Publications of the Astronomical Society of the Pacific, **116**:84–96, 2004 January

Correlation Statistics of Quantized Noiselike Signals

Carl Gwinn

Department of Physics, University of California, Santa Barbara, CA 93106; cgwinn@condor.physics.ucsb.edu

Received 2003 July 28; accepted 2003 November 4; published 2003 December 12

ABSTRACT. I calculate the statistics of correlation of two digitized noiselike signals, which are drawn from complex Gaussian distributions, sampled, quantized, correlated, and averaged. Averaged over many such samples, the correlation r approaches a Gaussian distribution. The mean and variance of r fully characterize the distribution of r. The mean corresponds to the reproducible part of the measurement, and the variance corresponds to the random part, or noise. I investigate the case of nonnegligible covariance ρ between the signals. Noise in the correlation can increase or decrease, depending on quantizer parameters, when ρ increases. This contrasts with the correlation of continuously valued or unquantized signals, for which the noise in phase with ρ increases with increasing ρ, and noise out of phase decreases. Indeed, for some quantizer parameters, I find that the correlation of quantized signals provides a more accurate estimate of ρ than would correlation without quantization. I present analytic results in exact form and as polynomial expansions, and compare these mathematical results with results of computer simulations.

1. INTRODUCTION

Nearly all signals from astrophysical sources can be represented as electric fields comprising Gaussian noise. These noiselike signals have zero mean. All information about the source is contained in the variance of the electric field and in covariances between different polarizations, positions, times, or frequency ranges. The intensity, for example, is simply the sum of the variances in two basis polarizations. More generally, all the Stokes parameters can be expressed in terms of the variances and covariances of these two polarizations. Similarly, in interferometry the covariance of electric fields at different positions is the visibility, the Fourier transform of source structure. In correlation spectroscopy the covariances of electric field at different time separations, expressed as the autocorrelation function, are the Fourier transform of the spectrum. Because the signals are drawn from Gaussian distributions, their variances and covariances completely characterize them. The single known exception to this rule of Gaussian statistics is radiation from pulsars, under certain observing conditions (Jenet, Anderson, & Prince 2001).

Particularly at wavelengths of a millimeter or more, covariances are usually estimated by correlation. (Actually, correlation is used at all wavelengths, but at wavelengths shortward of a millimeter quantum mechanical processes come into the picture, complicating it). Correlation involves forming products of samples of the two signals. The average of many such products approximates the covariance. In mathematical terms, for two signals x and y, the covariance is $\rho = \langle xy \rangle$, where the angular brackets $\langle \ldots \rangle$ represent a statistical average over an ensemble of all statistically identical signals. Correlation approximates this enormously infinite average with a finite average over N_q samples of x and y: $r_\infty = 1/N_q \sum_{i=1}^{N_q} x_i y_i$. Here, the subscript ∞ reflects the fact that x and y are unquantized; their accuracy is not limited to a finite number of quantized levels. The subscript q indicates sampling at the Nyquist rate, as I assume (see Thompson, Moran, & Swenson 1986).

Because the number of samples in most measurements of correlation is large, the results of a finite correlation follow a Gaussian distribution. This is a consequence of the central-limit theorem. Thus, one expects a set of identical measurements of r_∞ to be fully characterized by their mean, $\langle r_\infty \rangle$, and their standard deviation, $\langle (r_\infty - \langle r_\infty \rangle)^2 \rangle^{1/2}$. The mean is the deterministic part of the measurement; it provides an estimate of ρ. The standard deviation characterizes the random part of the measurement and is often called "noise" (but is to be distinguished from the noiselike signals x and y that are being correlated). In principle, the best measurement minimizes the random part while preserving the relation between the deterministic part and ρ. The signal-to-noise ratio (S/N) of the correlation, $\mathcal{R}_\infty = \langle r_\infty \rangle / \langle (r_\infty - \langle r_\infty \rangle)^2 \rangle^{1/2}$, provides a figure of merit that quantifies the relative sizes of deterministic and random parts (see, e.g., Thompson et al. 1986).

The electric field is commonly digitized before correlation. Digitization includes sampling and quantization. Sampling involves averaging the signal over short time windows; it thus restricts the range of frequencies that can be uniquely represented. For simplicity, in this paper I restrict discussion to "video" or "baseband" signals, for which a frequency range of 0 up to some maximum frequency is present; and I assume that they are sampled at the Nyquist rate, or at half the shortest period represented. I also assume that the signals are "white," in the sense that samples x_i and y_j are correlated only if $i = j$; and that the signals are stationary, so that the correlation of x_i and y_j is independent of i. These assumptions limit the influence of sampling. I will discuss spectrally varying signals elsewhere (C. Gwinn 2004, in preparation).

Quantization limits the values that can be represented so that the digitized signal imperfectly represents the actual signal. Quantization thus introduces changes in both the mean correlation $\langle r_M \rangle$ and its standard deviation $\langle (r_M - \langle r_M \rangle)^2 \rangle^{1/2}$. Here the subscript "$M$" represents the fact that the quantized signal can take on M discrete values. The mean and standard deviation of r_M can be calculated from the statistics of the quantized signals $\hat{x}_i$ and $\hat{y}_i$, and the details of the quantization scheme.

A number of previous authors have addressed the effects of quantization on correlation of noiselike signals (see, e.g., Thompson et al. 1986, chap. 8, and references therein). Notably, Van Vleck & Middleton (1966) found the average correlation and the standard deviation for two-level quantization—the case in which quantization reduces the signals to only their signs. Cooper (1970) found the average correlation and its standard deviation for small normalized covariance $\rho \ll 1$, for four-level correlation: quantization reduces the signals to signs and whether they lie above or below some threshold v_0. He found the optimal values of v_0 and the relative weighting of points above and below the threshold n, as quantified by the S/N $\mathcal{R}_4$. Hagen & Farley (1973) generalized this to a broader range of quantization schemes and studied effects of oversampling. Bowers & Klingler (1974) examined Gaussian noise and a signal of general form in the small-signal limit. They devise a criterion for the accuracy with which a quantization scheme represents a signal and show that this yields the highest S/N for $\rho \ll 1$. Most recently, Jenet & Anderson (1998) examined the case of many-level correlators. They use a criterion similar to that of Bowers & Klingler (1974) to calculate the optimal level locations for various numbers of levels. Jenet & Anderson (1998) also find the mean spectrum for a spectrally varying source, as measured by an autocorrelation spectrometer; they find that quantization introduces a uniform offset to the spectrum and scales the spectrum by a factor, and they calculate this offset and factor. D'Addario et al. (1984) provide an extensive analysis of errors in a three-level correlator.

In this paper, I calculate the average correlation and its standard deviation for nonvanishing covariance ρ. I provide exact expressions for these quantities, and approximations valid through fourth order in ρ. Interestingly, noise actually declines for large ρ for correlation of quantized signals. Indeed, the S/N for larger ρ can actually exceed that for correlation of an unquantized signal. In other words, correlation of a quantized signal can provide a more accurate measure of ρ than would correlation of the signal before quantization. This fact is perhaps surprising; it reflects that correlation is not always the most accurate way to determine the covariance of two signals.

The organization of this paper is as follows. In § 2, I review the statistics of the correlation of unquantized, or continuously variable, complex signals. In § 3, I present expressions that give the average correlation and the standard deviation, in terms of integrals involving the characteristic curve. I include statistics of real and imaginary parts, which are different. I present expansions of these integrals as a power series in ρ. In § 4 I discuss computer simulations of correlation to illustrate these mathematical results. I summarize the results in § 5.

2. CORRELATION OF UNQUANTIZED VARIABLES

2.1. Bivariate Gaussian Distribution

Consider two random, complex signals, x and y. Suppose that each of these signals is a random variable drawn from a Gaussian distribution. Suppose that the signals x and y are correlated so that they are, in fact, drawn from a Gaussian joint probability density function $P(x,y)$ (Meyer 1975). Without loss of generality, I assume that each signal has variance of 2:

$$\langle xx^* \rangle = \langle yy^* \rangle \equiv 2. \tag{1}$$

In this expression, the angle brackets $\langle \ldots \rangle$ denote a statistical average; in other words, an average over all systems with the specified statistics. This choice for a variance of 2 for x and y is consistent with the literature on this subject, much of which treats real signals (rather than complex ones) drawn from Gaussian distributions with unit variance (Cooper 1970; Thompson et al. 1986). Of course, the results presented here are easily scaled to other variances for the input signal. I require that the signals themselves have no intrinsic phase; in other words, the statistics remain invariant under the transformation $x \to xe^{i\phi_0}$ and $y \to ye^{i\phi_0}$, where ϕ_0 is an arbitrary overall phase. It then follows that

$$\langle xx \rangle = \langle yy \rangle = 0. \tag{2}$$

From these facts, one finds:

$$\langle \mathrm{Re}(x)\mathrm{Re}(x) \rangle = \langle \mathrm{Im}(x)\mathrm{Im}(x) \rangle = 1$$
$$\langle \mathrm{Re}(x)\mathrm{Im}(x) \rangle = 0, \tag{3}$$

and similarly for y. Thus, real and imaginary parts are drawn from Gaussian distributions with unit variance. The distributions are circular in the complex plane for both x and y.

