PUBLICATIONS OF THE ASTRONOMICAL SOCIETY OF THE PACIFIC, **119**: 805–827, 2007 July

Radio Frequency Interference Excision Using Spectral-Domain Statistics

GELU M. NITA, DALE E. GARY, AND ZHIWEI LIU

Center for Solar-Terrestrial Research, New Jersey Institute of Technology, Newark, NJ; gnita@njit.edu

GORDON J. HURFORD

Space Science Laboratory, University of California, Berkeley, CA

AND

STEPHEN M. WHITE

Department of Astronomy, University of Maryland, College Park, MD

Received 2007 March 26; accepted 2007 June 18; published 2007 July 24

ABSTRACT. A radio frequency interference (RFI) excision algorithm based on spectral kurtosis, a spectral variant of time-domain kurtosis, is proposed and implemented in software. The algorithm works by providing a robust estimator for Gaussian noise that, when violated, indicates the presence of non-Gaussian RFI. A theoretical formalism is used that unifies the well-known time-domain kurtosis estimator with past work related to spectral kurtosis, and leads naturally to a single expression encompassing both. The algorithm accumulates the first two powers of M power spectral density (PSD) estimates, obtained via Fourier transform, to form a spectral kurtosis (SK) estimator whose expected statistical variance is used to define an RFI detection threshold. The performance of the algorithm is theoretically evaluated for different time-domain RFI characteristics and signal-to-noise ratios η. The theoretical performance of the algorithm for intermittent RFI (RFI present in R out of M PSD estimates) is evaluated and shown to depend greatly on the duty cycle, $d = R/M$. The algorithm is most effective for $d = 1/(4 + \eta)$, but cannot distinguish RFI from Gaussian noise at any η when $d = 0.5$. The expected efficiency and robustness of the algorithm are tested using data from the newly designed FASR Subsystem Testbed radio interferometer operating at the Owens Valley Solar Array. The ability of the algorithm to discriminate RFI against the temporally and spectrally complex radio emission produced during solar radio bursts is demonstrated.

1. INTRODUCTION

The quality of radio astronomy scientific data can be greatly affected by radio frequency interference (RFI) contaminating the natural radio emission produced by astrophysical objects. With the increasing demand for wireless communication systems operating in frequency bands of scientific interest, the RFI environment becomes ever more hostile. At the same time, astronomers seek to observe over an ever broader part of the radio spectrum, both for increased sensitivity and to exploit spectral continuum diagnostics. These trends make it necessary to find efficient and robust methods that are able to discriminate and excise the RFI contamination, while preserving as much of the underlying useful information as possible.

There is no universal solution for the problem of RFI mitigation, and one must find or develop methods that are most suitable for a particular instrument and the RFI environment in which it operates (Fridman & Baan 2001). Among these methods, RFI mitigation algorithms based on higher moments of the distribution function in both the time domain (Ruf et al. 2006) and the frequency domain (Fridman 2001; Fridman & Baan 2001) have become increasingly popular in recent years, due to the feasibility of their implementation using field-programmable gate arrays (FPGAs), equipped on many modern digital instruments.

Although relatively new in the context of RFI mitigation for astrophysical radio instruments, the idea of using higher order statistics in the spectral domain was originally proposed more than two decades ago by Dwyer (1983) in the context of underwater acoustic signals. Pointing out that the real and imaginary parts of the digital Fourier transform (DFT) for a Gaussian signal separately obey Gaussian distributions, he proposed using the kurtosis of the real part of the DFT as a tool for non–Gaussian signal detection. He named this estimator "frequency-domain kurtosis" to distinguish it from the traditional time-domain kurtosis estimator. However, as Dwyer (1983) himself pointed out, a similar estimator can be defined for the imaginary part of the DFT, which should be considered for obtaining a complete description of the signal. Since the work of Dwyer (1983), others have investigated the statistical properties of power spectral density (PSD) estimates by treating both the real and imaginary parts of the DFT, and a theory of circular random complex variables has even emerged from such efforts (see Amblard et al. 1996a, 1996b, and references therein). In

a more recent work, Vrabie et al. (2003) provided a comprehensive study, from both the theoretical and practical points of view, of the newly emerged concept that they call spectral kurtosis (SK), and provided a cumulant-based SK estimator (Kenney & Keeping 1962) and its expected variance, as well as the theoretical expectations for the SK parameter corresponding to some particular types of stationary time-domain signals. In a subsequent study, Vrabie et al. (2004) revisited the subject and presented a practical application of the SK estimator in detecting bearing faults in induction motors. More recently, Antoni (2006) refined the theoretical definition of the SK parameter, provided a direct link of this parameter with the classical time-domain kurtosis parameter, and pointed out its value as an estimator of non-stationary time-domain signals. Subsequently, Antoni (2007) proposed an algorithm for SK estimation based on a multirate filter-bank structure, and presented some experimental results related to the detection of transient faults in mechanical systems.

Despite its general applicability, the extent to which the frequency-domain statistical analysis has been implemented as a RFI detection tool in astrophysical radio instruments is rather modest, possibly due to the lack of a clear and unifying picture of the subject. In our opinion, such a picture should be theoretically rigorous, but free of unnecessary mathematical sophistication, and focused on the practical aspects of its implementation. This study attempts to provide a general framework that should be useful across disciplines, before applying it to the specific problem of RFI mitigation.

Another reason that there has not been more rapid progress in radio astrophysics may be the lack of existing instruments that are technically able to perform broadband spectral analysis with the temporal and frequency resolution needed for efficient implementation and testing of such RFI excision algorithms. Such instruments are useful in setting design requirements for the new generation of radio instruments, such as the Frequency-Agile Solar Radiotelescope (FASR; Bastian 2003). From this perspective, the newly designed FASR Subsystem Testbed (FST; Liu et al. 2007), operating in parallel with three of the radio antennas of the Owens Valley Solar Array (OVSA; Gary & Hurford 1999), offers a unique test bed capability. Its design allows the direct comparison of various RFI excision algorithms by playing back in a simulated real-time mode the data recorded with full 1 ns temporal resolution, which provides a 500 MHz instantaneous bandwidth that can be tuned anywhere in the 1–9 GHz frequency band.

This paper brings together what is known about spectral-domain statistical methods from the broader literature, places it in context with the somewhat better known time-domain statistical methods, and develops a unitary framework of practical use in RFI excision. In § 2, the statistical properties of the PSD estimate are derived in a transparent and rigorous manner, a statistical RFI discriminator is derived, and its equivalence with both the preexisting spectral and classical (time domain) kurtosis statistical parameters is demonstrated. In § 3, a practical estimator is defined, and its statistical fluctuations, important for establishing a reliable detection threshold, are rigorously determined. The expected performance of the proposed estimator as a detection tool for RFI signals is also presented. In § 4, we propose an efficient real-time RFI excision algorithm that is suitable for radio spectrographs or interferometers. A software implementation of the algorithm is used to demonstrate its efficiency and robustness using RFI-contaminated data recorded by the FST instrument while observing the solar radio emission in both quiet-Sun and solar radio burst conditions. We summarize our conclusions in § 5.

2. STATISTICAL PROPERTIES OF POWER SPECTRA ESTIMATES

One of the commonly used methods for obtaining PSD estimates for a real-time signal sampled at a constant rate $1/\Delta t$ is the discreet Fourier transform, which is almost invariably implemented using the fast Fourier transform (FFT) algorithm.

A legitimate question that always arises with respect to this method is "How good is such a PSD estimate?" More-or-less mathematically involved discussions on this subject are part of almost any textbook devoted to the frequency-domain analysis of digital signals. The common answer to the above question, which may sound striking when heard for the first time, is that "the uncertainty of a DFT-based PSD estimate at any given frequency bin is as high as 100% of its computed value" (Press et al. 1989). However, this is strictly correct only for a particular type of time-domain signal, stationary Gaussian noise. The uncertainty of the PSD estimate may have a different behavior for time-domain signals governed by other probability distributions. Although the uncertainty is significantly reduced in practice by means of averaging, either in time or frequency, it is its magnitude before averaging that proves to be a very useful property in the context of this study, and therefore worthy of more detailed analysis in this section.

If the nonaveraged PSD estimate[1] $\hat{P}_k$ corresponding to the frequency bin f_k is characterized by its mean $\mu_k = \langle \hat{P}_k \rangle$ and variance $\sigma_k^2 = \langle \hat{P}_k^2 \rangle - \langle \hat{P}_k \rangle^2$, we can quantify its relative statistical fluctuation by defining the dimensionless magnitude

$$V_k^2 = \frac{\sigma_k^2}{\mu_k^2}, \tag{1}$$

where V_k is the coefficient of dispersion (Kenney & Keeping 1962). We further refer to V_k^2, at least for the moment, by the more descriptive term "spectral variability."

The expectation value of the spectral variability depends on the underlying probability distribution of the time-domain sig-

[1] Throughout this paper, we use a "hat" notation to indicate an estimator of a quantity to distinguish it from the quantity itself.

nal from which the PSD estimate is derived. Since our goal is to distinguish departures of the signal from natural noise, we first examine the expected spectral variability in the particular case of a time-domain signal governed by a Gaussian distribution.

2.1. Spectral Variability of a Gaussian Time-Domain Signal

We assume that the digitized time-domain signal $\{x(t_n)\} \equiv \{x_n\}$ is characterized by the following properties:

1. The time-domain data samples are random variables drawn from the same Gaussian parent population characterized by a zero mean, $\langle x_n \rangle = 0$, and a constant variance $\sigma_x^2 = \langle x_n^2 \rangle$.
2. The time-domain data samples are uncorrelated random variables; i.e., they have null covariance, $\mathrm{Cov}\,(x_n, x_m) = 0$.
3. The signal has a natural or artificial bandwidth limitation exactly bounded by the frequency interval $[0, f_c]$, where $f_c = 1/2\Delta t$ is the Nyquist frequency associated with the sampling rate Δt.
4. The test signal does not carry any spectral information; i.e., it is a band-limited Gaussian white noise.

To obtain a PSD estimate of a time-domain signal, one has to compute first the complex DFT coefficients given by

$$X_k = \sum_{n=0}^{N-1} w_n x_n e^{-2\pi i k n/N}, \quad k = 0, \ \ldots, N-1, \tag{2}$$

where $\{w_n\}$ is a set of N real coefficients of an arbitrary time window such as usually employed in order to reduce the spectral leakage among the PSD frequency bins (Press et al. 1989). Although the use of a time-domain window is a common practice, its potential influence on the statistical properties of the PSD estimates is ignored in most theoretical studies addressing this subject, yet we show below that the concept of an arbitrary window plays a central role in our generalized results.

The DFT coefficients obtained from a windowed stationary time-domain Gaussian signal $\{x_n\}$ can be regarded as corresponding to a nonwindowed DFT of a modified time sequence $\{w_n x_n\}$, which, being a linear combination of Gaussian random variables, should also obey a Gaussian distribution whose time-domain variance is (Kenney & Keeping 1962)

$$\sigma_{wx}^2 = \frac{1}{N} \sum_{i=0}^{N-1} w_n^2 \sigma_x^2. \tag{3}$$

The fact that a time-domain window cannot change the Gaussian nature of the original signal seems a good reason for not investigating its influence in more detail. However, as we show presently, there is a subtle connection between the statistical properties of the PSD estimate at each frequency bin and the spectral properties of the time-domain window, which results in a frequency-dependent correction of the PSD statistical properties that has been overlooked in previous studies. Although this correction will eventually prove negligible for the symmetrical time-domain windows (e.g., Hanning and Hamming) usually employed in standard spectral analysis, an evolving Gaussian background, such as that presented by a solar radio burst, can be mathematically modeled as the product of a stationary Gaussian background and a time-domain window of arbitrary shape. The effect of such an asymmetric window is generally not negligible.

The actual expressions for the spectral powers corresponding to each of the $N/2 + 1$ discrete frequency bins,

$$f_k = \frac{2k}{N} f_c, \quad k = 0, 1, \ \ldots, \frac{N}{2}, \tag{4}$$

of the PSD estimate are subject to different normalization conventions. The one adopted here is

$$\hat{P}_k = \frac{2}{N \sum w_n^2} |X_k|^2, \tag{5}$$

where X_0 and $X_{N/2}$ are real magnitudes, and we have taken into account that for a real-time signal, the other DFT coefficients and their complex conjugates are related as $X_k = X_{N-k}^*$. Note that for the power spectral densities corresponding to f_0(DC component) and $f_{N/2}$(Nyquist frequency), the above normalization convention is different by a factor of 2 from the classical half–spectrum periodogram normalization convention (Press et al. 1989, p. 464). This is because $\hat{P}_k$ should be regarded as a true power spectral density (i.e., power divided by the unit frequency interval), and thus we have taken into account the fact that the frequency bins corresponding to these two special frequencies are half as large as the others.

The statistical properties of the PSD estimate defined in equation (5) can be derived from the statistical properties of the real and imaginary parts of the DFT coefficients, $A_k = \Re(X_k)$ and $B_k = \Im(X_k)$, which, being linear combinations of zero-mean Gaussian variables, are themselves zero-mean Gaussian variables characterized by their individual variances $\sigma_{A_k}^2$ and $\sigma_{B_k}^2$. The rigorous derivation of the mean and variance of the PSD estimate, $\hat{P}_k = A_k^2 + B_k^2$, requires rather involving mathematical manipulations, which are presented for the interested reader in Appendix A. The remarkably simple results of this detailed analysis are summarized below.

Because the DC and Nyquist frequency bins are purely real, their probability distribution can be expressed in terms of the variance $\sigma_{A_k}^2 \equiv \mu_k$ as

$$p\left(\hat{P}_k\right) = \left(2\pi \mu_k \hat{P}_k\right)^{-1/2} \exp\left(-\frac{\hat{P}_k}{2\mu_k}\right), \quad k = 0, N/2 \tag{6}$$

(Appendix A), which is a χ^2 distribution with 1 degree of freedom. Equation (6) can be employed to evaluate the expected values

$$\langle \hat{P}_k \rangle = \mu_k, \quad \langle \hat{P}_k^2 \rangle = 3\mu_k^2, \quad k = 0, N/2, \tag{7}$$

which, employed in equation (1), provide

$$V_k^2 = 2, \quad k = 0, N/2, \tag{8}$$

which is a result that is independent of the shape of the time-domain window w_n.

