IEEE
Signal Processing
MAGAZINE
APRIL 1991

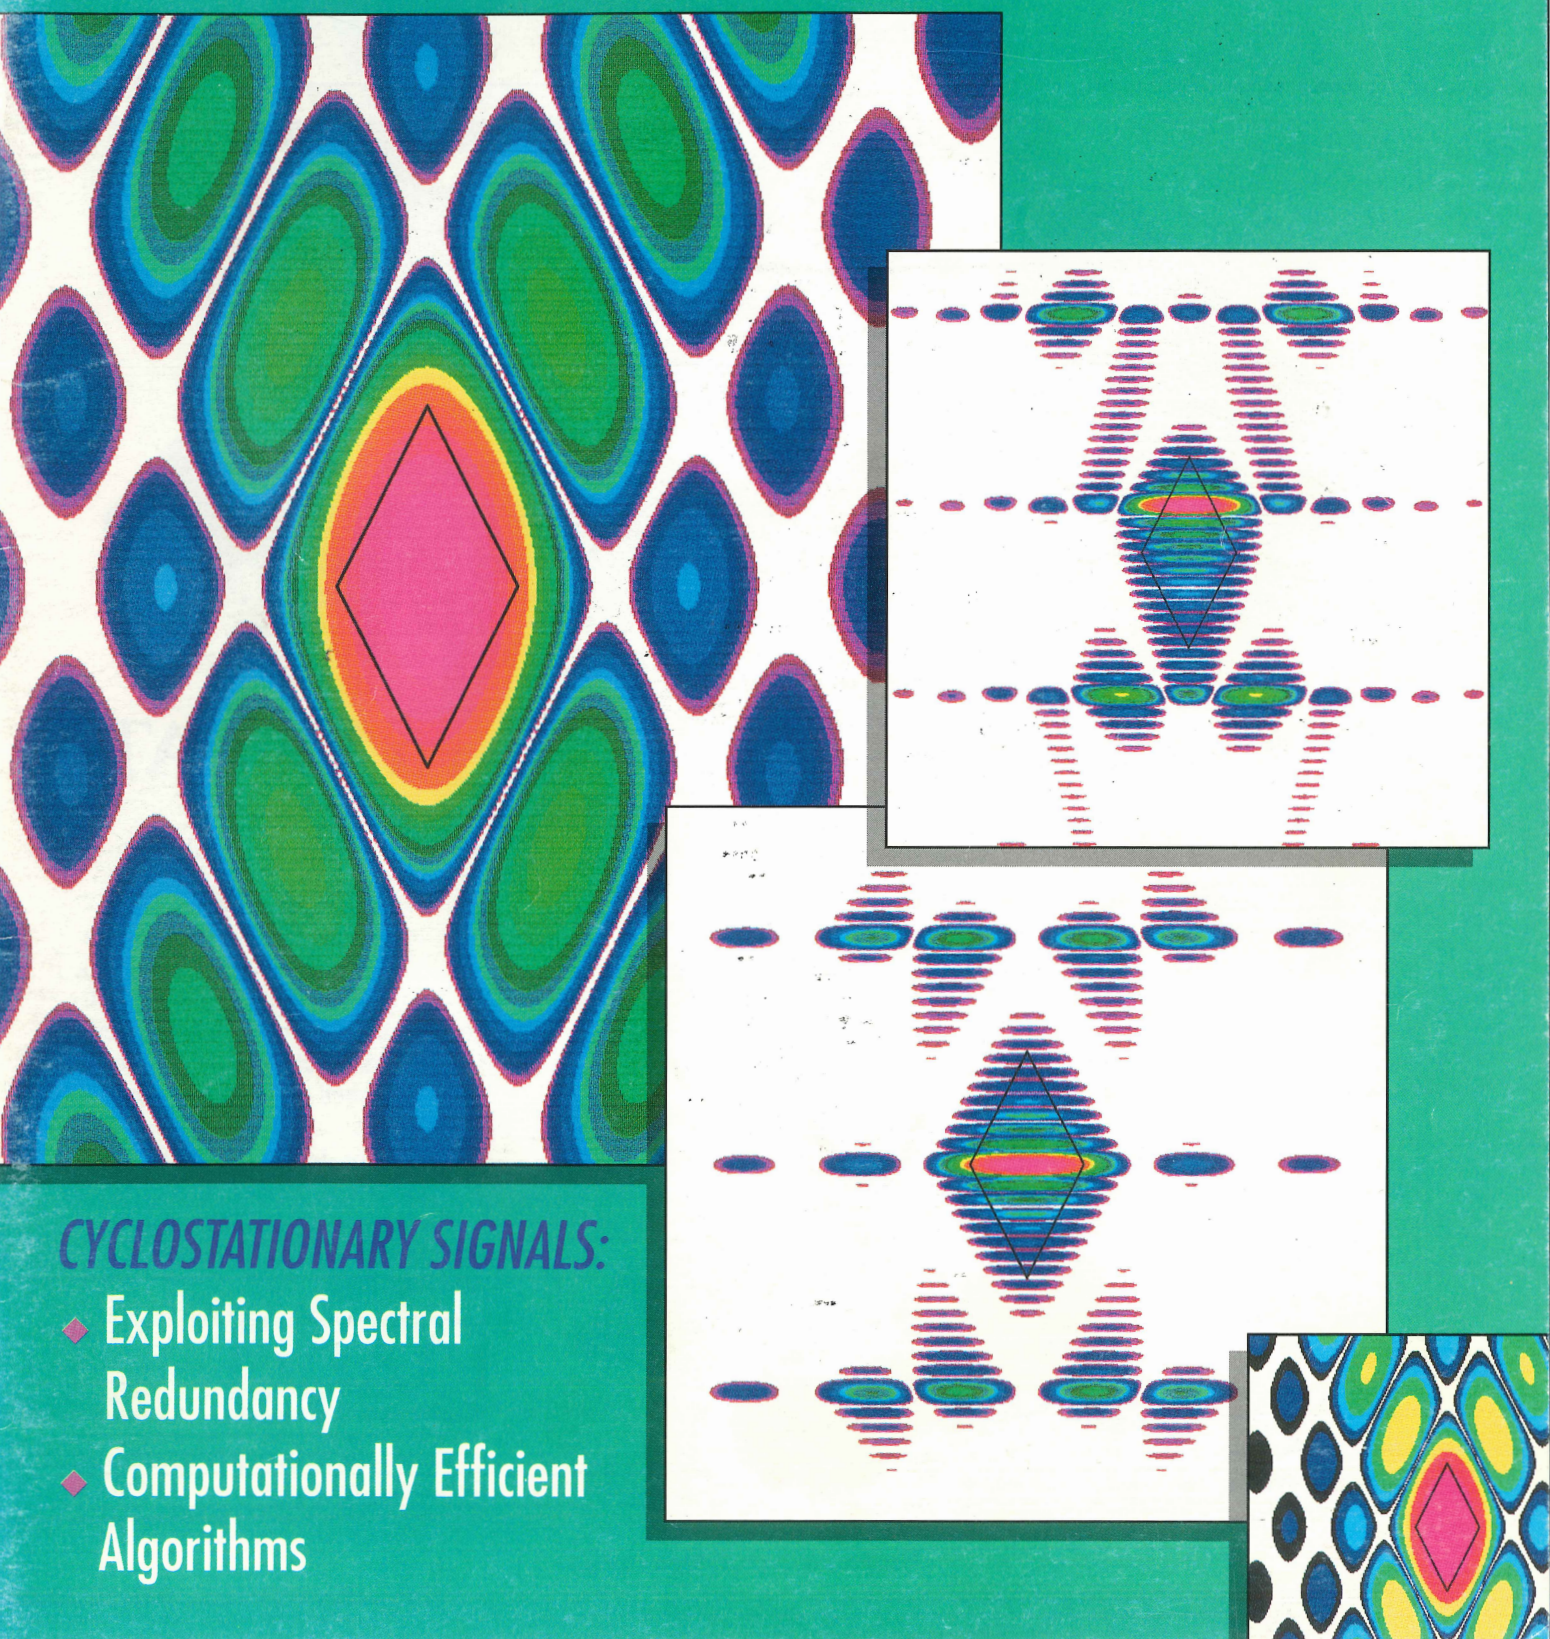

CYCLOSTATIONARY SIGNALS:

- Exploiting Spectral Redundancy
- Computationally Efficient Algorithms

IEEE
Signal Processing
MAGAZINE

VOLUME 8 NUMBER 2 ISSN 1053-5888 APRIL 1991
A Publication of the IEEE Signal Processing Society

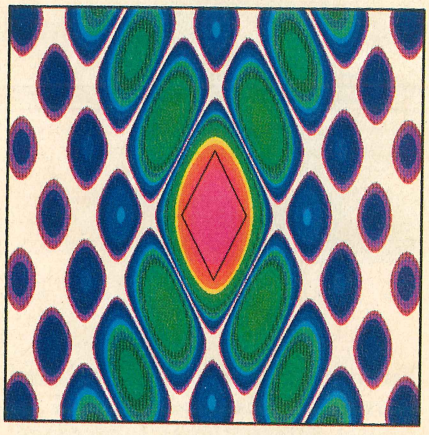

FEATURES

DEPARTMENTS

In this issue ...

Guest Editor William A. Gardner composed a magnificent issue on the important topic of cyclostationarity. Professor Gardner's research group at the University of California at Davis, and colleagues in government and industry, are internationally recognized for their work on this subject. His guest editorial to follow will introduce the issue and survey the papers.

Jack Deller

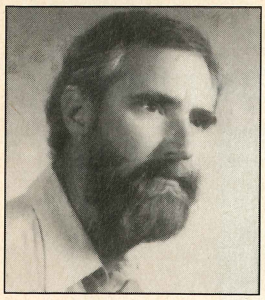

Guest Editorial

Exploitation of Spectral Redundancy in Cyclostationary Signals

The Signal Processing Group at the University of California, Davis, in collaboration with colleagues in government and industrial research and development laboratories, has been developing and promoting a type of signal processing that exploits a commonly occurring but often unrecognized property of manmade signals. This property, called spectral redundancy or more specifically spectral correlation, is a direct result of periodic (cyclic) structure present in otherwise stationary random signals. These cyclostationary signals are encountered in many signal processing systems including especially communications and telemetry, but also control, radar, sonar, and others. By designing signal processors that intentionally exploit spectral redundancy, new levels of performance can be achieved. This is particularly so for signal processing tasks such as detection and estimation of highly corrupted signals, that is, signals subjected to temporally and spectrally overlapping interference, or relatively strong noise, or severe distortion.

This special issue of the *IEEE Signal Processing Magazine* contains two articles. The first and longer of the two provides the reader with a concise introduction to the concept and theory of spectral redundancy, and a brief survey of some of its applications in signal processing. The second article introduces the reader to the problem of designing computationally efficient algorithms and architectures for digital measurement of spectral correlation. Although some of the applications of spectral redundancy described in the first article do not require actual measurement of spectral correlation, or require relatively modest measurement (e.g., equivalent to the measurement of a power spectrum), other applications require exhaustive spectral-correlation analysis. It is this latter situation that is addressed in the second article.

It is our hope that this special issue will attract the attention of those uninitiated in spectral redundancy and will assist them in determining if exploitation of spectral redundancy holds some promise for their signal processing problems.

A few words about the first article are in order here. The subject is inherently statistical and many of the concepts are, relatively speaking, advanced. Many engineers do not learn about the basic theory of random signals in their undergraduate training, and quite a few engineers with the M.S. degree, who are working in signal processing, have not had a formal course in random signals. Consequently, this article, which treats a subject that extends and generalizes concepts, theory, and methods that some readers have not yet learned about, except perhaps at a superficial level, is destined to be only partially understood by many. The only way this could be avoided would be to expand the article into a full-blown treatise that would be far too long for publication in this magazine. As a result of this situation I have had to strike a balance in the first article between the use of detailed low-level explanations with plenty of pictures, more technical explanations with more equations and fewer pictures, and the use of references for more expansive discussions and more results on the implementation and performance of specific signal processing methods. The required compromise here is reflected in the contradictory suggestions from the reviewers of this article. One reviewer wanted more equations and more technical and/or mathematical detail, while another reviewer wanted more lightweight discussion and fewer equations. I sympathized with both of these reviewers, but it was impossible to completely satisfy both of them with the limited number of pages available, given the broad scope of the article. In fact, the problems stemming from the inherently technical nature of the subject are exacerbated by the broad scope of the article. Unlike typical articles on technical subjects in this magazine, this article attempts to introduce the reader to a broad area of study that incorporates numerous topics (any one of which could easily occupy the entire article), such as signal detection, modulation recognition, signal parameter estimation, direction finding, time-difference estimation, interference suppression, distortion reduction, and linear prediction, as well as the underlying theory which generalizes the theory of autocorrelation and power spectral density from stationary random signals to cyclostationary random signals.

Nevertheless, I am optimistic that the article, in spite of its unavoidable drawbacks, will serve its purpose of introducing readers to a valuable new area of study, motivating them to learn more about the subject, and providing them with a bibliography that will help them to learn more.

William A. Gardner
University of California, Davis

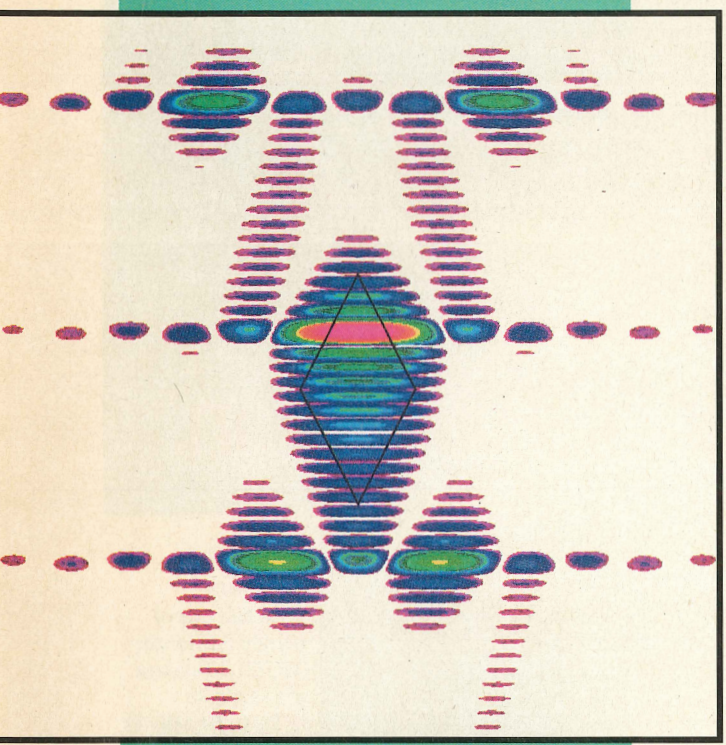

Exploitation of Spectral Redundancy in Cyclostationary Signals

WILLIAM A. GARDNER

Three properties shared by many manmade signals are described: (1) the property that enables generation of spectral lines by quadratically transforming the signal; (2) the statistical property called second-order ***cyclostationarity****, which means that the autocorrelation function fluctuates periodically with time; and (3) the correlation property for signal components in distinct spectral bands. It is shown that these three properties are different manifestations of a single attribute called* ***spectral redundancy****. This intriguing attribute is studied, and a variety of ways for exploiting it to perform signal processing tasks involving detection and estimation of highly corrupted manmade signals are explained.*

1053-5888/91/0400-0014$1.00 © 1991 IEEE

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 can be 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.

This article 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

This material is based upon work supported by the National Science Foundation under Grant No. MIP-88-12902.

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 article 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 and references are given in [1, Chapters 12 and 14], [2, Chapter 12].

Finally, and most importantly, this article 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.

In the next section, the possibility of generating spectral lines by simply squaring the signal is illustrated for two types of signals: the random-amplitude modulated sine wave and the random-amplitude modulated periodic pulse train. Then it is explained that the property that

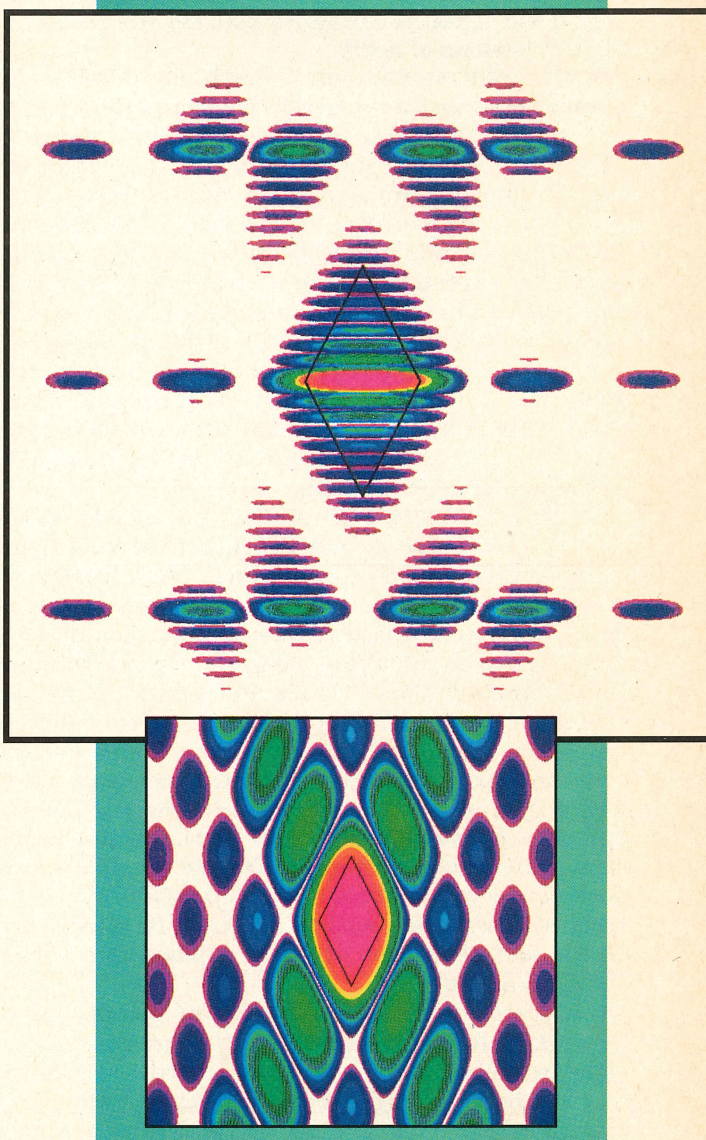

enables spectral-line generation with some type of quadratic time-invariant transformation is called *cyclostationarity* and is characterized by the *cyclic autocorrelation function*, which is a generalization of the conventional autocorrelation function. Following this, it is shown that a signal exhibits cyclostationarity if and only if the signal is correlated with certain frequency-shifted versions of itself.

In the third section, the correlation of frequency-shifted versions of a signal is localized in the frequency domain and this leads to the definition of a *spectral-correlation density function* (SCD). It is then explained that this function is the Fourier transform of the cyclic autocorrelation function. This Fourier-transform relation between these two functions includes as a special case the well-known *Wiener relation* between the power spectral density function and the conventional autocorrelation function. A normalization of the spectral-correlation density function that converts it into a spectral correlation coefficient, whose magnitude is between zero and unity, is then introduced as a convenient measure of the degree of spectral redundancy in a signal.

Continuing in this SCD section, the effects on the spectral-correlation density function of several signal processing operations are described. These include filtering and waveform multiplication, which in turn include the special cases of time delay and multipath propagation, bandlimiting, frequency conversion, and time sampling. These results are used to derive the spectral-correlation density function for the random-amplitude modulated sine wave, the random-amplitude modulated pulse train, and the binary phase-shift keyed sine wave. The spectral-correlation density functions for some other types of phase-shift keyed signals are also described graphically. Finally in this section, the measurement of (estimation of the ideal) spectral-correlation density function is briefly discussed and illustrated with a simulation of a phase-shift keyed signal.

The fourth section of this article contains the payoff for working through the preceding two sections. It provides explanations of how the spectral redundancy that is inherent in signals that exhibit cyclostationarity can be exploited in a variety of statistical signal processing tasks. The spectral redundancy can generally be exploited to enhance the accuracy and reliability of information gleaned from measurements of corrupted signals. Such information includes the following:

1) A decision as to the presence or absence of a random signal with a particular modulation type in a background of noise and other modulated signals,
2) A classification of multiple received signals in noise according to their modulation types,
3) An estimate of a signal parameter, such as carrier phase, pulse timing, or direction of arrival, or of the number of signals being received simultaneously, based on noise-and-interference-corrupted measurements,
4) An estimate of a message being communicated by a signal over a channel corrupted by noise, interference, and distortion,
5) A prediction of a future value of a random signal, and
6) An estimate of the input-output relation of a linear or nonlinear system based on measurements of the system's response to random excitation.

The article concludes with a brief section indicating how the theory of second-order cyclostationarity surveyed in the preceding sections can be generalized to higher-order cyclostationarity, which corresponds to the property that enables spectral-line generation using higher-order nonlinearities, such as cubic, quartic, and so on. References to more in-depth treatments of the theory and its applications are given throughout the article.

CYCLOSTATIONARITY

Spectral line generation

A signal $x(t)$ contains a *finite-strength additive sine-wave component* (an ac component) with frequency α, say

$$a \cos(2\pi\alpha t + \theta) \qquad \text{with } \alpha \neq 0 \tag{1}$$

if the Fourier coefficient

$$M_x^\alpha = \langle x(t) e^{-i2\pi\alpha t} \rangle \tag{2}$$

is not zero, in which case (1) gives

$$M_x^\alpha = \tfrac{1}{2} a e^{i\theta}$$

In (2), the operation $\langle \cdot \rangle$ is the time-averaging operation

$$\langle \cdot \rangle \triangleq \lim_{T\to\infty} \frac{1}{T} \int_{-T/2}^{T/2} (\cdot)\, dt$$

In this case, the power spectral density (PSD) of $x(t)$ includes a spectral line at frequency $f = \alpha$ and its image $f = -\alpha$. (The PSD is defined in the next section.) That is, the PSD contains the additive term[1]

$$|M_x^\alpha|^2 [\delta(f - \alpha) + \delta(f + \alpha)] \tag{3}$$

where $\delta(\cdot)$ is the Dirac delta, or impulse, function. For convenience in the sequel, it is said that such a signal exhibits *first-order periodicity*, with frequency α.

Let $x(t)$ be decomposed into the sum of its finite-strength sine-wave component, with frequency α, and its residual, say $n(t)$,

$$x(t) = a \cos(2\pi\alpha t + \theta) + n(t) \tag{4}$$

where $n(t)$ is defined to be that which is left after subtraction of (1) from $x(t)$. It is assumed that $n(t)$ is random. Here, the term *random* is used to denote nothing

[1]The strength of the spectral line is $|M_x^\alpha|^2$ as indicated in (3) if and only if the limit (2) exists in the temporal mean square sense with respect to the time parameter u obtained by replacing t with $t + u$ in (2) [1, Chapter 15, exc. 6].

more than the vague notion of erratic or unpredictable behavior. If the sine wave is weak relative to the random residual, it might not be evident from visual inspection of the signal that $x(t)$ contains a periodic component. Hence, it is said to contain *hidden periodicity*. However, because of the associated spectral lines, hidden periodicity can be detected and in some applications exploited through techniques of spectral analysis.

This article is concerned with signals that contain more subtle types of hidden periodicity that, unlike first-order periodicity, do not give rise to spectral lines in the PSD, but that can be converted into first-order periodicity by a nonlinear time-invariant transformation of the signal. In particular, we shall focus on the type of hidden periodicity that can be converted by a quadratic transformation to yield spectral lines in the PSD.

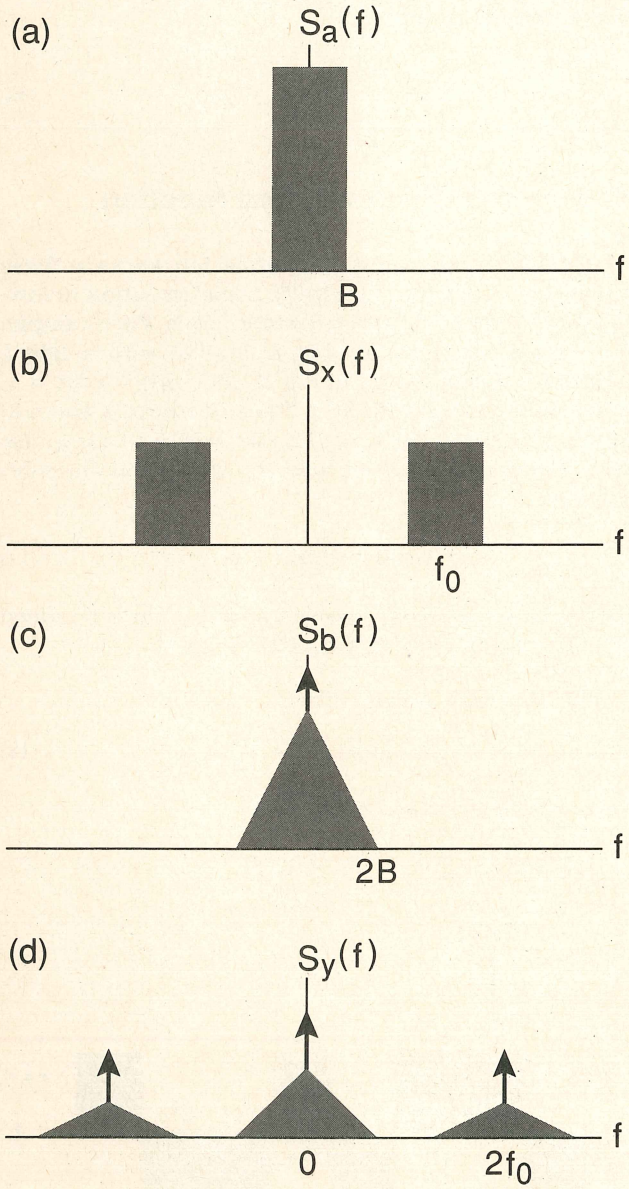

Fig. 1. a) Power spectral density (PSD) of a lowpass signal. b) PSD of an amplitude-modulated (AM) signal. c) PSD of a squared lowpass signal. d) PSD of a squared AM signal.