Without loss of generality, I assume that the normalized covariance of the signals ρ is purely real:

$$\rho \equiv \frac{\langle xy^* \rangle}{\sqrt{\langle |x|^2 \rangle \langle |y|^2 \rangle}} = \frac{1}{2} \langle xy^* \rangle. \tag{4}$$

One can always make ρ purely real by rotating x (or y) in the complex plane: $x \to xe^{i\phi_x}$, and $y \to y$. Note that because of the absence of any intrinsic phase,

$$\langle \mathrm{Re}(x)\mathrm{Re}(y)\rangle = \langle \mathrm{Im}(x)\mathrm{Im}(y)\rangle = \rho,$$
$$\langle \mathrm{Re}(x)\mathrm{Im}(y)\rangle = \langle \mathrm{Im}(x)\mathrm{Re}(y)\rangle = 0, \tag{5}$$

so that

$$\langle xy\rangle = 0. \tag{6}$$

In other words, the real parts of x and y are correlated, and the imaginary parts of x and y are correlated, but real and imaginary parts are uncorrelated. In mathematical terms, real parts (or imaginary parts) are drawn from the bivariate Gaussian distribution

$$P_2(X, Y) = \frac{1}{2\pi\sqrt{1-\rho^2}} \times \exp\left[\frac{-1}{2(1-\rho^2)}(X^2+Y^2-2\rho\, XY)\right], \tag{7}$$

where (X, Y) stands for either $(\mathrm{Re}(x), \mathrm{Re}(y))$ or $(\mathrm{Im}(x), \mathrm{Im}(y))$. The distributions for real and imaginary parts are identical, but real and imaginary parts are uncorrelated. Therefore,

$$P\,(x,y) = P_2\,(\mathrm{Re}(x), \mathrm{Re}(y)) \times P_2\,(\mathrm{Im}(x), \mathrm{Im}(y))\,. \tag{8}$$

2.2. Correlation of Unquantized Signals

2.2.1. Correlation

Consider the product of two random complex signals, drawn from Gaussian distributions as in the previous section, that are sampled in time. Suppose that the signals are not quantized; they can take on any complex value. The product of a pair of samples $c_i = x_i y_i^*$ does not follow a Gaussian distribution. Rather, the distribution of c_i is the product of an exponential of the real part, multiplied by the modified Bessel function of the second kind of order zero of the magnitude of c_i (Gwinn 2001). However, the average of many such products, over a large number of pairs of samples, approaches a Gaussian distribution, as the central limit theorem implies. In such a large but finite sum,

$$r_\infty = \frac{1}{2N_q}\sum_{i=1}^{N_q} x_i y_i^* \tag{9}$$

provides an estimate of the covariance ρ. Here the index i runs over the samples, which are commonly taken at different times. The total number of samples correlated is N_q. Hereafter, I assume that in all summations, indices run from 1 to N_q. The subscript ∞ on the correlation r_∞ again indicates that the correlation has been formed for variables x_i and y_i, which can take on any of an infinite number of values; in other words, it indicates that x_i and y_i have not been quantized.

2.2.2. Mean Correlation for Unquantized Signals

The mean correlation is equal to the covariance, in a statistical average

$$\langle r_\infty\rangle = \langle \mathrm{Re}(r_\infty)\rangle = \tfrac{1}{2}\,[\langle \mathrm{Re}(x_i)\mathrm{Re}(y_i)\rangle + \langle \mathrm{Im}(x_i)\mathrm{Im}(y_i)\rangle] = \rho, \tag{10}$$

where I used the assumption that the phase of ρ is zero, equation (5). This can also be seen from the mean of the distribution of the products $c_i = x_i y_i^*$ (Gwinn 2001), or simply by integrating over the joint distribution of x_i and y_i (eqs. [7] and [8]).

2.2.3. Noise for Correlation of Unquantized Signals

Because the distribution of r_∞ is Gaussian, the distribution of r_∞ is completely characterized by its mean (eq. [10]) and by the variances $\langle r_\infty r_\infty^*\rangle$ and $\langle r_\infty r_\infty\rangle$. Suppose that the samples x_i and y_i are independent; in mathematical terms, suppose that

$$x_i y_j = 0 \quad \text{for } i \neq j. \tag{11}$$

If they are not independent, the results will depend on the correlations among samples; this case is important if, for example, the signal has significant spectral structure. Jenet & Anderson (1998) find average spectrum in this case; I will discuss the noise in future work. Here I consider only independent samples. In that case,

$$\begin{aligned}\langle r_\infty r_\infty^*\rangle &= \frac{1}{4N_q^2}\sum_{ij}\langle x_i y_i^* x_j^* y_j\rangle \\ &= \frac{1}{4N_q^2}\sum_{i}\langle x_i y_i^* x_i^* y_i\rangle + \frac{1}{4N_q^2}\sum_{i\neq j}\langle x_i y_i^*\rangle\langle x_j y_j^*\rangle \\ &= \frac{1}{4N_q}\langle x_i y_i^* x_i^* y_i\rangle + \frac{N_q-1}{4N_q}(2\rho)^2,\end{aligned} \tag{12}$$

where I have separated the terms with $i = j$ from those with $i \neq j$, and appealed to the facts that x_i and y_j are covariant only if $i = j$ ("white" signals), and that for $i = j$ the statistics are stationary in i. For Gaussian variables with zero mean a, b, c, and d, all moments are related to the second moments,

$$\langle abcd\rangle = \langle ab\rangle\langle cd\rangle + \langle ac\rangle\langle bd\rangle + \langle ad\rangle\langle bc\rangle, \tag{13}$$

so that

$$\langle x_i y_i^* x_i^* y_i\rangle = (2\rho)^2 + 4. \tag{14}$$

Therefore,

$$\langle r_\infty r_\infty^*\rangle = \frac{1}{N_q} + \rho^2,$$

$$\langle r_\infty r_\infty^*\rangle - \langle r_\infty\rangle\langle r_\infty^*\rangle = \frac{1}{N_q}. \tag{15}$$

An analogous calculation yields

$$\langle x_i y_i^* x_i y_i^*\rangle = 2(2\rho)^2, \tag{16}$$

and so

$$\langle r_\infty^2\rangle = \frac{1}{N_q}\rho^2 + \rho^2,$$

$$\langle r_\infty^2\rangle - \langle r_\infty\rangle^2 = \frac{1}{N_q}\rho^2. \tag{17}$$

I combine these facts to find the means and standard deviations of the real and imaginary parts of the measured correlation r_∞:

$$\langle \mathrm{Re}(r_\infty)\rangle = \langle r_\infty\rangle = \rho,$$

$$\langle \mathrm{Im}(r_\infty)\rangle = 0,$$

$$\langle \mathrm{Re}(r_\infty)^2\rangle - \langle \mathrm{Re}(r_\infty)\rangle^2 = \frac{1}{2}(\langle r_\infty r_\infty^*\rangle + \langle r_\infty r_\infty\rangle) - \langle r_\infty\rangle^2$$

$$= \frac{1}{2N_q}(1+\rho^2),$$

$$\langle \mathrm{Im}(r_\infty)^2\rangle = \frac{1}{2}(\langle r_\infty r_\infty^*\rangle - \langle r_\infty r_\infty\rangle)$$

$$= \frac{1}{2N_q}(1-\rho^2). \tag{18}$$

If the number of independent samples N_q is large, the central-limit theorem implies that $\mathrm{Re}(r_\infty)$ and $\mathrm{Im}(r_\infty)$ are drawn from Gaussian distributions. The means and variances of these distributions, as given by equation (18), completely characterize r_∞. The fact that the real part of r_∞ has greater standard deviation than the imaginary part reflects the presence of self-noise or source noise. Sometimes this is described as the contribution of the noiselike signal to the noise in the result.

Commonly, and often realistically, astrophysicists suppose that ρ measures the intensity of one signal that has been superposed with two uncorrelated noise signals to produce x and y (see Bowers & Klingler 1974; Kulkarni 1989; Anantharamaiah et al. 1991). A change in ρ then corresponds to a change in the intensities $\langle |x|^2\rangle$ and $\langle |y|^2\rangle$ as well. Here, I suppose that $\langle |x|^2\rangle = \langle |y|^2\rangle \equiv 1$ while ρ varies. The results presented here can be scaled to those for the alternative interpretation.

2.2.4. S/N for Correlation of Unquantized Signals

The S/N for $\mathrm{Re}(r_\infty)$ is

$$\mathcal{R}_\infty(\mathrm{Re}(r_\infty)) = \frac{\langle \mathrm{Re}(r_\infty)\rangle}{\sqrt{\langle \mathrm{Re}(r_\infty)^2\rangle - \langle \mathrm{Re}(r_\infty)\rangle^2}}$$

$$= \sqrt{2N_q}\,\frac{\rho}{\sqrt{1+\rho^2}}. \tag{19}$$

Note that for a given number of observations N_q, S/N increases with ρ; the increase is proportional for $\rho \ll 1$.