If no special time-domain windowing function is applied (i.e., $w_n \equiv 1$), it can be shown that for $k \neq 0, N/2$, the real and imaginary parts of the DFT coefficients are independent zero-mean Gaussian variables characterized by equal variances $\sigma_{A_k}^2 = \sigma_{B_k}^2 \equiv \mu_k/2$. The probability distribution associated with the PSD estimate in this case is an exponential distribution given by

$$p(\hat{P}_k) = \frac{1}{\mu_k} \exp\left(-\frac{\hat{P}_k}{\mu_k}\right), \tag{9}$$

which provides the expected values

$$\langle \hat{P}_k \rangle = \mu_k, \quad \langle \hat{P}_k^2 \rangle = 2\mu_k^2, \quad k = 0, N/2; \tag{10}$$

hence,

$$V_k^2 = 1, \quad k \neq 0, N/2. \tag{11}$$

However, as shown in Appendix A, in the general case of a windowed time-domain signal, the variances $\sigma_{A_k}^2$ and $\sigma_{B_k}^2$ are no longer identical, and if the windowing function is asymmetric, the zero-mean Gaussian random variables A_k and B_k are no longer independent, being characterized by a nonzero correlation coefficient ρ_k. Although in this general case the probability distribution of the PSD estimate cannot be written in a closed form, the joint probability distribution of two correlated Gaussian variables,

$$p(A_k, B_k) = \exp\left[\frac{1}{1-\rho_k^2}\left(-\frac{A_k^2}{2\sigma_{A_k}^2} - \frac{B_k^2}{2\sigma_{B_k}^2} + \rho_k \frac{A_k B_k}{\sigma_{A_k}\sigma_{B_k}}\right)\right] \times \left(2\pi\sigma_{A_k}\sigma_{B_k}\sqrt{1-\rho_k^2}\right)^{-1}, \tag{12}$$

(Davenport & Root 1958) can be employed to obtain a generally valid and remarkably compact expression of the PSD spectral variability,

$$V_k^2 = 1 + |W_{2k}|^2, \tag{13}$$

where

$$W_{2k} = \frac{1}{\sum w_n^2} \sum_{n=0}^{N-1} w_n^2 e^{-4\pi i k n/N} \tag{14}$$

is the normalized DFT of the squared time-domain window evaluated at the even-indexed discreet frequencies $f_{2k} = 2f_k$.

Since by definition, the normalized DFT of the squared time-domain window, $|W_{2k}|^2$, is bounded by the [0, 1] interval, the spectral variability of a windowed Gaussian signal can range, at least in principle, anywhere between 1 and 2, depending on the particular shape of the time-domain window. However, because windows (e.g., Parzen, Hanning, Hamming, etc.) are chosen to reduce spectral leakage, they are invariably well-behaved, symmetric time-domain windows that gradually approach zero at both ends of the time-domain interval. This has the effect of drastically reducing $|W_{2k}|^2$ in most frequency bins, so that $V_k^2 = 1$ in all but a few of them—those near the DC and Nyquist frequency bins. However, we consider it worthwhile to investigate the influence of the time-domain window not only in the case of these standard windowing functions, but also in the more general case of an arbitrary time-domain window, which as we already mentioned can be used to model an evolving Gaussian background.

For this purpose, we present in Figure 1 a comparison between the $|W_{2k}|^2$ spectra corresponding to a Hanning window, $w_n = [1 - \cos(2\pi n/N)]/2$, and a Hanning window multiplied by a linearly increasing factor, $w_n' = (1 + n)[1 - \cos(2\pi n/N)]/2$. The rather extreme change rate, one per sample, is chosen just for illustration purposes. All spectra have been obtained by performing an $N = 8192$ point FFT, resulting in 4097 frequency bins. The lower line in Figure 1*a* displays the normalized $|W_{2k}|^2 = |\mathrm{FFT}(w_{2k}^2)/\sum w_{2k}^2|^2$ spectrum corresponding to a Hanning window, while the upper line represents the spectrum $|W_{2k}'|^2 = |\mathrm{FFT}(w_{2k}'^2)/\sum w_{2k}^2|^2$, normalized for the window function w_n, but not for the linear increase. Notice that except for a few frequencies at both ends of the frequency interval, the $|W_{2k}|^2$ and $|W_{2k}'|^2$ spectra have overall flat shapes that are many orders of magnitude below 1, which translates to a negligible contribution to the expected spectral variability expressed by equation (13) at most of the PSD frequency bins. To better quantize this behavior, we display in Figures 1*a* and 1*b* a close-up view of the same spectra at both ends of the frequency interval. Except for a couple of bins at both ends of the frequency interval, the $|W_{2k}|^2$ contribution to the spectral variability corresponding to the Hanning window is negligible. However, the partially normalized $|W_{2k}'|^2$ spectrum is several orders of magnitude above 1 for a few tens of frequency bins. However, by fully normalizing the spectrum, $|W_{2k}|^2 =$

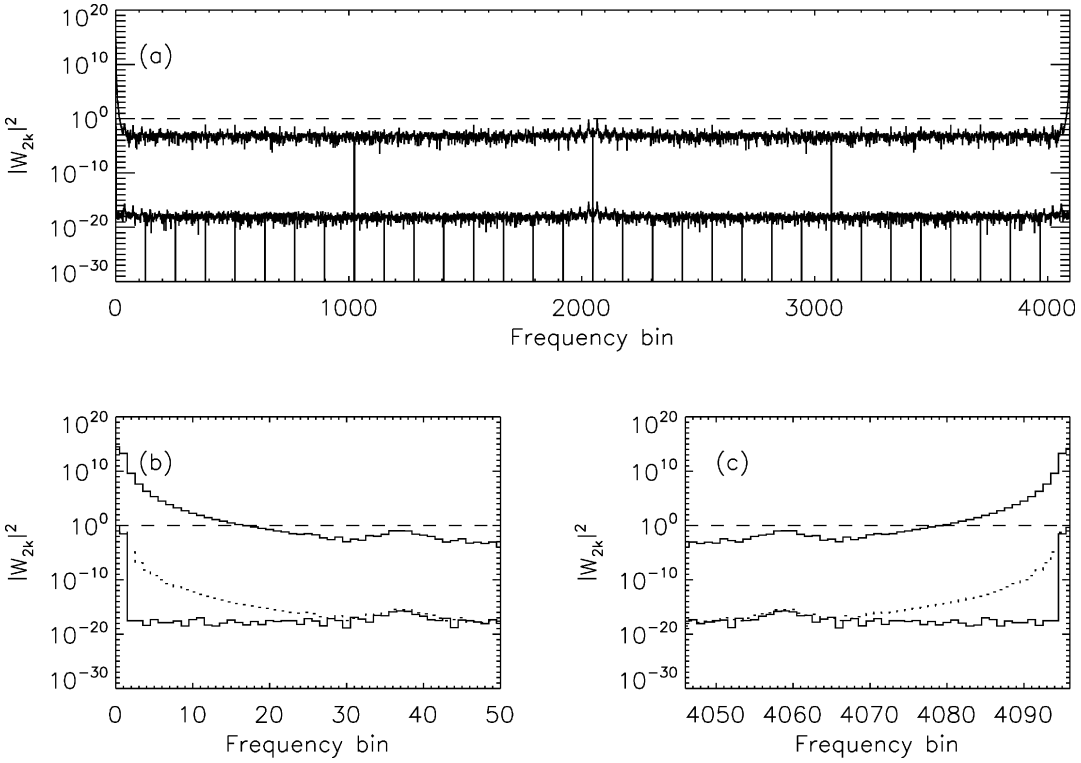

FIG. 1.—Comparison of normalized spectra $|W_{2k}|^2$ and unnormalized spectra $|W'_{2k}|^2$ corresponding to a Hanning window w_n, and to the same window multiplied by a linearly increasing function $c_n = 1 + n$ obtained for $N = 8192$ points: (*a*) $|W_{2k}|^2$ (*lower spectrum*) and $|W'_{2k}|^2$ (*upper spectrum*). (*b*) and (*c*) Expanded view of the low- and high-frequency ends of the spectra presented in (*a*). The dotted lines show the normalized version of $|W'_{2k}|^2$ at both ends of the spectrum. For comparison purposes, a horizontal dashed line in all three panels indicates the unity level.

$|\mathrm{FFT}(w'^2_{2k})/\sum w'^2_{2k}|^2$, as shown by the dotted line in the bottom panels, the spectral variability is practically indistinguishable from that of the original Hanning window.

2.2. Equivalence of Spectral Variability and Spectral Kurtosis

We now compare our definition of spectral variability, V_k^2, to that of similar statistical parameters defined by other authors as a means of discriminating Gaussian and non-Gaussian signals. Historically, as already mentioned in the introduction, Dwyer (1983) was the first one to employ frequency-domain analysis with the goal of detecting non-Gaussian signals in underwater acoustics experiments. Pointing out that the real and imaginary parts of the DFT are zero-mean random variables, Dwyer (1983) used their moments about the origin to define

$$K(f_k) = \frac{\langle A_k^4 \rangle}{\langle A_k^2 \rangle^2}, \tag{15}$$

which is expected to be exactly 3 for a Gaussian process. In this respect, the spectral variability as defined in equation (1) is, up to a constant, mathematically equivalent to Dwyer's definition for frequency-domain kurtosis, for both the DC (assuming a zero-mean time-domain signal) and Nyquist frequencies, where $B_{N/2} = 0$, and $V_{N/2}^2 = K(f_{N/2}) - 1$ has an expected value of 2.

More recently, Vrabie et al. (2003) provided a cumulant-based definition of the newly emerged concept of "spectral kurtosis" (SK) as

$$\mathrm{SK}(f_k) = \frac{\langle |X_k|^4 \rangle - 2\langle |X_k|^2 \rangle^2}{\langle |X_k|^2 \rangle^2}, \tag{16}$$

which is expected to be exactly 0 for a Gaussian time-domain signal. Taking into account that $|X_k|^2 \sim \hat{P}_k$, it immediately follows that

$$V_k^2 = \mathrm{SK}(f_k) + 1, \tag{17}$$

which proves, up to a constant, the mathematical equivalence of these parameters at all but the DC and Nyquist frequencies. However, Antoni (2006) has already pointed out that the SK parameter should be redefined to reflect the purely real nature

of the X_0 and $X_{N/2}$ DFT coefficients, and when this is done, the equivalence becomes exact.

Therefore, the spectral variability V_k^2, as derived in Appendix A, completely characterizes the statistical nature of a time-domain signal at all its spectral components, including the DC and Nyquist frequency bins, hence unifying previously published, incomplete results. Although one may argue whether V_k^2 is a second- or fourth-order statistical parameter, the definition of SK given by equation (16) lacks an immediate physical interpretation, as recognized by Antoni (2006). This is because the only directly *observable* physical quantity in the frequency domain is $\hat{P}_k \sim |X_k|^2$, whose expected value $\langle \hat{P}_k \rangle$ can be used as a meaningful physical reference from which its higher moments can be always computed. In contrast, $|X_k|$ is an artificially defined parameter whose statistical expectation $\langle |X_k| \rangle$ (which is not the same as $\langle \hat{P}_k \rangle^{1/2}$), aside from being physically meaningless, cannot be analytically computed, even for most of the simple statistical distributions. Hence, in terms of observable quantities, V_k^2 can be regarded as a second-order statistical parameter.

However, from the perspective of its practical application, the spectral variability V_k^2 plays a role equivalent to that of the more familiar time-domain kurtosis. Since "spectral kurtosis" is an appropriate name for such a spectral tracer of non-Gaussian processes, and since we have proven their practical equivalence, we see no reason not to refer to V_k^2 by this already familiar name, keeping in mind that it differs by a constant from the original definition, equation (16). This parallels the common practice of using the term "kurtosis" for what one may, in a more formal way, refer to as "kurtosis excess" (Kenney & Keeping 1962). Henceforth in this paper, we refer to V_k^2 as spectral kurtosis, and $\hat{V}_k^2$, defined below, as the spectral kurtosis estimator.

3. SPECTRAL KURTOSIS ESTIMATOR

3.1. Definition and Variance

Our goal is to decide with some level of confidence whether or not the observed PSD estimate $\hat{P}_k$ corresponding to a physical process may be assumed to be driven by an underlying Gaussian distribution. The statistical parameter defined by equation (1) is to be estimated from the data and subsequently compared with the expected value, equation (13), for an ideal Gaussian process. The ideal estimator would be an unbiased one, in other words an estimator having an expectation value identical to the expected value of the magnitude being estimated. Due to its linearity, the PSD estimate $\hat{P}_k$ is an unbiased estimator of the true spectral power density of the signal. However, the statistical parameter defined by equation (1) as the ratio of the variance and the squared mean of the PSD estimate is no longer a linear combination of its components. Therefore, the best option we have is to define the spectral kurtosis estimator as

$$\hat{V}_k^2 = \frac{\hat{\sigma}_k^2}{\hat{\mu}_k^2}, \tag{18}$$

where $\hat{\sigma}_k^2$ and $\hat{\mu}_k$ are unbiased estimators of their associated parameters; i.e., $\langle \hat{\sigma}_k^2 \rangle = \sigma_k^2$ and $\langle \hat{\mu}_k \rangle = \mu_k$. In practice, these two unbiased estimators will be computed from a number of M adjacent blocks of N time-domain samples used to obtain M spectral estimates $\hat{P}_{ki}$ ($k = 0, \ldots, N/2$; $i = 1, \ldots, M$).

Introducing short-hand notations for the sums,

$$S_1 = \sum_{i=1}^{M} \hat{P}_{ki}, \quad S_2 = \sum_{i=1}^{M} \left(\hat{P}_{ki}\right)^2, \tag{19}$$

the unbiased estimators of the true σ_k^2 and μ_k are given by the k-statistics unbiased estimators of the first two cumulants of the $\hat{P}_k$ distribution as:

$$\hat{\mu}_k = \frac{1}{M} S_1, \quad \hat{\sigma}_k^2 = \frac{MS_2 - S_1^2}{M(M-1)}, \tag{20}$$

(Kendall & Stuart 1958, p. 280), where S_1 and S_2 should be understood as having an implicit dependence on f_k. Using these expressions, the spectral kurtosis estimator becomes

$$\hat{V}_k^2 = \frac{M}{M-1}\left(M\frac{S_2}{S_1^2} - 1\right), \tag{21}$$

which is similar in form, but only asymptotically equivalent, to the SK estimator obtained by Vrabie et al. (2003) in terms of the second- and fourth-order cumulants of the complex X_k distribution.

In Appendix B, we present a detailed derivation of the variance of the spectral kurtosis estimator defined in equation (21). The result, which is dependent on the particular shape of the time-domain windowing function, is

$$\mathrm{Var}\left(\hat{V}_k^2\right) = \frac{4M}{(M-1)^2}\left(1 + 5|W_{2k}|^2 - |W_{2k}|^4 + |W_{2k}|^6\right). \tag{22}$$

However, as discussed in § 2 and shown in Figure 1, at the DC and Nyquist frequency bins $|W_{2k}| = 1$, and at most of the other bins $|W_{2k}| \simeq 0$; hence, the $1/M$-order approximation of equation (22) simplifies to

$$\mathrm{Var}\left(\hat{V}_k^2\right) \simeq \begin{cases} 24/M, & k = 0, N/2, \\ 4/M, & k = 1, \ldots, (N/2 - 1), \end{cases} \tag{23}$$

which can be used as a statistical detection threshold for practical applications.