The discussion begins with a couple of examples.

Example 1: *AM.* Let $a(t)$ be a random lowpass signal (say lowpass filtered thermal noise) with the PSD $S_a(f)$ shown in Fig. 1a, which contains no spectral lines. If $a(t)$ is used to modulate the amplitude of a sine wave, we obtain the amplitude modulated (AM) signal

$$x(t) = a(t) \cos (2\pi f_0 t) \tag{5}$$

whose PSD $S_x(f)$ is given by [1, Chapter 3, Sec. D]

$$S_x(f) = \tfrac{1}{4} S_a(f + f_0) + \tfrac{1}{4} S_a(f - f_0) \tag{6}$$

as shown in Fig. 1b. Although the PSD is centered about $f = f_0$ and $f = -f_0$, there is no spectral line at f_0 or $-f_0$. The reason for this is that, as shown in Fig. 1a, there is no spectral line in $S_a(f)$ at $f = 0$. This means that the dc component

$$M_a^0 \triangleq \langle a(t) \rangle \tag{7}$$

is zero, since the strength of any spectral line at $f = 0$ is $|M_a^0|^2$.

Let us now square $x(t)$ to obtain

$$y(t) = x^2(t) = a^2(t) \cos^2 (2\pi f_0 t)$$

$$= \tfrac{1}{2} [b(t) + b(t) \cos (4\pi f_0 t)] \tag{8}$$

where

$$b(t) = a^2(t) \tag{9}$$

Since $b(t)$ is nonnegative, its dc value must be positive: $M_b^0 > 0$. Consequently, the PSD of $b(t)$ contains a spectral line at $f = 0$, as shown in Fig. 1c. The PSD for $y(t)$ is given by

$$S_y(f) = \tfrac{1}{4} [S_b(f) + \tfrac{1}{4} S_b(f + 2f_0) + \tfrac{1}{4} S_b(f - 2f_0)] \tag{10}$$

and, as shown in Fig. 1d, it contains spectral lines at $f = \pm 2f_0$ as well as at $f = 0$. Thus, by putting $x(t)$ through a quadratic transformation (a squarer in this case) we have converted the hidden periodicity resulting from the sine-wave factor $\cos (2\pi f_0 t)$ in (5) into first-order periodicity with associated spectral lines. This is particularly easy to see if $a(t)$ is the asynchronous random telegraph signal, which switches back and forth at random times between $+1$ and -1, because then $b(t) \equiv 1$ and $y(t)$ in (8) is therefore a periodic signal

$$y(t) = \tfrac{1}{2} + \tfrac{1}{2} \cos (4\pi f_0 t)$$

Example 2: *PAM.* As another example, we consider the pulse-amplitude modulated (PAM) signal

$$x(t) = \sum_{n=-\infty}^{\infty} a(nT_0) p(t - nT_0) \tag{11}$$

where the pulse $p(t)$ is confined within the interval

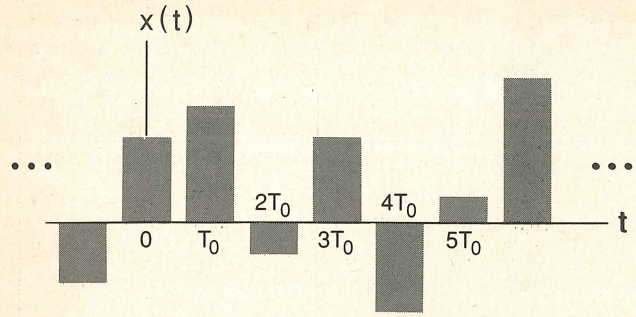

Fig. 2. A pulse-amplitude-modulated (PAM) signal with pulse width less than interpulse time.

$(-T_0/2, T_0/2)$ so that the pulse translates do not overlap, as shown in Fig. 2. The PSD of $x(t)$ is given by [1, Chapter 3, Sec. D]

$$S_x(f) = \frac{1}{T_0}|P(f)|^2 \sum_{m=-\infty}^{\infty} S_a(f - m/T_0) \tag{12}$$

where $S_a(f)$ is shown in Fig. 1a, which contains no spectral lines, and where $P(f)$ is the Fourier transform of $p(t)$. Since there are no spectral lines in $S_a(f)$ (or $P(f)$ since $p(t)$ has finite duration), there are none in $S_x(f)$, as shown in Fig. 3a, regardless of the periodic repetition of pulses in $x(t)$. But, let us look at the square of $x(t)$:

$$y(t) = x^2(t) = \sum_{n=-\infty}^{\infty} b(nT_0)q(t - nT_0) \tag{13}$$

where

$$b(nT_0) = a^2(nT_0) \tag{14a}$$

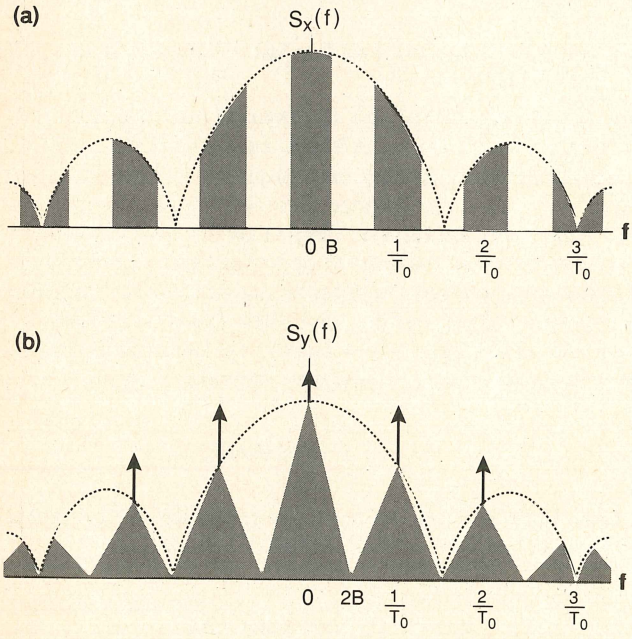

Fig. 3. a) Power spectral density (PSD) of a pulse-amplitude-modulated (PAM) signal with 67% duty-cycle pulses. b) PSD of the squared PAM signal.

and

$$q(t) = p^2(t) \tag{14b}$$

The PSD for $y(t)$ is given by

$$S_y(f) = \frac{1}{T_0}|Q(f)|^2 \sum_{m=-\infty}^{\infty} S_b(f - m/T_0) \tag{15}$$

where $Q(f)$ is the Fourier transform of $q(t)$. Because of the spectral line at $f = 0$ in $S_b(f)$, which is shown in Figure 1c, we have spectral lines in $S_y(f)$ at the harmonics m/T_0 (for some integer values of m) of the pulse rate $1/T_0$, as shown in Fig. 3b. Thus, again, we have converted the hidden periodicity in $x(t)$ into first-order periodicity with associated spectral lines by using a quadratic transformation. This is particularly easy to see if $a(nT_0)$ is a random binary sequence with values ± 1, because then $b(nT_0) \equiv 1$ and $y(t)$ in (13) is therefore a periodic signal

$$y(t) = \sum_{n=-\infty}^{\infty} q(t - n/T_0) \tag{16}$$

The cyclic autocorrelation function

Although the squaring transformation works in these examples, a different quadratic transformation involving delays can be required in some cases. For example, if $a(nT_0)$ is again binary, but $p(t)$ is flat with height 1 and width T_0, as shown in Fig. 4, then $y(t) = x^2(t) = 1$, which is a constant for all t. Thus, we have a spectral line at $f = 0$ but none at the harmonics of the pulse rate. Nevertheless, if we use the quadratic transformation

$$y(t) = x(t)x(t - \tau) \tag{17}$$

for any of a number of nonzero delays τ, we will indeed obtain spectral lines at $f = m/T_0$. That is,

$$\begin{aligned} M_y^{\alpha} &= \langle y(t)e^{-i2\pi\alpha t}\rangle \\ &= \langle x(t)x(t - \tau)e^{-i2\pi\alpha t}\rangle \neq 0 \end{aligned} \tag{18}$$

for $\alpha = m/T_0$ for some integers m.

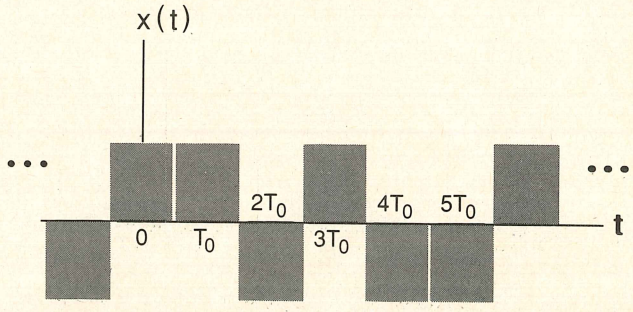

Fig. 4. A binary pulse-amplitude-modulated (PAM) signal with full duty-cycle pulses.

The most general time-invariant quadratic transformation is simply a linear combination of delay products of the form (17), viz.,

$$y(t) = \int h(\tau)x(t)x(t - \tau)\, d\tau$$

for some weighting function $h(\tau)$ that is analogous to the impulse-response function for a linear transformation. This motivates us to define the property of *second-order periodicity* as follows: The signal $x(t)$ contains second-order periodicity if and only if the PSD of the delay-product signal (17) for some delays τ contains spectral lines at some nonzero frequencies $\alpha \neq 0$, that is, if and only if (18) is satisfied.

In developing the theory of second-order periodicity it is more convenient to work with the symmetric delay product

$$y_\tau(t) = x(t + \tau/2)x^*(t - \tau/2) \tag{19}$$

(The complex conjugate * is introduced here for generality to accommodate complex-valued signals, but it is mentioned that for some complex-valued signals, the quadratic transformation without the conjugate can also be useful [1, Chapter 10, Sec. C].) Thus, the fundamental parameter (18) of second-order periodicity becomes

$$R_x^\alpha(\tau) \triangleq \langle x(t + \tau/2)x^*(t - \tau/2)e^{-i2\pi\alpha t}\rangle \tag{20}$$

which is the Fourier coefficient $M_{y_\tau}^\alpha$ of the additive sine-wave component with frequency α contained in the delay-product signal $y_\tau(t)$.

The notation $R_x^\alpha(\tau)$ is introduced for this Fourier coefficient because, for $\alpha = 0$, (20) reduces to the conventional autocorrelation function

$$R_x^0(\tau) = \langle x(t + \tau/2)x^*(t - \tau/2)\rangle \tag{21}$$

for which the notation $R_x(\tau)$ is commonly used. Furthermore, since $R_x^\alpha(\tau)$ is a generalization of the autocorrelation function, in which a cyclic (sinusoidal) weighting factor $e^{-i2\pi\alpha t}$ is included before the time averaging is carried out, $R_x^\alpha(\tau)$ is called the *cyclic autocorrelation function*.[2] Thus, we have two distinct interpretations of $R_x^\alpha(\tau) = M_{y_\tau}^\alpha$. In fact, we have yet a third distinct interpretation, which can be obtained by simply factoring $e^{-i2\pi\alpha t}$ in order to reexpress (20) as

$$R_x^\alpha(\tau) = \langle [x(t + \tau/2)e^{-i\pi\alpha(t+\tau/2)}][x(t - \tau/2)e^{+i\pi\alpha(t-\tau/2)}]^*\rangle \tag{22}$$

That is, $R_x^\alpha(\tau)$ is actually a conventional crosscorrelation function

$$R_{uv}(\tau) \triangleq \langle u(t + \tau/2)v^*(t - \tau/2)\rangle = R_x^\alpha(\tau) \tag{23}$$

where

$$u(t) = x(t)e^{-i\pi\alpha t} \tag{24a}$$

and

$$v(t) = x(t)e^{+i\pi\alpha t} \tag{24b}$$

Recall that multiplying a signal by $e^{\pm i\pi\alpha t}$ shifts the spectral content of the signal by $\pm\alpha/2$. For example, the Fourier transforms of $u(t)$ and $v(t)$ (if they exist) are

$$U(f) = X(f + \alpha/2) \tag{25a}$$

and

$$V(f) = X(f - \alpha/2) \tag{25b}$$

Similarly, their PSDs are

$$S_u(f) = S_x(f + \alpha/2) \tag{26a}$$

and

$$S_v(f) = S_x(f - \alpha/2) \tag{26b}$$

It follows from (23)–(24) that $x(t)$ exhibits second-order periodicity ((20) is not identically zero as a function of τ for some $\alpha \neq 0$) if and only if *frequency translates* (frequency-shifted versions) of $x(t)$ (namely $u(t)$ and $v(t)$) are correlated with each other ((23) is not identically zero as a function of τ for some $\alpha \neq 0$ in (24)). This third interpretation of $R_x^\alpha(\tau)$ suggests an appropriate way to normalize $R_x^\alpha(\tau)$ as explained next.

As long as the mean values of the frequency translates $u(t)$ and $v(t)$ are zero (which means that $x(t)$ does not contain finite-strength[3] additive sine-wave components at frequencies $\pm\alpha/2$ and, therefore, that $S_x(f)$ has no spectral lines at $f = \pm\alpha/2$), the crosscorrelation $R_{uv}(\tau) \equiv R_x^\alpha(\tau)$ is actually a temporal crosscovariance $K_{uv}(\tau)$. That is,

$$\begin{aligned} K_{uv}(\tau) &\triangleq \langle [u(t + \tau/2) - \langle u(t + \tau/2)\rangle] \\ &\quad \cdot [v(t - \tau/2) - \langle v(t - \tau/2)\rangle]^*\rangle \\ &= \langle u(t + \tau/2)v^*(t - \tau/2)\rangle = R_{uv}(\tau) \end{aligned} \tag{27}$$

An appropriate normalization for the temporal crosscovariance is the geometric mean of the two corresponding temporal variances. This yields a temporal correlation coefficient, the magnitude of which is upper bounded by unity. It follows from (24) that the two variances are given by

$$K_u(0) = R_u(0) = \langle |u(t)|^2\rangle = R_x(0) \tag{28a}$$

and

$$K_v(0) = R_v(0) = \langle |v(t)|^2\rangle = R_x(0) \tag{28b}$$

Therefore, the temporal correlation coefficient for frequency translates is given by

$$\frac{K_{uv}(\tau)}{[K_u(0)K_v(0)]^{1/2}} = \frac{R_x^\alpha(\tau)}{R_x(0)} \triangleq \gamma_x^\alpha(\tau) \tag{29}$$

[2]Although some readers will recognize the similarity between the cyclic autocorrelation function and the radar ambiguity function, the relationship between these two functions is only superficial. The concepts and theory underlying the cyclic autocorrelation function, as summarized in this article, have little in common with the concepts and theory of radar ambiguity (cf. [1, Chapter 10, Sec. C]).

[3]It does contain infinitesimal sine-wave components.

Hence, the appropriate normalization factor for $R_x^\alpha(\tau)$ is simply $1/R_x(0)$.

This is a good point at which to introduce some more terminology. A signal $x(t)$ for which the autocorrelation $R_x(\tau)$ exists (e.g., remains finite as the averaging time T goes to infinity) and is not identically zero (as it is for transient signals) is commonly said to be *stationary* (in the wide sense or of second order). But we need to refine the terminology to distinguish between those stationary signals that exhibit second-order periodicity ($R_x^\alpha(\tau) \not\equiv 0$ for some $\alpha \neq 0$) and those stationary signals that do not ($R_x^\alpha(\tau) \equiv 0$ for all $\alpha \neq 0$). Consequently, we shall call the latter for which $R_x^\alpha(\tau) \equiv 0$ *purely stationary* (of second order) and the former for which $R_x^\alpha(\tau) \not\equiv 0$ *cyclostationary* (of second order). We shall also call any nonzero value of the frequency parameter α for which $R_x^\alpha(\tau) \not\equiv 0$ a *cycle frequency*. The case $\alpha = 0$ can be considered to be a degenerate cycle frequency, since $e^{i2\pi\alpha t} = e^0 = 1$ is a degenerate sinusoid. Thus, some stationary signals are also cyclostationary or *exhibit cyclostationarity*; only stationary signals that do not exhibit cyclostationarity are purely stationary. The discrete set of cycle frequencies is called the *cycle spectrum*. For example, if a signal exhibits only one fundamental period of second-order periodicity, the cycle spectrum contains only harmonics (integer multiples) of the fundamental cycle frequency, which is the reciprocal of the fundamental period. But if there are multiple incommensurate periods, then the cycle spectrum contains harmonics of each of the incommensurate fundamental cycle frequencies.

We conclude this section by reconsidering the AM example and determining the cyclic autocorrelation function for the AM signal.

***Example 1** continued: AM.* Let $a(t)$ be a real random purely stationary signal with zero mean:

$$\langle a(t) \rangle = 0 \tag{30a}$$

$$\langle a(t + \tau/2)a^*(t - \tau/2) \rangle \not\equiv 0 \tag{30b}$$

$$\langle a(t + \tau/2)a^*(t - \tau/2)e^{-i2\pi\alpha t} \rangle \equiv 0 \text{ for all } \alpha \neq 0 \tag{30c}$$

Equation (30c) guarantees that

$$\langle a(t)e^{-i2\pi\alpha t} \rangle \equiv 0 \text{ for all } \alpha \neq 0 \tag{30d}$$

We consider the amplitude-modulated sinewave

$$\begin{aligned} x(t) &= a(t) \cos(2\pi f_0 t + \theta) \\ &= \tfrac{1}{2} a(t) [e^{i(2\pi f_0 t + \theta)} + e^{-i(2\pi f_0 t + \theta)}] \end{aligned} \tag{31}$$

Because of (30d), $a(t)$ contains no finite-strength additive sine-wave components and, therefore (together with (30a)), $x(t)$ contains no finite-strength additive sine-wave components. This means that its power spectral density contains no spectral lines. However, the quadratic transformation

$$\begin{aligned} y_\tau(t) &= x(t + \tau/2)x^*(t - \tau/2) \\ &= a(t + \tau/2)a^*(t - \tau/2) \tfrac{1}{4} [e^{i2\pi f_0 \tau} + e^{-i2\pi f_0 \tau} \\ &\quad + e^{i(4\pi f_0 t + 2\theta)} + e^{-i(4\pi f_0 t + 2\theta)}] \end{aligned} \tag{32}$$

does contain finite-strength additive sine-wave components with frequencies $\alpha = \pm 2f_0$, since (30b) renders one or the other of the last two terms in the quantity

$$\begin{aligned} \langle y_\tau(t)e^{-i2\pi\alpha t} \rangle &= \tfrac{1}{4} e^{i2\pi f_0 \tau} \langle a(t + \tau/2)a^*(t - \tau/2)e^{-i2\pi\alpha t} \rangle \\ &\quad + \tfrac{1}{4} e^{-i2\pi f_0 \tau} \langle a(t + \tau/2)a^*(t - \tau/2)e^{-i2\pi\alpha t} \rangle \\ &\quad + \tfrac{1}{4} e^{i2\theta} \langle a(t + \tau/2)a^*(t - \tau/2)e^{-i2\pi(\alpha - 2f_0)t} \rangle \\ &\quad + \tfrac{1}{4} e^{-i2\theta} \langle a(t + \tau/2)a^*(t - \tau/2)e^{-i2\pi(\alpha + 2f_0)t} \rangle \end{aligned} \tag{33}$$

nonzero for $\alpha = \pm 2f_0$. That these are the only two cycle frequencies $\alpha \neq 0$ follows from the fact that (30c) renders (33) equal to zero for all α except $\alpha = 0$ and $\alpha = \pm 2f_0$. Thus, the cycle spectrum consists of only the two cycle frequencies $\alpha = \pm 2f_0$ and the degenerate cycle frequency $\alpha = 0$.

Hence, the versions $u(t)$ and $v(t)$ of $x(t)$ obtained by frequency shifting $x(t)$ up and down by $\alpha/2 = f_0$ are correlated. This is not surprising since (31) reveals that $x(t)$ is obtained from $a(t)$ by frequency shifting up and down by f_0 and then adding. In conclusion, from (33) we have the cyclic autocorrelation function

$$R_x^\alpha(\tau) = \begin{cases} \frac{1}{4} e^{\pm i2\theta} R_a(\tau) & \text{for } \alpha = \pm 2f_0 \\ \frac{1}{2} R_a(\tau) \cos(2\pi f_0 \tau) & \text{for } \alpha = 0 \\ 0 & \text{otherwise} \end{cases} \tag{34}$$

from which it follows that the temporal correlation coefficient is given by

$$\gamma_x^\alpha(\tau) = \begin{cases} \frac{1}{2} e^{\pm i2\theta} \gamma_a^0(\tau) & \text{for } \alpha = \pm 2f_0 \\ 0 & \text{otherwise} \end{cases} \tag{35a}$$

Thus, the strength of correlation between $x(t + \tau/2)e^{-i\pi\alpha(t + \tau/2)}$ and $x(t - \tau/2)e^{i\pi\alpha(t - \tau/2)}$, which is given by

$$|\gamma_x^\alpha(\tau)| = \tfrac{1}{2} |\gamma_a^0(\tau)| \tag{35b}$$

can be substantial for this amplitude-modulated signal, e.g., $|\gamma_x^\alpha(0)| = \frac{1}{2}$

As an especially simple specific example of $a(t)$, we consider as before a stationary random telegraph wave, which switches back and forth between +1 and −1 at random (Poisson distributed) switching times [2, Chapter 6, exc. 12]. If we consider $\tau = 0$ in (32), we obtain

$$\begin{aligned} y_0(t) &= |x(t)|^2 = |a(t)|^2 \cos^2(2\pi f_0 t + \theta) \\ &= \tfrac{1}{2} + \tfrac{1}{2} \cos(4\pi f_0 t + 2\theta) \end{aligned}$$

which clearly contains finite-strength additive sine-wave components with frequencies $\alpha = \pm 2f_0$. In fact, in this very special case, there is no random component in $y_0(t)$. On the other hand, for $\tau \neq 0$, $y_\tau(t)$ contains both a sine-wave component and a random component. ■

Other examples of cyclostationary signals can be similarly viewed as mixtures of stationarity and periodicity. Examples are cited in the introductory Section. Typical cycle spectra include harmonics of pulse rates, keying rates, spreading-code chipping rates, frequency hopping rates, code repetition rates, doubled carrier frequencies, and sums and differences of these [1, Chapter 12].

THE SPECTRAL-CORRELATION DENSITY FUNCTION

Definition of the SCD

In the same way that it is beneficial for some purposes to localize in the frequency domain the average power $\langle |x(t)|^2 \rangle = R_x(0)$ in a stationary random signal, it can be very helpful to localize in frequency the correlation $\langle u(t)v^*(t) \rangle = \langle |x(t)|^2 \, e^{-i2\pi\alpha t} \rangle \equiv R_x^\alpha(0)$ of frequency-shifted signals $u(t)$ and $v(t)$ for a cyclostationary random signal $x(t)$. In the former case of localizing the power, we simply pass the signal of interest $x(t)$ through a narrowband bandpass filter and then measure the average power at the output of the filter. By doing this with many filters whose center frequencies are separated by the bandwidth of the filters, we can partition any spectral band of interest into a set of contiguous narrow disjoint bands. In the limit as the bandwidths approach zero, the corresponding set of measurements of average power, normalized by the bandwidth, constitute the power spectral density (PSD) function. That is, at any particular frequency f, the PSD for $x(t)$ is given by

$$S_x(f) \triangleq \lim_{B \to 0} \frac{1}{B} \langle |h_B^f(t) \otimes x(t)|^2 \rangle \qquad (36)$$

where $\otimes$ denotes convolution and $h_B^f(t)$ is the impulse response of a one-sided bandpass filter with center frequency f, bandwidth B, and unity gain at the band center, Fig. 5.

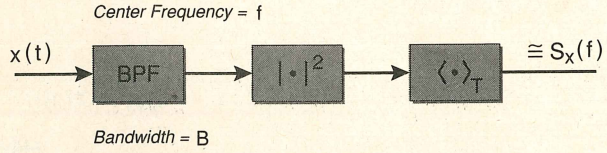

Fig. 5. One channel of a spectrum analyzer for measuring the power spectral density (PSD). (The symbol ≅ indicates that the output only approximates the ideal function $S_x(f)$ for finite T and B.)

In the latter case of localizing the correlation, we simply pass both of the two frequency translates $u(t)$ and $v(t)$ of $x(t)$ through the same set of bandpass filters as are used for the PSD and then measure the temporal correlation of the filtered signals to obtain

$$S_x^\alpha(f) \triangleq \lim_{B \to 0} \frac{1}{B} \langle [h_B^f(t) \otimes u(t)] \, [h_B^f(t) \otimes v(t)]^* \rangle \qquad (37)$$

which is called the *spectral-correlation density (SCD) function*, Fig. 6. This yields the spectral density of correlation in $u(t)$ and $v(t)$ at frequency f, which is identical to the spectral density of correlation in $x(t)$ at frequencies $f + \alpha/2$ and $f - \alpha/2$, Fig. 7. That is, $S_x^\alpha(\tau)$ is the bandwidth-normalized (i.e., divided by B) correlation of the amplitude and phase fluctuations of the

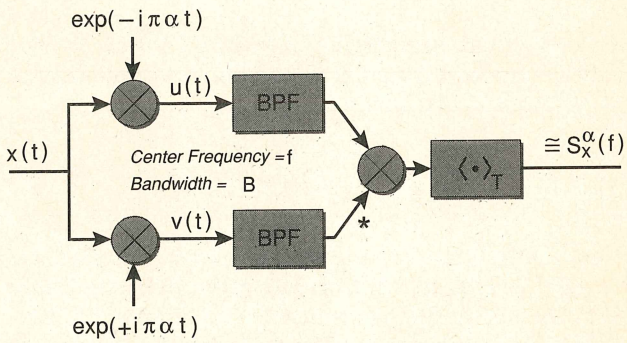

Fig. 6. One channel-pair of a spectral-correlation analyzer (or a cyclic-spectrum analyzer) for measuring the spectral-correlation density (or cyclic spectral density).