A related quantity to S/N is the rms phase, statistically averaged over many measurements. The phase is $\phi(r_\infty) = \tan^{-1}[\mathrm{Im}(r_\infty)/\mathrm{Re}(r_\infty)]$. When the number of observations is sufficiently large, and the true value of the phase is 0, as assumed here, the standard deviation of the phase is $\langle\phi(r_\infty)\rangle = \langle \mathrm{Im}(r_\infty)^2\rangle^{1/2}/\langle \mathrm{Re}(r_\infty)\rangle$. The inverse of the standard deviation of the phase (in radians) is analogous to the S/N for the real part (eq. [19]). This S/N for phase is

$$\mathcal{R}_\infty(\phi(r_\infty)) = \frac{\langle \mathrm{Re}(r_\infty)\rangle}{\sqrt{\langle \mathrm{Im}(r_\infty)^2\rangle}} = \sqrt{2N_q}\,\frac{\rho}{\sqrt{1-\rho^2}}. \tag{20}$$

For constant N_q the S/N of the phase increases with ρ and increases much faster than proportionately for $\rho \to 1$.

3. QUANTIZED SIGNALS

Quantization converts the variables x, y to discrete variables $\hat{x}$, $\hat{y}$. These discrete variables depend on x and y through a multiple-step function known as a characteristic curve. Each step extends over some range $[v_i, v_{i+1}]$ of the unquantized signal and is given some weight n_i in correlation. The function $\hat{X}(X)$ denotes the characteristic curve, where again X stands for either $\mathrm{Re}(x)$ or $\mathrm{Im}(x)$. The complex quantized variable $\hat{x}$ is thus given by $\hat{x} = \hat{X}(\mathrm{Re}(x)) + i\hat{X}(\mathrm{Im}(x))$. The same characteristic curve is applied to real and imaginary parts. I hold open the possibility that the characteristic curves $\hat{X}(X)$ and $\hat{Y}(Y)$ are different. I assume in this paper that the characteristic curves are antisymmetric: $\hat{X}(-X) = -\hat{X}(X)$, and similarly for Y. D'Addario et al. (1984) describe effects of departures from antisymmetry and how antisymmetry can be enforced. Figure 1 shows a typical characteristic curve for four-level (or two-bit) sampling.

Systems with M levels of quantization can be described by

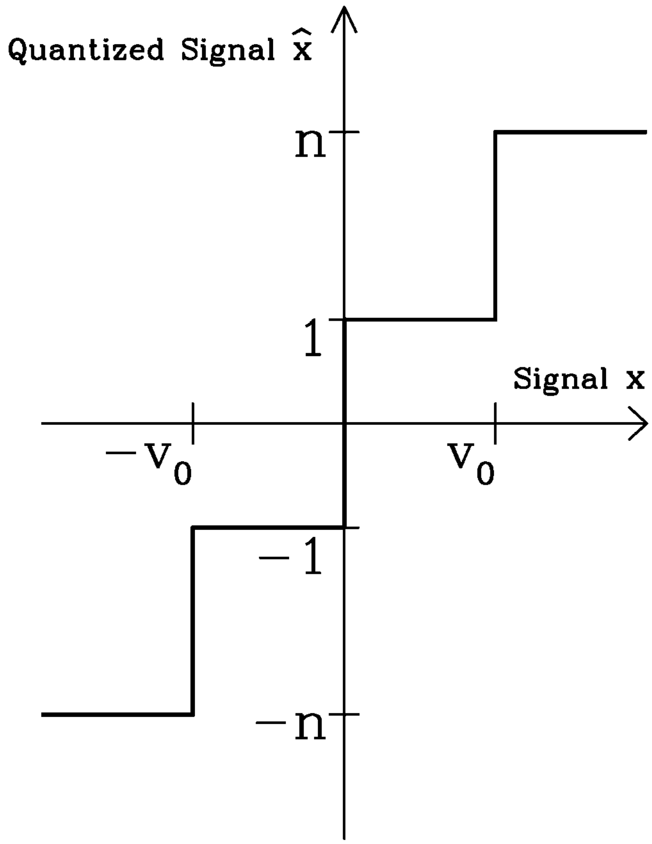

FIG. 1.—Characteristic curve $\hat{X}(X)$ for four-level quantization.

analogous, more complicated characteristic curves, and corresponding sets of weights $\{n_i\}$ and levels $\{v_{ix}\}$ and $\{v_{iy}\}$ (Jenet & Anderson 1998). For M-level sampling, the correlation of N_q quantized samples is

$$\hat{r}_M = \frac{1}{2N_q}\sum_{\imath=1}^{N_q} \hat{x}_\imath \hat{y}_\imath^*. \quad (21)$$

In practical correlators, deviations from theoretical performance can often be expressed as deviations of the characteristic curve from its desired form. In principle, these can be measured by counting the number of samples in the various quantization ranges and using the assumed Gaussian distributions of the input signals to determine the actual levels v_{ix} and v_{iy}. D'Addario et al. (1984) present an extensive discussion of such errors, and techniques to control and correct them.

Although the results presented below are applicable to more complicated systems, the four-level correlator will be used as a specific example in this paper, with correlation $\hat{r}_4$. For four-level sampling, commonly a sign bit gives the sign of X, and an amplitude bit assigns weight 1 if $|X|$ is less than some threshold v_{0x}, and weight n if $|X|$ is greater than v_{0x}. Together, sign and amplitude bits describe the four values possible for $\hat{X}(X)$. Other types of correlators, including two-, three-, or "reduced" four-level (in which case the smallest product for $|\hat{X}| = 1$, and $|\hat{Y}| = 1$ is ignored), can be formed as special cases or sums of four-level correlators (Hagen & Farley 1973).

3.1. Correlation of Quantized Signals: Exact Results

3.1.1. Mean Correlation: Exact Result

Ideally, from measurement of the quantized correlation $\hat{r}_M$ one can estimate the true covariance ρ. The statistical mean of $\hat{r}_M$ is

$$\begin{aligned}\langle \hat{r}_M \rangle &= \frac{1}{2N_q}\sum_i \langle \hat{x}_i \hat{y}_i^* \rangle \\ &= \frac{1}{2}\langle [\mathrm{Re}(\hat{x})\mathrm{Re}(\hat{y}) + \mathrm{Im}(\hat{x})\mathrm{Im}(\hat{y})] \\ &\quad + i\,[-\mathrm{Re}(\hat{x})\mathrm{Im}(\hat{y}) + \mathrm{Im}(\hat{x})\mathrm{Re}(\hat{y})]\rangle. \quad (22)\end{aligned}$$

Because $\mathrm{Re}(\hat{x})$ depends only on $\mathrm{Re}(x)$ and $\mathrm{Im}(\hat{y})$ depends only on $\mathrm{Im}(y)$, and $\mathrm{Re}(x)$ and $\mathrm{Im}(y)$ are completely independent [and similarly for $\mathrm{Im}(\hat{x})$ and $\mathrm{Re}(\hat{y})$],

$$\langle \mathrm{Re}(\hat{x})\mathrm{Im}(\hat{y})\rangle = \langle \mathrm{Im}(\hat{x})\mathrm{Re}(\hat{y})\rangle = 0. \quad (23)$$

Thus, the imaginary part of $\langle \hat{r} \rangle$, which involves products of these statistically independent terms, has an average of zero (eq. [3]). For the real part,

$$\langle \mathrm{Re}(\hat{x})\mathrm{Re}(\hat{y})\rangle = \langle \mathrm{Im}(\hat{x})\mathrm{Im}(\hat{y})\rangle, \quad (24)$$

where I use the assumption that the characteristic curves are identical for real and imaginary parts, and that real and imaginary parts of x and y have identical statistics (eqs. [3] and [5]). I use the bivariate Gaussian distribution for real and imaginary parts to find a formal expression for the statistical average of the correlation:

$$\begin{aligned}\langle \hat{r}_M \rangle &= \langle \mathrm{Re}(\hat{x})\mathrm{Re}(\hat{y})\rangle \\ &= \Upsilon_{XY} \equiv \int dX\, dY\, P_2(X,Y)\hat{X}(X)\hat{Y}(Y). \quad (25)\end{aligned}$$

This integral defines Υ_{XY}. For the assumed antisymmetric characteristic curves $\hat{X}(X)$ and $\hat{Y}(Y)$, one can easily show that $\partial\Upsilon_{XY}/\partial\rho > 0$. In other words, the ensemble-averaged quantized correlation is an increasing function of the covariance ρ for completely arbitrary quantizer settings (so long as the characteristic curves are antisymmetric).

The discussion of this section reduces the calculation of the average quantized correlation to that of integrating $P_2(X,Y)$ over each rectangle in a grid, with the edges of the rectangles given by the thresholds in the characteristic curves (Kokkeler, Fridman, & van Ardenne 2001). The function Υ_{XY} depends on ρ through $P_2(X,Y)$. This function is usually expanded through first order in ρ, because ρ is small in most astrophysical observations.