Remarkably, equation (23) reveals that the variance of the SK estimator at the DC and Nyquist frequencies, where it

equals the variance of a kurtosis estimator, is about 6 times larger than the variance at other frequency bins. We consider it worthwhile to discuss this difference in more detail, due to its importance from a practical point of view, as related to the implementation of a classical kurtosis estimator $\hat{K}$ for non-Gaussian signal detection using filter-bank spectrometers (Ruf et al. 2006).

To correctly quantify the difference between the variances of the $\hat{V}_k^2$ and $\hat{K}$ estimators, we have to take into account the fact that in the case of a finite impulse response (FIR) filter-bank spectrometer, the amplitudes of the filtered signal at $N/2 + 1$ frequencies are obtained directly as real linear combinations of $N/2 + 1$ time-domain samples. Therefore, from the same $N \times M$ data samples as used by the DFT-based PSD estimator, one would obtain in each filter-bank frequency channel a set of $N \times M/(N/2 + 1) \simeq 2M$ independent PSD estimates. Therefore, for a given amount of time-domain data, the $\hat{V}_k^2$ variance of $\sim 4/M$ has to be compared with a $\hat{K}$ variance of about $\sim 12/M$. Therefore, for a given time-domain data set, the SK estimator relying on DFT-based PSD estimates may be more efficient by a factor of 3, in terms of its variance, than a kurtosis estimator based on filter-bank PSD estimates.

3.2. Performance and Limitations of the SK Estimator

In general, non-Gaussian signals should cause the SK estimator $\hat{V}_k^2$ to deviate statistically (with respect to its variance, eq. [23]) from its expected value, equation (13). Thus, the SK estimator provides a potentially useful tool for discriminating Gaussian and non-Gaussian signals. We now use simulations to examine the performance and limitations of this detection tool for various non-Gaussian time-domain signals.

3.2.1. Detection of Stationary Deterministic Signals

We consider in this section the simplest case of a time-domain signal composed of background Gaussian noise of zero mean and variance σ_x^2 mixed with a deterministic sinusoidal signal of modulated amplitude A_n and random phase ϕ_n uniformly distributed on the $[0, 2\pi]$ interval. Moreover, we consider that the frequency of the deterministic signal exactly matches one of the non-DC or Nyquist discrete DFT frequencies f_k, and that no windowing function is applied. A rigorous mathematical treatment under these assumptions can be found in McDonough & Whalen (1995), where the probability distribution of the PSD estimate is given as

$$p(\hat{P}_k) = \frac{1}{2\,\sigma^2} \exp\left(-\frac{P_k + A^2}{2\,\sigma^2}\right) I_0\left(\frac{\sqrt{P_k}A}{\sigma^2}\right), \qquad (24)$$

which is a noncentral χ^2 distribution with 2 degrees of freedom. Equation (24) can be straightforwardly employed to obtain the mean and variance of the PSD estimate $\hat{P}_k$ as

$$\mu_k = 2\,\sigma^2(1 + \eta_k), \quad \sigma_k^2 = 4\,\sigma^4(1 + 2\eta_k), \qquad (25)$$

where $\eta_k = A^2/2\,\sigma^2$ denotes the signal-to-noise ratio (S/N) at the frequency bin f_k.

Using the results above, the corresponding V_k^2 statistical parameter can be written in terms of S/N as

$$V_k^2 = 1 - \frac{\eta_k^2}{(1 + \eta_k)^2}, \qquad (26)$$

which is the same, up to a constant of 1, as the result obtained by Vrabie et al. (2003) based on similar considerations.

Equation (26) shows that in the case of a monochromatic sine wave mixed with Gaussian noise, the V_k^2 monotonically goes from 1 to 0 when η_k goes from zero to infinity. If we consider a conservative $\hat{V}_k^2$ detection threshold of about $3\,\sigma$ (i.e., $6/\sqrt{M}$), we can determine an expression for the number M of PSD estimates needed for a reliable detection from the background Gaussian noise of a monochromatic RFI signal, which is given by

$$M(\eta_k) \geq 36(1 + 1/\eta_k)^4. \qquad (27)$$

Equation (27) reveals that, for example, $M \simeq 600$ should be sufficient for a reliable detection of a signal having an average power equal to the background noise, while only about 60 PSD estimates would be needed to reliably detect a monochromatic signal about 10 times stronger than the background noise. However, the same relationship shows that for $M \leq 36$, the SK estimator would not be able to detect a monochromatic signal, even if its S/N goes to infinity.

Before continuing to evaluate the performance and limitations of the SK estimator, it should be mentioned that as shown by Fridman (2001), the system of equations (25) may in principle be solved for σ^2 and A^2 in terms of the observed quantities μ_k and σ_k, leading to a cleaned background spectrum and an averaged power estimate of the deterministic signal. Although attractive, we want to point out that the method proposed by Fridman (2001) can be valid only for the particular case under discussion. However, the method could still be used with some degree of reliably as long as $V_k^2 < 1$–$3\,\sigma$. As shown later in the text, such situations are less frequently observed than the case of an SK estimator having values well above unity.

3.2.2. Detection of Transient Signals

Although derived for the assumption of time-domain stationarity of the signal, the SK estimator can be very effective in detecting transient signals mixed with the stationary background noise, as noted by Vrabie et al. (2003). To support this assertion, which may be of particular practical importance, we consider below the case of a transient monochromatic time-domain signal that is switched on and off during the observed

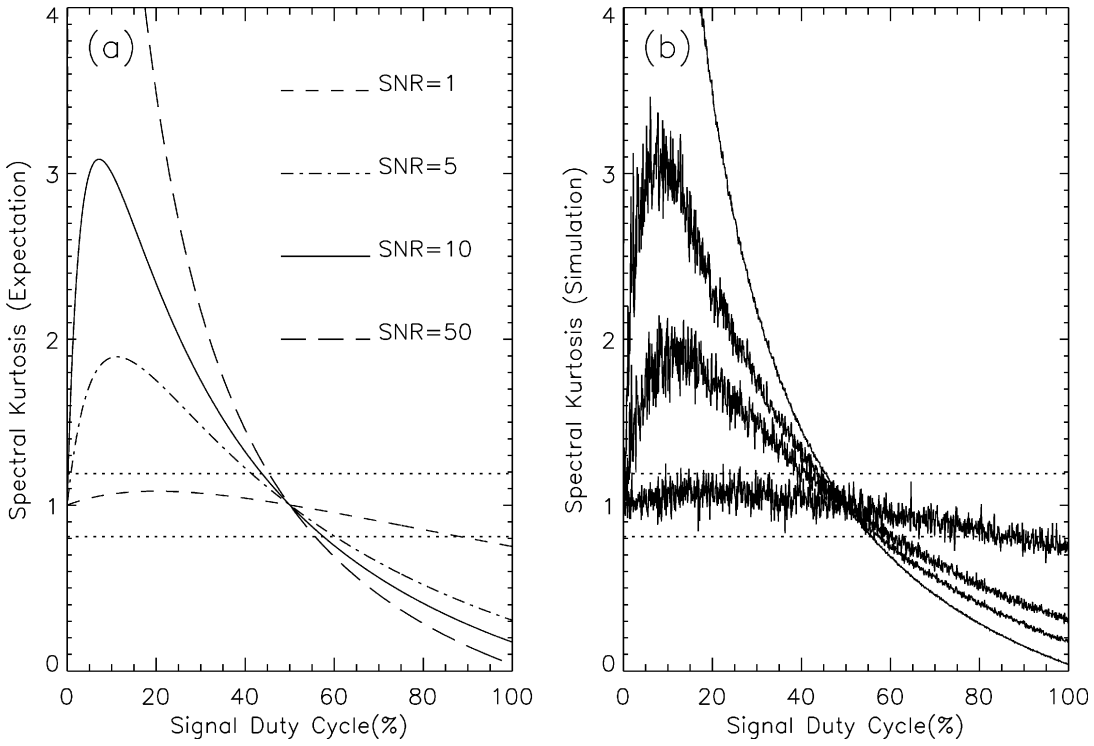

FIG. 2.—SK estimator vs. duty cycle of transient RFI for different S/Ns. (*a*) Theoretical expectation for S/Ns ranging from 1 to 50. The horizontal dotted lines represent the $\pm 3\ \sigma$ detection thresholds for $M = 1000$ PSD estimates. (*b*) Simulated SK results corresponding to the theoretical parameters presented in (*a*). The duty-cycle ranges for which the transient signal cannot be statistically distinguished from the Gaussian background lie between the intersections of the SK curves and the horizontal dotted lines.

$N \times M$ time-domain samples. To simplify the discussion, we further assume that the randomly occurring on/off transitions exactly coincide with some of the edges of the M nonoverlapping blocks of lengths N. Although this assumption is rather unrealistic, we believe that it will not significantly affect the validity of the result.

Therefore, measurements under the assumptions above result in M PSD estimates at a given frequency f_k, out of which, let us say, R estimates have a finite η_k, while the remaining $(M - R)$ have $\eta_k = 0$. Hence, for each individual data block, equation (24) provides the expectation values

$$\begin{aligned} \langle P_k \rangle &= \mu_k = 2\,\sigma^2(1 + \eta_k), \\ \langle P_k^2 \rangle &= \mu_k^2 + \sigma_k^2 = 4\,\sigma^4(2 + 4\eta_k + \eta_k^2), \end{aligned} \tag{28}$$

which, with the appropriate values of η_k, can be used to estimate the power sums entering the SK estimator,

$$\begin{aligned} \langle S_1 \rangle &= 2R(1 + \eta_k)\sigma^2 + 2(M - R)\sigma^2, \\ \langle S_2 \rangle &= 4R(2 + 4\eta_k + \eta_k^2)\sigma^4 + 8(M - R)\sigma^4. \end{aligned} \tag{29}$$

Hence, for the particular case being discussed, the SK estimator, equation (21), is

$$\langle \hat{V}_k^2 \rangle = \frac{M}{M - 1}\left[1 + \frac{(M/R - 2)\eta_k^2}{(M/R + \eta_k)^2}\right], \tag{30}$$

which, as expected, reduces to equation (26) for $R = M$ and reasonably large M.

Equation (30) allows us to evaluate the performance of the SK estimator for different duty cycles, $d = R/M$, of the randomly occurring transient signal at different values of the S/N parameter η_k. First of all, the duty cycle dependence of the SK estimator is not monotonic, having a local maximum at $d = 1/(4 + \eta_k)$, given by

$$\hat{V}_k^2\left(\frac{1}{4 + \eta_k}\right) = \frac{M}{M - 1}\left(1 + \frac{\eta_k^2}{8 + 4\eta_k}\right). \tag{31}$$

This result indicates that the SK estimator may indeed be very efficient in detecting low–duty-cycle transients, which may result in a value for the estimator well above 1, depending on the S/N parameter. However, it should be mentioned at the same time that equation (30) also reveals an important limitation of the estimator. Indeed, if the duty cycle is exactly 50%

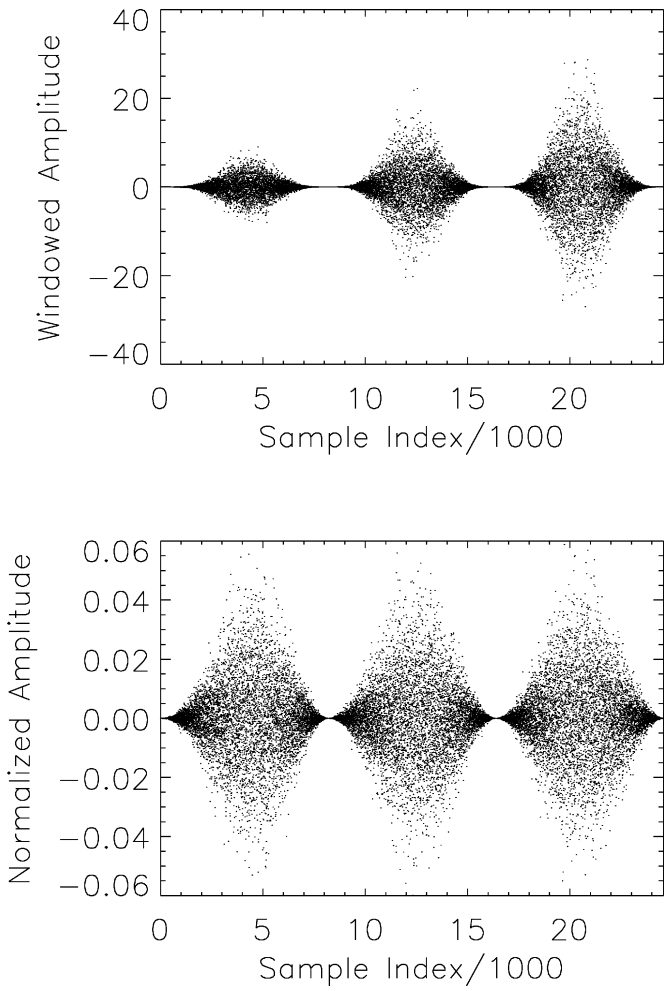

FIG. 3.—Windowed time-domain signal obtained from a stationary Gaussian signal $\{x_i\}$ multiplied by a linear factor $c_i = 1 + \alpha i$. Each of the three data blocks displayed contains $N = 8192$ data samples; the sample index i runs from 0 to $3N$; and $\alpha = 10/(3N)$. (*a*) Original time-domain signal $\{c_i x_i\}$ multiplied by three successive Hanning windows, w_i. (*b*) Signal displayed in (*a*) normalized to the total power contained in each data block, $\{w_i c_i x_i / [\sum_{\text{block}} (w_i c_i x_i)^2]^{1/2}\}$.

(i.e., $R = M/2$), the value of the SK estimator becomes practically 1. This makes the transient signal statistically indistinguishable from a Gaussian signal, regardless of the S/N—behavior that is identical to that known in the case of the time-domain kurtosis estimator (De Roo et al. 2007).