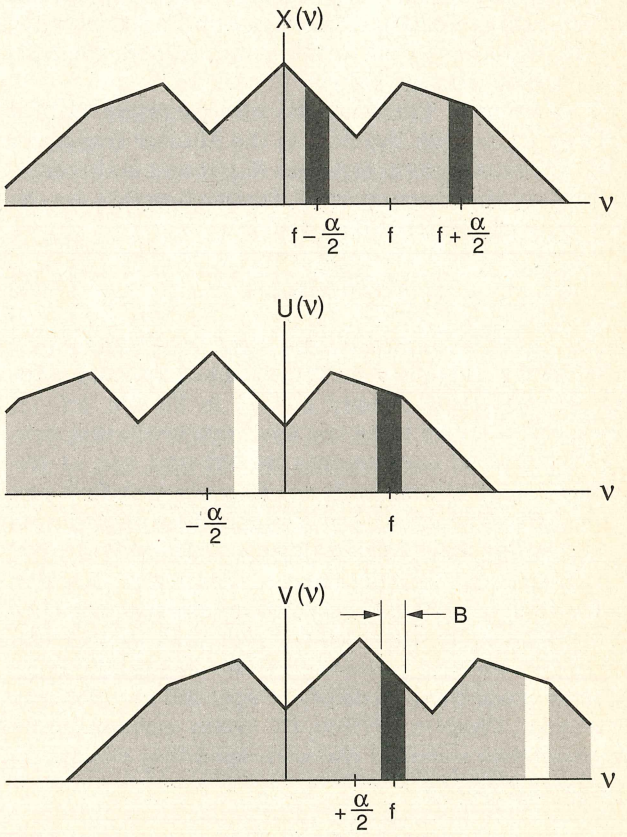

Fig. 7. Illustration of spectral bands used in the measurement of the spectral-correlation density $S_x^\alpha(f)$. (ν is a dummy frequency variable; the light and dark bands are the bands selected by the BPFs.)

narrowband spectral components in $x(t)$ centered at frequencies $f + \alpha/2$ and $f - \alpha/2$, in the limit as the bandwidth B of these narrowband components approaches zero.

It is well known (see, for example, [1, Chapter 3, Sec. C] for a proof) that the PSD obtained from (36) is equal to the Fourier transform of the autocorrelation function,

$$S_x(f) = \int_{-\infty}^{\infty} R_x(\tau) e^{-i2\pi f\tau}\, d\tau \tag{38}$$

Similarly, it can be shown [1, Chapter 11, Sec. C] that the SCD obtained from (37) is the Fourier transform of the cyclic autocorrelation function,

$$S_x^\alpha(f) = \int_{-\infty}^{\infty} R_x^\alpha(\tau) e^{-i2\pi f\tau}\, d\tau \tag{39}$$

Relation (38) is known as the *Wiener relation* (see, for example [1, Chapter 3, Sec. C]), and (39) is therefore called the *cyclic Wiener relation* [1, Chapter 11, Sec. C]. The cyclic Wiener relation includes the Wiener relation as the special case of $\alpha = 0$. (In the probabilistic framework of stochastic processes, which is based on expected values [ensemble averages] instead of time averages, the probabilistic counterpart of (38) is known as the *Wiener-Khinchin relation*, and therefore the probabilistic counterpart of (39) is called the *cyclic Wiener-Khinchin relation* [2, Chapter 12, Sec. 12.2].) Because of the relation (39), the SCD is also called the *cyclic spectral density* (CSD) function [1, Chapter 10, Sec. B]. Unlike the PSD which is real valued, the SCD is in general complex valued.

It follows from (39) and the interpretation (23) of $R_x^\alpha(\tau)$ as $R_{uv}(\tau)$ that the SCD is the Fourier transform of the crosscorrelation function $R_{uv}(\tau)$ and is therefore identical to the cross-spectral density function for the frequency translates $u(t)$ and $v(t)$:

$$S_x^\alpha(f) \equiv S_{uv}(f) \tag{40}$$

where $S_{uv}(f)$ is defined by the right hand side of (37) for arbitrary $u(t)$ and $v(t)$. This is to be expected since the cross-spectral density $S_{uv}(f)$ is known (cf. [1, Chapter 7, Sec. A]) to be the spectral-correlation density for spectral components in $u(t)$ and $v(t)$ at frequency f, and $u(t)$ and $v(t)$ are frequency-shifted versions of $x(t)$. The identity (40) suggests an appropriate normalization for $S_x^\alpha(f)$: as long as the PSDs of $u(t)$ and $v(t)$ contain no spectral lines at frequency f, which means that the PSD of $x(t)$ contains no spectral lines at either of the frequencies $f \pm \alpha/2$, then the correlation of the spectral components (40) is actually a covariance since the means of the spectral components are zero [1, Chapter 11, Sec. C]. When normalized by the geometric mean of the corresponding variances, which are given by

$$S_u(f) = S_x(f + \alpha/2) \tag{41a}$$

and

$$S_v(f) = S_x(f - \alpha/2) \tag{41b}$$

the covariance becomes a correlation coefficient:

$$\frac{S_{uv}(f)}{[S_u(f)S_v(f)]^{1/2}} = \frac{S_x^\alpha(f)}{[S_x(f + \alpha/2)S_x(f - \alpha/2)]^{1/2}} \triangleq \rho_x^\alpha(f) \tag{42}$$

Since $|\rho_x^\alpha(f)|$ is bounded to the interval $[0,1]$, it is a convenient measure of the degree of local spectral redundancy that results from spectral correlation. For example, for $|\rho_x^\alpha(f)| = 1$, we have complete spectral redundancy at $f + \alpha/2$ and $f - \alpha/2$.

Let us now return to the AM example considered in the preceding Section.

***Example 1** continued: AM.* By Fourier transforming (34) and invoking the cyclic Wiener relation (39), we obtain the following SCD function for the amplitude-modulated signal (31):

$$S_x^\alpha(f) = \begin{cases} \frac{1}{4} e^{\pm i2\theta} S_a(f) & \text{for } \alpha = \pm 2f_o \\ \frac{1}{4} S_a(f + f_o) + \frac{1}{4} S_a(f - f_o) & \text{for } \alpha = 0 \\ 0 & \text{otherwise} \end{cases} \tag{43}$$

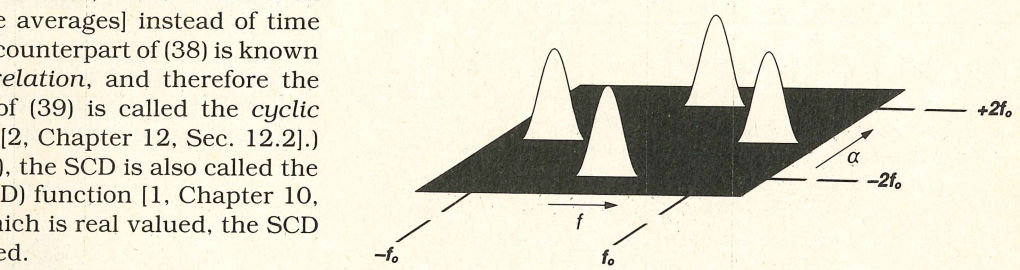

Fig. 8. Magnitude of the spectral-correlation density function for an AM signal graphed as a height above the bifrequency plane with coordinates f and α.

The magnitude of this SCD is graphed in Fig. 8 as the height of a surface above the bifrequency plane with coordinates f and α. For purposes of illustration, $a(t)$ is assumed to have an arbitrary low-pass PSD for this graph. Observe that, although the argument f of the SCD is continuous, as it always will be for a random signal, the argument α is discrete, as it always will be since it represents the harmonic frequencies of periodicities underlying the random process (the sine-wave carrier in this example).

It follows from (43) that the spectral correlation coefficient is given by

$$\rho_x^\alpha(f) = \frac{S_a(f)e^{\pm i2\theta}}{\{[S_a(f + 2f_o) + S_a(f)]\,[S_a(f) + S_a(f - 2f_o)]\}^{1/2}} \quad \text{for } \alpha = \pm 2f_o \tag{44a}$$

Thus, the strength of correlation between spectral component in $x(t)$ at frequencies $f + \alpha/2$ and $f - \alpha/2$ is unity:

$$|\rho_x^\alpha(f)| = 1 \quad \text{for } |f| < f_o \text{ and } \alpha = \pm 2f_o \tag{44b}$$

provided that $a(t)$ is bandlimited to $|f| \le f_o$,

$$S_a(f) = 0 \qquad \text{for } |f| \ge f_o \tag{45}$$

This is not surprising since the two spectral components in $x(t)$ at frequencies $f \pm \alpha/2 = f \pm f_o$ are obtained from the single spectral component in $a(t)$ at frequency f simply by shifting and scaling. Thus, they are perfectly correlated. That is, the upper (lower) sideband for $f > 0$ carries exactly the same information as the lower (upper) sideband for $f < 0$. Techniques for exploiting this *spectral redundancy* are described in a later section. ■

Before considering other examples of the SCD, let us first gain an understanding of the effects of some basic signal processing operations on the SCD. This greatly facilitates the determination of the SCD for commonly encountered manmade signals.

Filtering

When a signal $x(t)$ undergoes a linear time-invariant (LTI) transformation, (i.e., a convolution or a filtering operation),

$$z(t) = h(t) \otimes x(t) \triangleq \int_{-\infty}^{\infty} h(u)x(t-u) \tag{46}$$

the spectral components in $x(t)$ are simply scaled by the complex-valued transfer function $H(f)$, which is the Fourier transform

$$H(f) = \int_{-\infty}^{\infty} h(t)e^{-i2\pi ft}\,dt \tag{47}$$

of the impulse-response function $h(t)$ of the transformation. As a result, the PSD gets scaled by the squared magnitude of $H(f)$ (see, for example, [1, Chapter 3, Sec. C] or [2, Chapter 10, Sec. 10.1])

$$S_z(f) = |H(f)|^2 S_x(f) \tag{48}$$

Equation (48) can be derived from the definition (36) of the PSD. Similarly, because the spectral components of $x(t)$ at frequencies $f \pm \alpha/2$ are scaled by $H(f \pm \alpha/2)$, the SCD gets scaled by the product $H(f + \alpha/2)H^*(f - \alpha/2)$:

$$S_z^\alpha(f) = H(f + \alpha/2)H^*(f - \alpha/2)S_x^\alpha(f) \tag{49}$$

This result, called the *input-output SCD relation for filtering*, which can be derived from the definition (37) of the SCD, includes (48) as the special case of $\alpha = 0$. Observe that it follows from (49) and the definition (42) that

$$|\rho_z^\alpha(f)| \equiv |\rho_x^\alpha(f)| \tag{50}$$

That is, the magnitude of the spectral correlation coefficient is unaffected by filtering (if $H(f \pm \alpha/2) \neq 0$).

Example 3: *Time Delay.* As our first example of (49), we consider a filter that simply delays the input by t_o; then $h(t) = \delta(t - t_o)$ and $H(f) = e^{-i2\pi ft_o}$. Therefore, for $z(t) = x(t - t_o)$, we obtain from the input-output SCD relation (49)

$$S_z^\alpha(f) = S_x^\alpha(f)e^{-i2\pi\alpha t_o} \tag{51}$$

which indicates that, unlike the PSD, the SCD of a cyclostationary signal is sensitive to the timing or phase of the signal.

Example 4: *Multipath Propagation.* As a second example of (49), if $x(t)$ undergoes multipath propagation during transmission to yield a received signal

$$z(t) = \sum_n a_n x(t - t_n)$$

where a_n and t_n are the attenuation factor and delay of the nth propagation path, we have

$$H(f) = \sum_n a_n e^{-i2\pi ft_n} \tag{52}$$

and therefore (49) yields

$$S_z^\alpha(f) = S_x^\alpha(f) \sum_{n,m} a_n a_m^* \cdot \exp\{-i2\pi[f(t_n - t_m) + \alpha(t_n + t_m)/2]\} \tag{53}$$

Example 5: *Bandpass Signals.* As a third example of the utility of the relation (49), let us determine the support region in the (f, α) plane for a bandpass signal with lowest frequency b and highest frequency B. To enforce such a spectrum, we can simply put any signal $x(t)$ through an ideal bandpass filter with transfer function

$$H(f) = \begin{cases} 1 & \text{for } b < |f| < B \\ 0 & \text{otherwise} \end{cases}$$

It then follows directly from the input-output SCD relation (49) that the SCD for the output of this filter can be nonzero only for $||f| - |\alpha|/2| > b$ and $|f| + |\alpha|/2 < B$:

$$S_z^\alpha(f) = \begin{cases} 0 & \text{for } ||f| - |\alpha|/2| \le b \text{ or } |f| + |\alpha|/2 \ge B \\ S_x^\alpha(f) & \text{otherwise} \end{cases} \tag{54}$$

This shows that the support region in the (f, α) plane for a bandpass signal is the four diamonds located at the vertices of a larger diamond, depicted in Fig. 9a. By letting $b \to 0$, we obtain the support region for a lowpass signal, and by letting $B \to \infty$, we obtain the support region for a highpass signal. This is shown in Figs. 9b and 9c.

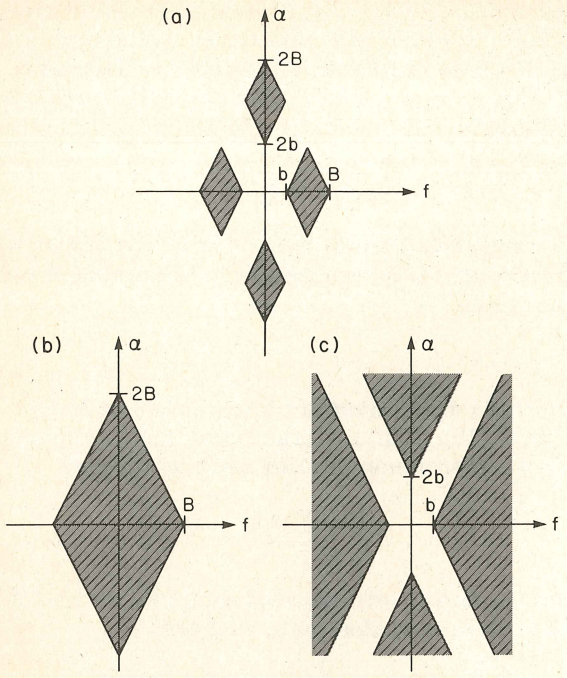

Fig. 9. a) Support region in the bifrequency plane for the spectral-correlation density function of a bandpass signal. b) Support region for a lowpass signal. c) Support region for a highpass signal.

Waveform multiplication

When finite segments of two signals are multiplied together, we know from the convolution theorem that their Fourier transforms get convolved. From this, we expect some sort of convolution relation to hold for the SCDs of signals passing through a product modulator. In fact, it can be shown [1, Chapter 11, Sec. C], [2, Chapter 12, exc. 41] that if $x(t)$ is obtained by multiplying together two statistically independent[4] time-series $r(t)$ and $s(t)$,

$$x(t) = r(t)s(t) \tag{55}$$

then the cyclic autocorrelation of $x(t)$ is given by the discrete convolution in cycle frequency of the cyclic autocorrelations of $r(t)$ and $s(t)$:

$$R_x^\alpha(\tau) = \sum_\beta R_r^\beta(\tau) R_s^{\alpha-\beta}(\tau) \tag{56}$$

where, for each α, β ranges over all values for which $R_r^\beta(\tau) \not\equiv 0$. By Fourier transforming (56), we obtain the *input-output SCD relation for waveform multiplication*:

$$S_x^\alpha(f) = \int_{-\infty}^{\infty} \sum_\beta S_r^\beta(\nu) S_s^{\alpha-\beta}(f-\nu)\, d\nu \tag{57}$$

which is a double convolution that is continuous in the variable f and discrete in the variable α.

Example 6: *Frequency Conversion.* As an example of (57), if $s(t)$ is simply a sinusoid,

$$s(t) = \cos(2\pi f_o t + \theta)$$

the product modulator becomes a frequency converter when followed by a filter to select either the up-converted version or the down-converted version of $r(t)$. By applying first the input-output SCD relation (57) for the product modulator (which applies since a sinusoid is statistically independent of all time-series [1, Chapter 15, Sec. A]), and then (49) for the filter, we can determine the up-converted or down-converted SCD. To illustrate, we first determine the SCD for the sinusoid $s(t)$. By substituting the sinusoid $s(t)$ into the definition of the cyclic autocorrelation, we obtain

$$R_s^\alpha(\tau) = \begin{cases} \frac{1}{2}\cos(2\pi f_o \tau) & \text{for } \alpha = 0 \\ \frac{1}{4} e^{\pm i 2\theta} & \text{for } \alpha = \pm 2f_o \\ 0 & \text{otherwise} \end{cases} \tag{58}$$

Fourier transforming then yields the SCD

$$S_s^\alpha(f) = \begin{cases} \frac{1}{4}\delta(f-f_o) + \frac{1}{4}\delta(f+f_o) & \text{for } \alpha = 0 \\ \frac{1}{4} e^{\pm i 2\theta}\,\delta(f) & \text{for } \alpha = \pm 2f_o \\ 0 & \text{otherwise} \end{cases} \tag{59}$$

which is illustrated in Fig. 10a. Using (57), we convolve this SCD with that of a stationary signal $r(t)$, for which

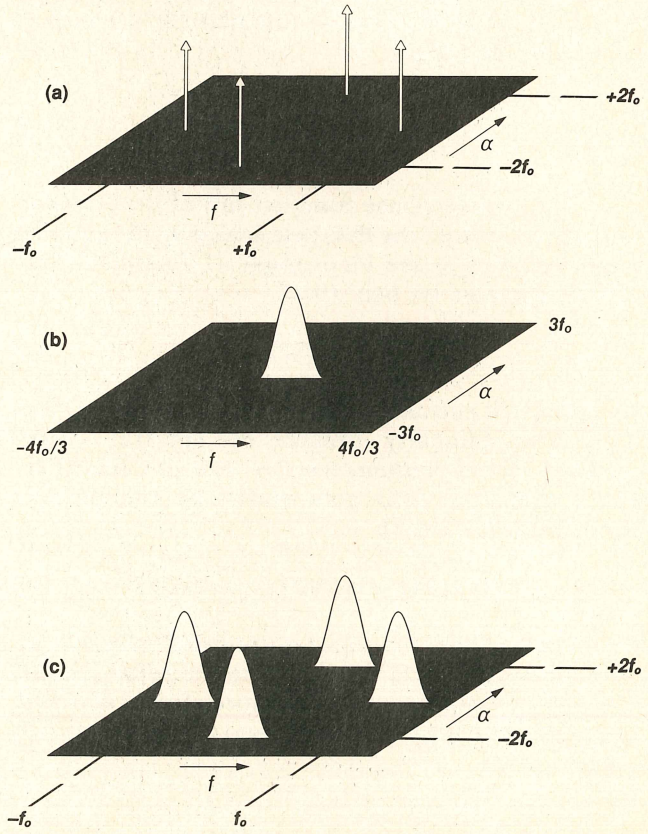

Fig. 10. a) Magnitude of the spectral-correlation density (SCD) for a sine wave of frequency f_0. b) SCD for a lowpass stationary signal. c) SCD magnitude for the product of signals corresponding to a) and b), obtained by convolving the SCDs in a) and b).

[4]Time-series are statistically independent if their joint fraction-of-time probability densities factor into products of individual fraction-of-time probability densities, as explained in [1, Chapter 15, Sec. A].

$$S_r^\alpha(f) = \begin{cases} S_r(f) & \text{for } \alpha = 0 \\ 0 & \text{for } \alpha \neq 0 \end{cases} \tag{60}$$

which is illustrated in Fig. 10b. The result is that the SCD of the stationary signal simply gets replicated and scaled at the four locations of the impulses in the SCD of the sinusoid, as illustrated in Fig. 10c.

Example 7: *Time Sampling.* Another important signal processing operation is periodic time sampling. It is known that for a purely stationary signal $x(t)$, the PSD $\tilde{S}_x(f)$ of the sequence of samples $\{x(nT_s)$: $n = 0, \pm 1, \pm 2, \cdots\}$ is related to the PSD $S_x(f)$ of the waveform by the aliasing formula [1, Chapter 3, Sec. E], [2, Chapter 11, Sec. 11.1]

$$\tilde{S}_x(f) = \frac{1}{T_s} \sum_{n=-\infty}^{\infty} S_x\left(f - \frac{n}{T_s}\right) \tag{61}$$

It is shown in [1, Chapter 11, Sec. C], [2, Chapter 12, Sec. 12.4] that this aliasing formula generalizes for the SCD to

$$\tilde{S}_x^\alpha(f) = \frac{1}{T_s} \sum_{m,n=-\infty}^{\infty} S_x^{\alpha + m/T_s}\left(f - \frac{m}{2T_s} - \frac{n}{T_s}\right) \tag{62}$$

We also can obtain essentially the same result (except for an additional factor of $1/T_s$) by applying the input-output SCD relation for waveform multiplication (57) to an impulse sampler, which is a product modulator with one input equal to a periodic train of impulses,

$$s(t) = \sum_{n=-\infty}^{\infty} \delta(t - nT_s) \tag{63}$$

This alternative approach can be carried out (cf. [1, Chapter 11, Sec. C]) by formally expanding (63) into the Fourier series

$$s(t) = \frac{1}{T_s} \sum_{m=-\infty}^{\infty} e^{i2\pi mt/T_s} \tag{64}$$

■

Observe that, when $x(t)$ is not purely stationary (i.e., when $S_x^\alpha(f) \not\equiv 0$ for $\alpha = m/T_s$ for some nonzero integers m), the conventional PSD aliasing formula (61) must be corrected according to (62) evaluated at $\alpha = 0$:

$$\tilde{S}_x(f) = \frac{1}{T_s} \sum_{m,n=-\infty}^{\infty} S_x^{m/T_s}\left(f - \frac{m}{2T_s} - \frac{n}{T_s}\right) \tag{65}$$

This reflects the fact that, when aliased overlapping spectral components add together, their PSD values add only if they are uncorrelated. When they are correlated, as in a cyclostationary signal, the PSD value of the sum of overlapping aliased components depends on the particular magnitudes and phases of their correlations. The SCD aliasing formula (62) is illustrated graphically in Fig. 11, where the support regions for the SCD $\tilde{S}_x^\alpha(f)$ for the sequence of samples $\{x(nT_s)\}$ is depicted in terms of the single diamond support region for a lowpass waveform $x(t)$, which is shown in Fig. 9b.

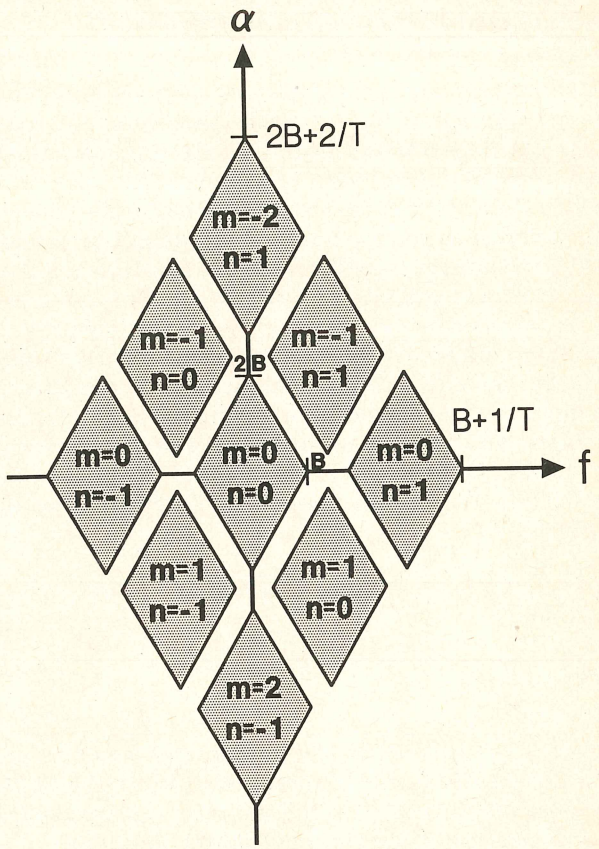

Fig. 11. Illustration of support regions in the bifrequency plane for the spectral-correlation densities that are aliased by periodic time-sampling.

Discrete time

Before leaving the topic of time sampling, it is explained that the discrete-time counterpart of the symmetric definition of the cyclic autocorrelation function (20), which uses delays of $+\tau/2$ and $-\tau/2$, is not appropriate since time samples midway between the given samples $\{x(nT_s)\}$ are not available. Therefore, the asymmetric definition

$$\tilde{R}_x^\alpha(kT_s) \triangleq \langle x(nT_s + kT_s)x^*(nT_s)e^{-i2\pi\alpha nT_s}\rangle\, e^{-i\pi\alpha kT_s} \tag{66}$$

(where $\langle\cdot\rangle$ denotes discrete-time averaging over n), which uses delays of τ and 0 and which includes the correction factor $e^{-i\pi\alpha\tau}$ for $\tau = kT_s$ that makes the asymmetric definition agree with the symmetric definition, has been adopted [1, Chapter 11, Sec. C], [2, Chapter 12, Sec. 12.4].