3.1.2. Simpler Form for Υ_{XY}

The integral Υ_{XY} and similar integrals can be converted into one-dimensional integrals for easier analysis. If one defines

$$Q(v_{0x},v_{0Y}) = \int_{v_{0X}}^{\infty} dX \int_{v_{0Y}}^{\infty} dY\, P_2(X,Y), \tag{26}$$

then the Fourier transform of $P_2(v_{0x},v_{0Y})$ is equal to that of $\partial Q(v_{0x},v_{0Y})/\partial\rho$, as one finds from integration by parts. Thus,

$$Q(v_{0x},v_{0Y}) = \int_0^{\rho} d\rho_1 P_2(v_{0x},v_{0Y}) + \frac{1}{4}\,\mathrm{erfc}\left(\frac{v_{0X}}{\sqrt{2}}\right)\mathrm{erfc}\left(\frac{v_{0Y}}{\sqrt{2}}\right). \tag{27}$$

The integral Υ_{XY} is the sum of one such integral and one such constant for each step in the characteristic curve. This one-dimensional form is useful for numerical evaluation and expansions. Kashlinsky, Hernández-Monteagudo, & Atrio-Barandela (2001) also present an interesting expansion of Υ_{XY} in Hermite polynomials, in v_0.

3.1.3. Noise: Exact Results

This section presents an exact expression for the variance of the correlation of a quantized signal when averaged over the ensemble of all statistically identical measurements. Real and imaginary parts of $\hat{r}_M$ have different variances. This requires calculation of both $\langle \hat{r}_M \hat{r}_M^*\rangle$ and $\langle \hat{r}_M \hat{r}_M\rangle$. Note that

$$\begin{aligned}\langle \hat{r}_M\, \hat{r}_M^*\rangle &= \frac{1}{4N_q^2}\left\langle \sum_i \sum_j \hat{x}_i\, \hat{y}_i^* \hat{x}_j^*\, \hat{y}_j\right\rangle \\ &= \frac{1}{4N_q^2}\sum_i \langle \hat{x}_i\, \hat{x}_i^* \hat{y}_i\, \hat{y}_i^*\rangle + \frac{1}{4N_q^2}\sum_{i\neq j}\langle \hat{x}_i\, \hat{y}_i^*\rangle\langle \hat{x}_j^*\, \hat{y}_j\rangle \\ &= \frac{1}{4N_q}\langle \hat{x}_i\, \hat{x}_i^* \hat{y}_i\, \hat{y}_i^*\rangle + \frac{N_q-1}{N_q}\langle \hat{r}_M\rangle^2, \end{aligned} \tag{28}$$

where the sum over $i \neq j$ is simplified by the fact that samples at different times are uncorrelated, and by equation (22). I expand the first average in the last line:

$$\begin{aligned}\langle \hat{x}_i\, \hat{x}_i^* \hat{y}_i\, \hat{y}_i^*\rangle &= \langle \mathrm{Re}(\hat{x}_i)^2 \mathrm{Re}(\hat{y}_i)^2 + \mathrm{Im}(\hat{x}_i)^2 \mathrm{Im}(\hat{y}_i)^2 \\ &\quad + \mathrm{Re}(\hat{x}_i)^2 \mathrm{Im}(\hat{y}_i)^2 + \mathrm{Im}(\hat{x}_i)^2 \mathrm{Re}(\hat{y}_i)^2\rangle \\ &= \langle \mathrm{Re}(\hat{x}_i)^2 \mathrm{Re}(\hat{y}_i)^2\rangle + \langle \mathrm{Im}(\hat{x}_i)^2 \mathrm{Im}(\hat{y}_i)^2\rangle \\ &\quad + \langle \mathrm{Re}(\hat{x}_i)^2\rangle\langle \mathrm{Im}(\hat{y}_i)^2\rangle + \langle \mathrm{Im}(\hat{x}_i)^2\rangle\langle \mathrm{Re}(\hat{y}_i)^2\rangle, \end{aligned} \tag{29}$$

where I have used the fact that the real part of $\hat{x}$ has zero covariance with the imaginary part of $\hat{y}$, and vice versa. Because the real and imaginary parts are identical, this sum can be expressed formally in terms of the integrals:

$$\begin{aligned}\langle \mathrm{Re}(\hat{x}_i)^2 \mathrm{Re}(\hat{y}_i)^2\rangle &= \langle \mathrm{Im}(\hat{x}_i)^2 \mathrm{Im}(\hat{y}_i)^2\rangle \\ &= \Upsilon_{X2Y2} \equiv \int dX\; dY\; P(X,Y)\; \hat{X}(X)^2\; \hat{Y}(Y)^2, \\ \langle \mathrm{Re}(\hat{x}_i)^2\rangle = \langle \mathrm{Im}(\hat{x}_i)^2\rangle &= A_{X2} \equiv \int dX \frac{1}{\sqrt{2\pi}} e^{-X^{2/2}} \hat{X}(X)^2, \\ \langle \mathrm{Re}(\hat{y}_i)^2\rangle = \langle \mathrm{Im}(\hat{y}_i)^2\rangle &= A_{Y2} \equiv \int dY \frac{1}{\sqrt{2\pi}} e^{-Y^{2/2}} \hat{Y}(Y)^2. \end{aligned} \tag{30}$$

These expressions define Υ_{X2Y2}, A_{X2}, and A_{Y2}. Thus,

$$\langle \hat{r}_M \hat{r}_M^*\rangle = \frac{1}{N_q}\left[\frac{1}{2}\Upsilon_{X2Y2} + \frac{1}{2}A_{X2}A_{Y2} - (\Upsilon_{XY})^2\right] + (\Upsilon_{XY})^2. \tag{31}$$

Note that in these expressions A_{X2} and A_{Y2} are constants that depend on the characteristic curve, but not on ρ, whereas Υ_{XY} and Υ_{X2Y2} depend on ρ in complicated ways as well as on the characteristic curve.

Similarly,

$$\langle \hat{r}_M \hat{r}_M\rangle = \frac{1}{4N_q}\langle \hat{x}_i \hat{y}_i^* \hat{x}_i \hat{y}_i^*\rangle + \frac{N_q-1}{N_q}\langle \hat{r}_M\rangle^2. \tag{32}$$

I again expand the first sum in the last line:

$$\langle \hat{x}_i \hat{y}_i^* \hat{x}_i \hat{y}_i^*\rangle = 2\Upsilon_{X2Y2} - 2A_{X2}A_{Y2} + 4\Upsilon_{XY}^2, \tag{33}$$

where I omit the imaginary terms, all of which average to zero. Therefore,

$$\langle \hat{r}_M \hat{r}_M\rangle = \frac{1}{N_q}\left(\frac{1}{2}\Upsilon_{X2Y2} - \frac{1}{2}A_{X2}A_{Y2}\right) + (\Upsilon_{XY})^2. \tag{34}$$

Using the same logic as in the derivation of equation (18), equations (30) and (33) can be used to find the means and

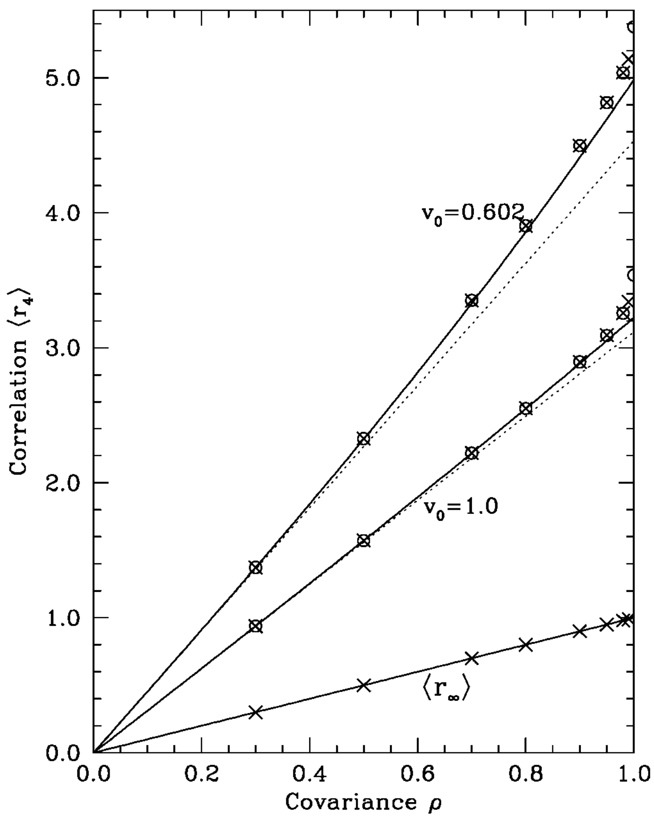

FIG. 2.—Average correlation plotted with covariance ρ. Curves show $\langle r_\infty \rangle$ and approximations to $\langle \hat{r}_4 \rangle$ for $v_0 = 1.0$ and 0.602, with $n = 3$. Heavy lines show the third-order approximation of eq. (39), and light lines show the linear approximation of Cooper (1970). Circles show true values as computed by direct integration of eq. (25). Crosses show results of simulations (§ 4). The sharp rise in $\langle \hat{r}_4 \rangle$ for $\rho \approx 1$, shown by the circles at far right, motivates the approximation of Jenet & Anderson (1998) that $\langle \hat{r}_4 \rangle$ varies proportionately with ρ for $\rho < 1$, with a spike at $\rho = 1$.

standard deviations of the real and imaginary parts of $\hat{r}_M$:

$$\langle \mathrm{Re}(\hat{r}_M) \rangle = \Upsilon_{XY},$$

$$\langle \mathrm{Im}(\hat{r}_M) \rangle = 0,$$

$$\langle \mathrm{Re}(\hat{r}_M)^2 \rangle - \langle \mathrm{Re}(\hat{r}_M) \rangle^2 = \frac{1}{2N_q}[\Upsilon_{X2Y2} - (\Upsilon_{XY})^2],$$

$$\langle \mathrm{Im}(\hat{r}_M)^2 \rangle = \frac{1}{2N_q}[A_{X2}A_{Y2} - (\Upsilon_{XY})^2]. \quad (35)$$

Again, note that A_{X2} and A_{Y2} are constants that depend on the characteristic curve, but not on the covariance ρ, whereas Υ_{XY} and Υ_{X2Y2} depend on the actual value of ρ as well as the characteristic curve. For particular characteristic curves and particular values of ρ, these expressions nevertheless yield the mean correlation and the standard deviations of real and imaginary parts about the mean. Figures 2–5 show examples and compare them with the approximate results from the following section.