To better quantify this limitation, one should compare the estimate given by equation (30) with the detection threshold limits above and below the expected value for a Gaussian signal (i.e., $1 \pm 6/\sqrt{M}$), which provides the range of duty cycle over which the transient signal is undetectable by means of the SK estimator. The dependence of equation (30) is shown graphically in Figure 2. Panel *a* presents the expected behavior of the SK estimator in the case of monochromatic transient signals having duty cycles ranging from 0% to 100% and $\eta_k = 1$, 5, 10, and 50. The dotted lines show the $\pm 3\ \sigma$ limits for $M = 1000$ PSD estimates. The figure shows that the $\eta_k = 1$ signal can be detected from the background noise only for a duty cycle exceeding about 90%, while even for $\eta_k = 50$, the estimator fails to detect the transient signals having a duty cycle between ~45% and ~55%. Figure 2*b* presents a simulation obtained for the same parameters as used in Figure 2*a*. Besides a very good agreement with the theoretical expectations, this figure also illustrates how the statistical variance of the estimator is influenced both by S/N level and duty cycle. It can be seen that this variance strongly decreases with the increase of the S/N level, and it also decreases to some extent with the increase of the duty cycle.

3.2.3. Quasistatic Gaussian Background

We now turn to the case of an evolving Gaussian background, such as that expected for a strong solar radio burst. To simulate this, we assume that each individual data sample is drawn from different Gaussian distributions with variances $\sigma_{x_n}^2$. This situation is mathematically equivalent to replacing the original data set $\{x_i\}$ with a linear combination $\{c_i x_i\}$, where the real coefficients c_i may have a continuous time evolution that extends beyond the range of the N samples used to produce one PSD estimate. However, we assume that the timescale of the analyzed physical process is slow compared with the time span of the $N \times M$ data samples, so that the c_i coefficients have only a smooth temporal variation. From the example illustrated in Figure 1, we have already learned that the $|W_{2k}|^2$ spectrum of a Hanning window multiplied by a smooth function, such as $c_n = 1 + \alpha n$, cannot significantly deviate from 1 in more than a few frequency bins of the SK estimator. However, this assumes a periodic sawtooth variation of the c_i coefficients, independent of PSD index M. Fortunately, we show that a similar behavior can be obtained by building an SK estimator based on normalized PSD estimates. Indeed, if after each realization of a $\hat{P}_k$ estimate its spectrum is normalized to the full-bandwidth total power $\sum \hat{P}_k$, the $\hat{V}_k^2$ statistics become insensitive to smooth temporal variations of the input signal. Note that this spectral-domain normalization is equivalent to a time-domain normalization obtained by dividing each windowed data block by $[\sum_{\text{block}} (w_i c_i x_i)^2]^{1/2}$. Figure 3 displays the original and normalized time sequences for three consecutive blocks of $N = 8192$ data samples generated by numerical means. The bottom panel of this figure suggests that the normalized time-domain sequence might be statistically equivalent to the case of a stationary Gaussian signal analyzed with a modified Hanning window similar to the one analyzed in Figure 1.

To validate our expectations, we present in Figure 4 a numerical simulation obtained for $M = 1000$ data blocks of $N = 8192$ data samples. The data set consists of a linearly increasing Gaussian background, $c_{mn} = 1 + \alpha[(m-1)N + n]$, mixed with two deterministic signals. The first one, corresponding to the frequency bin $k = 1000$, is a continuous mono-

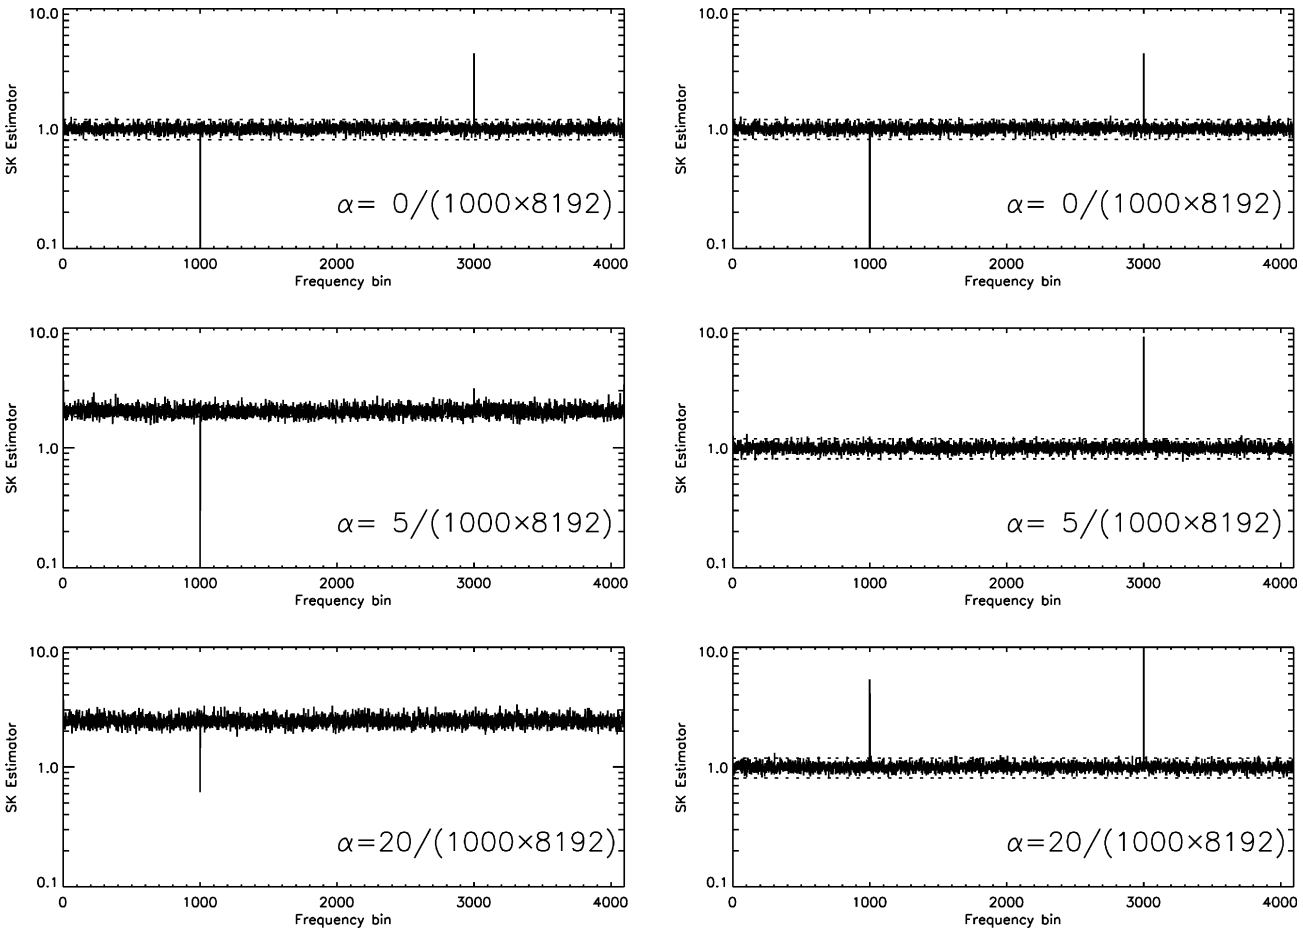

FIG. 4.—SK estimator spectra for a time-domain signal obtained for a stationary Gaussian background $\{x_i\}$ multiplied by a linear factor $c_i = 1 + \alpha i$, with α indicated in each panel. The background signal is mixed with a 100% duty cycle random-phase sinusoidal signal at $f_k = 1000$ and a 20% duty cycle random-phase sinusoidal signal at $f_k = 3000$. Both signals have S/Ns of 1000. Each data block contains $N = 8192$ samples, and $M = 1000$ data blocks were used to generate the statistics. *Left column*: $\hat{V}_k^2$ spectra obtained from PSD estimates corresponding to the original time-domain data blocks multiplied by a Hanning window. The horizontal dashed lines indicate the $\pm 6/\sqrt{M}$ ranges about the medians of each spectrum. *Right column*: $\hat{V}_k^2$ spectra obtained from PSD estimates normalized to their full-bandwidth total powers. The horizontal dashed lines indicate the $\pm 6/\sqrt{M}$ ranges about 1.

chromatic signal with S/N = 1000, while the second one, corresponding to the frequency bin $k = 3000$, has the same S/N but is randomly occurring with a 20% duty cycle. The figure is organized in three rows and two columns. Each row of the left column, from top to bottom, corresponds to a different change-rate parameter α; i.e., $0/(M \times N)$, $5/(M \times N)$, and $20/(M \times N)$, respectively, corresponding to 0%, 0.5%, and 2% variation over N samples. The right column presents the results corresponding to the same parameters, but for an SK estimator built on normalized spectra. Two important effects are apparent. The first is that without normalization (*left column*), the $\hat{V}_k^2$ spectrum is shifted above unity, while its variance obviously becomes larger than for the stationary case (*dotted lines*). This problem is avoided by applying the normalization (*right column*). We actually observed this phenomenon on real solar data during solar bursts, where the estimator failed until we applied the normalization.

The second effect is that the expected duty-cycle behavior described in Figure 2 is altered for both normalized and unnormalized spectra. Indeed, the left column shows that in the case of unnormalized spectra, the SK estimator level of the 100% duty-cycle line decreases when α increases, while the 20% duty-cycle line becomes practically indistinguishable from the background. On the other hand, when normalized spectra are used (*right column*), the low–duty-cycle signals can be easily detected, while for particular values of the change rate, $\alpha = 5/(M \times N)$ in this case, the continuous signal may become undistinguishable from the Gaussian background. However, for $\alpha = 20/(M \times N)$, the SK estimator for the continuous signal becomes larger than unity and is well above the threshold level. In this respect, the change in the background power level appears to have a similar effect on the SK estimator for continuous signals as the change of the duty cycle has for the transient ones.

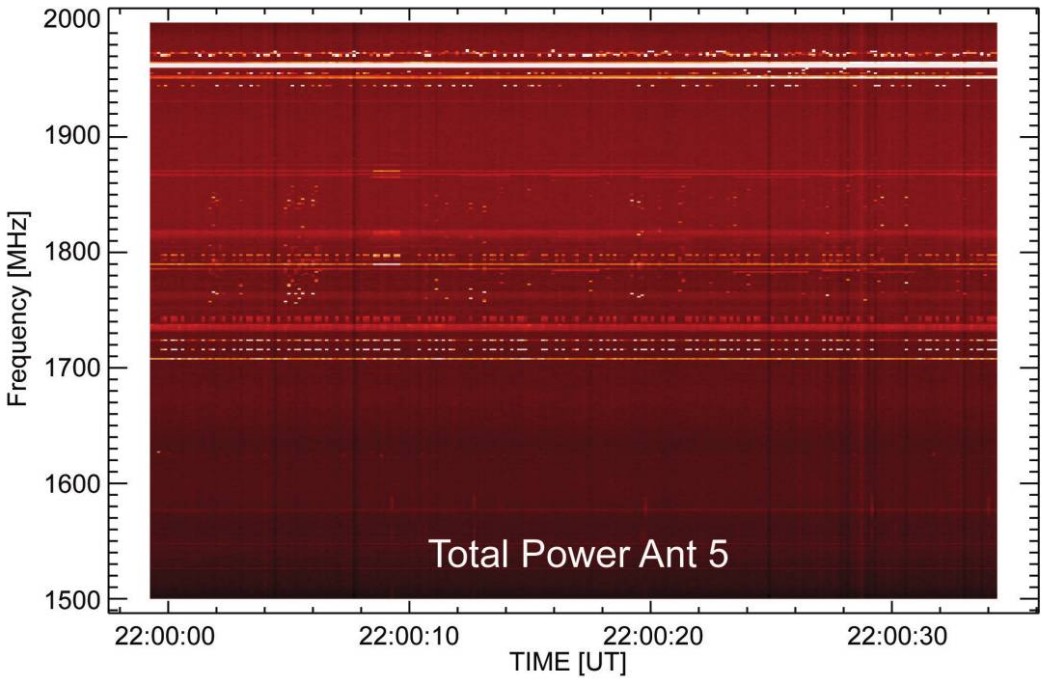

FIG. 5.—RFI-contaminated dynamic spectrum produced from FST radio data obtained during a quiet-Sun observation. The time and frequency resolution of the averaged dynamic spectrum are ~142 ms and ~1 MHz, respectively. Various types of continuous, clustered, and scattered RFI contamination are evident.

4. IMPLEMENTATION AND RESULTS

Having established the theoretical properties of the SK estimator, our aim is to suggest a general approach for its practical implementation as an RFI detection and excision tool. A straightforward implementation of an RFI excision algorithm for a one-channel DFT-based radio spectrograph would require accumulating the sums S_1 and S_2 defined in equation (19) over the desired M spectra, each derived from N blocks of time-domain data and normalized by the band-integrated total power, if necessary. From these sums, the SK estimator, equation (21), is calculated and compared with a threshold value based on equation (22) (or eq. [23]) to define a flag array to be applied to S_1 to remove bad channels. Because of the narrow bandwidth of most RFI, one can maximize retention of RFI-free data channels by accumulating the SK statistics at a finer frequency resolution than the desired (science-based) bandwidth and combine frequency channels after RFI excision.

In the case of an interferometric radio array, the logical AND of the RFI flag arrays (assuming channels with RFI are flagged as zero) yield an RFI flag array for each pair of antennas. Note that this approach avoids the theoretical difficulties of defining a statistical estimator for the correlated spectra, which would have to take into account the fact that the data sets recorded by two different antennas are not statistically independent (Fridman 2001).