The discrete-time counterpart of the SCD can be defined just as is done at the beginning of this Section (but with a discrete-time bandpass filter instead of the continuous-time filter used there). The SCD can then be shown to be the discrete-time Fourier transform of the cyclic autocorrelation [1, Chapter 11, Sec. C], [2, Chapter 12, Sec. 12.4]:

$$\tilde{S}_x^{\alpha}(f) = \sum_{k=-\infty}^{\infty} \tilde{R}_x^{\alpha}(kT_s)\, e^{-i2\pi kT_s f} \tag{67}$$

Periodically time-variant filtering

Many signal processing devices such as pulse and carrier modulators, multiplexers, samplers, and scanners, can be modeled as periodically time-variant filters, especially if multiple incommensurate periodicities (i.e., periodicities that are not harmonically related) are included in the model. By expanding the periodically time-variant impulse-response function in a Fourier series as explained shortly, any such system can be represented by a parallel bank of sinusoidal product modulators followed by time-invariant filters. Consequently, the effect of any such system on the SCD of its input can be determined by using the SCD relations for filters and product modulators. In particular, it can be shown [1, Chapter 11, Sec. D] that the SCD of the output $z(t)$ of a multiply-periodic system with input $x(t)$ is given by

$$S_z^{\alpha}(f) = \sum_{\beta,\gamma \in A} G_{\beta}(f + \alpha/2) G_{\gamma}^{*}(f - \alpha/2) \cdot S_x^{\alpha-\beta+\gamma}\left(f - \frac{\beta+\gamma}{2}\right) \tag{68}$$

where $G_{\beta}(f)$ are the transfer functions of the filters and A is the set of sinusoid frequencies associated with the product modulators in the system representation. More specifically, for the input-output equation

$$z(t) = \int_{-\infty}^{\infty} h(t, u)x(u)\, du \tag{69}$$

the multiply-periodic impulse-response function $h(t, u)$ can be expanded in the Fourier series

$$h(t + \tau, t) = \sum_{\beta \in A} g_{\beta}(\tau) e^{i2\pi\beta t} \tag{70}$$

where the Fourier coefficients (for each τ) are given by

$$g_{\beta}(\tau) = \langle h(t + \tau, t) e^{-i2\pi\beta t} \rangle \tag{71}$$

It follows from (69)–(70) that the filter output can be expressed as

$$z(t) = \sum_{\beta \in A} [x(t) e^{i2\pi\beta t}] \otimes g_{\beta}(t) \tag{72}$$

where $g_{\beta}(t)$ are the impulse-response functions of the filters with corresponding transfer functions $G_{\beta}(f)$. Thus, periodically time-variant filters perform time-invariant filtering on frequency-shifted versions $x(t)e^{i2\pi\beta t}$ of the input. This results in summing scaled, frequency-shifted, cycle-frequency-shifted versions of the SCD for the input $x(t)$ to obtain the SCD for the output $z(t)$, as indicated in (68).

Let us now consider a couple of additional examples of modulation types, making use of the results obtained in the preceding paragraphs to determine SCDs.

***Example 2** continued: PAM.* Let $\{a_n\}$ be a purely stationary random sequence, and let us interpret these random variables as the time-samples of a random waveform, $a_n = a(nT_o)$, with PSD $S_a(f)$. We consider the PAM signal

$$x(t) = \sum_{n=-\infty}^{\infty} a_n p(t - nT_o + \epsilon) \tag{73}$$

where $p(t)$ is a deterministic finite-energy pulse and ϵ is a fixed pulse-timing phase parameter. To determine the SCD of $x(t)$, we can recognize that $x(t)$ is the output of a periodically time-variant linear system with input $a(t)$, and impulse response

$$h(t, u) = \sum_{n=-\infty}^{\infty} p(t - nT_o + \epsilon)\delta(u - nT_o)$$

We can then use the input-output SCD relation (68). Or we can recognize that this particular periodically time-variant system is composed of a product modular, that implements an impulse sampler, followed by a linear time-invariant pulse-shaping filter with impulse-response function $h(t) = p(t)$, as shown in Fig. 12. We can then use the input-output SCD relation (57), as it applies to impulse sampling, together with the relation (49) for filtering. The result is

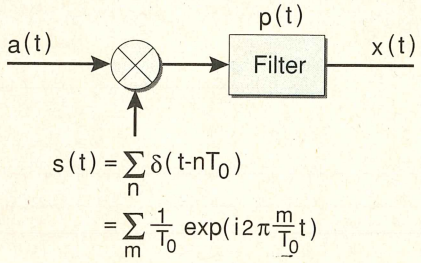

Fig. 12. Interpretation of PAM signal generator as the cascade of an impulse sampler and a pulse-shaping filter.

$$S_x^{\alpha}(f) = \frac{1}{T_o^2} P(f + \alpha/2)P^{*}(f - \alpha/2) \cdot \sum_{n,m=-\infty}^{\infty} S_a^{\alpha + m/T_o}\left(f - \frac{n}{T_o} - \frac{m}{2T_o}\right) e^{i2\pi\alpha\epsilon} \tag{74}$$

Using the SCD aliasing formula (62) for $a(t)$ we can re-express (74) as

$$S_x^{\alpha}(f) = \frac{1}{T_o} P(f + \alpha/2)P^{*}(f - \alpha/2)\tilde{S}_a^{\alpha}(f) e^{i2\pi\alpha\epsilon} \tag{75}$$

where $\tilde{S}_a^{\alpha}(f)$ is the SCD for the pulse-amplitude sequence $\{a_n\}$. Having assumed that $\{a_n\}$ is purely stationary, and using the periodicity property (exhibited by all SCDs for discrete time-series [1, Chapter 11, Sec. C]

$$\tilde{S}_a^{\alpha}(f) = \begin{cases} \tilde{S}_a(f + \alpha/2) & \text{for } \alpha = k/T_o \\ 0 & \text{otherwise} \end{cases} \tag{76}$$

for $k = 0, \pm 1, \pm 2, \cdots$, we can express (75) as

$$S_x^{\alpha}(f) = \begin{cases} \dfrac{1}{T_o} P(f + \alpha/2)P^*(f - \alpha/2)\tilde{S}_a(f + \alpha/2)e^{i2\pi\alpha\epsilon} & \text{for } \alpha = k/T_o \\ 0 & \text{otherwise} \end{cases} \tag{77}$$

A graph of the magnitude of this SCD for the full-duty-cycle rectangular pulse

$$p(t) = \begin{cases} 1 & \text{for } |t| \le T_o/2 \\ 0 & \text{otherwise} \end{cases} \tag{78}$$

and a white-noise amplitude sequence with PSD

$$\tilde{S}_a(f) = 1 \tag{79}$$

is shown in Fig. 13.

It follows from (77) that for all $\alpha = k/T_o$ for which $\tilde{S}_a(f \pm \alpha/2) \neq 0$ and $P(f + \alpha/2)P^*(f - \alpha/2) \neq 0$, the spectral correlation coefficient $\rho_x^{\alpha}(f)$ is unity in magnitude:

$$|\rho_x^{\alpha}(f)| = 1 \tag{80}$$

Thus, all spectral components outside the band $|f| < 1/2T_o$ are completely redundant with respect to those inside this band. Techniques for exploiting this spectral redundancy are described in the next section.

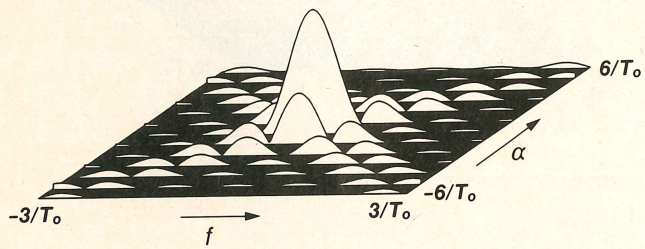

Fig. 13. Magnitude of the spectral-correlation density for a PAM signal with full duty-cycle rectangular pulses.

By inverse Fourier transforming the SCD (77), we obtain the cyclic autocorrelation function

$$R_x^{\alpha}(\tau) = \begin{cases} \dfrac{1}{T_o} \displaystyle\sum_{n=-\infty}^{\infty} R_a(nT_o) r_p^{\alpha}(\tau - nT_o) e^{i2\pi\alpha\epsilon} & \text{for } \alpha = k/T_o \\ 0 & \text{otherwise} \end{cases} \tag{81}$$

where

$$r_p^{\alpha}(\tau) \triangleq \int_{-\infty}^{\infty} p(t + \tau/2)p^*(t - \tau/2)e^{-i2\pi\alpha t}\, dt \tag{82}$$

For a white-noise amplitude-sequence as in (79), (81) reduces to

$$R_x^{\alpha}(\tau) = \frac{1}{T_o} r_p^{\alpha}(\tau) e^{i2\pi\alpha\epsilon} \quad \text{for } \alpha = k/T_o \tag{83}$$

and, for a rectangular pulse as in (78), this yields the temporal correlation coefficient

$$\gamma_x^{\alpha}(\tau) = \frac{\sin[\pi\alpha(T_o - |\tau|)]}{\pi\alpha T_o} e^{i2\pi\alpha\epsilon} \quad \text{for } |\tau| \le T_o \tag{84a}$$

which peaks for $\alpha = 1/T_o$ at $\tau = T_o/2$, where it takes on the value

$$|\gamma_x^{\alpha}(T_o/2)| = 1/\pi \quad \text{for } \alpha = 1/T_o \tag{84b}$$

That is, the strongest possible spectral line that can be regenerated in a delay-product signal (cf. remark made following (20)) for this particular PAM signal occurs when the delay equals half the pulse period. In contrast to this, when the more bandwidth-efficient pulse whose transform is a raised cosine is used, the optimal delay for sine-wave regeneration is zero.

An especially simple example of a sequence of pulse amplitudes $\{a_n\}$ is a binary sequence with values ± 1. If we consider $\tau = 0$ in the delay-product signal, then we obtain

$$y_0(t) = |x(t)|^2 = \sum_{n,m=-\infty}^{\infty} a_n a_m p(t - nT_o + \epsilon)p(t - mT_o + \epsilon)$$

If the pulses do not overlap (i.e., if $p(t) = 0$ for $|t| \ge T_o/2$), this reduces to

$$y_0(t) = \sum_{n=-\infty}^{\infty} a_n^2 p^2(t - nT_o + \epsilon)$$

$$= \sum_{n=-\infty}^{\infty} p^2(t - nT_o + \epsilon)$$

which is periodic with period T_o and therefore contains finite-strength additive sine-wave components with frequencies k/T_o (except when $p(t)$ is flat as in (78)). In this very special case where $\{a_n\}$ is binary and the pulses do not overlap, there is no random component in $y_0(t)$; but, for $\tau \neq 0$, $y_\tau(t)$ contains both sine-wave components and random components (even when $p(t)$ is flat).

Example 8: *ASK and PSK.* By combining the amplitude-modulated sine wave and the digital amplitude-modulated pulse train, we obtain the amplitude-shift-keyed (ASK) signal

$$x(t) = a(t) \cos(2\pi f_o t + \theta) \tag{85a}$$

where

$$a(t) = \sum_{n=-\infty}^{\infty} a_n p(t - nT_o + \epsilon) \tag{85b}$$

and $\{a_n\}$ are digital amplitudes. By using the SCD relation (57) for waveform multiplication and the result

(77) for the SCD of $a(t)$, we can obtain the SCD for the signal (85) by simply convolving the SCD functions shown in Figs. 10a and 13. The result is shown in Fig. 14a, where the cycle frequencies shown are $\alpha = \pm 2f_o + m/T_o$ and $\alpha = m/T_o$ for integers m, and $f_o = 3.3/T_o$.

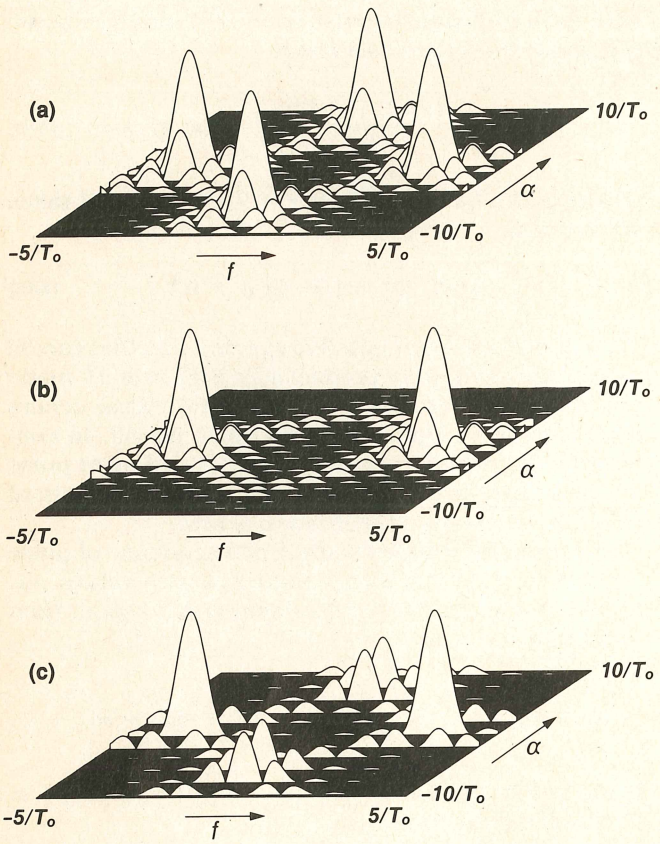

Fig. 14. Magnitude of spectral-correlation densities. a) BPSK, b) QPSK, and c) SQPSK. (Each signal has a rectangular keying envelope.)

For a binary sequence with $a_n = \pm 1$, this amplitude-shift keyed signal, with the pulse (78), is identical to the binary phase-shift keyed (BPSK) signal

$$x(t) = \sin\left[2\pi f_o t + \theta + \frac{\pi}{2}\sum_{n=-\infty}^{\infty} a_n p(t - nT_o)\right] \qquad (86)$$

since shifting the phase of a sine wave by $\pm\pi/2$ is the same as shifting it by $\pi/2$ and multiplying its amplitude by ± 1. Other commonly used types of phase-shift-keyed signals include quaternary phase-shift keying (QPSK) and staggered QPSK (SQPSK). The details of these signal types are available in the literature (see, for example [1, Chapter 12, Sec. E], [2, Chapter 12, Sec. 12.5]). Only their SCD-magnitude surfaces are shown here in Figs. 14b and 14c, where again $f_o = 3.3/T_o$.

It is emphasized that the three signals BPSK, QPSK, and SQPSK differ only in their carrier phase shifts and pulse timing and, as a result, they have identical PSDs, as shown in Fig. 14 (consider $\alpha = 0$). However, as also shown in Fig. 14, these differences in phase and timing result in substantially different SCDs (consider $\alpha \neq 0$). That is, the phase-quadrature component present in QPSK but absent in BPSK results in cancellation of the SCD at cycle frequencies associated with the carrier frequency (viz., $\alpha = \pm 2f_o + m/T_o$ for all integers m) in QPSK. Similarly, the pulse staggering by $T_o/2$ (between the in-phase and quadrature components) present in SQPSK but absent in QPSK results in the SCDs being cancelled at $\alpha = \pm 2f_o + m/T_o$ only for even integers m, and at $\alpha = m/T_o$ only for odd integers m in SQPSK. This again illustrates the fact that the SCD contains phase and timing information not available in the PSD. In fact, as formulas (43) and (77) reveal, the carrier phase θ in (31) and the pulse timing ϵ in (73) are contained explicitly in the SCDs for these carrier- and pulse-modulated signals.

Measurement of the SCD

The ideal SCD function (37) is derived by idealizing the practical spectral correlation measurement depicted in Fig. 6, by letting the averaging time T in the correlation measurement approach infinity and then letting the spectral resolving bandwidth B approach zero. Consequently, the practical measurement with finite parameters T and B can be interpreted as an estimate of the ideal SCD. This estimate will be statistically reliable only if $TB \gg 1$. Numerous alternative implementations of this practical measurement are described in [1, Chapter 13], and computationally efficient digital architectures for some of these, which are developed in [3] and [4], are presented in this issue in [5]. The statistical behavior (bias and variance) of

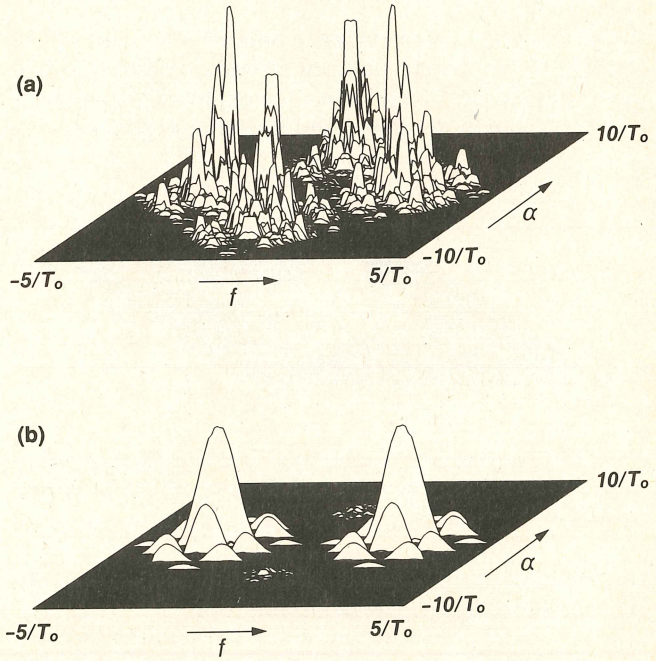

Fig. 15. Magnitude of measured spectral-correlation density (SCD) estimated from a finite data record for the QPSK signal whose ideal SCD is shown in Fig. 14b. a) Record length is 128 time samples, and four adjacent frequency (f) bins are averaged together. b) Record length is 32,768 and 1,024 adjacent frequency (f) bins are averaged together. (The sampling rate in both a) and b) is $10/T_0$, where $1/T_0$ is the keying rate of the QPSK signal.)

such estimates is analyzed in detail in [1, Chapter 15, Sec. B]. For the purpose of making the applications described in the next section more concrete, it suffices here to simply point out that because the SCD $S_x^\alpha(f)$ is equivalent to a particular case of the conventional cross spectral density $S_{uv}(f)$ (cf. (40)), one can envision any of the conventional methods of cross spectral analysis as being used in the applications.

Example 9: *QPSK.* As an example, the result of using the Wiener-Daniell method [1], based on frequency smoothing of the cross-periodogram of $u(t)$ and $v(t)$ (the conjugate product of their FFTs), is illustrated in Fig. 15 for a QPSK signal with carrier frequency $f_o = 1/4T_s$ and keying rate $1/T_o = 1/8T_s$, where $1/T_s$ is the sampling rate. An FFT of length 128 ($T = 128T_s$) was used in Fig. 15a, and only four frequency bins were averaged together ($B = 4/T$), whereas in Fig. 15b, the FFT length used was 32,768 ($T = 32{,}768T_s$) and 1,024 bins were averaged together ($B = 1{,}024/T$). It is easily seen by comparing with the ideal SCD in Fig. 14b that unless $TB \gg 1$, the variability of the SCD estimate can be very large.

EXPLOITATION OF SPECTRAL REDUNDANCY

Spectral redundancy

The existence of correlation between widely separated spectral components (separation equal to α) can be interpreted as *spectral redundancy*. The meaning of the term *redundancy* that is intended here is essentially the same as that used in the field of information theory and coding. Specifically, multiple randomly fluctuating quantities (random variables) are redundant if they are statistically dependent, for example, correlated. In coding, undesired redundancy is removed from data to increase the efficiency with which it represents information, and redundancy is introduced in a controlled manner to increase the reliability of storage and transmission of information in the presence of noise by enabling error detection and correction.

Here, redundancy is to be exploited to enhance the accuracy and reliability of information gleaned from the measurements of corrupted signals, but the term *information* is interpreted in a broad sense. For instance, it includes the six examples outlined in the introductory section. In all these examples, the performance of the signal processors that make the decisions and/or produce the estimates can be substantially improved by suitably exploiting spectral redundancy. The degree of improvement relative to the performance of more commonly used signal processors that ignore spectral redundancy depends on both the severity of the signal corruption (noise, interference, distortion) and the degree of redundancy in the signal $x(t)$, as measured by the magnitude of the spectral correlation coefficient $|\rho_x^\alpha(f)|$ defined in the preceding section.

The primary feature of spectral redundancy that enables it to be readily exploited is its distinctive character. That is, most manmade signals exhibit spectral redundancy, but most noise (all noise that is not cyclostationary) does not. Furthermore, in many practical situations where multiple signals of interest, as well as signals not of interest (interference), overlap in both time and frequency, their spectral redundancy functions are nonoverlapping because their cycle frequencies α are distinct. This is a result of signals having distinct carrier frequencies and/or pulse rates or keying rates, even when occupying the same spectral band.

The distinctive character of spectral redundancy makes signal selectivity possible. Specifically, for the received signal

$$x(t) = \sum_{\ell=}^{L} s_\ell(t) + n(t) \tag{87}$$

where the $\{s_\ell(t)\}_1^L$ include both signals of interest and interference—all of which are statistically independent of each other—and where $n(t)$ is background noise, we have the SCD

$$S_x^\alpha(f) = \sum_{\ell=1}^{L} S_{s_\ell}^\alpha(f) + S_n^\alpha(f) \tag{88}$$

But if the only signal with the particular cycle frequency α_k is $s_k(t)$, then (for measurement time $T \to \infty$) we have

$$S_x^\alpha(f) = S_{s_k}^\alpha(f) \tag{89}$$

regardless of the temporal or spectral overlap among $\{s_\ell(t)\}_1^L$ and also $n(t)$.

Example 10: *BPSK Signal in Multiple AM Interference and Noise.* To illustrate the concept of signal selectivity, let us consider the situation in which a broadband BPSK signal of interest is received in the presence of white noise and five interfering AM signals with narrower bandwidths that together cover the entire band of the BPSK signal. The noise and each of the five interfering signals have equal average power. Therefore, the total signal-to-interference-and-noise rate (SINR) is approximately −8 dB. The BPSK signal has carrier frequency $f_o = 0.25/T_s$ and keying rate $\alpha_o = 0.0625/T_s$. It has full-duty-cycle half-cosine envelope, which results in an approximate bandwidth of $B_o = 0.1875/T_s$. The five AM signals have carrier frequencies $f_1 = 0.156/T_s$, $f_2 = 0.203/T_s$, $f_3 = 0.266/T_s$, $f_4 = 0.313/T_s$, $f_5 = 0.375/T_s$, and bandwidths $B_1 = 0.04/T_s$, $B_2 = 0.05/T_s$, $B_3 = 0.045/T_s$, $B_4 = 0.04/T_s$, $B_5 = 0.08/T_s$. With the use of the same measurement parameters (FFT length = 32,768) as in the preceding *Example 9* for the measurement of the SCD of the QPSK, the SCD for these six signals in noise was measured. The resultant SCD magnitude is shown in Fig. 16a. Also shown in Figs. 16b and 16c are the SCD magnitudes for the BPSK signal alone and for the five AM interferences plus noise alone. Although all six signals exhibit strong spectral redundancy ($|\rho_{s_\ell}^\alpha(f)| = 1$), the cycle frequencies α at which this redundancy exists

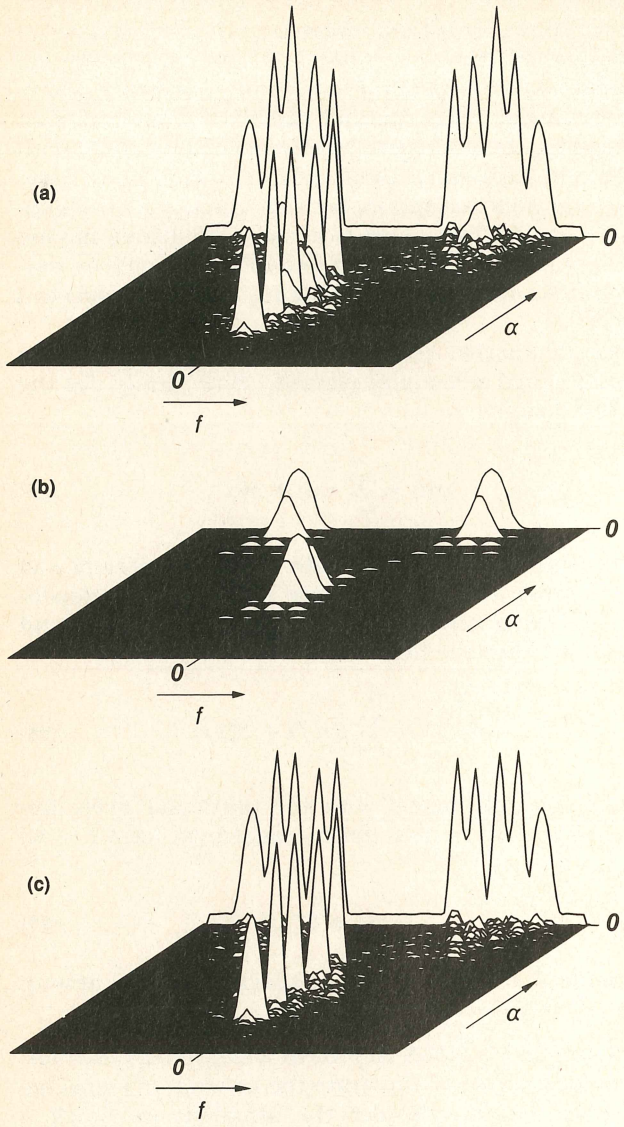

Fig. 16. Magnitudes of estimated spectral-correlation densities (SCDs). a) SCD magnitude for a BPSK signal corrupted by white noise and five AM interferences. b) SCD magnitude for the BPSK signal alone. c) SCD magnitude for the white noise and five AM interferences. (The power levels, center frequencies, and bandwidths for the signals and noise are specified in the text; the record length used is 32,768 time-samples and 1,024 adjacent frequency (f) bins are averaged together.)

are distinct because the carrier frequencies are all distinct. Thus, an accurate estimate of the SCD for the BPSK signal is easily extracted from the SCD for the corrupted measurements. Similarly, accurate estimates of the SCDs for each of the five AM signals can be extracted. Consequently, any information contained in these SCDs can be reliably extracted.