Note that, because Υ_{XY} is an increasing function of ρ, the standard deviation of the imaginary part decreases with increasing covariance ρ. This holds for arbitrary quantizer parameters so long as the characteristic curves are antisymmetric. In other words, the noise in the imaginary part always decreases when the correlation increases.

For uncorrelated signals, $\rho = 0$. One finds then that $\Upsilon_{XY} = 0$ and $\Upsilon_{X2Y2} = A_{X2}A_{Y2}$. In this case both real and imaginary parts have identical variances, as they must:

$$\langle \mathrm{Re}(\hat{r}_M)^2 \rangle - \langle \mathrm{Re}(\hat{r}_M) \rangle^2 = \langle \mathrm{Im}(\hat{r}_M)^2 \rangle = \frac{1}{2N_q} A_{X2}A_{Y2} \quad \text{for } \rho = 0. \quad (36)$$

This recovers the result of Cooper (1970) and others for the noise in this limit.

If the characteristic curves are identical so that $\hat{X}(X) = \hat{Y}(Y)$, then if the signals are identical ($\rho = 1$), one finds that $\Upsilon_{X2Y2} = A_{X2} = A_{Y2}$. Under these assumptions then $\langle \mathrm{Im}(\hat{r}_M)^2 \rangle = 0$.

3.2. Correlation of Quantized Signals: Approximate Results

3.2.1. Mean Correlation: Approximate Result

Unfortunately, the expressions for the mean correlation and the noise, for quantized signals, both depend, in a complicated way, on the covariance ρ, the quantity one seeks to measure. Often the covariance ρ is small. Various authors discuss the correlation $\hat{r}_M$ of quantized signals $\hat{x}$ and $\hat{y}$ to first order in ρ, as is appropriate in the limit $\rho \to 0$ (Van Vleck & Middleton 1966; Cooper 1970; Hagen & Farley 1973; Thompson et al. 1986; Jenet & Anderson 1998). Jenet & Anderson (1998) also calculate $\langle r_4 \rangle$ for $\rho = 1$. As they point out, this case is important for autocorrelation spectroscopy. D'Addario et al. (1984) present an expression for $\langle \hat{r} \rangle$ for a three-level correlator, and present several useful approximate expressions for the inverse relationship $\rho(\langle \hat{r}_M \rangle)$. Here I find the mean correlation $\langle \hat{r}_M \rangle$ through fourth order in ρ.

For small covariance ρ, one can expand $P(X,Y)$ in equation

(7) as a power series in ρ:

$$P(X,Y) = \frac{1}{2\pi\sqrt{1-\rho^2}} \times \exp\left[\frac{-1}{2(1-\rho^2)}(X^2+Y^2-2\rho XY)\right]$$
$$\approx \frac{1}{2\pi} e^{-X^{2/2}} e^{-Y^{2/2}} \times \left[1 + XY\rho + \frac{1}{2}(1-X^2)(1-Y^2)\rho^2 + \frac{1}{6}(3X-X^3)(3Y-Y^3)\rho^3 + \frac{1}{24}(3-6X^2+X^4)(3-6Y^2+Y^4)\rho^4 \ldots\right]. \quad (37)$$

Note that the coefficient of each term in this expansion over ρ can be separated into two factors that depend on either X alone or Y alone. The extension to higher powers of ρ is straightforward, and the higher-order coefficients have this property as well.

As noted above, I assume that the characteristic curve is antisymmetric: $\hat{X}(-X) = -\hat{X}(X)$. The integral Υ_{XY} (eq. [25]) involves first powers of the functions $\hat{X}(X)$ and $\hat{Y}(Y)$. In this integral, only terms odd in both X and Y match the antisymmetry of the characteristic curve and yield a nonzero result. Such terms are also odd in ρ, as is seen from inspection of equation (36). The first-order terms thus involve the integrals

$$B_X \equiv \int dX \frac{1}{\sqrt{2\pi}} e^{-X^{2/2}} \hat{X}(X) X,$$
$$B_Y \equiv \int dY \frac{1}{\sqrt{2\pi}} e^{-Y^{2/2}} \hat{Y}(Y) Y, \quad (38)$$

where again X and Y can stand for either real or imaginary parts of x and y. Here I consider terms up to order 3 in ρ. One thus encounters the further integrals

$$D_X \equiv \int dX \frac{1}{\sqrt{2\pi}} e^{-X^{2/2}} \hat{X}(X) X^3, \quad (39)$$

and the analogous expression for D_Y. Therefore, through fourth order in ρ,

$$\langle \hat{r}_M \rangle = \Upsilon_{XY} \approx \int dX\, dY\, \frac{1}{2\pi} e^{-X^{2/2}} e^{-Y^{2/2}} \hat{X}(X)\hat{Y}(Y) \times \left[(XY)\rho + \tfrac{1}{6}(3X-X^3)(3Y-Y^3)\rho^3 + \ldots\right]$$
$$\approx B_X B_Y\, \rho + \tfrac{1}{6}(3B_X - D_X)(3B_Y - D_Y)\, \rho^3 + \ldots . \quad (40)$$

For thresholds $v_0 \approx 1$, the linear approximation is quite accurate (Cooper 1970; Thompson et al. 1986; Jenet & Anderson 1998). However, for other values of v_0 the higher order terms can become important. My notation differs from that of previous authors; my $B_X B_Y$ is equal to $[(n-1)E+1]$ of Cooper (1970) and Thompson et al. (1986). It is equal to the $A\,(\hat{\sigma}^2/\sigma^2)$ of Jenet & Anderson (1998).

Figure 2 shows typical results of the expansion of equation (39) for a four-level correlator, and compares this estimate for $\hat{r}_4(\rho)$ with the results from direct integration of equation (25) over rectangles in the x-y plane. In this example, $v_{0x} = v_{0y} \equiv v_0$. For both $v_0 = 1$ and $v_0 = 0.602$, $\langle \hat{r}_M \rangle$ is relatively flat, with a sharp upturn very close to $\rho = 1$. However, in both cases, but especially for $v_0 = 0.602$, the curve of $\hat{r}_4$ bends upward well before $\rho = 1$ so that the linear approximation is good only for relatively small ρ.

3.2.2. Noise: Approximate Results

Expressions for the noise in the integral involve the integral Υ_{X2Y2} (eq. [34]). This integral involves only the squares of the characteristic curves $\hat{X}(X)^2$ and $\hat{Y}(Y)^2$. Because the characteristic curves are antisymmetric about 0, their squares are symmetric: $\hat{X}(X)^2 = \hat{X}(-X)^2$ and $\hat{Y}(Y)^2 = \hat{Y}(-Y)^2$. Therefore, the only contributions come from terms in the expansion of $P(x,y)$ (eq. [36]) that are even in X and Y. One thus encounters the integrals

$$A_{X2} \equiv \int dX\, \frac{1}{\sqrt{2\pi}} e^{-X^{2/2}}\, \hat{X}(X)^2,$$
$$C_{X2} \equiv \int dX\, \frac{1}{\sqrt{2\pi}} e^{-X^{2/2}}\, \hat{X}(X)^2 X^2,$$
$$E_{X2} \equiv \int dX\, \frac{1}{\sqrt{2\pi}} e^{-X^{2/2}}\, \hat{X}(X)^2 X^4, \quad (41)$$

and analogously for A_{Y2}, C_{Y2}, and E_{Y2}. Then, through fourth order in ρ,

$$\begin{aligned}\Upsilon_{X2Y2} &= \int dX\ dY\ P(X,Y)\ \hat{X}(X)^2\ \hat{Y}(Y)^2 \\ &\approx \int dX\ dY\ \frac{1}{2\pi} e^{-X^2/2} e^{-Y^2/2} \hat{X}(X)^2\ \hat{Y}(Y)^2 \\ &\quad \times \left[1 + \frac{1}{2}(1-X^2)(1-Y^2)\rho^2 \right. \\ &\quad \left. + \frac{1}{24}(3-6X^2+X^4)(3-6Y^2+Y^4)\rho^4 \ldots\right] \\ &\approx A_{X2}A_{Y2} + \frac{1}{2}(A_{X2}-C_{X2})(A_{Y2}-C_{Y2})\rho^2 \\ &\quad + \frac{1}{24}(3A_{X2}-6C_{X2}+E_{X2}) \\ &\quad \times (3A_{Y2}-6C_{Y2}+E_{Y2})\rho^4 \ldots . \end{aligned} \tag{42}$$