4.1. Software Implementation of an RFI Excision Algorithm for the FST Interferometric System

One of the main goals of the newly developed FASR Subsystem Testbed (FST; Liu et al. 2007) is to provide a test bed for studying methods of RFI excision. The design of this instrument, which records full time-resolution time-domain data at a sampling rate of 1 ns, allows playback of data in simulated real time while experimenting with software-implemented RFI excision algorithms in either time or frequency domains over an instantaneous 500 MHz bandwidth. This 500 MHz band can be tuned anywhere in the 1–9 GHz frequency range. Since the FST system is composed of three independent data channels linked to three of the seven antennas of the Owens Valley Solar Array (OVSA; Gary & Hurford 1999), the performance of the RFI excision algorithms on correlated data can also be investigated. However, this flexibility does not come without cost. Due to the limitation to 2 Mbytes per channel of on-board memory of the 8 bit resolution digitizing system (Acqiris DC271), no more than about 2 ms of contiguous time-domain data segments can be recorded at a time, with time gaps between data segments that are due to the time needed to download the recorded data and store it on the host computer hard drive. The duty cycle for data acquisition is roughly 1%. The data we use to illustrate the SK algorithm were taken in modes 1 and 3, as defined in Liu et al. (2007), and a standard Hanning window has been applied to reduce the FFT spectral leakage. For three-

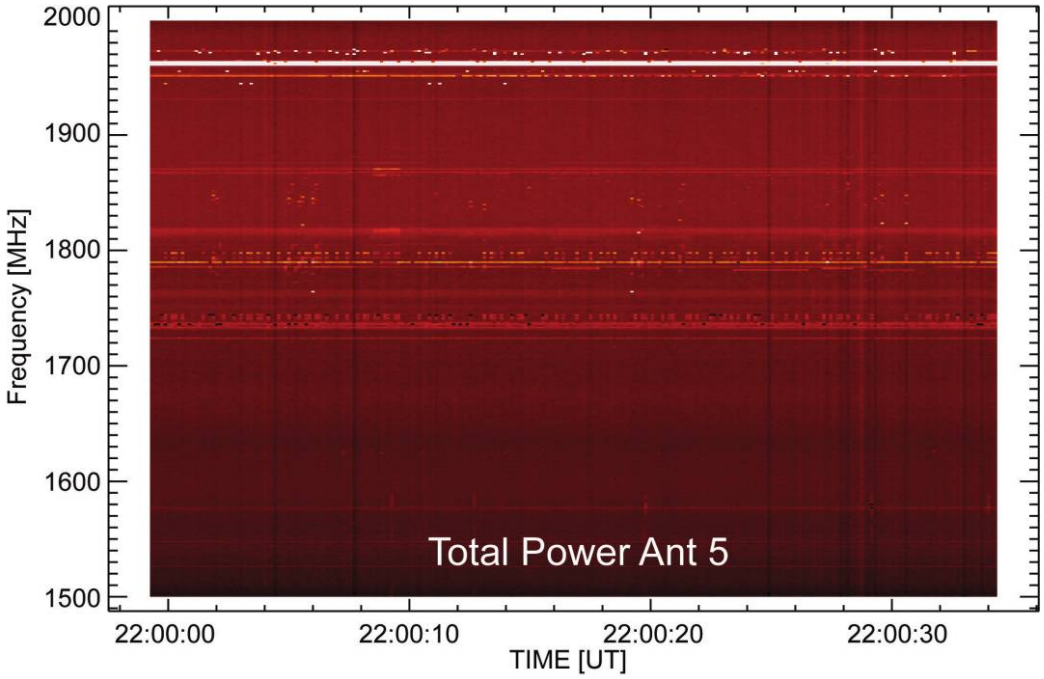

FIG. 6.—RFI-cleaned dynamic spectrum obtained from the same data displayed in Fig. 5. For each averaged spectrum, the RFI flags have been generated based on the $\hat{V}_k^2$ statistics obtained from $M = 244$ contiguous PSD estimates. Most of the RFI spikes have been removed, and the continuous RFI spectral lines are reduced.

channel data taken in mode 1, 2 ms of contiguous time-domain data are recorded, with time gaps between segments on the order of 160 ms. In mode 3, 100 μs of data are separated by 20 ms gaps. As a consequence of these hardware constraints, the contiguous SK statistics cannot be extended beyond an upper limit (for mode 1 data) of $N \times M = 2 \times 10^6$ contiguous data samples, which, depending on the frequency resolution desired, puts an upper limit on the maximum number of PSD

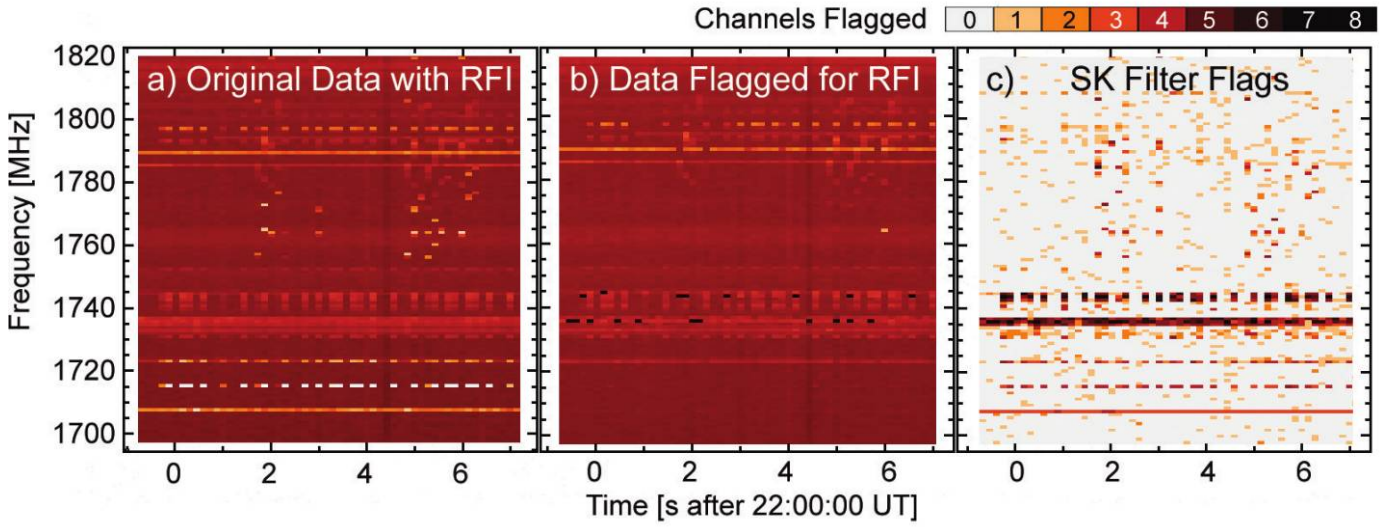

FIG. 7.—(*a*) Magnified portion of Fig. 5 containing RFI. (*b*) Same portion of Fig. 6, with some of the RFI removed. (*c*) Number of subchannels flagged for RFI, in dynamic spectrum format, corresponding to the same portion of the dynamic spectrum as in (*a*) and (*b*). The number of subchannels is color-coded according to the scale above (*c*).

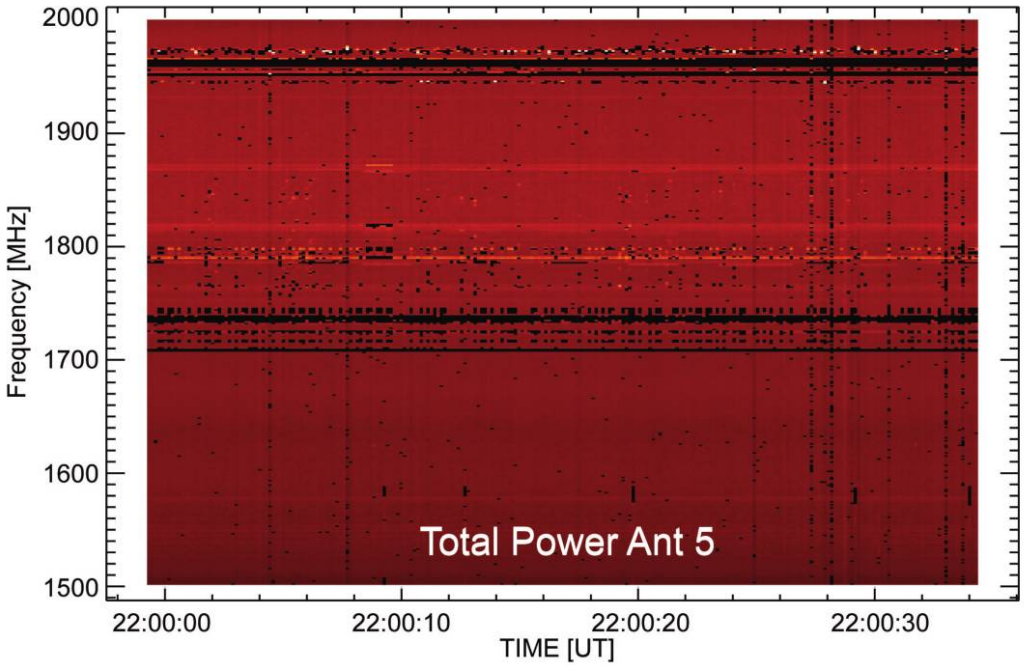

FIG. 8.—RFI-cleaned dynamic spectrum obtained from the same data displayed in Fig. 5, but with frequency resolution reduced by a factor of 8 in order to improve the SK statistics ($M = 1952$ contiguous PSD estimates). To allow a better comparison with Figs. 5 and 6, no subchannel integration has been performed.

estimates, M, that can be obtained from the same contiguous data set. Nevertheless, as shown later in the text, this limitation can be overcome by ignoring the time gaps and performing SK statistics based on more-or-less evenly distributed, noncontiguous PSD estimates. But one should keep in mind that the efficiency of the estimator using this operating mode may be affected by the temporal behavior of nonstationary RFI signals.

The FST system also provides the ability to investigate the validity and efficiency of the RFI excision algorithms under extreme conditions of rapidly varying power levels of the Gaussian signal, which occur during the impulsive phases of some solar radio bursts recorded by the instrument. As shown in the previous section, it is necessary in such situations to obtain SK statistics based on normalized PSD estimates. However, the FST system does not need special treatment in such situations, due to its built-in automatic gain control (AGC) loop, which dynamically maintains a fixed power level of 6 dBm over the 500 MHz bandwidth data channel inputs of the digitizer. Therefore, the unscaled PSD estimates obtained directly from the raw data are already normalized by the default behavior of the system. However, it is one of the main goals of this study to check whether or not the SK estimator really behaves as expected under the extreme conditions of fast-varying radio bursts, a situation we consider of great practical interest for the design of future solar radio telescopes with real-time RFI excision capability. Moreover, studying the behavior of the SK estimator under radio burst conditions should help us evaluate the extent to which the assumption of an underlying Gaussian parent population still holds in the presence of various coherent or incoherent emission mechanisms that might be responsible for the natural radio emission produced by the solar radio bursts. In particular, we want to investigate whether or not a dynamic SK-based RFI excision algorithm can be safely used during a radio burst, without the risk of confusion with artificial RFI signals.

4.2. Performance of the Algorithm in the Case of Strong RFI

In the 1–9 GHz observational range of the FST instrument, the 1.5–2.0 GHz frequency band is one that is strongly contaminated by RFI. Figure 5 displays a dynamic spectrum of the quiet Sun that is heavily contaminated by RFI signals with a variety of time-domain characteristics. The dynamic spectrum has a frequency resolution of ~1 MHz and a time resolution of ~142 ms, each spectrum being obtained by averaging data obtained from 2 ms segments of contiguous data corresponding to one mode 1 acquisition, separated by 140 ms of dead time. Each segment has been divided into $M = 244$ adjacent blocks of $N = 8192$ data samples each, resulting in 244 PSD estimates in each of 4097 frequency channels, having a frequency resolution of ~122 kHz. We then integrate over eight frequency bins (subchannels) and accumulate in time to produce a time-averaged spectrum with a frequency resolution of ~1 MHz.

Figure 6 illustrates the performance of the RFI excision al-

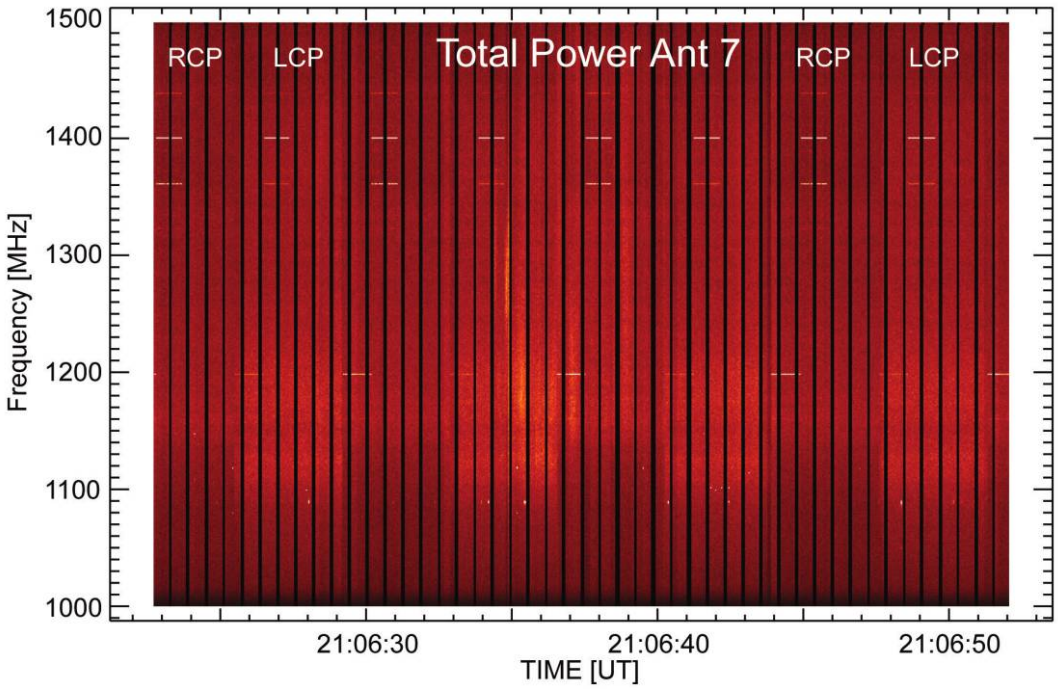

FIG. 9.—Total power dynamic spectrum of a weak solar radio burst recorded on 2006 April 5, around 21:06:30 UT, in the 1–1.5 GHz frequency band. The time resolution of the dynamic spectrum is 25 ms, and the frequency resolution is 1 MHz. The vertical black stripes in the plot represent variable time gaps needed to dump each separate acquisition on the storage media. Four on/off periodic RFI signals can be seen at 1190, 1360, 1400, and 1440 MHz, in addition to a few narrowband features scattered below 1200 MHz, most likely due to aircraft distance-measuring equipment. RCP and LCP labels indicate a few of the periodic changes in the background intensity, about every 4 s, due to periodic switching of the observed circular polarization, which is the default mode of operation of OVSA.

gorithm in the case of $\hat{V}_k^2$ statistics based on a single averaged spectrum. A detection threshold of $\sim 1 \pm 3\,\sigma = 1 \pm 6\sqrt{M}$ has been used to flag the RFI-contaminated channels and remove them prior to subchannel integration. Channels for which subchannels have been removed are proportionally scaled in amplitude to correct for the missing subchannels. Given the small number ($M = 244$) of PSD estimates we have available as a result of the low duty cycle of the FST, we consider the performance illustrated by Figure 6 to be promising. Even in this case of low-number statistics, Figure 6 shows that the algorithm successfully removes most of the transient RFI without affecting the continuum background. The removal of the strong, apparently continuous RFI is less effective, perhaps due to a combination of low-number statistics, time gaps forced by the low duty cycle of the instrument, or residual nonstationary power level, which can reduce the effectiveness of the algorithm for continuous RFI as shown in Figure 4. We anticipate that the algorithm will perform better for continuous RFI in an instrument with a 100% duty cycle. Indeed, a detailed inspection of the RFI flags generated by the algorithm reveals that it is sensitive to the continuous RFI. Figure 7 zooms in on two areas in Figures 5 and 6 and shows the number (zero to eight) of surviving (unflagged) subchannels in each frequency and time bin. It is evident that in the case of continuous RFI components, for much of the time, at least one out of eight frequency channels are flagged as being contaminated by RFI, and we anticipate that if better statistics (larger M) were possible, the results would have been better. The general salt-and-pepper aspect of the flag spectrum is explained by normal statistical fluctuations of the SK estimator, expected to be on an order of $\sim 1\%$ for the 3 σ detection threshold used. The total number of bins flagged as RFI represent $\sim 2.5\%$ of the original data, which suggests that $\sim 1.5\%$ of the data have been successfully identified as contaminated with RFI, and removed.

