal Processing*, Vol. 40, 1992 (in press).

[16] W. A. Gardner and C. M. Spooner, "Weak-signal detection and source location: Simplifications of the maximum-likelihood receiver," in press (same as item 4).

[17] C. M. Spooner and W. A. Gardner, "On the performance of cyclic TDOA algorithms," *SSPI Technical Report*, Sept. 1991.

[18] W. A. Gardner and C. M. Spooner, "Comparison of auto- and cross-correlation methods for signal-selective TDOA estimation," *IEEE Transactions on Signal Processing* (in press).

[19] W. A. Gardner, "Cyclic autocorrelation algorithms for joint FDOA/TDOA estimation," *SSPI Technical Report*, Sept. 1991.

[20] W. A. Gardner, S. V. Schell, and C. M. Spooner, "A comparison of array-based AOA estimation and TDOA estimation," *SSPI Technical Report*, May 1991.

[21] W. A. Gardner, C. M. Spooner, and C. K. Chen, "Study of cyclic features of dehopped FH signals," *SSPI Technical Report*, 1989.

[22] S. V. Schell and W. A. Gardner, "High resolution direction finding," Chapter in *Handbook of Statistics*, N. K. Bose and C. R. Rao (Eds.), Elsevier, 1992 (in press).

[23] S. V. Schell and W. A. Gardner, "Estimating the directions of arrival of cyclostationary signals. Part I: Theory and methods," *IEEE Transactions on Signal Processing* (submitted).

[24] S. V. Schell, "Estimating the directions of arrival of cyclostationary signals. Part II: Performance," *IEEE Transactions on Signal Processing* (submitted).

[25] S. V. Schell and W. A. Gardner, "Cramer-Rao lower bound for directions of arrival of Gaussian cyclostationary signals," *IEEE Transactions on Information Theory* (in press).

[26] B. G. Agee, S. V. Schell, and W. A. Gardner, "Spectral self-coherence restoral: A new approach to blind adaptive signal extraction using antenna arrays," *Proceedings of the IEEE*, Vol. 78, pp. 753-767, 1990.

[27] W. A. Gardner, "Cyclic Wiener filtering: Theory and method," *IEEE Transactions on Communications*, Vol. 40, 1992 (in press).

[28] C. M. Spooner and W. A. Gardner, "An overview of the theory of higher-order cyclostationarity," *Proceedings of the NASA/Hampton Workshop on Nonstationary Random Processes*, Hampton University, VA, August 1-2, 1991 (in press).

[29] C. M. Spooner and W. A. Gardner, "Estimation of cyclic polyspectra," *Proceedings of the Twenty-fifth Annual Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, November 4-6, 1991 (Invited paper).