I find the standard deviations of real and imaginary parts from equations (34), (39), and (41):

$$\begin{aligned}\langle \mathrm{Re}(\hat{r}_M)^2\rangle - \langle \mathrm{Re}(\hat{r}_M)\rangle^2 &\approx \frac{1}{2N_q}\left\{(A_{X2}A_{Y2}) \right. \\ &\quad + \left[\frac{1}{2}(A_{X2}-C_{X2})(A_{Y2}-C_{Y2}) - B_X^2B_Y^2\right]\rho^2 \\ &\quad + \left[\frac{1}{24}(3A_{X2}-6C_{X2}+E_{X2}) \right. \\ &\quad \times (3A_{Y2}-6C_{Y2}+E_{Y2}) \\ &\quad \left.\left. - \frac{1}{3}B_X(3B_X-D_X)B_Y(3B_Y-D_Y)\right]\rho^4 \ldots\right\}, \\ \langle \mathrm{Im}(\hat{r}_M)^2\rangle &\approx \frac{1}{2N_q}\left\{(A_{X2}A_{Y2}) - (B_X^2B_Y^2)\,\rho^2 \right. \\ &\quad \left. - \left[\frac{1}{3}B_X(3B_X-D_X)B_Y(3B_Y-D_Y)\right]\rho^4 \ldots\right\}. \end{aligned} \tag{43}$$

I have used the fact that $\langle \hat{r}_M\rangle$ is purely real; this is a consequence of the assumption that ρ is purely real.

Figure 3 shows examples of the standard deviations of $\mathrm{Re}(\hat{r}_4)$ and $\mathrm{Im}(\hat{r}_4)$ for two choices of v_0. These are the noise in estimates of the correlation. Note that the noise varies with ρ. The quadratic variation of these quantities with ρ is readily apparent. The higher order variation is more subtle, although it does lead to an upturn of the standard deviation of $\mathrm{Re}(\hat{r}_4)$ near $\rho \approx 0.7$ for $v_0 = 1$. The series expansions become inaccurate near $\rho = 1$, as expected. The standard deviation of $\mathrm{Re}(\hat{r}_4)$ can also increase, instead of decrease, for large ρ. Such an increase is more common for parameter choices with $v_0 > 1$. Again, note that the standard deviation of the imaginary part always decreases with increasing ρ.

3.2.3. S/N for Quantized Correlation

The S/N for a quantizing correlator is the quotient of the mean and standard deviation of $\hat{r}_M$, the results of §§ 3.2.1 and 3.2.2. I recover the results of Cooper (1970) for the S/N for a quantizing correlator by using our approximate expressions through first order in ρ (see also Hagen & Farley 1973; Thompson et al. 1986; Jenet & Anderson 1998):

$$\begin{aligned}\mathcal{R}_M(\mathrm{Re}(\hat{r}_M)) &\approx \frac{\langle \mathrm{Re}(\hat{r}_M)\rangle}{\sqrt{\langle \mathrm{Re}(\hat{r}_M)^2\rangle - \langle \mathrm{Re}(\hat{r}_M)\rangle^2}} \\ &\approx \sqrt{2N_q}\,\frac{B_XB_Y}{\sqrt{A_{X2}A_{Y2}}}\,\rho. \end{aligned} \tag{44}$$

For a four-level correlator, a S/N of $\mathcal{R}_4 = 0.88115\rho$ is attained for $n = 3, v_0 = 1$, in the limit $\rho \rightarrow 0$. Many four-level correlators use these values. The maximum value for $\mathcal{R}_4$ is actually obtained for $n = 3.3359$, $v_0 = 0.9815$, for which $\mathcal{R}_4 = 0.88252\rho$. This adjustment of quantization constants provides a very minor improvement in S/N.

For nonvanishing ρ, the optimum-level settings depend upon the covariance ρ. The S/N is

$$\mathcal{R}_M(\mathrm{Re}(\hat{r}_M)) = \sqrt{2N_q}\,\frac{\Upsilon_{XY}}{\sqrt{\Upsilon_{X2Y2} - (\Upsilon_{XY})^2}}. \tag{45}$$

This can be approximated using the expansions for Υ_{XY} and Υ_{X2Y2}; Figure 4 shows the results. Note that in the examples in the figure, the S/N for the quantized correlations actually curve above that for the unquantized correlation beyond $\rho \approx 0.5$. This indicates that correlation of quantized signals can actually yield higher S/N than would be obtained from correlating the same signals before quantization. This results from the decline in noise with increasing ρ, visible in Figure 3.

For a proper comparison of S/Ns, one must compare with the S/N obtained for nonquantized correlation, $\mathcal{R}_\infty(\mathrm{Re}(\hat{r}_M))$ (eq. [19]). One finds

$$\frac{\mathcal{R}_M(\mathrm{Re}(\hat{r}_M))}{\mathcal{R}_\infty(\mathrm{Re}(\hat{r}_M))} = \frac{\Upsilon_{XY}}{\sqrt{\Upsilon_{X2Y2} - (\Upsilon_{XY})^2}}\,\frac{\sqrt{1-\rho^2}}{\rho}. \tag{46}$$

Figure 5 shows this ratio for two choices of v_0. The ratio can exceed 1, again indicating that a quantized correlation provides a more accurate result than would correlation of an unquantized signal.

The S/N for a measurement of phase for a quantized correlation is the inverse of the standard deviation of the phase,

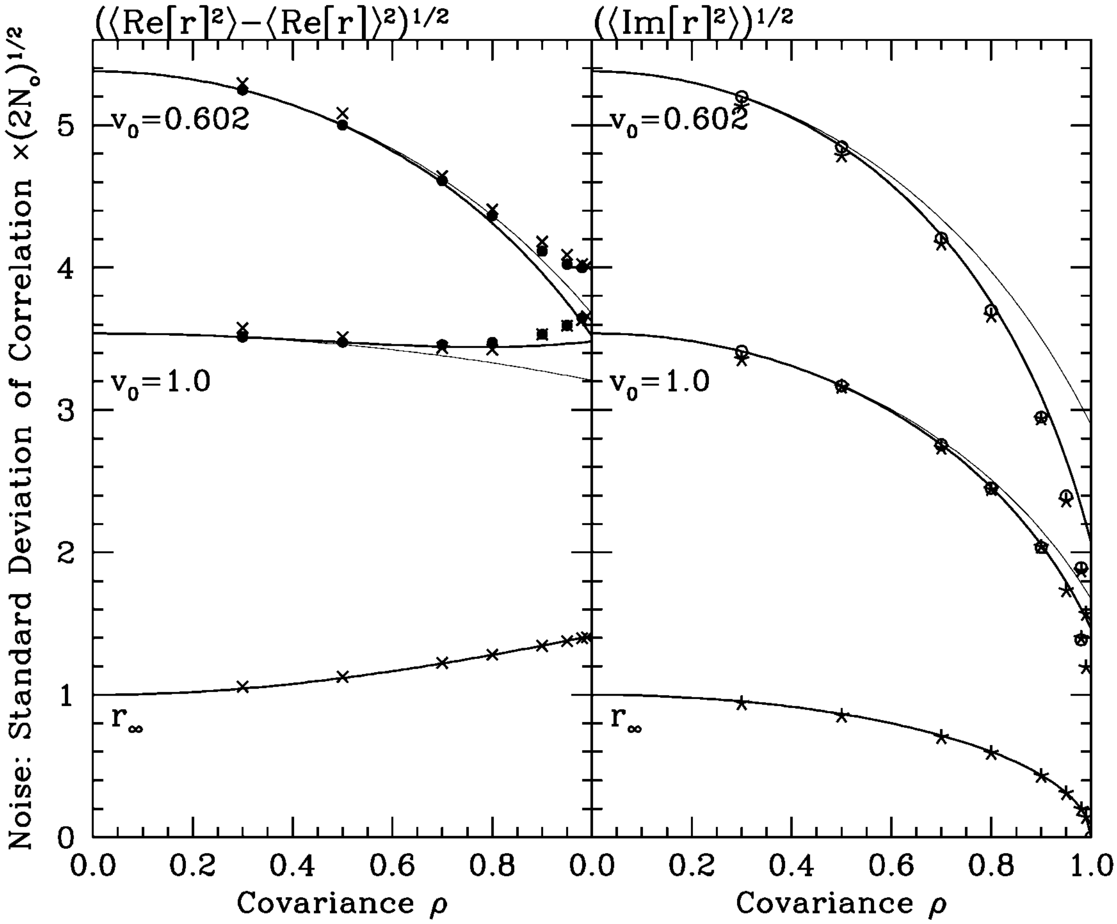

FIG. 3.—Standard deviation of the real part of correlation (*left*) and of the imaginary part (*right*), plotted with covariance ρ. Both are normalized for the number of samples by the factor $(2N_q)^{1/2}$. Curves show $\langle \hat{r}_\infty \rangle$ and approximations to $\langle \hat{r}_4 \rangle$ for quantization with $v_0 = 1.0$ and 0.602, with $n = 3$. Heavy lines show the fourth-order approximation of eq. (42), and light lines show the approximation to second order. The values found by Cooper (1970) are the y-intercepts ($\rho = 0$). Filled circles show true values as computed by direct integration of eq. (34). Crosses show results of simulations (§ 4).

as discussed in § 19. For a quantized correlation, this is

$$\mathcal{R}_M(\phi(r_M)) = \frac{\mathrm{Re}(r_M)}{\sqrt{\langle \mathrm{Im}(r_M)^2 \rangle}}$$

$$= \sqrt{2N_q}\,\frac{\Upsilon_{XY}}{\sqrt{A_X A_Y - (\Upsilon_{XY})^2}}. \quad (47)$$

Again, because Υ_{XY} always increases with ρ, the S/N of the phase always increases with increasing covariance.