To check the extent to which the performance of the SK estimator would be improved by better statistics, we present in Figure 8 the results obtained by reducing the estimator frequency resolution by a factor of 8, yielding $M = 1952$ PSD estimates in each frequency-time bin. To allow a better comparison with the results presented in Figure 6, we have not performed in this case any frequency integration, so that the RFI-flagged bins appear black in the plot. Although the results have improved, we note that the algorithm performance is still not optimal. This is because by reducing the frequency resolution of the estimator, the S/N of narrowband RFI is also reduced, due to the wider frequency bins. The total data loss

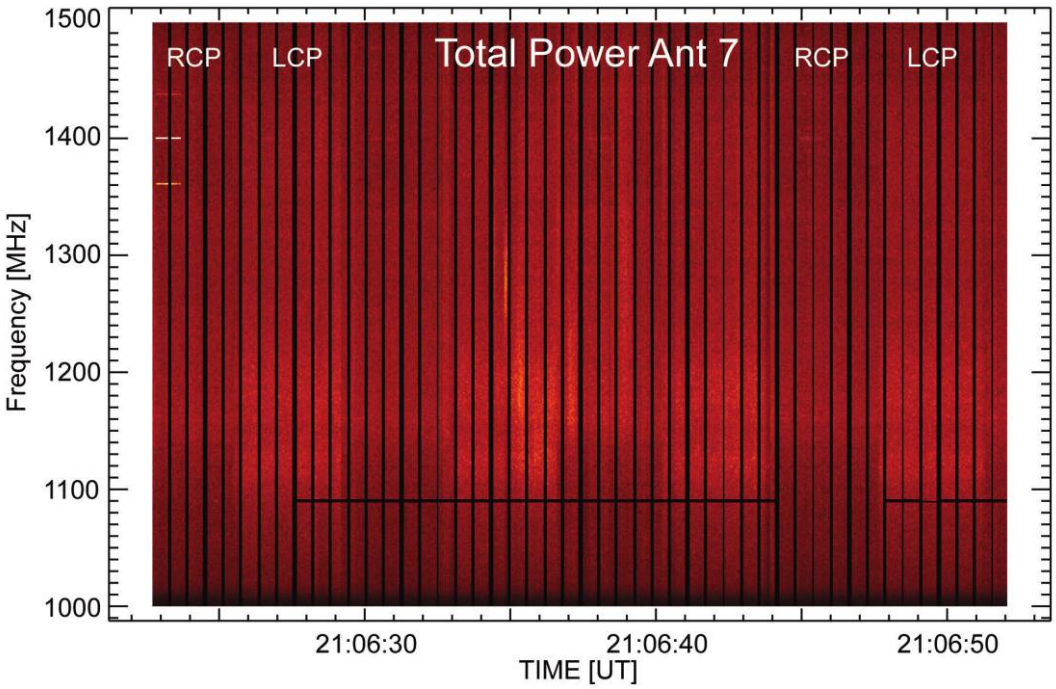

FIG. 10.—RFI-cleaned version of the spectrum displayed in Fig. 9, using a running buffer as explained in the text. A total of $M = 2880$ PSD estimates have been used to evaluate the SK estimator and generate the RFI flags. In comparison with the dynamic spectrum displayed in Fig. 9, the periodic fixed-frequency RFI signals, and nearly all of the scattered RFI spikes, have been removed, while the spectrum of the solar radio burst is not affected by the excision algorithm. The algorithm fails to detect the RFI spectral lines at the earliest times, until a statistically significant number of PSD estimates is accumulated in the running buffers. The RFI spikes in a few time bins at about 1090 MHz, having a very low duty cycle at the timescale of the running buffers, results in the removal of most of the data for the affected frequency channel.

in this case is ~7%, which may be partially due to better performance, but is also due to the fact that the narrowband RFI contaminates larger frequency bins.

4.3. Weak Solar Radio Burst Spectrum Contaminated by Moderate RFI

Figure 9 displays a dynamic spectrum corresponding to a solar radio burst recorded by the FST instrument on 2006 April 5 in the 1.0–1.5 GHz range. The burst consists of several weak emissions, the most intense being observed around 21:06:30 UT. Four on/off periodic RFI signals can be best observed at about 1200, 1360, 1400, and 1440 MHz, in addition to a few transient RFI spikes near 1100 MHz. A periodic change in the background intensity every 4 s is due to feed-switching between right-hand and left-hand circular polarization. The spectrum in each time bin results from $N \times M = 10^5$ time samples of 1 ns each ($N = 8192$, $M = 12$, $\Delta f = 0.122$ MHz) separated by 25 ms dead time, while ~100 ms time gaps occur between 20 consecutive spectra (*vertical black stripes*), due to additional dead time needed for writing the data to disk.

Although the number of PSD estimates, $M = 12$, is far too small for useful statistics for an individual bin, the efficiency of the SK algorithm may be greatly improved by accumulating the sums S_1 and S_2, on which the SK estimator is based, in running buffers $\{S_1\}$ and $\{S_2\}$ that span a much longer time. Sums over the $M = 12$ PSD estimates for each time bin are added to limited-length arrays holding a predefined number of partial sums. When the length limit of these arrays is reached, the oldest $\{S_1, S_2\}$ pair is discarded to make room for the incoming one. This procedure allows us to generate, based on the total sums $S_1 = \sum S_{1k}$ and $S_2 = \sum S_{2k}$, a statistically significant SK estimator that is used to flag the RFI-affected spectral components of the incoming time bin.

Figure 10 presents the cleaned version of the spectrum presented in Figure 9, based on running history buffers maintaining 240 partial sums of 12 contiguous PSD estimates. Thus, the statistics are generated from $M = 2880$ PSD estimates. Comparing Figs. 10 and 9, we conclude that the approach successfully removes most of the RFI spectral components without affecting the spectrum of the solar radio burst. Note that the algorithm fails to detect the RFI in the first few seconds, before a statistically significant number of PSD estimates is accumulated in the running SK buffers. While the "learning curve" of the estimator should not be of concern, given the

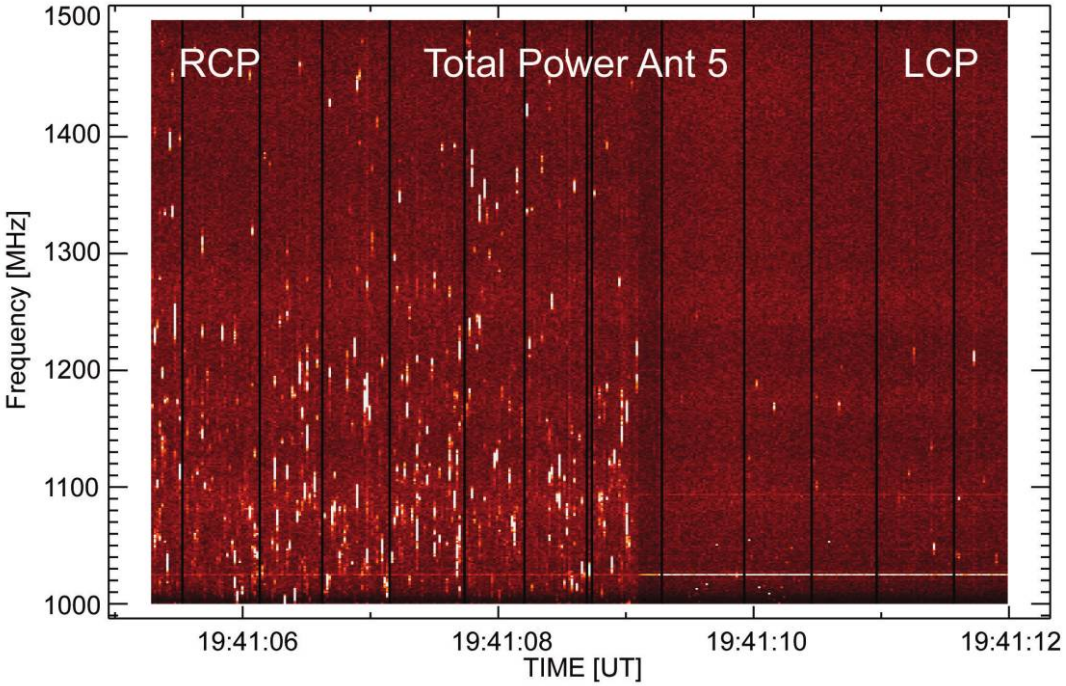

FIG. 11.—Portion of the dynamic spectrum of a solar radio burst recorded by FST on 2006 December 6, 19:41:00 UT, when strong radio spikes where observed. The frequency and time resolutions of the dynamic spectrum are ~0.122 MHz and 25 ms, respectively. However, each averaged spectrum corresponding to a 25 ms time bin was obtained only from 0.1 ms of contiguous data samples, resulting in $M = 12$ consecutive PSD estimates. The first 4 s of data, taken with a right-hand–polarized (RCP) feed, reveal intense radio spikes, with only weak RFI contamination. The rest of the data, taken with a left-hand–polarized (LCP) feed, reveal just a few radio spikes and clear RFI contamination. However, some of the spiky structures present in both polarizations might also be due to low–duty-cycle RFI.

fact that this is a time-limited "start up" effect, the high sensitivity of the estimator to transient RFI spikes might result in an undesired complete blanking of some spectral regions that are only transiently affected by RFI, as shown by the missing channel near 1090 MHz. RFI spikes in a few time bins can result in removal of all data for the affected frequency until the contribution to the sum "ages off" and leaves the buffer. A balance must be struck between increasing buffer size for improved statistics and the corresponding loss of data due to overlong response to transient RFI.

4.4. Strong Solar Radio Burst with Fine Time-Frequency Structures

One of the solar radio bursts recorded by FST that occurred on 2006 December 6 displayed extremely strong, narrowband (few MHz), short-duration (<20 ms) radio spikes that nominally look a great deal like RFI. Figure 11 displays ~7 s of the dynamic spectrum, where the highly right-circularly polarized (RCP) spike bursts are apparent during the first ~4 s, when the system was measuring RCP, but are nearly absent when the system was switched to left-circular polarization (LCP) in the last ~3 s. At the same time, RFI is weak during the first ~4 s, but more readily visible in the last ~3 s. As in the previous burst (Fig. 9), each time bin of the spectrum is obtained from $N \times M = 10^5$ time samples of 1 ns each ($N = 8192$, $M = 12$, $\Delta f = 0.122$ MHz), separated in this case by 20 ms dead time.

A legitimate question arising from the displayed spectral structures in Figure 11 is whether or not the $\hat{V}_k^2$ statistics will remain insensitive to the natural spiky structures while still discriminating the true RFI contamination. The answer to this question is important from both practical and theoretical points of views. From a practical standpoint of a real-time RFI excision algorithm, it would be highly undesirable for the natural solar emission to be irreversibly removed due to its being mistaken for low–duty-cycle RFI contamination, to which we already proved the high sensitivity of the SK estimator. From a theoretical point of view, the $\hat{V}_k^2$ analysis of this data set might reveal whether or not the intensity of such natural radio spikes, generally believed to be produced by the electron cyclotron maser emission mechanism (see Fleishman et al. 2003 and references therein), nevertheless obeys a Gaussian statistical distribution.

The number $M = 12$ of contiguous PSD estimates within a single time bin is too small to be statistically significant, and when we applied the running buffer approach of the previous

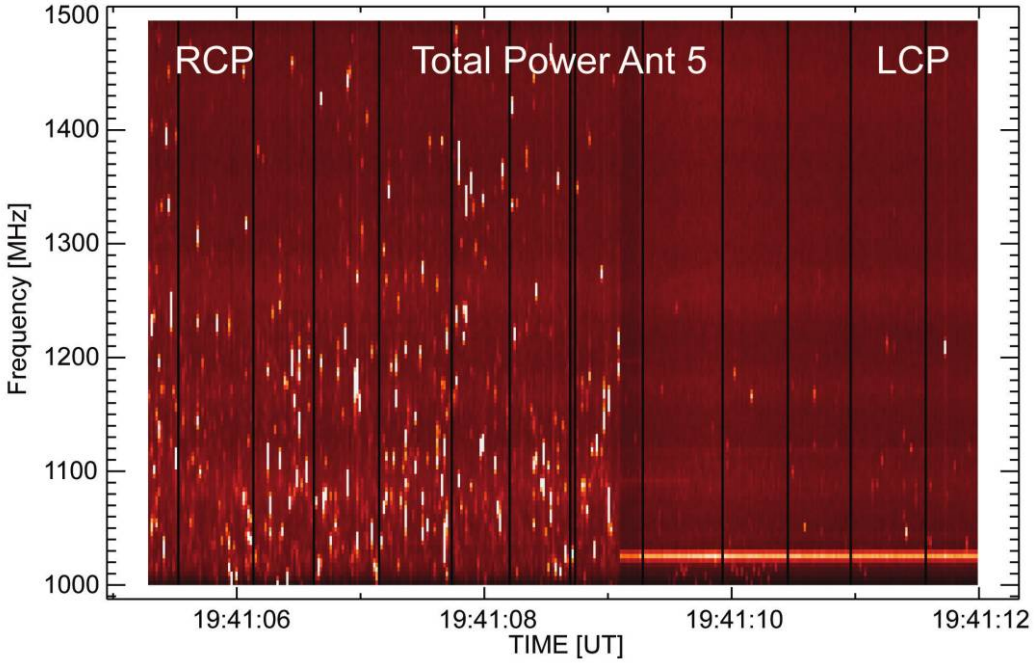

FIG. 12.—Same dynamic spectrum displayed in Fig. 11, but with much lower frequency resolution. Each of the contiguous 0.1 ms data segments have been divided in $M = 390$ adjacent blocks of $N = 256$ data samples, for an ~3.9 MHz frequency resolution. No integration in the frequency direction has been performed. Most of the spike clusters are still partially resolved in the RCP segment of the spectrum, while only one RFI spectral line at about 1020 MHz remains clearly visible.

section, the algorithm did indeed incorrectly flag the spikes as RFI. However, as mentioned above, we can increase M by decreasing N (i.e., decreasing spectral resolution), which is likely to be fruitful, since the spikes have an intrinsic spectral width of several MHz. Figure 12 displays the dynamic spectrum obtained for the same data set, but with a frequency resolution of 3.9 MHz. As expected, most of the individual spikes are still resolved, while some of the weaker RFI spectral lines present in Figure 11 are no longer visible, due to the fact that their S/N has been decreased by increasing the width of the frequency bins. While the temporal resolution of the dynamic spectrum remains the same, each of the individual spectrograms is now obtained from $M = 390$ contiguous PSD estimates, which according to equation (27) should be sufficient to discriminate continuous RFI with a S/N as low as ~1.3.