In connection with this example, let us briefly consider some of the signal processing tasks outlined in the introductory section.

Detection and classification

We can see from Fig. 16 that knowing the particular pattern of the SCDs for BPSK and AM signals (see Figs. 8 and 14) enables us to detect the presence of six signals and to classify them according to modulation type. This would be impossible if only PSD (SCD at $\alpha = 0$) measurements were used. One approach to exploiting the spectral redundancy of a signal to detect its presence is to generate a spectral line at one of its cycle frequencies and then detect the presence of the spectral line (cf. the earlier section on SCD). It has been shown that the maximum-SNR spectral-line generator for a signal $s(t)$ is additive Gaussian noise and interference with PSD $S_n(f)$ produces the detection statistic [1, Chapter 14, Sec. E]

$$z = \int_{-\infty}^{\infty} \hat{S}_x^{\alpha}(f) \frac{S_s^{\alpha}(f)^*}{S_n(f + \alpha/2) S_n(f - \alpha/2)}\, df \qquad (90)$$

for comparison to a threshold. In (90), $\hat{S}_x^{\alpha}(f)$ is a crude estimate of $S_x^{\alpha}(f)$ obtained by deleting the time-averaging operation $\langle \cdot \rangle$ and the limiting operation from (37) and choosing B equal to the reciprocal of the record length of $x(t)$. It can be shown that (90) is equivalent to whitening the noise and interference using a filter with transfer function $1/[S_n(f)]^{1/2}$, and then correlating the measured SCD for the noise-and-interference-whitened data with the ideal SCD of the signal (transformed by the whitener) to be detected [1, Chapter 14, Sec. E].

A detailed study of both optimum (e.g., maximum-likelihood and maximum-SNR) and more practical suboptimum detection on the basis of SCD measurement is reported in [6], and receiver operating characteristics for these detectors obtained by simulation are presented in [7].

Parameter estimation

Once the six signals have been detected and classified, their carrier frequencies and phases and the keying rate and phase of the BPSK signal can—with sufficiently long signal duration—be accurately estimated from the magnitude and phase of the SCD (cf., f_o, θ in (43) and T_o, ϵ in (77)). It is clear from the theory discussed in preceding sections that SCD measurement is intimately related to the measurement of the amplitudes and phases of sine waves generated by quadratic transformations of the data. Thus, the fact that an SCD feature occurs at $\alpha = 2f_o$ for each carrier frequency f_o is a direct result of the fact that a sine wave (spectral line) with frequency $\alpha = 2f_o$ and phase 2θ can be generated by putting the data through a quadratic transformation. Similarly, for the SCD feature at $\alpha = 1/T_o$, where $1/T_o$ is the keying rate, a spectral line with frequency $\alpha = 1/T_o$ and phase ϵ can be quadratically generated. Consequently, SCD measurement is useful either directly or indirectly for estimation of synchronization parameters (frequencies and phases) required for the operation of synchronized receivers. The link between synchronization problems and spectral redundancy is pursued in [8].

Time-difference-of-arrival estimation

The cross SCD $S_{wx}^{\alpha}(f)$ for two signals $x(t)$ and $w(t)$ is defined in a way that is analogous to the definition (37) and (24) of the auto SCD $S_x^{\alpha}(f)$. That is, $x(t)$ in (24a) is simply replaced with $w(t)$. If we were to compute the cross SCD for two sets of corrupted measurements obtained from two reception platforms, then the cross SCD magnitude would look very similar to that in Fig. 16 (except that the low flat feature at $\alpha = 0$, which represents the PSD of the receiver noise, would be absent), but the phase of the cross SCD would contain a term linear in f at each value of α where the auto SCD of one of the six signals is nonzero. The slope of this linear phase equals the time-difference-of-arrival (TDOA) of the wavefront at the two platforms for the particular signal with that feature. That is, for $x(t)$ from one platform given by (87) and $w(t)$ from the other platform given by

$$w(t) = \sum_{\ell=1}^{L} s_\ell(t - t_\ell) + m(t) \tag{91}$$

where $\{t_\ell\}$ are the TDOAs, we have

$$S_{wx}^{\alpha}(f) = S_{s_\ell}^{\alpha}(f)\, e^{-i2\pi(f+\alpha/2)t_\ell} \tag{92}$$

provided that $s_\ell(t)$ is the only signal with cycle frequency α. Consequently, accurate estimates of the TDOAs of each of these signals can be obtained from the cross SCD measurement, regardless of temporal and spectral overlap or of the closeness of the individual TDOAs. In other words, the signal selectivity in the α domain eliminates the problem of resolving TDOAs of overlapping signals. Detailed studies of signal-selective TDOA estimation are reported in [9] and [10], where various algorithms are introduced and their mean-squared-error performance is evaluated.

Spatial filtering

Continuing in the same vein, we consider receiving these same six signals in noise with an antenna array. Then we can use the signal selectivity in α to blindly (without any training information other than knowledge of the cycle frequencies α of the signals) adapt a linear combiner of the outputs from the elements in the array to perform spatial filtering. Specifically, by directing the linear combiner to enhance or restore spectral redundancy in its output at a particular cycle frequency α, the combiner will adapt to null out all other signals (if there are enough elements in the array to make this nulling possible). This behavior of the combiner can be seen from the fact that the spectral correlation coefficient for $x(t)$ in (87) is (from (89))

$$\rho_x^{\alpha}(f) = \frac{S_{s_\ell}^{\alpha}(f)}{[S_x(f + \alpha/2)S_x(f - \alpha/2)]^{1/2}} \tag{93}$$

and, similarly, the temporal correlation coefficient for the frequency-shifted versions of $x(t)$ is

$$\gamma_x^{\alpha}(\tau) = \frac{R_{s_\ell}^{\alpha}(\tau)}{R_x(0)} \tag{94}$$

Thus, nulling signals other than $s_\ell(t)$ in the output $x(t)$ of the linear combiner reduces the denominators in (93) and (94) but not the numerators. Hence, $|\rho_x^{\alpha}(f)|$ and $|\gamma_x^{\alpha}(\tau)|$ can be increased by nulling any of the signals other than $s_\ell(t)$. Moreover, the linear combiner needs no knowledge of the reception characteristics of the array (no calibration) to accomplish this nulling. A thorough study of spectral-coherence-restoral algorithms that perform this blind adaptive spatial filtering is reported in [11].

Direction finding

We can take this approach one step further if we do indeed have calibration data for the reception characteristics of an antenna array because we can then also exploit signal selectivity in α to perform high-resolution direction finding (DF) without some of the drawbacks (described below) of conventional methods for high-resolution DF, such as subspace fitting methods [12], that do not exploit spectral redundancy. In particular, let us consider the narrowband model

$$\boldsymbol{x}(t) = \sum_{\ell=1}^{L} \boldsymbol{a}(\theta_\ell)\, s_\ell(t) + \boldsymbol{n}(t) \tag{95}$$

for the analytic signal (or complex envelope) $\boldsymbol{x}(t)$ of the received data vector of dimension r, where $\boldsymbol{a}(\theta_\ell)$ is the direction vector associated with the ℓ-th received signal $s_\ell(t)$, and the function $\boldsymbol{a}(\cdot)$ is specified by the calibration data for the array. Then, by working with the magnitude and phase information contained in the $r \times r$ cyclic correlation matrix

$$\boldsymbol{R}_{\boldsymbol{x}}^{\alpha}(\tau) = \boldsymbol{R}_{s_k}^{\alpha}(\tau) = \boldsymbol{a}(\theta_k)\, R_{s_k}^{\alpha}(\tau)\, \boldsymbol{a}^{\dagger}(\theta_k) \tag{96}$$

for some fixed τ (where $\dagger$ denotes conjugate transpose), instead of working with the information contained in the conventional correlation matrix

$$\boldsymbol{R}_{\boldsymbol{x}}(0) = \sum_{\ell=1}^{L} \boldsymbol{R}_{s_\ell}(0) + \boldsymbol{R}_{\boldsymbol{n}}(0) = \sum_{\ell=1}^{L} \boldsymbol{a}(\theta_\ell)\, \boldsymbol{R}_{s_\ell}(0)\, \boldsymbol{a}^{\dagger}(\theta_\ell) + \boldsymbol{R}_{\boldsymbol{n}}(0) \tag{97}$$

we can avoid the need for advance knowledge of the correlation properties of the noise $\boldsymbol{R}_{\boldsymbol{n}}(0)$ and interference $\boldsymbol{R}_{\boldsymbol{s}_\ell}(0)$ for $\ell \neq k$, and we can avoid the constraint imposed by conventional methods that the number of elements in the array exceed the total number L of signals impinging on the array. Also, by resolving signals in α, we need not resolve them in direction of arrival. Consequently, superior effective spatial resolution is another advantage available through the exploitation of spectral redundancy. As an example of a cyclic DF method, we can exploit the fact that the $r \times r$ matrix in (96) has a rank of unity and the $(r - 1)$-dimensional null space of this matrix is orthogonal to $\boldsymbol{a}(\theta_k)$. Therefore, we can choose as our estimate of θ_k that value $\hat{\theta}_k$

which renders $\boldsymbol{a}(\hat{\theta}_k)$ most nearly orthogonal to the null space of an estimate of the matrix $\boldsymbol{R}_x^{\alpha}(\tau)$ obtained from finite-time averaging. A thorough study of this approach to signal-selective DF is reported in [13], where various algorithms are introduced and their performances are evaluated.

In the preceding paragraphs of this section, the signal processing tasks (with the exception of spatial filtering) involve decisions or parameter estimation, but do not involve estimating (or extracting) an entire signal or an information-bearing message carried by the signal. Nevertheless, for the signal-extraction problem, the utility of spectral redundancy is just as apparent, as explained in the following paragraphs.

Frequency-shift filtering for signal extraction

Spectrally redundant signals that are corrupted by other interfering signals can be more effectively extracted in some applications by exploiting spectral correlation through the use of periodic or multiply-periodic linear time-variant filters, instead of the more conventional time-invariant filters. These time-variant filters enable spectral redundancy to be exploited for signal extraction, because such filters perform frequency-shifting operations (cf. (72)) as well as the frequency-dependent magnitude-weighting and phase-shifting operations performed by time-invariant filters. The utility of this is easily seen for the simple example in which interference in some portions of the spectral band of the signal is so strong that it overpowers the signal in those partial bands. In this case, a time-invariant filter can only reject both the signal and the interference in those highly corrupted bands, whereas a time-variant filter can replace the rejected spectral components of the signal of interest with spectral components from other uncorrupted (or less corrupted) bands that are highly correlated with the rejected components from the signal.

AM is an obvious example of this because of the complete redundancy that exists between its upper sideband (above the carrier frequency) and its lower sideband (below the carrier frequency). Although this redundancy is exploited in the conventional double sideband demodulator to obtain a 3 dB gain in SNR performance, it is seldom exploited properly when partial-band interference is present. The proper exploitation in this case is illustrated in Fig. 17. Figure 17a shows the spectral content (Fourier transform magnitude of a finite segment of data) for an AM signal with partial-band interference in the upper sideband. Figure 17b shows the spectral content after the interference has been rejected by time-invariant filtering. The signal distortion caused by rejection of the signal components along with the interference can be completely removed by simply shifting replicas of perfectly correlated components from the lower sideband into the upper sideband, and then properly adjusting their magnitudes and phases, as suggested in Fig. 17c. A less easily explained example involves two spectrally overlapping linearly modulated signals such as AM,

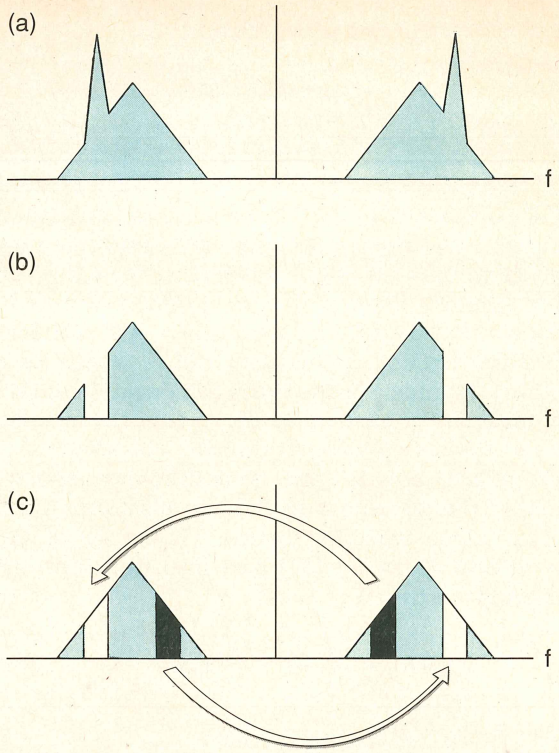

Fig. 17. Illustration of power spectral densities (PSDs) for cochannel-interference removal with minimal signal distortion. a) PSD for AM signal plus interference. b) PSD after interference removal by time-invariant filtering. c) PSD after distortion removal by frequency-shifting.

PAM, ASK, PSK, or digital QAM (quadrature AM). It can be shown that, regardless of the degree of spectral and temporal overlap, each of the two interfering signals can be perfectly extracted by using frequency shifting and complex weighting, provided only that they have either different carrier frequencies or phases (AM, ASK, BPSK) or different keying rates or phases (PAM, ASK, PSK, digital QAM) and at least 100% excess bandwidth (bandwidth in excess of the minimum Nyquist bandwidth for zero intersymbol interference). In addition, when the excess bandwidth is $(L - 1)$ 100%, L spectrally overlapping signals can be separated if they have the same keying rate but different keying phases or carrier frequencies. Also, when broadband noise is present, extraction of each of the signals can in many cases be accomplished without substantial noise amplification.

To illustrate how spectrally overlapping signals can be separated, we consider the case of two QPSK signals with unequal carrier frequencies and unequal keying rates and 100% excess bandwidth. The graphs in Fig. 18 show the overlapping spectra for these two signals. Starting from the top of this figure, each pair of graphs illustrates the result of one filtering and frequency-shifting stage. The sub-band shaded with a single set of parallel lines represents spectral components from one signal that are not corrupted by the other signal. These components are selected and complex-weighted by a filter and then frequency-shifted to cancel the components in another sub-band, which is identified by crosshatched shading. The result of this

cancellation is shown in the second graph (which contains no shading) of each pair. After five such stages, a full sideband of each of the two QPSK signals has been completely separated. In each stage the complex spectral redundancy between components separated by the keying rate is being exploited, and this same spectral redundancy can be used to reconstruct the entire QPSK from either one of its sidebands.

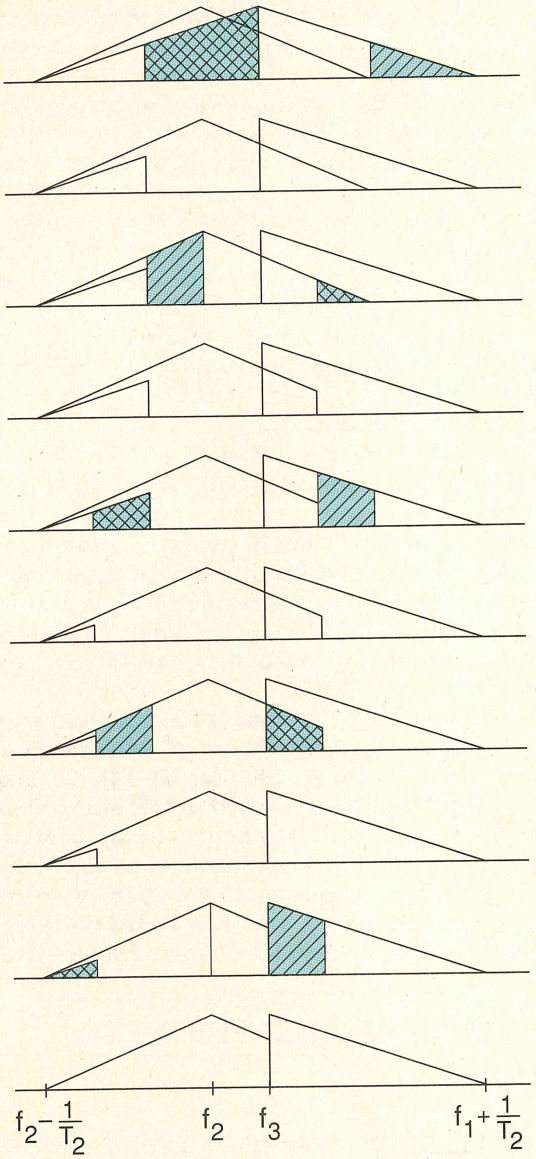

Fig. 18. Illustration of power spectral densities for cochannel-QPSK-signal separation. The keying rates of the two signals are different and the carrier frequencies also are different. Each QPSK signal has a positive-frequency bandwidth equal to twice its keying rate.

The five cascaded stages of filtering, frequency-shifting, and adding operations can be converted into one parallel connection of frequency-shifters, each followed by a filter, simply by using standard system-transformations to move all frequency shifters to the input.

A final example involves the reduction of the signal distortion due to frequency-selective fading caused by multipath propagation. Straightforward amplification in faded portions of the spectrum using a time-invariant filter suffers from the resultant amplification of noise. In contrast to this, a periodically time-variant filter can replace the faded spectral components with stronger highly correlated components from other bands. If these correlated spectral components are weaker than the original components before fading there will be some noise enhancement when they are amplified. But the amount of noise enhancement can be much less than that which would result from the time-invariant filter, which can only amplify the very weak faded components.

Detailed studies of the principles of operation and the mean-squared-error performance of both optimum and adaptive frequency-shift filters are reported in [1, Chapter 14, Secs. A, B], [2, Chapter 12, Sec. 12.8], [3], [14]–[17].

Frequency-shift prediction

If a signal is correlated with time-shifted versions of itself (i.e., if it is not a white-noise signal), then its past can be used to predict its future. The higher the degree of temporal coherence $|\gamma_x^0(\tau)|$, the better the prediction. A signal that exhibits cyclostationarity is also correlated with frequency-shifted versions of itself. Consequently, its future can be better predicted if frequency-shifted versions of its past are also used, so that its spectral coherence as well as its temporal coherence can be exploited. For example, if $x(nT_s)$ has cycle frequencies $\{\alpha_1, \cdots, \alpha_{N-1}\}$ then we can estimate the future value $x[(n+k)T_s]$ for some $k > 0$ using a linear combination of the pasts of the N signals

$$x_q(nT_s) = x(nT_s)\, e^{i2\pi\alpha_q nT_s} \quad \text{for } q = 0, \cdots, N-1 \tag{98}$$

That is, the predicted value is given by

$$\hat{x}[(n+k)T_s] = \sum_{m=0}^{M-1} \sum_{q=0}^{N-1} h_q(m)\, x_q[(n-m)T_s] \tag{99}$$

where M is the memory-length of the predictor. The set of MN prediction coefficients that minimize the time-averaged (over n) squared magnitude of the prediction error $\hat{x}[(n+k)T_s] - x[(n+k)T_s]$ can be shown to be fully specified by the cyclic correlation functions for the N cycle frequencies. Specifically, the set of MN coefficients $\{h_q(m)\}$ is the solution to the set of MN simultaneous linear equations

$$\sum_{m=k}^{M+k-1} \sum_{q=0}^{N-1} h_q(m)\, \tilde{R}_x^{\alpha_q - \alpha_p}[(n-m)T_s] = \tilde{R}_x^{-\alpha_p}(nT_s) \tag{100}$$

for $n = k, \cdots, M+k-1$ and $p = 0, \cdots, N-1$. Also, the percent accuracy of prediction is determined solely by the temporal coherence functions for the frequency translates, which are the discrete-time analogues of (29). It can be shown that for each cycle frequency α_q exploited, there is a corresponding increase in the percent accuracy of the prediction.

In the same way that time-invariant autoregressive model-fitting of stationary time-series data is mathematically equivalent to time-invariant linear prediction [1, Chapter 9, Sec. B], it can be shown that frequency-shift (or multiply-periodic time-variant) autoregressive model-fitting is mathematically equivalent to frequency-shift linear prediction. Studies of this problem are reported in [18]–[24].

HIGHER-ORDER CYCLOSTATIONARITY

Some types of modulated signals like QPSK and digital QAM (quadrature AM) exhibit second-order cyclostationarity associated with the carrier only after the signal has gone through a nonlinear transformation, like a signal squarer or signal quadrupler. In other words, fourth-order (or higher-order) time-delay products, such as (for real $x(t)$)

$$y_{\tau_1 \cdots \tau_4}(t) \triangleq x(t+\tau_1)\, x(t+\tau_2)\, x(t+\tau_3)\, x(t+\tau_4) \quad (101)$$

exhibit a spectral line at the fourth (or higher) harmonic α of the carrier frequency,

$$\langle y_{\tau_1 \cdots \tau_4}(t)\, e^{-i2\pi\alpha t} \rangle \not\equiv 0 \quad (102)$$

even though the second order time-delay product

$$y_{\tau_1 \tau_2}(t) \triangleq x(t+\tau_1)\, x(t+\tau_2) \quad (103)$$

(which includes $y_\tau(t) \triangleq x(t+\tau/2)\, x(t-\tau/2)$ from the earlier section headed "CYCLOSTATIONARITY" as a special case) does not exhibit a spectral line at the second harmonic,

$$\langle y_{\tau_1 \tau_2}(t)\, e^{-i2\pi\alpha t} \rangle \equiv 0 \quad (104)$$

As a result, such signals exhibit nonzero nth order moments of spectral components only for $n \geq 4$; that is, the fourth-order (or higher-order) spectral moment function

$$S_x^{f_1 \cdots f_4}(f) \triangleq \lim_{B \to 0} \langle [h_B^f(t) \otimes x_1(t)]^{(*)}\, [h_B^f(t) \otimes x_2(t)]^{(*)} \cdot [h_B^f(t) \otimes x_3(t)]^{(*)}\, [h_B^f(t) \otimes x_4(t)]^{(*)} \rangle \quad (105)$$

where

$$x_q(t) \triangleq x(t)\, e^{-i2\pi f_q t} \text{ for } q = 1, 2, 3, 4 \quad (106)$$

and

$$\alpha = (-) f_1 + (-) f_2 + (-) f_3 + (-) f_4 \quad (107)$$

can be nonzero for some $\alpha \neq 0$, even though the second order spectral moment function

$$S_x^{f_1 f_2}(f) \triangleq \lim_{B \to 0} \langle [h_B^f(t) \otimes x_1(t)]^{(*)}\, [h_B^f(t) \otimes x_2(t)]^{(*)} \rangle \quad (108)$$

(which includes $S_x^\alpha(f)$ in the earlier section on SCD, by choosing $x_1(t) = u(t)$ and $x_2(t) = v(t)$, as a special case) is zero for all $\alpha = (-) f_1 + (-) f_2 \neq 0$. (In (105) (*) denotes independent optional conjugation and in (107) $(-)$ denotes corresponding independent optional minus sign. That is, depending on the particular spectral moment of interest, one can either include or exclude any of the optional conjugations and corresponding minus signs.)

The same is true for cyclostationarity and spectral moments associated with pulse and keying rates of some severely bandlimited digital signals such as partial-response-coded digital PAM with positive-frequency pulse-bandwidth less than half the pulse rate.

For such signals, exploitation of the higher-order cyclostationarity or spectral redundancy is, in principle, possible, but little progress has yet been made. A notable exception is the long established work on synchronization of communication receivers, where higher-order nonlinearities are commonly used to generate sine waves at harmonics of carrier frequencies to be used to synchronize local oscillators for demodulation of received signals. Recent progress on this problem is reported in [25].

A much more recent application of higher-order cyclostationarity is reported in [26], where new methods for identifying the input-output relations of nonlinear systems with memory are developed. These methods make use of second-order, third-order, fourth-order, and so on, cyclostationary inputs to identify the second-order (quadratic), third-order (cubic), fourth-order (quartic), etc., nonlinear components of the overall nonlinear system.

The foundation for developing the generalization of the spectral redundancy theory of cylcostationarity from the second order [1, Chapter 10–14], [2, Chapter 12] to higher orders is presented in [27] and [28]. Also, the foundation for the *strict sense* theory of cyclostationarity, based on fraction-of-time probability (or temporal probability), in contrast to the *wide-sense* theory described in this article, which is based on temporal moments, is developed in [1, Chapter 15] and [29].

CONCLUSION AND FURTHER READING

Spectral correlation and more general spectral redundancy associated with higher-order spectral moments are common in manmade signals. Almost all types of modulated signals encountered in communications and telemetry systems and also in some control, radar, and sonar systems exhibit spectral redundancy as a direct result of underlying periodicity associated with the modulation. With the substantial increase in the sophistication of signal processors that can be built by using modern digital technology, newly developing techniques for exploiting spectral redundancy promise significant improvements in the capability of signal processors for extracting information from corrupted signals for such purposes as detection and estimation.

This article provides a concise introduction to the rel-

atively new spectral-correlation theory of modulated signals and briefly illustrates the utility of this theory for a variety of signal processing problems.

Readers interested in more technical but still concise introductions to this theory and its applications are referred to [30] and [31]. A comprehensive treatment of the theory is given in [1, Part II], and some of its probabilistic counterpart formulated in terms of stochastic processes is given in [2, Chapter 12]. Details on the various types of modulated signals and the applications briefly described in this paper and additional associated theoretical development can be found in [1, Part II], [2, Chapter 12], and [3]–[31] and references cited therein. Much more work is needed for development of algorithms to exploit spectral redundancy, for evaluation of the performance of such algorithms, and for further development of the theory of higher-order spectral redundancy.