The ratio of the S/N for phase to that for correlation of an unquantized signal, $\mathcal{R}_M(\phi(\hat{r}_M))/\mathcal{R}_\infty(\phi(r_\infty))$ (see eq. [20]), provides an interesting comparison. For $\rho \to 0$, the statistics for the imaginary part of the correlation are identical to those for the real part (as they must be), and the highest S/N for the phase is given by the quantizer parameters that are optimal for the real part, traditionally $v_0 = 1$, $n = 3$. This ratio is approximately constant with ρ up to $\rho \approx 0.6$, and then decreases rather rapidly. Simulations suggest that quantized correlation is less efficient than unquantized for measuring phase. However, I have not proved this in general.

4. SIMULATIONS

Simulation of a four-level correlator provides a useful perspective. I simulated such a correlator by generating two sequences of random, complex numbers, x_i and y_i. The real parts of x_i and y_i are drawn from one bivariate Gaussian distribution, and their imaginary parts from another independent one (see eq. [8]). These bivariate Gaussian distributions can be described equivalently as elliptical Gaussian distributions, with major and minor axes inclined to the coordinate axes $X = \mathrm{Re}(x_i)$ and $Y = \mathrm{Re}(y_i)$ and the corresponding axes for the imaginary parts.

Meyer (1975) gives expressions that relate the semimajor and semiminor axes and angle of inclination of an elliptical Gaussian distribution to the normalized covariance ρ and var-

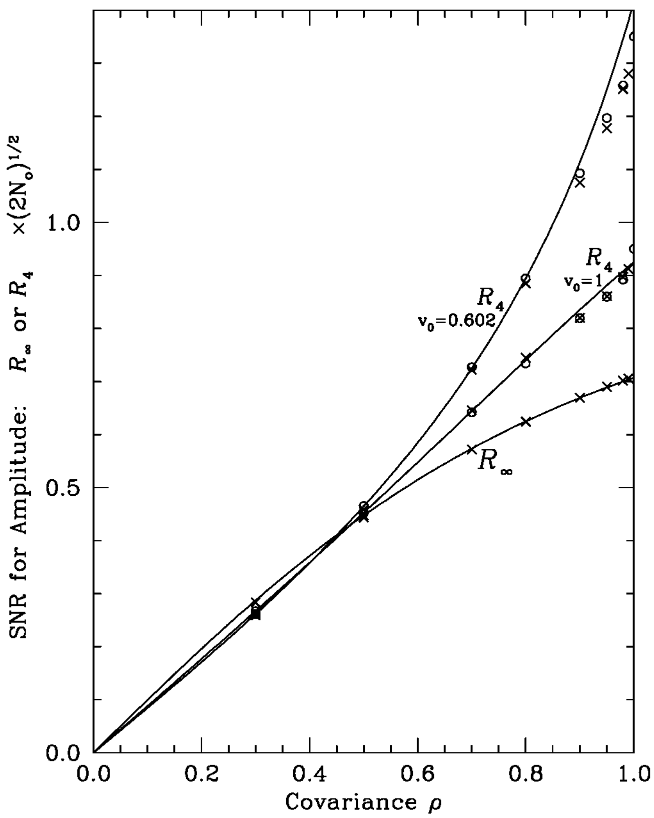

FIG. 4.—S/N $\mathcal{R}_\infty$ or $\mathcal{R}_4$, normalized for number of samples by $(2N_q)^{-1/2}$, and plotted with covariance ρ. Curves mark the exact expression given by eq. (19) for $\mathcal{R}_\infty$, or the approximate expression given by the expansions of eqs. (39) and (42) through fourth order in eq. (44). Filled circles give exact values as found from direct integration of eq. (34) in eq. (44). Crosses show results of computer simulations (§ 4). Quantization uses the characteristic curve in Fig. 1 with values of $v_0 = 1.0$ or 0.602, and $n = 3$.

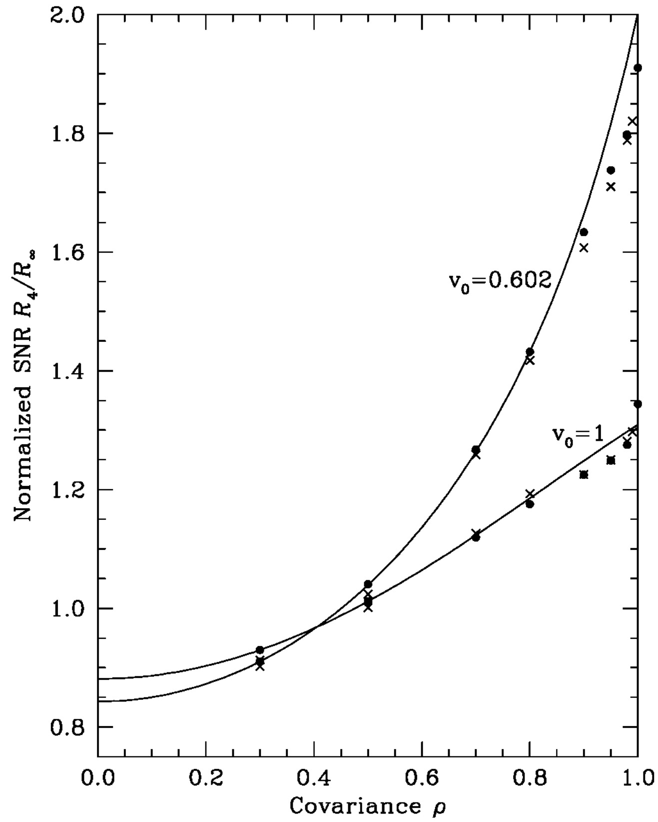

FIG. 5.—S/N for quantized signals, normalized to the S/N of the unquantized signal: $\mathcal{R}_4/\mathcal{R}_\infty$, plotted with covariance ρ. Curves mark the approximate expressions given by eqs. (39) and (42) for $\mathcal{R}_4$, normalized by $\mathcal{R}_\infty$ as given by eq. (18). Filled circles give exact expressions as found from direct integration of eq. (25). The y-intercept for $v_0 = 1.0$ is the standard S/N for a four-level correlator, $\mathcal{R}_4 = 0.88115$, with $v_0 = 1.0$ and $n = 3$, which is optimal for $\rho = 0$. Note that the ratio can be greater than 1. This indicates that quantized correlation can be more efficient than correlation of the original, unquantized signals.

iances σ_X^2 and σ_Y^2. For the special case of $\sigma_X = \sigma_Y = 1$ used in this work, the major axis always lies at angle $\pi/4$ to both coordinate axes, along the line $X = Y$. The semimajor axis b_1 and semiminor axis b_2 are then given by

$$b_1 = \sqrt{1+\rho},$$
$$b_2 = \sqrt{1-\rho}. \quad (48)$$

To form the required elliptical distributions, I drew pairs of elements from a circular Gaussian distribution, using the Box-Muller method (see Press et al. 1989). I scaled these random elements so that their standard deviations were b_1 and b_2. I then rotated the resulting two-element vector by $\pi/4$ to express the results in terms of X and Y. I repeated the procedure for the imaginary part.

I quantized the sequences x_i and y_i according to the four-level characteristic curve shown in Figure 1, to yield $\hat{x}_i$ and $\hat{y}_i$. Both the unquantized and the quantized sequences were correlated by forming the products $x_i y_i^*$ and $\hat{x}_i \hat{y}_i^*$, respectively, and results were averaged over $N_q = 10^5$ instances of the index i. This procedure yields one realization each of r and $\hat{r}_4$. I found that values for N_q that were smaller than about 10^5 could produce significant departures from Gaussian statistics for $\hat{r}_4$, particularly for larger values of ρ.

I repeated the process to obtain 4096 different realizations of r and $\hat{r}_4$. I found the averages and standard deviations for the real and imaginary parts for this set of realizations. Figures 2–5 show these statistical results of the simulations and compare them with the mathematical results of the preceding sections. Clearly, the agreement is good.

In graphical form, samples of the correlation form an elliptical Gaussian distribution in the complex plane, centered at the mean value of correlation $\langle r_\infty \rangle$ or $\langle \hat{r}_M \rangle$, as the case may be. The principal axes of the distribution lie along the real and imaginary directions (or, more generally, the directions in phase with ρ and out of phase with ρ). The lengths of these principal axes are the variances of real and imaginary parts.