Figure 13 presents the results of the RFI excision algorithm applied to the data displayed in Figure 12, with no frequency integration. Remarkably, almost all of the natural radio spikes survived the application of the $\hat{V}_k^2$ discriminator. This indicates that in the limit of the 390 PSD estimates available, a Gaussian distribution is at least a good approximation for the probability distribution function of spike intensity. On the other hand, two of the three adjacent RFI spectral lines have been completely removed, as have a few other bright features believed to be RFI. The low-frequency edge of the main RFI-affected bins is not removed, possibly due to a too low S/N or a duty cycle near 50%. The random removal of spectral bins is about equally numerous in both the RCP and LCP sections and is consistent with the 1% removal expected for a 3 σ discriminator.

Therefore, we can conclude that the implementation of an SK-based RFI excision algorithm in a real-time instrument, which can be designed to ensure sufficient statistics in the 20 ms lifetime of a spike, would be unlikely to trigger more than statistically expected false RFI alarms, even in the presence of intense and highly fluctuating natural radio emission.

4.5. Rapidly Varying Gaussian Transients

The success described in the previous section may be artificial, since the natural radio spikes do not vary over the 100 μs of contiguous data, yet on longer timescales (100 $\mu s < \tau < 20$ ms) they must rise from a low background to some peak level and then subside. To quantify the SK estimator sensitivity to such natural transient signals that obey Gaussian statistics, we follow similar steps to those in § 3.2.2, considering this time that a stationary Gaussian background, characterized by a constant variance σ^2, is mixed for R out of the M analyzed data blocks, with a narrowband Gaussian transient signal characterized by the variance $\eta_k\sigma^2$, where the ratio of the two variances, η_k, is still referred as a signal-to-noise ratio. Under these considerations, the power sums entering the SK estimator can be ex-

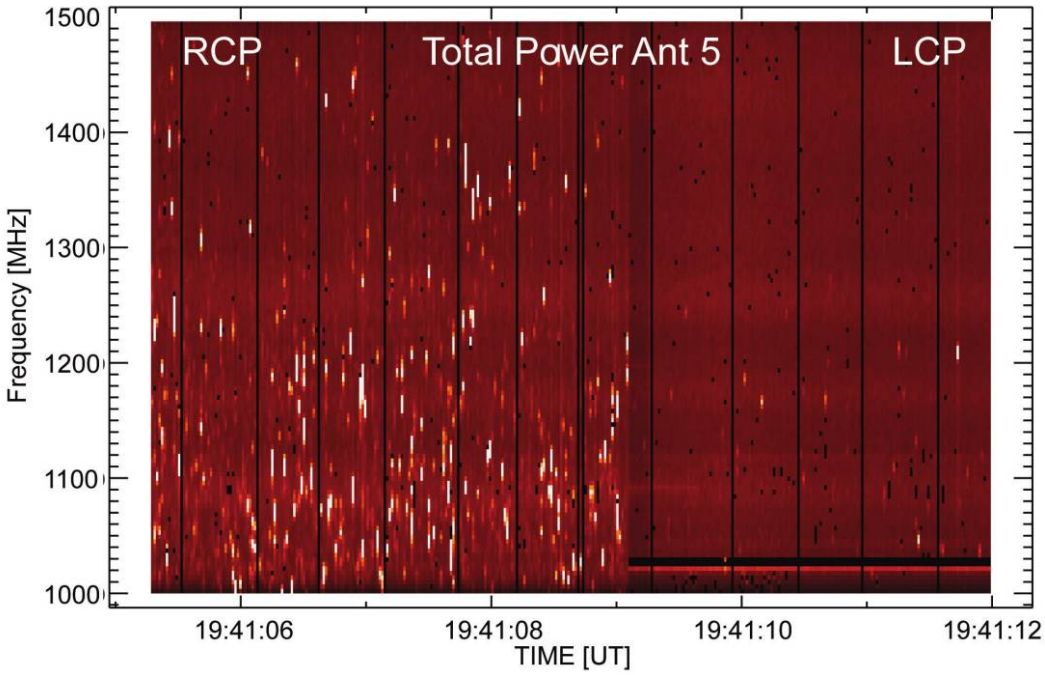

FIG. 13.—Results of the RFI excision algorithm applied to the data displayed in Fig. 12. Since no frequency integration has been performed, the RFI-flagged spectrogram bins appear black in the plot. Remarkably, the SK estimator proves to be insensitive to the natural spiky emission, while it successfully detects most of the continuous RFI spectral line at about 1020 MHz. The spectrogram also reveals a few scattered bins flagged by the algorithm, most likely false RFI alarms. However, the total data losses, ~0.7% in the RCP section, are close to the ~1% losses due to statistical fluctuations expected for the $\pm 3\,\sigma$ RFI detection threshold used.

pressed as

$$\langle S_1 \rangle = 2R(1+\eta_k)\sigma^2 + 2(M-R)\sigma^2,$$
$$\langle S_2 \rangle = 8R(1+\eta_k)^2\sigma^4 + 8(M-R)\sigma^4, \qquad (32)$$

and subsequently

$$\langle \hat{V}_k^2 \rangle = \frac{M}{M-1}\left[1 + \frac{2(M/R-1)\eta_k^2}{(M/R+\eta_k)^2}\right], \qquad (33)$$

which is similar in form to the expected SK estimator for a transient RFI signal given by equation (30).

Figure 14 displays the expected SK estimator as given by equation (33), along with the corresponding simulated results, for the same S/N values as in Figure 2. Comparing these two figures, we can conclude that the SK estimator is indeed highly sensitive not only to RFI transients, but also to Gaussian transients, the only difference being that the SK estimator in the case of a Gaussian transient is always greater than unity for any particular duty cycle other than exactly 0% or 100%. Thus, if we had longer contiguous data for the event in Figure 11, so that the Gaussian noise level varied significantly over the accumulation period, it is likely that the algorithm would have "failed" and the spikes would be flagged as RFI. This effectively limits the accumulation time to times shorter than the intrinsic variation of the natural emission, which is normally not a problem, except for such extreme events as solar spike bursts. Whether the difference in behavior demonstrated in Figures 2 and 14 can be exploited for reliable discrimination between RFI and Gaussian transients is an issue that may be worth addressing in a future study.

5. CONCLUSION

The statistical properties of the PSD estimates obtained in the spectral domain (e.g., via FFT) have been analyzed, and an RFI excision algorithm suitable for real–time implementation in radio telescopes has been proposed. However, the theoretical analysis is general and should be directly adaptable for many applications in addition to RFI mitigation in radio astronomy. The main findings are as follows:

1. There exists a discriminator that we refer to as the spectral kurtosis, which for the special case of Gaussian noise has the value $V_k^2 = 1 + |W_{2k}|^2$, where W_{2k} is given by equation (14).

2. Using sums S_1 and S_2, which are sums of the first two powers of M PSD estimates, we define a discriminator for Gaussian noise that we call the spectral kurtosis estimator, $\hat{V}_k^2 = M(MS_2/S_1^2 - 1)/(M-1)$.

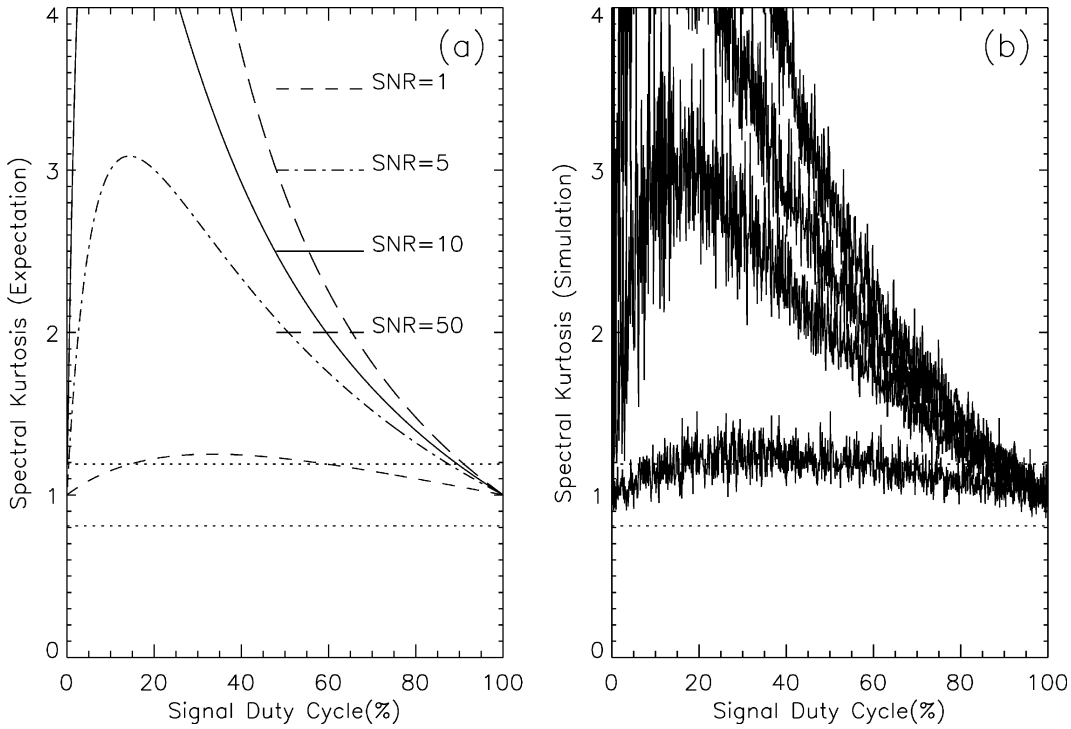

FIG. 14.—SK estimator vs. duty cycle of a transient Gaussian signal for different S/Ns. (*a*) Theoretical expectation for S/Ns ranging from 1 to 50. The horizontal dotted lines represent the $\pm 3\ \sigma$ detection thresholds for $M = 1000$ PSD estimates. (*b*) Simulated SK results corresponding to the theoretical parameters presented in (*a*).

3. The value of this SK estimator is unity in most spectral bins, with a variance given by equation (22), which for practical purposes (see § 3) reduces to

$$\mathrm{Var}\left(\hat{V}_k^2\right) \simeq \begin{cases} 24/M, & k = 0, N/2, \\ 4/M, & k = 1, \ \ldots\ , (N/2 - 1). \end{cases} \tag{34}$$

The variance in the DC and Nyquist bins ($k = 0, N/2$) is the well-known value $24/M$ obtained from time-domain analysis (e.g., Ruf et al. 2006), which unifies the time- and spectral-domain approaches.

Using the above expectation values of the estimator and its variance, one can determine whether any time-domain signal obeys Gaussian statistics by performing a Fourier transform, accumulating sums, and examining whether the result in any frequency bin k agrees with the expected value within a predefined deviation [e.g., $3\ \sigma = 3(4/M)^{1/2}$]. However, the extent to which the presence of non-Gaussian signals can be distinguished from noise depends on both the S/N η_k and the duty cycle $d = R/M$ of the signal, according to equation (30). In particular, the algorithm is "blind" to RFI with a duty cycle near 50%, as has been noted previously for the case of time-domain kurtosis (De Roo et al. 2007). In fact, we have shown that spectral kurtosis and time-domain kurtosis are completely equivalent, with the same strengths and weaknesses. Hence, either can be used interchangeably, with the choice being driven by whether spectral channels are derived from an FFT-based approach (where SK would be used) or from a finite impulse response (FIR) or polyphase filter approach (where time-domain kurtosis would be used).

The above results have been applied to the problem of RFI identification and excision using full-resolution time-domain data recorded by the FST instrument (Liu et al. 2007) during RFI surveys and solar observations. Taking into account the limitation on the number of PSD estimates allowed by the FST instrument, we consider the results to be a successful experimental validation of the use of the SK estimator for RFI excision. Moreover, the algorithm appears to be most effective in the case of transient, low–duty-cycle RFI. If it proves to be less effective for channels with continuous RFI, such channels are easy to detect from simple power-level considerations, and it is always possible to flag them using a predefined static frequency mask.

We have also demonstrated the robustness of the SK estimator for identifying RFI in conditions of considerable temporal and spectral complexity presented by solar radio bursts. The algorithm successfully identifies the non-Gaussian artificial signals for flagging while passing the natural solar radio emission, which recommends it as a promising and valuable tool to be considered for a real-time implementation in future radio

instruments, such as FASR. However, Gaussian transients such as spike bursts may limit the time window over which to gather statistics. We agree with Fridman & Baan (2001) that no single RFI excision technique will be 100% successful in identifying all possible artificial signals, but the SK estimator is a powerful tool when used in combination with other techniques.

We thank the anonymous referee for pointing out the effect of Gaussian transients, and for other comments that improved the paper. This work is supported by NSF grant AST 03-52915 to the New Jersey Institute of Technology. We acknowledge additional support for the Owens Valley Solar Array through NSF grant AST 06-07544 and NASA grant NNG06GJ40G.

APPENDIX A

SPECTRAL VARIABILITY OF A WINDOWED TIME-DOMAIN GAUSSIAN SIGNAL

To derive the spectral variability of a windowed Gaussian signal, we start by conveniently rewriting equation (5) as

$$\hat{P}_k = A_k^2 + B_k^2, \tag{A1}$$

where

$$A_k = \sqrt{\frac{2}{N \sum w_n^2}} \sum_{n=0}^{N-1} w_n x_n \cos\left(\frac{2\pi kn}{N}\right),$$

$$B_k = \sqrt{\frac{2}{N \sum w_n^2}} \sum_{n=0}^{N-1} w_n x_n \sin\left(\frac{2\pi kn}{N}\right). \tag{A2}$$

Since A_k and B_k are linear combinations of zero-mean Gaussian variables, they are also random variables belonging to normal parent distributions of zero mean, and have variances $\sigma_{A_k}^2$ and $\sigma_{B_k}^2$ that can be *exactly* computed from the formula, generally providing only the first-order approximation for the variance of a function of N random variables,

$$\sigma^2_{g(x_1,\ \ldots,\ x_n)} = \sum_{i,j=0}^{N-1} \frac{\partial g}{\partial m_i} \frac{\partial g}{\partial m_j} \mathrm{Cov}\,(x_i, x_j) \tag{A3}$$

(Kendall & Stuart 1958, p. 232), where $\partial g/\partial m_i$ are the partial derivatives evaluated at the mean values $m_i = \langle x_i \rangle$, and $\mathrm{Cov}\,(x_i, x_j)$ is the covariance of the two random variables, which is zero for independent variables (as is the case here) and reduces to the variance σ_x^2 for $i = j$.