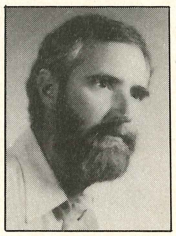

William A. Gardner was born in Palo Alto, California on November 4, 1942. He received the M.S. degree from Stanford University in 1967, and the Ph.D. degree from the University of Massachusetts, Amherst, in 1972, both in Electrical Engineering. He was a Member of the Technical Staff at Bell Laboratories in Massachusetts from 1967 to 1969. He has been a faculty member at the University of California, Davis, since 1972, where he is Professor of Electrical Engineering and Computer Science. Since 1982, he has also been President of the engineering consulting firm Statistical Signal Processing, Inc., Yountville, California. His research interests are in the general area of statistical signal processing, with primary emphasis on the theories of time-series analysis, stochastic processes, and signal detection and estimation. Professor Gardner is the author of *Introduction to Random Processes with Applications to Signals and Systems*, Macmillan, 1985, second edition, McGraw-Hill, 1990; *The Random Process Tutor: A Comprehensive Solutions Manual for Independent Study*, McGraw-Hill, 1990; and *Statistical Spectral Analysis: A Nonprobabilistic Theory*, Prentice-Hall, 1987. He holds several patents and is the author of many research-journal papers. He received the Best Paper of the Year Award from the European Association for Signal Processing in 1986 for the paper entitled "The spectral correlation theory of cyclostationary signals;" the 1987 Distinguished Engineering Alumnus Award from the University of Massachusetts; and the Stephen O. Rice Prize Paper Award in the Field of Communication Theory from the IEEE Communications Society in 1988 for the paper entitled "Signal interception: A unifying theoretical framework for feature detection." He is a member of the American Mathematical Society, the Mathematical Association of America, the American Association for the Advancement of Science, and the European Association for Signal Processing; and a member of the honor societies Sigma Xi, Tau Beta Pi, Eta Kappa Nu, and Alpha Gamma Sigma; and he is a Fellow of the Institute of Electrical and Electronics Engineers.

REFERENCES

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

[2] W. A. Gardner, *Introduction to Random Processes with Applications to Signals and Systems*, Second Edition, McGraw-Hill, NY, 1990.

[3] W. A. Brown, "On the theory of cyclostationary signals," Ph.D. Dissertation, Department of Electrical Engineering and Computer Science, University of California, Davis, September 1987.

[4] R. S. Roberts, "Architectures for digital cyclic spectral analysis," Ph.D. Dissertation, Department of Electrical Engineering and Computer Science, University of California, Davis, September 1989.

[5] R. S. Roberts, W. A. Brown, H. H. Loomis, Jr., "Computationally efficient algorithms for cyclic spectral analysis," *IEEE Signal Processing Magazine* (this issue).

[6] 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.

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

[8] W. A. Gardner, "The role of spectral correlation in design and performance analysis of synchronizers," *IEEE Transactions on Communications*, Vol. COM-34, No. 11, pp. 1089–1095, 1986.

[9] W. A. Gardner and C. K. Chen, "Interference-tolerant time-difference-of-arrival estimation for modulated signals," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, Vol. ASSP-36, No. 9, pp. 1385–1395, 1988.

[10] 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," *IEEE Transactions on Signal Processing* (in press).

[11] B. G. Agee, S. V. Schell, and W. A. Gardner, "Spectral self-coherence restoral: A new approach to blind adaptive signal extraction," *Proceedings of the IEEE*, Vol. 78, pp. 756–767, 1990.

[12] S. V. Schell and W. A. Gardner, "High-resolution direction-finding," *IEEE Signal Processing Magazine* (submitted).

[13] S. V. Schell, "Exploitation of spectral correlation for signal-selective direction finding," Ph.D. Dissertation, Department of Electrical Engineering and Computer Science, University of California, Davis, December 1990.

[14] W. A. Gardner and L. E. Franks, "Characterization of cyclostationary random signal processes," *IEEE Transactions on Information Theory*, Vol. IT-21, No. 1, pp. 4–14, 1975.

[15] W. A. Gardner and W. A. Brown, "Frequency-shift filtering theory for adaptive co-channel interference removal," *Proceedings of the Twenty-Third Annual Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, CA, Oct. 30–Nov. 1, 1989, pp. 562–567.

[16] W. A. Gardner and S. Venkataraman, "Performance of optimum and adaptive frequency-shift filters for co-channel interference and fading," *Proceedings of the Twenty-Fourth Annual Asilomar Conference of Signals, Systems, and Computers*, Pacific Grove, CA, Nov. 5–6, 1990.

[17] J. H. Reed and T. C. Hsia, "The performance of time-dependent adaptive filters for interference rejection," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, Vol. 38, pp. 1373–1385, August 1990.

[18] M. Pagano, "Periodic and multiple autoregression," *Annals of Statistics*, Vol. 6, pp. 1310–1317, 1978.

[19] B. M. Troutman, "Some results in periodic autoregression," *Biometrika*, Vol. 66, No. 2, pp. 219–228, 1979.

[20] G. C. Tiao and M. R. Grupe, "Hidden autoregressive-

moving average models in time series data," *Biometrika*, Vol. 67, No. 2, pp. 365–373, 1980.

[21] A. G. Miamee and H. Salehi, "On the prediction of periodically correlated stochastic processes," pp. 167–179 in *Multivariate Analysis, V*, P. R. Krishnaiah, ed., New York: North Holland Publ. Co., 1980.

[22] K. Hasselman and T. P. Barnett, "Techniques of linear prediction for systems with periodic statistics," *Journal of Atmospheric Science*, Vol. 38, pp. 2275–2283, 1981.

[23] H. J. Newton, "Using periodic autoregressions for multiple spectral estimation," *Technometrics*, Vol. 24, pp. 109–116, May 1982.

[24] A. V. Vecchia, "Maximum likelihood estimation for periodic autoregressive moving average models," *Technometrics*, Vol. 27, pp. 375–384, November 1985.

[25] D. E. Reed and M. A. Wickert, "Nonstationary moments of a random binary pulse train," *IEEE Transactions on Information Theory*, Vol. 35, pp. 700–703, 1989.

[26] T. L. Archer and W. A. Gardner, "New methods for identifying the Volterra kernels of a nonlinear system," *Proceedings of the Twenty-Fourth Annual Asilomar Conference of Signals, Systems, and Computers*, Pacific Grove, CA, Nov. 5–7, 1990.

[27] W. A. Gardner, "Spectral Characterization of *N*-th order cyclostationarity," *Proceedings of the IEEE Fifth ASSP Workshop on Spectrum Estimation and Modeling*, Rochester, NY, Oct. 10–12, 1990, pp. 251–255.

[28] W. A. Gardner and C. M. Spooner, "Higher order cyclostationarity, cyclic cumulants, and cyclic polyspectra," *Proceedings of the 1990 International Symposium on Information Theory and its Applications*, Hawaii, Nov. 27–30, 1990.

[29] W. A. Gardner and W. A. Brown, "Fraction-of-time probability for time-series that exhibit cyclostationarity," *Signal Processing* (in press) [EURASIP].

[30] W. A. Gardner, "The spectral correlation theory of cyclostationary time-series," *Signal Processing*, Vol. 11, pp. 13–36, 1986 [EURASIP].

[31] W. A. Gardner, "A unifying view of coherence in signal processing," *Signal Processing* (in press) [EURASIP].

Glossary of Acronyms and Statistical Terms

Numbers in parentheses following a term listed below are keyed to Equation numbers within this article, and indicate where that term is defined or discussed.

AM - Amplitude modulation - (5)
ASK - Amplitude-shift keying - (85)
Autocorrelation - (21)
BPSK - Binary phase-shift keying - ff. (86)
Cycle frequency - ff. (29)
Cycle spectrum - ff. (29)
Cyclic autocorrelation - (20)
CSD - Cyclic spectral density - ff. (39)
Cyclostationarity - ff. (29)
Periodicity, first-order - ff. (3)
Periodicity, second-order - ff. (18)
PSK - Phase-shift keying - ff. (86)
PSD - Power spectral density - (36)
PAM - Pulse-amplitude modulation - (11)
Pure stationarity - ff. (29)
QAM - Quadrature amplitude modulation - ff. (97)
QPSK - Quaternary phase-shift keying - ff. (86)
Spectral-correlation coefficient - (42)
SCD - Spectral-correlation density - (37)
SQPSK - Staggered quaternary phase-shift keying - ff. (86)
Stationarity - ff. (29)
Temporal correlation coefficient of frequency-translates - (29)

Two computationally efficient algorithms for digital cyclic spectral analysis, the FFT Accumulation Method (FAM) and the Strip Spectral Correlation Algorithm (SSCA), are developed from a series of modifications on a simple time smoothing algorithm. The signal processing, computational, and structural attributes of time smoothing algorithms are presented with emphasis on the FAM and SSCA.

Computationally Efficient Algorithms for Cyclic Spectral Analysis

RANDY S. ROBERTS, WILLIAM A. BROWN, and HERSCHEL H. LOOMIS, JR.

The need for computationally efficient cyclic spectral analysis algorithms becomes increasingly evident as cyclic spectral analysis grows in importance as a signal analysis tool [1]. 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. In the simplest case the spectral components of a real-valued signal are correlated, whereas in the most general case spectral components of two complex-valued signals are correlated. It is the potentially large number of correlation computations, rather than computing the spectral components, that makes cyclic spectral analysis computationally complex. Over the last six years several computationally efficient cyclic spectral analysis algorithms have evolved from the original methods introduced in [2] and [3]. The objective of this paper is to present these algorithms and describe their signal processing, computational, and structural properties.

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 considered to be 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 estimating the cyclic spectrum for a few values of cycle frequency or estimating the cyclic spectrum for small time-frequency resolution product [8]. With computational efficiency for general cyclic spectral analysis as our primary motivation, we focus our attention in this paper on time smoothing algorithms.

Our discussion of time smoothing algorithms begins by describing an algorithm based on the time smoothed cyclic cross periodogram. This algorithm is considered to be

1053-5888/91/0400-0038 $1.00 ©1991 IEEE

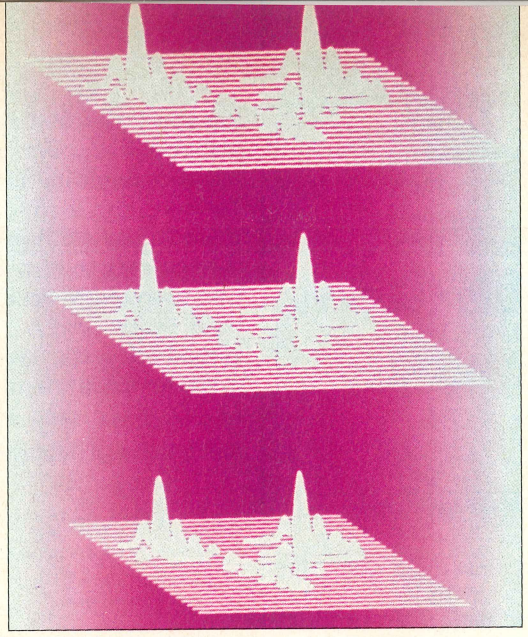

the most fundamental time smoothing algorithm and serves to illustrate the general characteristics of time smoothing algorithms. A mathematical description of the basic algorithm that lends itself to the study of algorithm attributes such as frequency and cycle frequency resolution, and also computational complexity, is developed. The mathematical description is quite general and we apply it to other algorithms as well. From there we develop successively more sophisticated (and substantially less computationally complex) algorithms by refining the basic algorithm. After several intermediate algorithms we arrive at two computationally efficient algorithms: the FFT Accumulation Method (FAM) and the Strip Spectral Correlation Algorithm (SSCA).

In order to fully examine the algorithms we consider the problem of estimating the cyclic cross spectrum of two complex-valued sequences. This problem is the most general and computationally demanding problem of digital cyclic spectral analysis. [1] Simplification of the resulting expressions to special cases of the cross cyclic spectrum of two complex-valued sequences, such as the cyclic spectrum of a single real-valued sequence, are easily found by replacing references to $y(n)$ with $x(n)$ where needed. Computational and structural simplifications arising from the specialization are described.

BASIC TIME SMOOTHING ALGORITHMS

An implementation of the time smoothed cyclic cross periodogram

All time smoothing algorithms are based on the time smoothed cyclic cross periodogram [2]:

$$S^{\alpha}_{xy_T}(n,f)_{\Delta t} = \frac{1}{T}\left\langle X_T(n,f+\alpha/2)Y^*_T(n,f-\alpha/2)\right\rangle_{\Delta t} \qquad (1)$$

[1] Complete cyclic spectral analysis of two complex sequences x(n) and y(n) requires the estimation of both $S^{\alpha}_{xy}(f)$ and $S^{\alpha}_{xy^*}(f)$.

The time smoothed cyclic cross periodogram has the physical interpretation of correlating spectral components of $x(n)$ with spectral components of $y(n)$ over a time span of Δt seconds. The spectral components $X_T(n,f+\alpha/2)$ and $Y_T(n,f-\alpha/2)$, also called complex demodulates, are the complex envelopes of narrowband, bandpass components of a signal. Figure 1 illustrates the processing for a special case of Equation (1), namely, the time smoothed cyclic periodogram of a real signal. A data tapering window of length T seconds slides over the data for a time span of Δt seconds. At each instant the complex demodulates of the data within the window are computed. (Details of computing the complex demodulates are considered later. For now, note that the demodulates are lowpass sequences and have bandwidths on the order of the reciprocal of the data tapering window, i.e., $\Delta f = 1/T$ Hz.) After the complex demodulates have been computed, they are

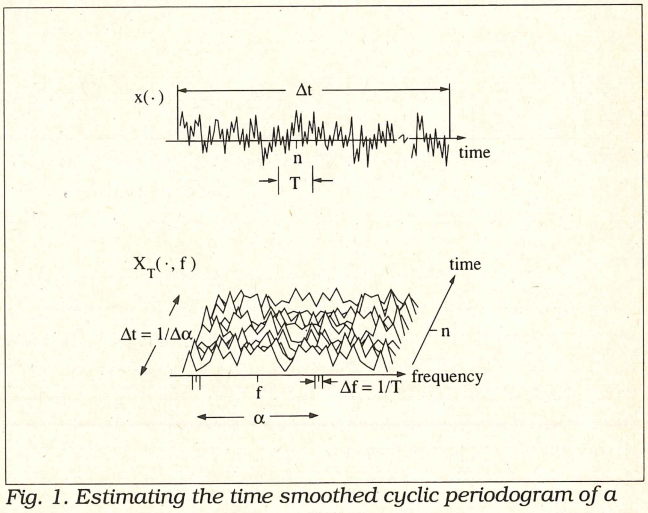

Fig. 1. Estimating the time smoothed cyclic periodogram of a single real-valued signal.

This work was supported in part by a grant from ESL Inc. with partial matching support from the California State MICRO Program (PI: W. A. Gardner).

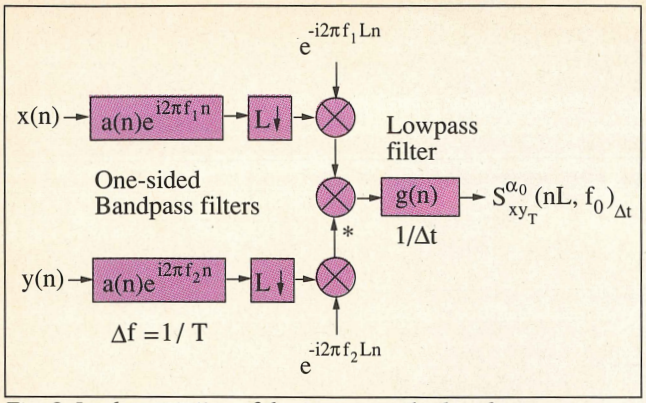

Fig. 2. Implementation of the time smoothed cyclic cross periodogram.

correlated by time averaging their conjugate product over an interval of Δt seconds to produce estimates of the cyclic spectrum. To estimate the cyclic spectrum at the point (f_0, α_0), demodulates separated in frequency by an amount α_0 and centered about a midpoint of f_0 are correlated. The quantities Δt and Δf are referred to as the time and frequency resolutions of the point estimate.

A system based on Equation (1) is illustrated in Fig. 2 (for now let the time decimation parameter $L = 1$). In this approach the complex demodulates are obtained by filtering the input sequences with one-sided bandpass filters and frequency shifting the filter outputs to baseband. To compute the point estimate at (f_0, α_0), the center frequencies of the filters are set to $f_1 = f_0 + \alpha_0 / 2$ and $f_2 = f_0 - \alpha_0/2$. Following the notation in [4], the input filters in Fig. 2 have memory lengths of $T = N' T_s$ seconds and therefore bandwidths on the order of $1/T = f_s/N'$ Hz. (T_s is the sampling interval and f_s is the sampling frequency.) Mathematically, computation of the complex demodulates is expressed as

$$X_{T(n,f)} = \sum_{r=-N/2}^{N/2} a(r)x(n-r)e^{-i2\pi f(n-r)T_s} \qquad (2)$$

where $a(r)$ is a data tapering window of length $T = N'T_s$ seconds. For convenience, the quantity Δa is defined to be the bandwidth of the input filters; hence, for this system the bandwidth of the input filters is also the frequency resolution of the estimate: $\Delta f = \Delta a = f_s/N'$.

After the complex demodulates have been computed they are correlated over a time span of Δt seconds. In Fig. 2 the correlation operation is performed by a complex multiplier followed by a lowpass filter. The lowpass filter has a memory length of $\Delta t = NT_s$ seconds and a bandwidth on the order of $1/\Delta t = f_s/N$ Hz. The correlation operation is expressed as

$$S^{\alpha_0}_{xy_T}(n, f_0)_{\Delta t} = \sum_r X_T(r, f_1) Y^*_T(r, f_2)\, g(n-r) \qquad (3)$$

where $g(n)$ is a data tapering window of width $\Delta t = NT_s$ seconds. The resulting output sequence $S^{\alpha_0}_{xy_T}(n, f_0)_{\Delta t}$ is the spectral cross correlation estimate at (f_0,α_0) and is composed of a bias (if spectral correlation is present) whose magnitude indicates the degree of spectral correlation, and a random component that can be interpreted as measurement noise. For a reliable estimate it is necessary to have $\Delta t \gg T$, or equivalently, to have the time-frequency resolution product be much greater than unity: $\Delta t \Delta f \gg 1$. It is shown in [2] that the time smoothed cyclic cross periodogram converges to the cyclic cross spectrum in the limit, as $\Delta t \to \infty$ followed by $\Delta f \to 0$, if the time windows $a(n)$ and $g(n)$ are properly normalized. Therefore, if $\sum_n a^2(n) = \sum_n g(n) = 1$ we have

$$\lim_{\Delta f \to 0} \lim_{\Delta t \to \infty} S^{\alpha_0}_{xy_T}(n, f_0)_{\Delta t} = S^{\alpha_0}_{xy}(f_0) \qquad (4)$$

where the cyclic cross spectrum $S^{\alpha_0}_{xy}(f_0)$ is defined by

$$S^{\alpha_0}_{xy}(f_0) = \sum_{k=-\infty}^{\infty} R^{\alpha_0}_{xy}(k) e^{-i2\pi f_0 kT_s} \qquad (5)$$

and the cyclic cross correlation $R^{\alpha_0}_{xy}(k)$is defined by

$$R^{\alpha_0}_{xy}(k) = \lim_{\Delta t \to \infty} \left\langle x(nT_s + kT_s) y^*(nT_s) e^{-i2\pi(n+k/2)T_s} \right\rangle_{\Delta t} \qquad (6)$$

See [2] and [5] for further details on the above relationships.

An interesting representation of the basic system is found by combining expressions for the complex demodulates with Equation (3) to get [4]

$$S^{\alpha_0}_{xy_T}(n, f_0)_{\Delta t} = \sum_q \sum_r m(q,r) x(n-q) y^*(n-r) e^{-i2\pi\alpha_0 nT_s} \qquad (7)$$

Equation (7) is a quadratic transformation of the input sequences with the kernel $m(q,r)$. By manipulating Equations (2), (3), and (7), the kernel for the basic system is found to be

$$m(q,r) = \sum_p g(p) a(q-p) a(r-p) e^{i2\pi f_0 (q-r)T_s} e^{i\pi\alpha_0 (q+r)T_s} \qquad (8)$$

In general, most cyclic spectral analysis algorithms can be mathematically described in the form of Equation (8), that is, as a quadratic transformation of the input sequences with a kernel that depends on the system parameters f_0, α_0, Δf, and Δt [2], [5]. The interesting feature of this representation is that the system parameters are contained solely within $m(q,r)$. Hence, Equation (8) describes how system parameters influence estimation of the cyclic cross spectrum. Although the system kernel provides a compact mathematical representation of cyclic spectral analysis algorithms, a transformed version of this representation provides greater insight into the algorithms.

In terms of the rotated Fourier Series Transform (FST) of the kernel $m(q,r)$

$$M(\alpha, f) = \sum_q \sum_r m(q,r) e^{-i2\pi(f+\alpha/2)qT_s} e^{i2\pi(f-\alpha/2)rT_s} \qquad (9)$$

the output is expressed as [4]

$$S^{\alpha_0}_{xy_T}(n,f_0)_{\Delta t}=T_s\sum_{\Re}\int_{-f_s/2}^{f_s/2}M(\beta,\nu)\,S^{\beta}_{x}(\nu)\,d\nu\,e^{i2\pi(\beta-\alpha_0)nT_s}+R(n) \quad (10)$$

where the $\Sigma_{\Re}$-indicated summation is over the region $-f_s/2 < \beta - \alpha_0 \le f_s/2$, and $R(n)$ is a random component that can be interpreted as measurement noise. Note that Equation (10) also describes how system parameters (represented by $M(\alpha,f)$) influence the estimation of $S^{\alpha}_{xy}(f)$, the underlying cyclic cross spectrum. Cycle features of $S^{\alpha}_{xy}(f)$ that are within the domain of $M(\alpha,f)$ are summed in cycle frequency and integrated in frequency to produce the output value of the estimate. Cycle features outside the domain of $M(\alpha,f)$ are suppressed. Thus, by understanding the properties of $M(\alpha,f)$ we can understand how system parameters affect the estimate. As a general rule, if $\Delta t\Delta f \rangle\rangle 1$ then the kernel transform $M(\alpha,f)$ can often be approximated by the separable form [2]

$$M(\alpha,f)\approx G_{1/\Delta t}(\alpha-\alpha_0)H_{1/T}(f-f_0) \quad (11)$$

The components of the kernel, $G_{1/\Delta t}(\alpha)$ and $H_{1/T}(f)$, are frequency windows that are simply related to the FSTs of data tapering windows used in the algorithm under consideration, and have approximate bandwidths of $1/\Delta t$ and $1/T$ Hz respectively. Typically, $M(\alpha,f)$ is a two dimensional pulse-like function surrounded by sidelobes. The sidelobes of $M(\alpha,f)$ are generally small compared to the mainlobe but can be troublesome in some measurements. The region of support of the main lobe of $M(\alpha, f)$ is called a Cyclic Spectrum Analyzer (CSA) cell. Equation (11) indicates that a CSA cell has a width on the order of $1/\Delta t$ in cycle frequency and a length on the order of $1/T$ in frequency. As described by Equation (10), features of $S^{\alpha}_{xy}(f)$ within a CSA cell are transmitted to the output while features outside of the region are suppressed. Thus, the width of a CSA cell determines the cycle frequency resolution $\Delta\alpha$ and the length of the cell determines the frequency resolution Δf. Figure 3 depicts an idealized (i.e., no sidelobes or skirts) CSA cell located at (f_0, α_0). For a proper measurement the cycle frequency resolution must be small enough to resolve the cycle features of $S^{\alpha}_{xy}(f)$ and the frequency resolution must be small enough to resolve $S^{\alpha}_{xy}(f)$ in frequency. Note that cycle

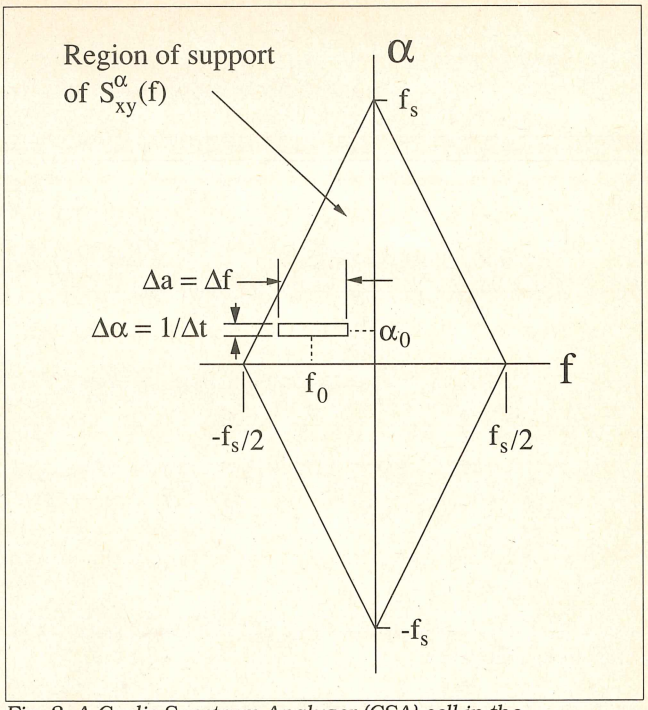

Fig. 3. A Cyclic Spectrum Analyzer (CSA) cell in the bifrequency plane.

features exterior to a CSA cell but within the sidelobes of $M(\alpha, f)$ can contribute to the estimate of $S^{\alpha}_{xy}(f)$. This undesirable effect is called cycle leakage and is minimized by designing $M(\alpha,f)$ to have sufficiently low sidelobes.