5. DISCUSSION AND SUMMARY

5.1. Change of Noise with Covariance

The fundamental result of this paper is that a change in covariance ρ affects quantized correlation $\hat{r}_M$ differently from unquantized correlation r_∞. For unquantized correlation, an increase in covariance ρ increases noise for estimation of signal amplitude. For quantized correlation, an increase in ρ can increase or decrease amplitude noise. For both quantized and unquantized correlation, an increase in ρ leads to a decrease in phase noise. In this work I arbitrarily set the phase of ρ to 0 so that amplitude corresponds to the real part, and phase to the imaginary part of the estimated correlation. The net noise (summed, squared standard deviations of real and imaginary parts) can decrease (or increase) with increasing ρ; noise is not conserved.

I present expressions for the noise as a function of quantization parameters, both as exact expressions that depend on ρ and on power-series expansions in ρ. These expressions, and a power-series expansion for the mean correlation, are given through fourth order in ρ.

The increase in noise with covariance ρ for analog correlation, sometimes called source noise or self-noise, is sometimes ascribed to the contribution of the original, noiselike signal to the noise of correlation. This idea is difficult to generalize to comparisons among quantized signals, because such comparisons require additional assumptions about changes in quantizer levels and the magnitude of the quantized signal when the covariance changes. These comparisons are simpler for multiple correlations derived from a single signal (as, for example, for the correlation function of a spectrally varying signal), and I will discuss them in that context elsewhere (C. Gwinn 2004, in preparation). The discussion in this paper is limited to "white" signals, without spectral variation and with only a single independent covariance.

5.1.1. Increase in S/N via Quantization

One interesting consequence of these results is that the S/N of correlation can actually be greater for quantized signals than it would be for correlation of the same signals before quantization. At small covariance ρ, S/N is always lower for quantized signals, but this need not be the case for covariance $\rho \gtrsim 0.4$. This appears to present a paradox, because the process of quantization intrinsically destroys information: the quantized signals $\hat{x}_i$, $\hat{y}_i$ contain less information than did the original signals x_i, y_i. However, correlation of unquantized signals also destroys information: it converts x_i and y_i to the single quantity $c_i = (x_i y_i^*)$. Different information is destroyed in the two cases.

Moreover, correlation does not always yield the most accurate estimate of the covariance ρ. As a simple example, consider the series $\{X_i\} = \{0.400, -0.800, 1.600\}$ and $\{Y_i\} = \{0.401, -0.799, 1.600\}$. Here $N_q = 3$. One easily sees that X and Y are highly correlated. If X and Y are known to be drawn from Gaussian distributions with unit standard deviation, equation (47) suggests that $\rho \approx 0.999$. However, the correlation is $r = \frac{1}{3}\sum_i X_i Y_i = 1.29$. Clearly r is not an optimal measurement of ρ. I will discuss strategies for optimal estimates of covariance elsewhere (C. Gwinn 2004, in preparation).

5.1.2. Quantization Noise

Sometimes effects of quantization are described as "quantization noise"; an additional source of noise that, like "sky noise" or "receiver noise," reduces the correlation of the desired signal. However, unlike other sources of noise, quantization destroys information in the signals, rather than adding unwanted information. The discussion in the preceding section suggests that the amount of information that quantization destroys (or, more loosely, the noise that it adds) depends on what information is desired; and that correlation removes information as well. Unless the covariance ρ is small, effects of quantization cannot be represented as a one additional, independent source of noise in general.

5.1.3. Applications

The primary result of this paper is that for quantized correlation, noise can increase or decrease when covariance increases, whereas for continuous signals it increases. This fact is important for applications requiring accurate knowledge of the noise level; as for example in studies of rapidly varying strong sources such as pulsars, where one wishes to know whether a change in correlation represents a significant change in the pulsar's emission; or for single-dish or interferometric observations of intraday variable sources, for which one wishes to know whether features that may appear and disappear are statistically significant.

A second result of this paper is that the S/N for quantized correlation can be quite different from that expected for a continuum source or for a continuum source with added noise. This effect is most important for large correlation, $\rho \gtrsim 0.5$. Correlation this large is often observed for strong sources, such as the strongest pulsars, or maser lines, and for the strongest continuum sources observed with sensitive antennas. For example, at Arecibo Observatory a strong continuum source easily dominates the system temperature at many observing frequencies. The effect will be even more common for some proposed designs of the Square Kilometer Array.

Many sources show high correlation that varies dramatically with frequency. Such sources include scintillating pulsars and maser lines. Typically, observations of these sources involve determination of the full correlation function and a Fourier transform to obtain an estimated cross-power or autocorrelation spectrum. I discuss the properties of noise for this analysis elsewhere.

5.1.4. S/N Enhancement?

An interesting question is whether one can take advantage of the higher S/N afforded by quantization at high ρ even for weakly correlated signals, perhaps by adding an identical signal to each of two weakly covariant signals and so increasing their covariance ρ, before quantizing them. The answer appears to be no. As a simple example, consider a pair of signals with covariance of $\rho \approx 0.01$. After correlation of $N_q = 2 \times 10^6$ instances of the signal using a four-level correlator with $v_0 = 1$ and $n = 3$, the S/N is 20, and one can determine at a level of 2 standard deviations whether $\rho = 0.010$ or 0.011. If a single signal with 4.6 times greater amplitude is added to both of the original signals, then these two cases correspond to $\rho = 0.7004$ or 0.7005. To distinguish them at 2 standard deviations requires a S/N of 1400, requiring $N_q = 4 \times 10^8$ samples of the quantized correlation. Thus, the increase in S/N is more than outweighed by the reduction in the influence on the observable.

5.2. Summary

In this paper I consider the result of quantizing and correlating two complex noiselike signals, x and y, with normalized covariance ρ. The signals are assumed to be statistically stationary, "white," and sampled at the Nyquist rate. The correlation r provides a measurement of ρ. The variation of r about that mean, characterized by its standard deviation, provides a measure of the random part of the measurement, or noise.

I suppose that the signals x and y are quantized to form $\hat{x}$ and $\hat{y}$. I suppose that the characteristic curves that govern quantization are antisymmetric, with real and imaginary parts subject to the same characteristic curve. I recover the classic results for the noise for $\rho = 0$ and for the mean correlation to first order in ρ in the limit $\rho \rightarrow 0$ (Van Vleck & Middleton 1966; Cooper 1970; Hagen & Farley 1973; Thompson et al. 1986). I find exact expressions for the mean correlation and the noise, and approximations that are valid through fourth order in ρ. I compare results with simulations. Agreement is excellent for the exact forms, and good for ρ that is not too close to 1, for the approximate expressions.

I find that for nonzero values of ρ, the noise varies, initially quadratically, with ρ. I find that the noise in an estimate of the amplitude of ρ can decrease with increasing ρ. This is opposite the behavior of noise for correlation of unquantized signals, for which noise always increases with ρ. The mean correlation can increase more rapidly than linearly with ρ. The S/N for correlation of quantized signals can be greater than that for correlation of unquantized signals, for $\rho \gtrsim 0.5$. In other words, correlation of quantized signals can be more efficient than correlation of the same signals before quantization, as a way of determining the covariance ρ.

I am grateful to the Dominion Radio Astrophysical Observatory (DRAO) for supporting this work with extensive correlator time. I also gratefully acknowledge the VLBI Space Observatory Programme (VSOP) Project, which is led by the Japanese Institute of Space and Astronautical Science in cooperation with many organizations and radio telescopes around the world. I thank an anonymous referee for useful comments. The National Science Foundation provided financial support.

REFERENCES

Anantharamaiah, K. R., Deshpande, A. A., Radhakrishnan, V., Ekers, R. D., Cornwell, T. J., Goss, W. M. 1991, in IAU Colloq. 131, Radio Interferometry: Theory, Techniques, and Applications (San Francisco: ASP), 6
Bowers, F. K., & Klingler, R. J. 1974, A&AS, 15, 373
Cooper, B. F. C. 1970, Australian J. Phys., 23, 521
D'Addario, L. R., Thompson, A. R., Schwab, F. R., & Granlund, J. 1984, Radio Sci., 19, 931
Gwinn, C. R. 2001, ApJ, 554, 1197
Hagen, J. B., & Farley, D. T. 1973, Radio Sci, 8, 775
Jenet, F. A., & Anderson, S. B. 1998, PASP, 110, 1467
Jenet, F. A., Anderson, S. B., & Prince, T. A. 2001, ApJ, 558, 302
Kashlinsky, A., Hernández-Monteagudo, C., & Atrio-Barandela, F. 2001, ApJ, 557, L1
Kokkeler, A. B. J., Fridman, P., & van Ardenne, A. 2001, Experimental Astronomy, 11, 33
Kulkarni, S. R. 1989, AJ, 98, 1112
Meyer, S. L. 1975, Data Analysis for Scientists and Engineers (New York: Wiley)
Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. 1989, Numerical Recipes (Cambridge: Cambridge Univ. Press)
Thompson, A. R., Moran, J. M., & Swenson, G. W. Jr. 1986, Interferometry and Synthesis in Radio Astronomy (New York: Wiley)
Van Vleck, J. H., & Middleton, D. 1966, Proc. IEEE, 54, 2