After some basic mathematical manipulations, equation (A3) leads to remarkably simple results:

$$\sigma_{A_k}^2 = [1 + \Re(W_{2k})] \frac{\sigma_x^2}{N}, \tag{A4}$$

$$\sigma_{B_k}^2 = [1 - \Re(W_{2k})] \frac{\sigma_x^2}{N},$$

where

$$\Re(W_{2k}) = \sum_{n=0}^{N-1} w_n^2 \cos\left(\frac{4\pi kn}{N}\right) \Big/ \sum_{n=0}^{N-1} w_n^2 \tag{A5}$$

is the real part of the normalized Fourier transform of the squared time-domain window,

$$W_k = \frac{1}{\sum w_n^2} \sum_{n=0}^{N-1} w_n^2 e^{-2\pi ikn/N}, \tag{A6}$$

evaluated at the discreet frequencies $f_{2k} = 2f_k$.

Note that by definition, $W_0 = W_N = 1$ for any windowing function, which implies that for the DC and Nyquist frequencies, $\sigma_{A_0}^2 = \sigma_{A_{N/2}}^2 = 2\sigma_x^2/N$. On the other hand, since A_k and B_k are zero-mean random variables, it immediately follows that $\langle A_k^2 \rangle = \sigma_{A_k}^2$ and $\langle B_k^2 \rangle = \sigma_{B_k}^2$; hence, the expected mean of the PSD estimate, $\mu_k \equiv \langle \hat{P}_k \rangle$, is given at all frequencies bins by

$$\mu_k = \langle A_k^2 \rangle + \langle B_k^2 \rangle = \frac{2\sigma_x^2}{N},$$

$$k = 0, 1, \ldots, N/2, \tag{A7}$$

which indicates that the power of the Gaussian signal (white noise) is, as expected, equally distributed among the frequency bins of the PSD estimate. However, contrary to the case of a square windowing function (generally an unstated assumption), equations (A4) and (A5) reveal that the variances of the real and imaginary parts of the PSD estimate are not generally equal. This means that at any given frequency bin, even a symmetrical time-domain window may produce an asymmetric distribution of the spectral power between the real and imaginary parts of the PSD estimate in a way that is dependent on the Fourier transform of the normalized, squared time-domain window $w_n^2/\sum w_n^2$.

Since for the particular DC and Nyquist frequency bins we have $B_0 = B_{N/2} = 0$, the PSD probability distribution function associated with these frequency bins can be derived directly from the Gaussian distribution associated with the

corresponding A_k coefficients,

$$p(A_k) = \frac{1}{\sqrt{2\pi}\sigma_{A_k}} \exp\left(-\frac{A_k^2}{2\sigma_{A_k}^2}\right), \tag{A8}$$

which after the variable change $A_k^2 \rightarrow \hat{P}_k$, becomes

$$p\left(\hat{P}_k\right) = \frac{1}{\sqrt{2\pi}\sigma_{A_k}} \hat{P}_k^{-1/2} \exp\left(-\frac{1}{2\sigma_{A_k}^2}\hat{P}_k\right), \quad k = 0, N/2. \tag{A9}$$

Equation (A9) can be used to compute any nth-order moment about the origin of the PSD estimate, $\mu_k'(n) \equiv \langle P_k^n \rangle$, which is given by

$$\mu_k'(n) = 1 \times 3 \ \ldots \times (2n-1)\sigma_{A_k}^{2n}, \quad k = 0, N/2, \tag{A10}$$

and the variance of the PSD estimate, $\sigma_k^2 = \mu_k'(2) - \mu_k'^2(1)$, at the DC and Nyquist frequency bins, as

$$\sigma_0^2 = \sigma_{N/2}^2 = \frac{8\sigma_x^4}{N^2}. \tag{A11}$$

Hence, taking into consideration equations (A7) and (A11), the spectral variabilities corresponding to the DC and Nyquist frequency bins turn out to be

$$V_0^2 = V_{N/2}^2 = 2, \tag{A12}$$

which is a generally valid result, independent of the actual shape of the time-domain windowing function.

However, for all other frequency bins, a complete statistical description of the PSD estimate cannot be obtained without evaluating the covariances of the A_k and B_k coefficients. Again, the linearity of these coefficients assures us that the general formula

$$\mathrm{Cov}\,(f, g) = \sum_{i,j=0}^{N-1} \frac{\partial f}{\partial m_i}\frac{\partial g}{\partial m_j}\mathrm{Cov}\,(x_i, x_j), \tag{A13}$$

provided by Kendall & Stuart (1958, p. 232) only as a first-order approximation, leads in this case to an exact result. Recalling the assumed zero correlation of the time-domain data samples $\{x_n\}$, we get

$$\sigma_{A_kB_k}^2 \equiv \mathrm{Cov}\,(A_k, B_k) = \Im(W_{2k})\frac{\sigma_x^2}{N}, \tag{A14}$$

where

$$\Im(W_{2k}) = \sum_{n=0}^{N-1} w_n^2 \sin\left(\frac{4\pi kn}{N}\right) \Big/ \sum_{n=0}^{N-1} w_n^2 \tag{A15}$$

is the imaginary part of the normalized Fourier transform of the squared time-domain window evaluated at the discreet frequencies $f_{2k} = 2f_k$. In the case of a symmetric time-domain window, it immediately follows that $\Im(W_{2k}) = 0$, which implies the statistical independence of A_k and B_k. However, since for the reasons mentioned in § 2.1 we do not want to impose any constraints on the time-domain window, we notice that according to equations (A14) and (A15), any asymmetric time-domain window would induce a statistical correlation of the DFT coefficients that has to be taken into consideration when computing the variance of the PSD estimate.

We have established so far that in the most general case, the DFT coefficients A_k and B_k corresponding to a given frequency bin are Gaussian random variables of zero mean and known variances. However, since for $k \neq 0, N/2$ the PSD estimate is a quadratic combination of two random variables that may be statistically correlated, its expected statistical moments should be derived from the joint distribution of two dependent Gaussian variables, which is

$$p(A_k, B_k) = \exp\left[\frac{1}{1-\rho_k^2}\left(-\frac{A_k^2}{2\sigma_{A_k}^2} - \frac{B_k^2}{2\sigma_{B_k}^2} + \rho_k\frac{A_kB_k}{\sigma_{A_k}\sigma_{B_k}}\right)\right] \times \left(2\pi\sigma_{A_k}\sigma_{B_k}\sqrt{1-\rho_k^2}\right)^{-1} \tag{A16}$$

(Davenport & Root 1958), where

$$\rho_k = \frac{\sigma_{A_kB_k}^2}{\sigma_{A_k}\sigma_{B_k}} \tag{A17}$$

is the correlation coefficient of the DFT coefficients.

The evaluation of the PSD variance in the most general case described by equation (A16) can be greatly simplified by using the moment-generating function

$$M(t) = \left\langle e^{(A_k^2+B_k^2)t}\right\rangle = \frac{1}{\sqrt{1-2(\sigma_{A_k}^2+\sigma_{B_k}^2)t + 4\sigma_{A_k}^2\sigma_{B_k}^2(1-\rho_k^2)t^2}}, \tag{A18}$$

which can be straightforwardly employed to compute the nth-

order moments about the origin,

$$\mu'_k(n) = \frac{d^n}{dt^n}\Big|_{t=0} M(t). \tag{A19}$$

Hence, the variance of the PSD estimate for $k \neq 0, N/2$ is

$$\sigma_k^2 = 2(\sigma_{A_k}^4 + \sigma_{B_k}^4 + 2\sigma_{A_kB_k}^4), \tag{A20}$$

which can be formally reduced to equation (A11) for $\sigma_{B_0}^2 = \sigma_{B_{N/2}}^2 = 0$, $\sigma_{A_0B_0}^2 = \sigma_{A_{N/2}B_{N/2}}^2 = 0$, and $\sigma_{A_0}^2 = \sigma_{A_{N/2}}^2 = 2\sigma_x^2/N$.

Therefore, using equations (A7) and (A20), as well as equations (A4) and (A14), the PSD spectral variability, as defined by equation (1), can be now expressed as

$$V_k^2 = \frac{2\,(\sigma_{A_k}^4 + \sigma_{B_k}^4 + 2\sigma_{A_kB_k}^2)}{(\sigma_{A_k}^2 + \sigma_{B_k}^2)^2}, \tag{A21}$$

which, making use of equations (A4), (A5), and (A15), can be rewritten in a remarkably simple, compact form as

$$V_k^2 = 1 + |W_{2k}|^2. \tag{A22}$$

Taking into consideration equation (A12), and that $W_0 = W_N = 1$, this is a result that is valid for all $k = 0, 1, \ldots N/2$.

APPENDIX B

VARIANCE OF THE SK ESTIMATOR

We rewrite equation (21) in terms of the first two sample moments about the origin of the random variable $\hat{P}_{ki}$; i.e., $m'_1 = S_1/M$ and $m'_2 = S_2/M$, which leads to

$$\hat{V}_k^2 = \frac{M}{M-1}\left(\frac{m'_2}{m'^2_1} - 1\right). \tag{B1}$$

Hence, from equation (A3), we obtain

$$\mathrm{Var}\left(\hat{V}_k^2\right) = \left(\frac{M}{M-1}\right)^2 \mathrm{Var}\left(\frac{m'_2}{m'^2_1}\right), \tag{B2}$$

which translates the problem of determining the variance of $\hat{V}_k^2$ to the more general problem of determining the variance of the ratio m'_2/m'^2_1 of any statistical distribution. Taking into account the fact that the expectation values of any sample moment of order q, m'_q, is the parent population moment of the same order, μ'_q, we can again employ equation (A3) to obtain

$$\mathrm{Var}\left(\frac{m'_2}{m'^2_1}\right) = \frac{4\mu'^2_2}{\mu'^6_1}\mathrm{Var}\,(m'_1) - \frac{4\mu'_2}{\mu'^5_1}\mathrm{Cov}\,(m'_1, m'_2) + \frac{1}{\mu'^4_1}\mathrm{Var}\,(m'_2). \tag{B3}$$

The exact variances and covariances entering equation (B3) are given by Kendall & Stuart (1958, p. 229) for any distribution as

$$\mathrm{Cov}\,(m'_q, m'_r) = \frac{1}{M}(\mu'_{q+r} - \mu'_q\mu'_r), \tag{B4}$$

where the covariance of any moment with itself should be understood as its variance. Using these relationships, equation (B3) becomes

$$\mathrm{Var}\left(\frac{m'_2}{m'^2_1}\right) = \frac{1}{M}\left\{\frac{\mu'_4 - \mu'^2_2}{\mu'^4_1} - \frac{4\mu'_2\mu'_3}{\mu'^5_1} + \frac{4\mu'^3_2}{\mu'^6_1}\right\}. \tag{B5}$$

For $k = 0, N/2$, equation (A10) leads to

$$\mathrm{Var}\left(\frac{m'_2}{m'^2_1}\right) = \frac{24}{M}, \tag{B6}$$

which, as expected, is the same as the variance of the classical kurtosis parameter computed from the fourth and second moments of the time-domain amplitude (Kendall & Stuart 1958).

For $k \neq 0, N/2$, equations (A19) and (A4) lead to

$$\mathrm{Var}\left(\frac{m'_2}{m'^2_1}\right) = \frac{4}{M}(1 + 5|W_{2k}|^2 - |W_{2k}|^4 + |W_{2k}|^6), \tag{B7}$$

which becomes equation (B6) for $k = 0, N/2$, since $|W_0| = |W_N| = 1$.

Therefore, the variance of the SK estimator (eq. [18]) is given at all PSD frequency bins by the unique expression

$$\mathrm{Var}\left(\hat{V}_k^2\right) = \frac{4M}{(M-1)^2} \times (1 + 5|W_{2k}|^2 - |W_{2k}|^4 + |W_{2k}|^6). \tag{B8}$$

REFERENCES

Amblard, P. O., Gaeta, M., & Lacoume, J.-L. 1996a, Signal Processing, 53, 1

———. 1996b, Signal Processing, 53, 15

Antoni, J. 2006, Mechanical Systems and Signal Processing, 20, 282

———. 2007, Mechanical Systems and Signal Processing 21, 108

Bastian, T. S. 2003, Proc. SPIE, 4853, 98

Davenport, W. R., & Root, W. L. 1958, An Introduction to Random Signals and Noise (New York: McGraw-Hill)

De Roo, R. D., Misra, S., & Ruf, C. S. 2007, IEEE Trans. Geoscience and Remote Sensing, 45, 1938

Dwyer, R. F. 1983, in Proc. IEEE Int. Conf. on Acoustic, Speech, and Signal Processing, Vol. 8 (New York: IEEE), 607

Fleishman, G. D., Gary, D. E., & Nita, G. M. 2003, ApJ, 593, 571

Fridman, P. A. 2001, A&A, 368, 369

Fridman, P. A., & Baan, W. A. 2001, A&A, 378, 327

Gary, D. E., & Hurford, G. J. 1999, in Proc. Nobeyama Symposium, ed. T. S. Bastian, N. Gopalswamy, & K. Shibasaki (NRO Rep. 479; Nagano: Nobeyama Radio Obs.), 429

Kendall, M. G., & Stuart, A. 1958, The Advanced Theory of Statistics, Vol. 1 (London: Griffin)

Kenney, J. F., & Keeping, E. S. 1962, Mathematics of Statistics (3rd. ed.; Princeton, NJ: Van Nostrand)

Liu, Z., Gary, D. E., Nita, G. M., White, S. M., & Hurford, G. J. 2007, PASP, 119, 303

McDonough, R. N., & Whalen, A. D. 1995, Detection of Signals in Noise (New York: Academic Press)

Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. 1989, Numerical Recipies in PASCAL (Cambridge: Cambridge Univ. Press)

Ruf, C. S., Gross, S. M., & Misra, S. 2006, IEEE Trans. Geoscience and Remote Sensing, 44, 694

Vrabie, V. D., Granjon, P., Maroni, C. S., & Leprettre, B. 2004, in 5th Int. Conf. on Acoustical and Vibratory Surveillance Methods and Diagnostic Techniques

Vrabie, V. D., Granjon, P., & Servière, C. 2003, in 2003 IEEE-EURASIP Workshop on Nonlinear Signal and Image Processing (NSIP-03)