The kernel transform is a useful means for studying the frequency resolution and cycle frequency resolution of an algorithm. As an example, the kernel transform for the basic system is

$$M(\alpha,f)=G(\alpha-\alpha_0)A\left(f-f_0+\frac{\alpha-\alpha_0}{2}\right)A^*\left(f-f_0-\frac{\alpha-\alpha_0}{2}\right) \quad (12)$$

where $A(f)$ and $G(\alpha)$ are the FSTs of $a(n)$ and $g(n)$. (Note that the bandwidth of $A(f)$ is on the order of $\Delta a = 1/T$ and that of $G(a)$ on the order of $1/\Delta t$.) Since $f_1 = f_0 + \alpha/2$ and $f_2 = f_0 - \alpha_0/2$, the transform kernel can also be expressed as

$$M(\alpha,f)=G(\alpha-\alpha_0)A\left(f-f_1+\frac{\alpha}{2}\right)A^*\left(f-f_2-\frac{\alpha}{2}\right) \quad (13)$$

Examination of Equations (12) or (13) reveals that most of the energy in $M(\alpha,f)$ is indeed confined to the

TABLE I

BASIC TIME SMOOTHING

Computation Section	Number of Complex Multiplications	
	Cyclic Cross Spectrum of Two Complex Signals	Cyclic Spectrum of a Single Real Signal
Filtering	$2N^2N'$	N^2N'
Frequency Shift	$2N^2$	N^2
Correlate	N^2N'	$N^2N'/4$

Computational complexity of the basic time smoothing algorithm in terms of the number of complex multiplications. Note that complete coverage of the bifrequency plane requires N N′ point estimates.

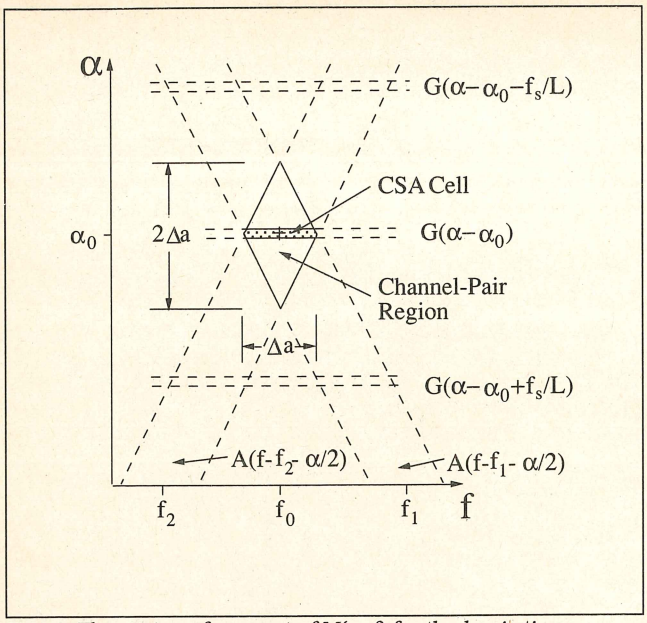

Fig. 4. The region of support of M(α, f) for the basic time smoothing algorithm.

region $|f-f_0| \le \Delta a/2$ and $|\alpha-\alpha_0| \le 1/2\Delta t$ (cf. Fig. 3). Hence, the frequency resolution of this algorithm is again seen to be $\Delta f = \Delta a = f_s/N'$, and cycle frequency resolution is $\Delta\alpha = 1/\Delta t = f_s/N$. Finally, the time-frequency resolution product is $\Delta t\Delta f = N/N'$.

In addition to determining the resolutions of a cyclic spectral analysis algorithm the CSA cell concept is useful for determining the computational complexity of an algorithm. Consider the problem of estimating the cyclic cross spectrum of two complex signals everywhere in the bifrequency plane at a given time instant. To perform the estimation approximately NN' point estimates must be calculated (this is the area of the region of support of $S_{xy}^{\alpha}(f)$ divided by the area of a CSA cell). All of the input data can be filtered with two banks of N one-sided bandpass filters for a total of N^2N' complex multiplications and the complex demodulates calculated with an additional $2N^2$ complex multiplications. Additionally, N^2N' complex multiplications are needed to compute the spectral correlations for all point estimates (assuming that $g(n)$ is a rectangular window with unity height). Table I summarizes the computational complexity of the algorithm in terms of the number of complex multiplications required to compute estimates of the cyclic cross spectrum of two complex signals and the cyclic spectrum of a single real signal.

To put these numbers into perspective a numerical example is considered. Say that we want to estimate the cyclic cross spectrum of two complex sequences for $\Delta f = 1/8$ ($T_s = 1$) and $\Delta t\,\Delta f = 16384$. For these values, $N'=8$ and $N=131072$. If the computations are performed on a uniprocessor computer (i.e., one ALU) capable of producing one complex product every 150 ns, then the computation would take approximately 18.6 hrs. For the case of estimating the cyclic spectrum of a real signal the symmetry relations $S_x^{\alpha}(-f) = S_x^{\alpha}(f)$ and $S_x^{-\alpha}(f) = S_x^{\alpha}(f)^*$ imply that $S_x^{\alpha}(f)$ need only be estimated in one quadrant of the bifrequency plane. Thus, the previous estimation problem with a single real signal would require approximately 7.9 hrs of computation. This example serves to illustrate the "brute force" approach to digital cyclic spectral analysis. Algorithms that are considerably more efficient are presented in the remainder of this paper.

Time smoothing with decimation

The computational efficiency of the previous algorithm can be improved by decimating the outputs of the bandpass filters by a suitable factor. Equivalently, data is shifted into the filters in blocks of L samples where $L<N'$; thus, only N/L samples are processed for each point estimate and the overall computational complexity of the algorithm is reduced by the factor L. Since the filter outputs are over sampled by a factor of N' the sampling rate can be reduced to f_s/L, $L \le N'$ before aliasing occurs. However, if the sampling rate is reduced by the maximum factor $L = N'$, then the effects of cycle leakage can be substantial. With decimation in effect the estimate by Equation (3) is modified to

$$S_{xy_T}^{\alpha_0}(nL,f_0)_{\Delta t} = \sum_r X_T(rL,f_1)Y_T^*(rL,f_2)g_c(n-r) \qquad (14)$$

Equation (10) still applies except that $M(\alpha, f)$ now also accounts for decimation. Shifting the input sequences into the system in blocks has the effect of mutating the lowpass filter G(.) in Equations (12) and (13) into the comb filter [2]

$$G_c(\alpha) = \sum_r g_c(r)e^{-i2\pi\alpha rLT_s} \qquad (15)$$

$$= \sum_n G\left(\alpha + \frac{nf_s}{L}\right) \qquad (16)$$

where $G(\alpha)$ consists of one period of $G_c(\alpha)$,

$$G(\alpha) = \begin{cases} G_c(\alpha) & |\alpha| < \dfrac{f_S}{2L} \\ 0 & \text{otherwise} \end{cases}$$

Consequentially, the kernel transform is now expressed as (cf. Equation (13))

$$M(\alpha,f) = \sum_n G\left(\alpha-\alpha_0+\frac{nf_s}{L}\right)A\left(f-f_1+\frac{\alpha}{2}\right)A^*\left(f-f_2-\frac{\alpha}{2}\right) \qquad (17)$$

It is convenient to consider Equation (17) as consisting of two components: the product of frequency shifted channelizer transfer functions

$$M_1(\alpha,f) = A\left(f-f_1+\frac{\alpha}{2}\right)A^*\left(f-f_2-\frac{\alpha}{2}\right) \qquad (18)$$

and the comb filter

$$G_c(\alpha-\alpha_0) = \sum_n G\left(\alpha-\alpha_0+\frac{nf_s}{L}\right) \qquad (19)$$

The bandwidth of $A(f)$ is on the order of Δa and $G_c(\alpha)$ has teeth with bandwidths on the order of $1/\Delta t$ spaced f_s/L Hz apart in cycle frequency. The approximate

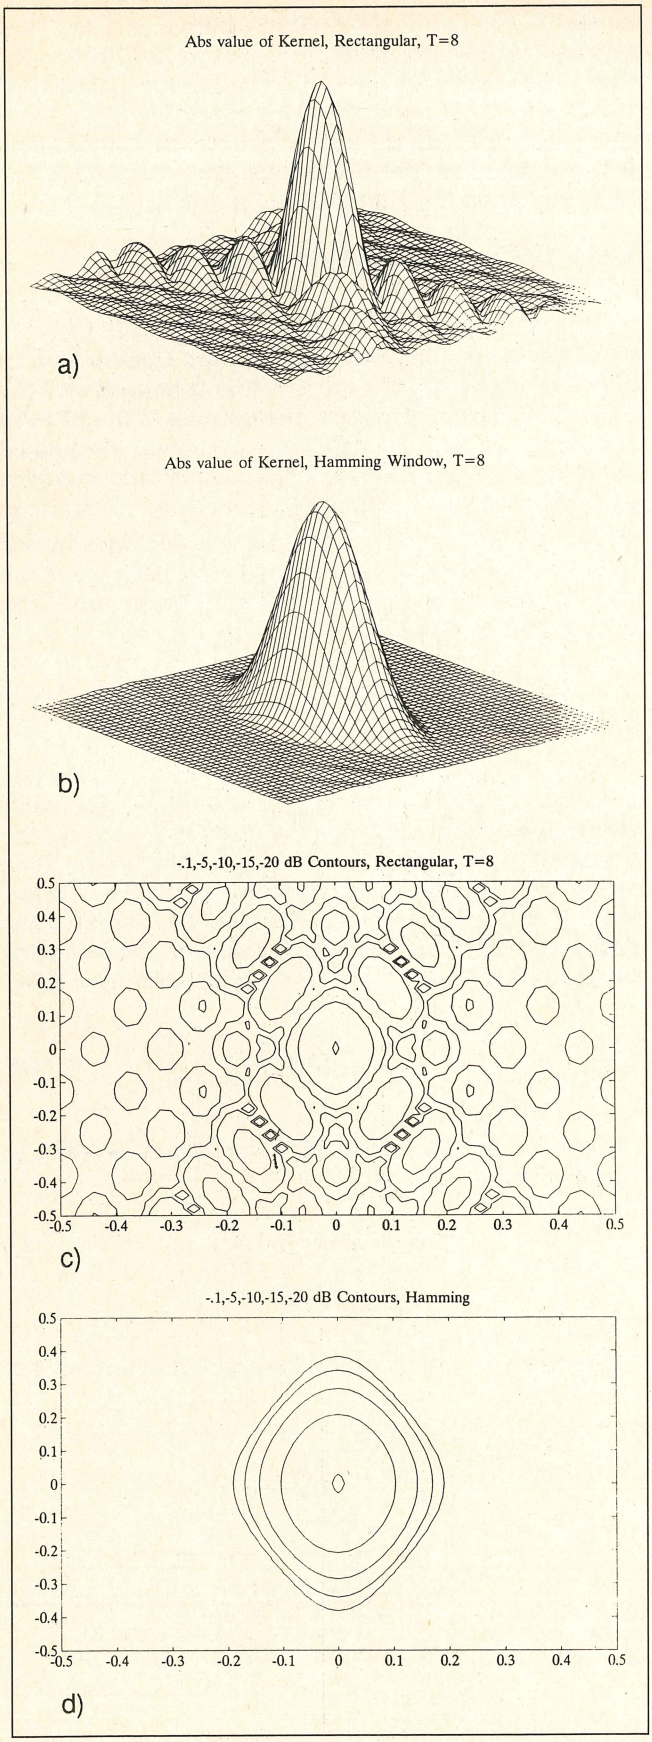

Fig. 5 Channelizer transfer function for a) a rectangular window and b) a Hamming window. Corresponding contour plots are c) rectangular and d) Hamming.

regions of support of $M_1(\alpha, f)$ and $G_c(\alpha - \alpha_{0})$ are illustrated in Fig. 4.

The diamond shaped region in Fig. 4 is called a channel-pair region, and it is the idealized (no sidelobes or skirts) region of support of $M_1(\alpha, f)$. A channel-pair region is approximately diamond shaped with length and width on the order of $2\Delta a$ and Δa respectively. The intersection of a channel-pair region with the region of support of the $n=0$ tooth of $G_c(\alpha - \alpha_{0})$ at (f_0, a_0) defines a CSA cell. Using the fixed ratio $P = N/L$ we can express the cycle frequency resolution of the decimated algorithm as $\Delta\alpha = f_s/PL$. Note that the frequency resolution of the algorithm remains the same as before: $\Delta f = f_s/N'$. Although the boundaries of the CSA cell and channel--pair region in Fig. 4 are depicted as well defined, this is generally not the case. Both $M_1(\alpha, f)$ and $G_c(\alpha - \alpha_{0})$ can have significant skirts and sidelobes beyond the regions indicated in Fig. 4. Figure 5a illustrates the magnitude of $M_1(\alpha, f)$ for $a(n)$ a rectangular data tapering window and Fig. 5b shows $M_1(\alpha, f)$ for $a(n)$ a Hamming window. Contour plots of Figs. 5a and 5b are shown in Figs. 5c and 5d respectively.

Recall that cycle leakage occurs if the sidelobes of $M(\alpha, f)$ are large enough to admit nearby cycle features. In the decimated algorithm, cycle leakage can also occur if teeth of $G_c(\alpha - \alpha_{0})$ other than the one at α_0 intersect the channel-pair region. Overlap of a channel-pair region with $n \neq 0$ teeth of $G_c(\alpha - \alpha_{0})$ occurs if L is too small or if the channel-pair region is larger than indicated in Fig. 4. An additional effect, known as cycle aliasing, occurs if the frequencies of leaked cycle features exceed the decimated sampling rate. We postpone determination of appropriate values for L until the discussion of the FAM.

FFT BASED TIME SMOOTHING ALGORITHMS

Time smoothing with a Fourier transform

Consider frequency shifting the product sequence in Fig. 2 by an amount ε from α_0 to $\alpha_0 + \varepsilon$ (see Fig. 6 for a block diagram of the processing). In this case, the output of the system is given by

$$S_{xy_T}^{\alpha_0+\varepsilon}(n, f_0)_{\Delta t} = \sum_r X_T(r, f_1)\, Y_T^*(r, f_2)\, g(n-r)\, e^{-i2\pi\varepsilon r T_s} \quad (20)$$

Following some manipulations the kernel transform is found to be

$$M(\alpha, f) = G(\alpha - \alpha_0 - \varepsilon) A(f - f_1 + \alpha/2) A^*(f - f_2 - \alpha/2) \quad (21)$$

The region of support for the kernel transform is depicted in Fig. 7. It is seen that $M(\alpha, f)$ is a pulse centered at $(f_0, \alpha_0 + \varepsilon)$ with a width of $\Delta\alpha = 1/\Delta t$ and length of $\Delta f = \Delta a - |\varepsilon|$, where $|\varepsilon| < \Delta a$. If several values of ε are desired, evaluation of the sum in Equation (20) can be simplified by discretizing the values of ε to be ε

$= q\Delta\alpha$. In this case the output of the algorithm is expressed as

$$S_{xy_T}^{\alpha_i+q\Delta\alpha}(n,f_j)_{\Delta t}=\sum_r X_T(r,f_1)Y_T^*(r,f_2)g(n-r)\,e^{-i2\pi rq/N} \quad (22)$$

in which the sum in Equation (22) can be evaluated with an N-point FFT. Thus, point estimates with constant cycle frequency can be computed in blocks by Fourier transforming the product sequence instead of averaging the product sequences individually as in Equation (1).

For complete coverage of the bifrequency plane a bank of bandpass filters is required to produce the necessary complex demodulates. An efficient method for producing the required complex demodulates is based on a sliding FFT [6]. In this approach the frequencies of the filter bank are discretized to $f_k = k(f_s/N'), k = -N'/2 \ldots (N'/2)-1$. The channel-pair regions associated with all pairs of complex demodulates are located at (f_j, α_i) where the frequency coordinates are

$$f_j=\frac{f_k+f_l}{2} \quad (23)$$

$$=\frac{k+l}{2}\left(\frac{f_s}{N'}\right) \quad (24)$$

and the cycle frequency coordinates are

$$\alpha_i=f_k-f_l \quad (25)$$

$$=(k-l)\left(\frac{f_s}{N'}\right) \quad (26)$$

Figure 8 shows the pattern of channel-pair regions for $N' = 8$. The ordered pair associated with each channel-pair region in Fig. 8 is the coordinate of the region written in terms of indices: $(f_j, \alpha_i) \rightarrow (j,i)$. For an N'-point channelizer there are $(N')^2$ possible combinations of channelizer streams; hence, there are at most $(N')^2$ channel-pair regions. Due to symmetry, estimation of the cyclic spectrum of a single real signal requires only $(N')^2/4$ channel-pair regions (one quadrant of the bifrequency plane).

A troubling aspect of this approach is the nonuniform frequency resolution of the point estimates (recall that $\Delta f = \Delta a - |\varepsilon|$). Near the top and bottom of a channel-pair region the frequency resolution approaches zero, and as $\Delta f \rightarrow 0$ so does the time-frequency resolution product. As a result, point estimates at the ends of channel-pair regions become increasingly unreliable (i.e., the estimates have high variability). Several methods have been developed to deal with this problem. One approach, developed in [7], combines estimates in adjacent channel-pair regions to obtain new estimates

$$S_{xy_T}^{\alpha_0}(f_0+\varepsilon_1/2)_{\Delta t}=(S_{xy_T}^{\alpha_0}(f_0)_{\Delta t}+S_{xy_T}^{\alpha_0}(f_0+\text{sign}(\varepsilon_1)/2N')_{\Delta t})/2 \quad (27)$$

The estimates in this method have a frequency resolution of $\Delta f = \Delta a$ provided that $A(f)$ is rectangular. Another method, emphasized here, is to retain only those estimates within $\pm\Delta a/2$ of the center of the channel-pair region [4], [5]. Although this approach is simple to implement, the frequency resolution and time-frequency resolution product still vary within a channel-pair region. Furthermore, discarding estimates in this manner leaves coverage gaps in the bifrequency plane. However, if the input signals are resolvable with $\Delta f \approx \Delta a$ then loss of coverage is not detrimental since cycle features extend to at least one channel-pair region. See [5] for other ways to eliminate coverage gaps.

The FFT accumulation method

The FFT Accumulation Method (FAM) incorporates all of the ideas so far discussed [4], [5]. Channelization is performed by an N'-point FFT that is hopped over the data in blocks of L samples. The outputs of the FFT are frequency shifted to baseband to obtain decimated complex demodulate sequences. After the complex demodulates are computed, product sequences $X_T(nL, f_k)\ Y_T^*(nL, f_l)$ are formed and Fourier transformed with a P-point FFT. (Recall that $P=N/L$ where N is the total number of samples in the input sequences.) The output of the FAM is therefore

$$S_{xy_T}^{\alpha_i+q\Delta\alpha}(nL,f_j)_{\Delta t}=\sum_r X_T(rL,f_k)Y_T^*(rL,f_l)g_c(n-r)e^{-i2\pi rq/P} \quad (28)$$

and the kernel transform is given by Equation (17) with $f_0 = f_j$ and $\alpha_0 = \alpha_i + q\Delta\alpha$. The cycle frequency resolution of the FAM is identical to that of the decimated algorithm described above (see "Time smoothing with decimation"); thus, $\Delta\alpha = f_s/PL$. The frequency resolution of the FAM is identical to that of the Fourier transform smoothing algorithm of the preceding Section; thus, $\Delta f(q) = \Delta a - |q|\Delta\alpha$ and since $\Delta a = f_s/N'$

$$\Delta f(q)=\left(1-|q|\frac{N'}{PL}\right)\frac{f_s}{N'} \quad (29)$$

Likewise, the time-frequency resolution product is variable and is given as

$$\Delta t\Delta f=\Delta t(\Delta a-|q|\Delta\alpha) \quad (30)$$

$$=\frac{N}{N'}-|q| \quad (31)$$

The time-frequency resolution product of the FAM is

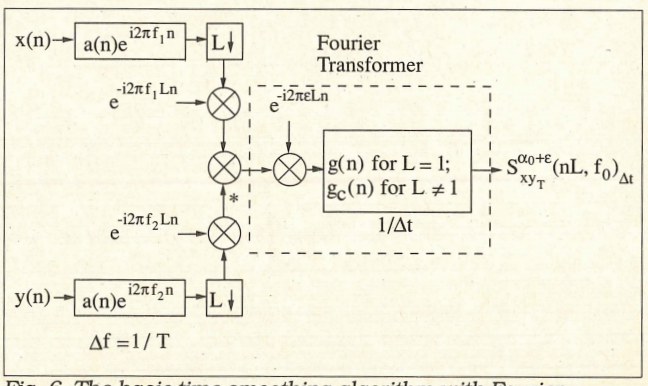

Fig. 6. The basic time smoothing algorithm with Fourier transformer output.

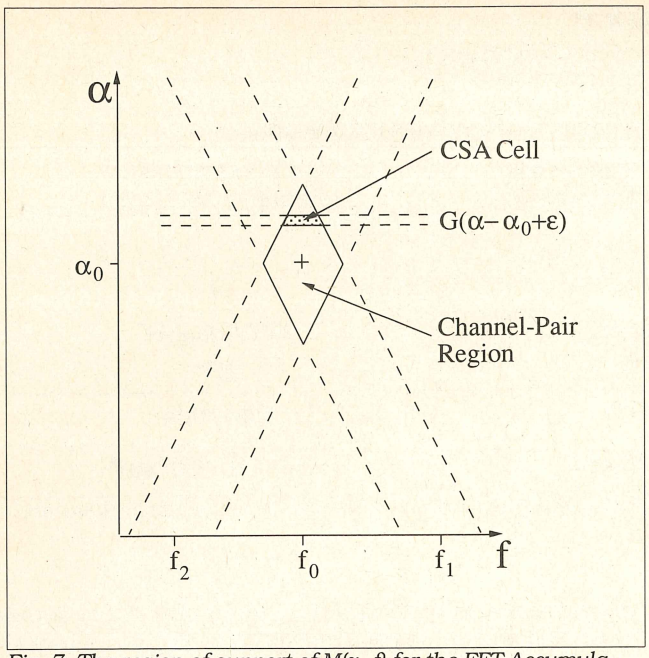

Fig. 7. The region of support of M(α, f) for the FFT Accumulation Method.

typically referenced to the center of a channel-pair region ($q=0$) so that $\Delta t\Delta f = N/N' = PL/N'$}. To minimize the variability of point estimates near the top and bottom of the channel-pair regions we can retain only those estimates within $\pm\Delta a/2$ of the center of the region and discard all others. In terms of the index parameter q, only estimates within the range

$$-\frac{\Delta a}{2} \le q\Delta\alpha < \frac{\Delta a}{2} \tag{32}$$

or

$$-\frac{PL}{2N'} \le q \le \frac{PL}{2N'} - 1 \tag{33}$$

in Equation (28) are retained.

Both cycle leakage and cycle aliasing can be minimized by restricting the amount of overlap of the $n\neq 0$ teeth of $G_c(\alpha - \alpha_0)$ with the channel-pair region. Diminishing the overlap is accomplished in two ways. First, L must be selected to ensure that the passbands of $G_c(\alpha - \alpha_0)$ are sufficiently removed from the channel-pair region. In practice, a decimation factor of $L=N'/4$ has been found to be a good compromise between maintaining computational efficiency and minimizing cycle leakage and cycle aliasing [4]. Second, cycle leakage can be reduced by minimizing the skirts and sidelobes of $M_1(\alpha,f)$ and $G_c(\alpha)$. From Equation (18) it is evident that a channelizer with steep transition bands and low sidelobes produces low skirts and low sidelobes in the channel-pair region. Thus, a data tapering window whose Fourier transform has low skirts and low sidelobes (e.g., the Hamming window) is desirable. A data tapering window for the output lowpass filter is not as crucial as a data tapering window for the input bandpass filters. For simplicity $g_c(p)$ is often taken to be a rectangular window.

The computational complexity of the FAM, in terms of the number of complex multiplications required to compute estimates of the cyclic cross spectrum of two complex signals and the cyclic spectrum of a single real signal, is given in Table II. In order to give some feel for the magnitude of calculations typically encountered when estimating the cyclic cross spectrum with the FAM, a numerical example is given. Consider the previous problem of estimating the cyclic cross spectrum of two complex signals where a time-frequency resolution product of $\Delta t\Delta f = 16384$ and a frequency resolution of $\Delta f = 1/8$ is desired. Channelization for this problem requires approximately 2×10^6 complex multiplications; computing and Fourier transforming all product sequences requires approximately 4×10^6 and 67×10^6 complex multiplications respectively. If a complex multiplication is performed every 150 ns, this computation would take approximately 11 seconds to perform. Estimating the cyclic spectrum of a real signal with the same processing parameters would take on the order of 2.75 seconds.

A close look at the FAM reveals that the algorithm has a high degree of parallelism [8]. The parallelism in the FAM is a direct result of the independence of the product sequences $X_T(p,f_k)Y^*_T(p,f_l)$ for each k and l during and after their computation. A convenient way to describe the parallelism in the FAM is through a series of time-sequenced signal flow graphs (SFGs). The time sequencing of the signal flow graphs indicates the ordering of computations in the algorithm, that is, the output of SFG(i) is the input to SFG(i+1). The ordered

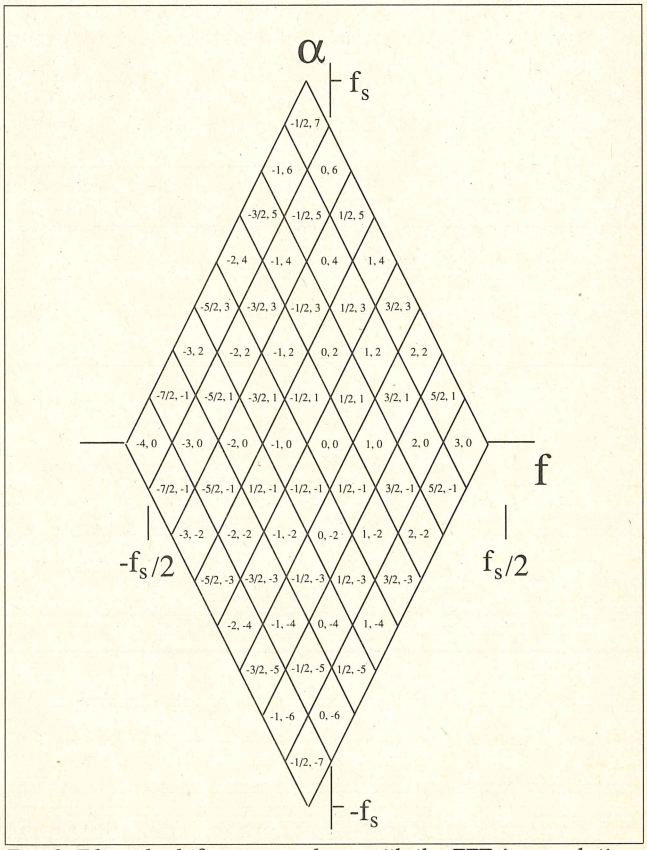

Fig. 8. Tiling the bifrequency plane with the FFT Accumulation Method for N'=8

TABLE II
FFT ACCUMULATION METHOD

Computation Section	Number of Complex Multiplications	
	Cyclic Cross Spectrum of Two Complex Signals	Cyclic Spectrum of a Single Real Signal
Channelizer:		
Data Tapering	2 N′P	†
N′-point FFT	PN′$\log_2$N′	P(N′/2)$\log_2$N′
Frequency Shift	2 N′P	N′P
Cross Multiply	P(N′)^2	P[(N′)^2]/4
FFT Product Sequences	(N′)^2(P/2)$\log_2$P	[(N′)^2/4](P/2)$\log_2$P

Computational complexity of the FAM in terms of the number of complex multiplications. † Note that N′P real multiplications are required for data tapering.

signal flow graph description also has the advantage of easy translation into digital architectures. For the FAM the signal flow graphs are:

SFG(1) Computation of the Complex Demodulates
SFG(2) Cross Multiplication, Complex Demodulates
SFG(3) P -point FFT array (including I/O buffers)

These three signal flow graphs are shown in Fig. 9.[2] For clarity the FFT signal flow graphs in SFG(1) and SFG(3) are shown as blocks. Additionally, only one node of SFG(3) is shown.

Operationally, data is input to SFG(1) in blocks of L samples and output in N'-point blocks. Each block of data contains $X_T(\cdot, f_k)$ and $Y^*_T(\cdot, f_k)$, $-N'/2 \le k \le (N'/2)-1$. As illustrated in Fig. 9, data from SFG(1) is input to SFG(2). SFG(2), which cross multiplies all complex demodulates, is detailed in Fig. 10 for $N' = 8$. The nodes of SFG(2) in Fig. 10 are labeled according to channel--pair regions associated with the node (cf. Fig. 8). In general, SFG(2) consists of 2N′ complex data paths and $(N')^2$ complex multiplier nodes configured in an $N' \times N'$ array. Complex demodulate streams flow through a node of SFG(2), are cross multiplied and fed into nodes of SFG(3). SFG(3) Fourier transforms the output product sequence streams of SFG(2). For each node in SFG(2) there is a corresponding node in SFG(3). The array structure of SFG(3) is similar to that of SFG(2) except that nodes of SFG(3) are not connected to one another – due to the independence of the product sequences. Each node of SFG(3) consists of an input buffer and a P-point FFT processor. The buffers convert the product sequence streams output from SFG(2) into P-point blocks suitable for Fourier transforming by the FFT processors.

The parallel description of the FAM can be used as a guide to mapping the FAM calculation onto a multiprocessor computer, or for designing application-specific architectures. In either case the algorithm is easily partitioned. The most natural partition for the FAM is by channel-pair region. If processing elements compute and Fourier transform all product sequences concurrently, then the amount of time required to estimate the cyclic cross spectrum decreases by a factor of $(N')^2$. Such an implementation significantly decreases the time required to perform the FAM computation since computing and Fourier transforming the product sequences represent the bulk of the computations. Whatever the implementation, it is important to maintain a smooth flow of data between the signal flow

[2] Note that the structure illustrated in Fig. 9 can be used to complete other quadratic surfaces, e.g. the cross ambiguity function, by modifying the pre- and post- quadratic tranform blocks (blocks (a) and (c) in Fig. 9).

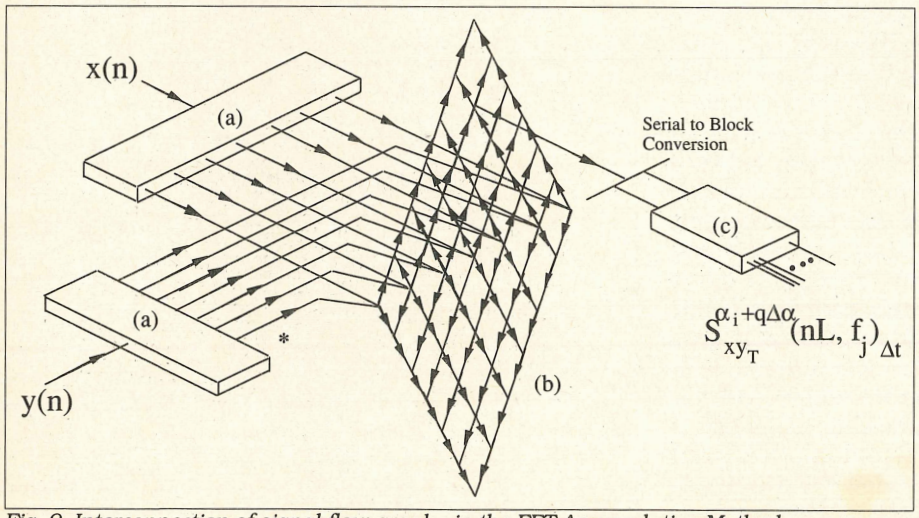

Fig. 9. Interconnection of signal flow graphs in the FFT Accumulation Method. a) N′- point FFT channelizers.b) SFG(2) for cross multiplication of complex demodulates. c) P-point FFT processors.

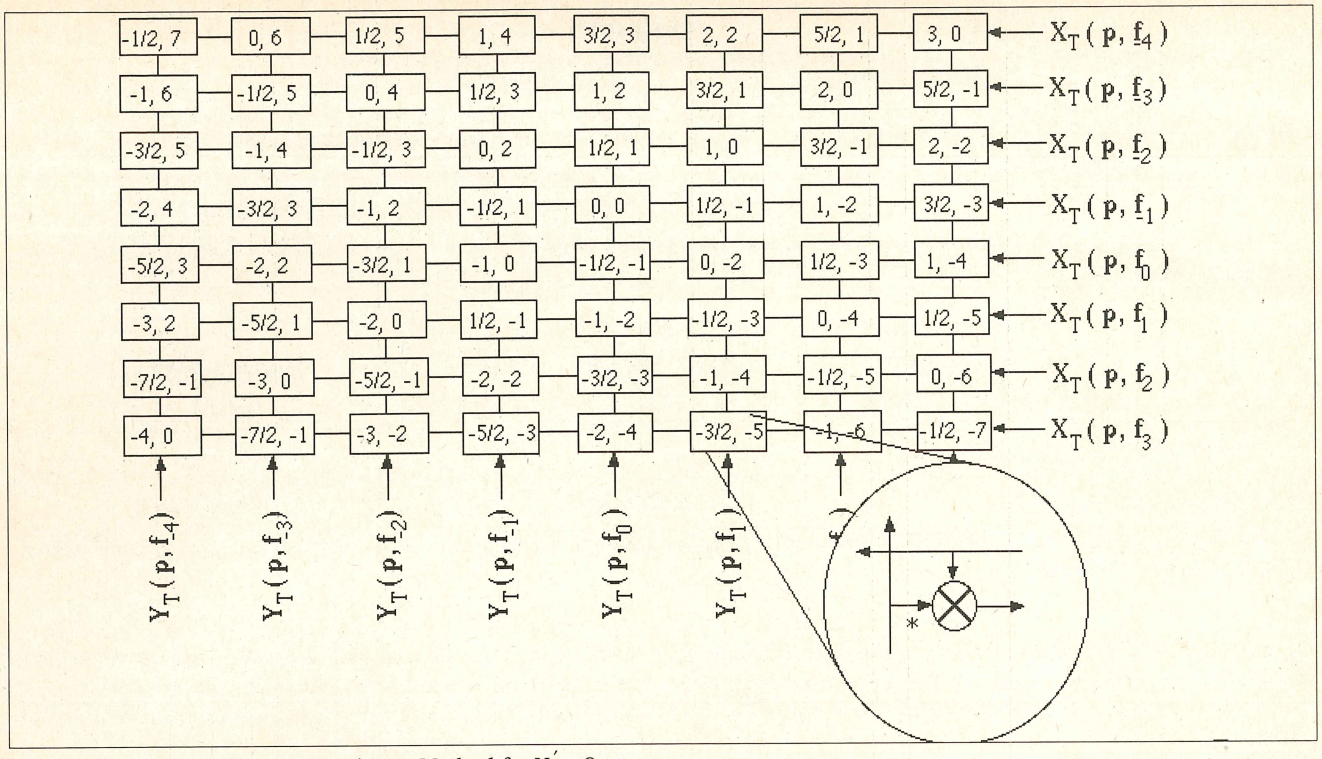

Fig. 10. SFG(2) in the FFT Accumulation Method for $N' = 8$.

graphs, i.e., to computationally balance the realization. As an example, in [8] it is shown that architectures based explicitly on the parallel description of the FAM contain redundant arithmetic units. However, by computationally balancing designs based explicitly on the parallel description, efficient architectures can be realized.

The strip spectral correlation algorithm

The second FFT-based time smoothing algorithm is the Strip Spectral Correlation Algorithm (SSCA) [4], [5]. In the SSCA the complex demodulates $X_T(n,f_k)$ directly multiply $y^*(n)$. As a result, the point estimates produced by the SSCA lie along the frequency-skewed family of lines $\alpha = 2f_k - 2f$ (see Fig. 11). A strip of point estimates is computed using the formula

$$S_{xy_T}^{f_k+q\Delta\alpha}\left(n,\frac{f_k}{2}-q\frac{\Delta\alpha}{2}\right)_{\Delta t}=\sum_r X_T(r,f_k)y^*(r)g(n-r)e^{-i2\pi qr/N} \quad (34)$$

(Recall that $f_k = k(f_s/N')$, $-N'/2 \le k \le (N'/2) - 1$). The SSCA uses an N-point FFT to compute the sum in

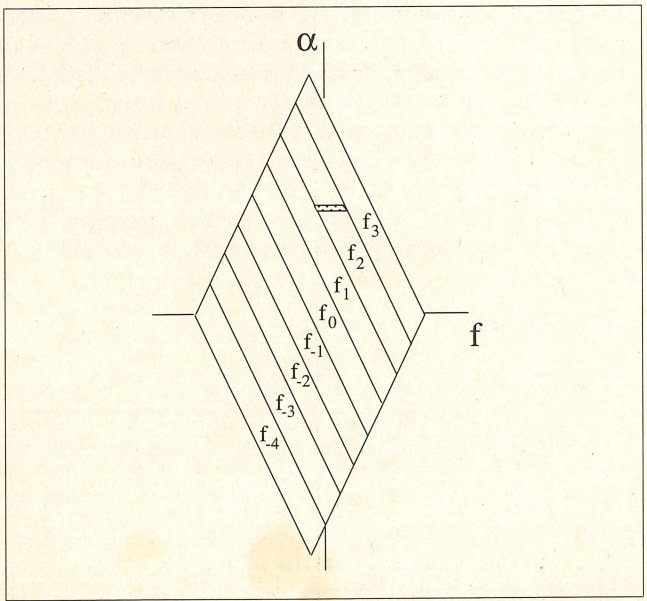

Fig. 11. Tiling the bifrequency plane with the Strip Spectral Correlation Algorithm for $N' = 8$. A single CSA cell is shown..

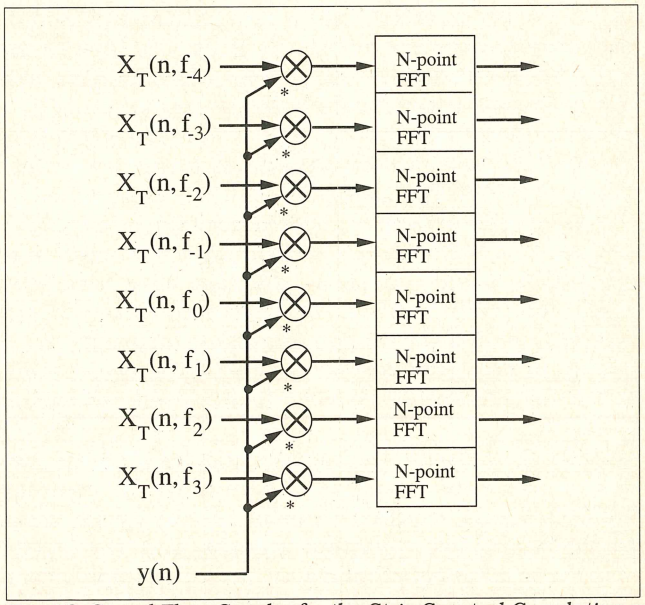

Fig. 12. Signal Flow Graphs for the Strip Spectral Correlation Algorithm excluding SFG(1) (the channelizer for x(n)).

TABLE III
STRIP SPECTRAL CORRELATION ALGORITHM

Computation Section	Number of Complex Multiplications	
	Cylic Cross Spectrum of Two Complex Signals	Cyclic Spectrum of a Single Real Signal
Channelizer:		
Data Tapering	NN′	†
N′-point FFT	$N(N'/2)\log_2 N'$	$N(N'/2)\log_2 N'$
Frequency Shift	NN′	NN′/2
Compute Product Sequences	NN′	NN′/2
FFTs Product Sequences	$N'(N/2)\log_2 N$	$(N'/2)(N/2)\log_2 N$

Computational complexity of the SSCA in terms of the number of complex multiplications. † Note that N N′ real multiplications are required for data tapering.

Equation (34). Note that the sampling rate of $X_T(n, f_k)$ cannot be decimated; in order for $X_T(n, f_k)$ to properly multiply $y^*(n)$ both sequences must be sampled at the same rate.

The properties of the SSCA are determined by studying the CSA cells produced by the algorithm. The SSCA kernel transform for a CSA cell located at (f_0,α_0) is [4]

$$M(\alpha,f)=G(\alpha-\alpha_0)A\left(f-f_0+\frac{\alpha-\alpha_0}{2}\right) \tag{35}$$

where the allowable values of f_0 and α_0 are

$$f_0=\frac{f_k}{2}-q\frac{\Delta\alpha}{2} \tag{36}$$

$$\alpha_0=f_k+q\Delta\alpha \tag{37}$$

In Equation (35), $G(\alpha)$ and $A(f)$ are the FSTs of $g(n)$ and $a(n)$. From Equations (35) – (37), the region of support of $A(f - f_0 + (\alpha - \alpha_0)/2)$ is found to be a strip approximately Δa wide along the line $\alpha = 2f_k - 2f$. Thus, CSA cells in the SSCA have a constant length of Δa which implies that the SSCA has a constant frequency resolution of $\Delta f = \Delta a$. The uniform frequency resolution of the SSCA is in marked contrast to the nonuniform frequency resolution of the FAM. The width of a CSA cell, determined by the bandwidth of $G(\alpha)$, is gleaned from Equation (35) to be $1/\Delta t = f_s/N$. Hence, the cycle frequency resolution of the SSCA is $\Delta\alpha = 1/\Delta t = f_s/N$. The time-frequency resolution product of the SSCA is therefore $\Delta t\Delta f = N/N'$. Figure 11 shows the bifrequency plane tiled with strips of $A(\cdot)$. In all, there are N' strips in the cyclic cross spectrum estimate and $N'/2$ in the cyclic spectrum estimate.

Table III summarizes the computational complexity of the SSCA in terms of the number of complex multiplications required to estimate the cyclic cross spectrum of two complex signals and the cyclic spectrum of a single real signal. As an example of the computational complexity, the previous problem where $\Delta f = 1/8$ and $\Delta t\Delta f = 16348$ which corresponds to $N' = 8$ and $N = 131072$, would require approximately two seconds to execute. Estimating the cyclic spectrum of a single real signal with the same parameters would take approximately one second.

Like the FAM, the SSCA is a highly parallel algorithm and the parallel description of the SSCA is similar to the parallel description of the FAM but less complex [8]. The parallel description of the SSCA consists of three time-sequenced signal flow graphs:

SFG(1) Computation of Complex Demodulates
SFG(2) Computation of Product Sequences
SFG(3) N-point FFT (including I/O buffers)

SFG(1) is identical to that of SFG(1) in the FAM parallel description except that no sample rate decimation occurs ($L=1$). SFG(2) is composed of N' complex multiplier nodes configured in a linear array. Figure 12 shows the SSCA signal flow graphs, excluding SFG(1), for $N' = 8$. Complex demodulate streams from SFG(1) enter SFG(2) on the left and product sequence streams exit SFG(2) on the right for SFG(3). The array structure of SFG(3) follows directly from the array structure of SFG(2). Additionally, the nodes of SFG(3) are identical to nodes of SFG(3) in the FAM parallel description (except N-point FFTs are used instead of P-point FFTs). The prior comments on the use of the parallel description to implement the FAM are also applicable to the SSCA.

SUMMARY

We began the discussion with a simple algorithm based on a direct implementation of the time smoothed cyclic cross periodogram. Although this algorithm is not computationally attractive for estimating the cyclic (cross) spectrum over the entire bifrequency plane, it illustrates the essence of time smoothing algorithms. Several representations of the basic time smoothing algorithm were developed next. The transformed kernel representation is perhaps the most important characterization of cyclic spectral analysis algorithms due to

its generality and conceptual utility. From this representation the signal processing attributes of an algorithm, such as frequency and cycle frequency resolution, and phenomena such as cycle leakage can be studied. Additionally, the computational complexity of cyclic spectral analysis algorithms is readily determined using this representation. From the basic time smoothing algorithm several improved algorithms were developed. Modifications to reduce the computational complexity of the basic algorithm included 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, we arrived at the computationally efficient FFT Accumulation Method. A second computationally efficient algorithm, the Strip Spectral Correlation Algorithm was developed as an alternative to the 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.

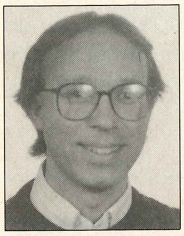

Randy S. Roberts (M'89) was born in Oildale, CA, on October 26, 1957. He received the B.S. degree in electrical engineering in 1980, the M.S. degree in electrical engineering in 1982, and the Ph.D. degree in electrical engineering in 1989, all from the University of California at Davis.

From 1982 to 1984 he was a member of the Technical Staff at Signal Science Inc., Santa Clara, CA, working in the area of digital signal processing. He is currently a staff member at the Los Alamos National Laboratory, Los Alamos, NM. His research interests include statistical signal processing, parallel algorithms and architectures, and image processing.

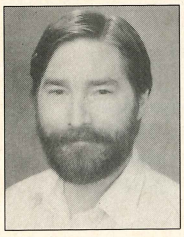

William A. Brown (S'73-M'78) was born in Richland, WA, on February 13, 1952. He received the B.S. degree from California State University, Chico, in 1974, the M.S. degree from Illinois Institute of Technology, Chicago, in 1975, and the Ph.D. degree from University of California, Davis, in 1987, all in electrical engineering.

From 1977 to 1981 he was a member of the technical staff at ARGOSystems, Sunnyvale, CA, working on passive sonar signal processing. From 1981 to 1987 he was engaged in doctoral research on the theory of cyclostationary signals. Since 1987, he has been a staff scientist for Mission Research Corporation, Monterey, CA, working on communication signal interception as well as mitigation of ionospheric propagation effects on radar tracking and imaging systems.

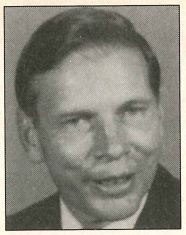

Herschel H. Loomis, Jr. (S'59-M'63-SM'75) was born in Wilmington, Delaware on May 31, 1934. He graduated from Wilmington Friends School in 1952. He received the B.E.E. Degree from Cornell University in 1957, the M.S. in Electrical Engineering from the University of Maryland in 1959 and the Ph.D. degree from Massachusetts Institute of Technology in 1963.

From 1957 until 1959 he served as an electronic engineer with the United States Navy. He was Assistant Professor, Associate Professor and Professor of Electrical and Computer Engineering at the University of California, Davis, California from 1963 until 1981. He served as Department Chairman at UC Davis from 1970 until 1975. From 1981 to 1983, he was Naval Electronics Systems Command Chair Professor at the Naval Postgraduate School in Monterey. He is currently Professor of Electrical and Computer Engineering at the Naval Postgraduate School.

He has served as a consultant to Lawrence Livermore Laboratory, Fairchild Semiconductor, and Signal Science, Inc.

Dr. Loomis is a Senior Member of the IEEE, and is a member of the Association for Computing Machinery, Eta Kappa Nu, Phi Kappa Phi, Sigma Xi and Tau Beta Pi. He is graduate scholarship chairman for the Special Interest Group on Design Automation of the ACM.

REFERENCES

[1] W. A. Gardner, "Exploitation of spectral redundancy in man-made signals," *IEEE SP Magazine,* this issue.

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

[3] W. A. Gardner, "Measurement of spectral correlation," *IEEE Trans. Acoust., Speech, Signal Proc.,* vol. ASSP-34, No. 5, pp. 1111-1123, October 1986.

[4] W. A. Brown and H. H. Loomis, Jr. "Digital implementations of spectral correlation analyzers," *Proceedings of the Fourth Annual ASSP Workshop on Spectrum Estimation and Modeling,* Minneapolis MN. August 1988. pp. 264-270.

[5] W. A. Brown,"On the Theory of Cyclostationary Signals," Ph.D Dissertation, Department of Electrical Engineering and Computer Science, University of California, Davis, 1987. (Thesis Advisors: W. A. Gardner and H. H. Loomis, Jr.)

[6] R. Crochiere and L. Rabiner, *Multirate Digital Signal Processing.* Englewood Cliffs, NJ: Prentice-Hall, 1983.

[7] B. G. Agee and W. A. Gardner, "Cyclic spectrum analysis study: executive summary and final briefing," ARGOSystems Technical Report No. B83-0014, June 1984.

[8] R. S. Roberts, "Architectures for Digital Cyclic Spectral Analysis," Ph.D Dissertation, Dept. of Electrical Engineering and Computer Science, University of California, Davis, 1989. (Thesis Advisors: W. A. Gardner and H. H. Loomis, Jr.)

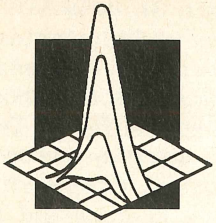

Ph. D. Abstracts

Title: **Exploitation of Spectral Correlation for Signal-Selective Direction Finding**

Author: **Stephan V. Schell**
Advisor: **William A. Gardner**
Granting Institution: University of California, Davis
Acceptance Date: Nov. 14, 1990
Further Information: Dr. Stephan V. Schell or Prof. William A. Gardner,
Dept. of Electrical Engineering and Computer Science, University of California, Davis, CA 95616.
Tel. (916) 752-1326, 1951, 0583.
Email: schell@iris.ucdavis.edu or {ucbvax,lll-crg}!ucdavis!iris!schell

The spectral-correlation theory of cyclostationary time-series is applied to the signal-selective direction-finding problem, in which desired signals, undesired signals or interference, and noise are received by an array of sensors, but the directions of arrival of only the desired signals are to be estimated. Several new direction-finding methods (including Cyclic MUSIC) that exploit spectral correlation to discriminate against the undesired signals or interference and noise are developed and analyzed. The performance of these methods is evaluated for finite numbers of data samples in an extensive set of Monte Carlo simulations, where it is shown that the new methods can greatly outperform conventional super-resolution methods, such as MUSIC and maximum likelihood, in some environments. Also, the Cramer-Rao Lower Bound (CRLB) for the directions of arrival (and other parameters) of Gaussian cyclostationary signals is developed and evaluated for several specific signal environments, where it is shown to differ substantially from the CRLB for Gaussian stationary signals.

Cyclic Spectral Analysis Software Package

The Cyclic Spectral Analysis Software package rapidly computes and displays spectral correlation surfaces (cyclic spectra) for internally generated signals or user-supplied data. As a basic signal analysis tool for communications, telemetry, reconnaissance and direction finding, the CSAS package can be used for:

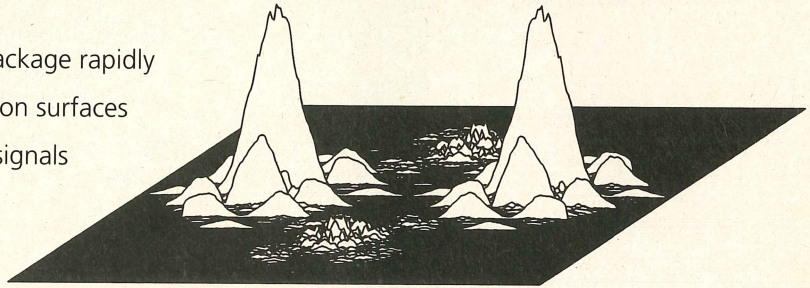

Signal Detection

Signal Classification

Modulation Recognition

Parameter Estimation

Synchronization

Time Delay Estimation

System Identification

Features

Graphical display software that drives Postscript laser printers and Tektronix graphics terminals

C source code and user's manuals

Runs on UNIX-like operating systems and VMS

Complex-PAM and Gaussian noise generators

Base price: $3495.

For information and ordering:
SSPI 6950 Yount Street Yountville, California 94599
707 944.0648