

The Effects of Digitization on Nonstationary Stochastic Signals with Applications to Pulsar Signal Baseband Recording

F. A. JENET

California Institute of Technology, Space Radiation Laboratory, Mail Code 206-49 Pasadena, CA 91125; merlyn@srl.caltech.edu

AND

S. B. ANDERSON

California Institute of Technology, Space Radiation Laboratory, Mail Code 206-49 Pasadena, CA 91125; sba@srl.caltech.edu

Received 1998 January 18; accepted 1998 August 26

ABSTRACT. The process of digitizing a stochastic signal introduces systematic distortions into the resulting digitized data. Further processing of these data may result in the appearance of unwanted artifacts, especially when the input signal was generated by a nonstationary stochastic process. In this paper the magnitude of these distortions are calculated analytically and the results are applied to a specific example found in pulsar signal processing. A pulsar signal is an excellent example of a nonstationary stochastic process. When analyzing pulsar data, the effects of interstellar medium (ISM) dispersion must be removed by digitally filtering the received signal. The distortions introduced through the digitization process cause unwanted artifacts to appear in the final “dedispersed” signal. These artifacts are demonstrated using actual 2-bit (4-level) digitized data of the pulsar PSR B0833-45 (Vela). Techniques are introduced that simultaneously minimize these artifacts and maximize the signal-to-noise ratio of the digitized data. The distortion analysis and artifact removal techniques described in this paper hold for an arbitrary number of input digitization thresholds (i.e., number of bits). Also presented are tables of the optimum digitizer thresholds for both uniform and nonuniform input threshold digitizers.

1. INTRODUCTION

Digital signal processing (DSP) techniques are used widely in the scientific community. Such techniques represent a continuous signal as a sequence of discrete time points with quantized values. Standard analyses of DSP algorithms are primarily concerned with the discrete time aspect of the digital signal. Relatively little attention has been given to the effects of quantization since many applications use a large number of discrete amplitude levels (i.e., many bits) to reduce any digitization artifacts. However, in applications that emphasize high time resolution over dynamic range and fidelity, a small number of discrete amplitude levels are often required.

One important example of such an application is baseband signal recording in radio pulsar astronomy. This technique records a representation of the electric field vector incident upon a radio telescope. Conventional techniques record only the total power received, hence losing vital information about the source of the field and the intervening medium. As recording technology continues to become faster and cheaper, baseband recording will be the method of choice for most pulsar observations. Baseband signal recording is also used in VLBI observations.

Due to finite recording rates, there is a trade-off between bandwidth, B , and the total number of bits used to represent

the recorded signal. The bandwidth is usually maximized to increase the signal-to-noise ratio. Therefore, baseband recording systems are designed with a small number of bits. Unfortunately, digitization distortion increases as the number of bits decreases. Thus, an accurate analysis of the distortion is necessary.

The effects of digitization on nonstationary stochastic signals are analyzed in this paper. A nonstationary stochastic signal is a stochastic signal whose statistical properties change with time. Throughout this paper, the existence of a small timescale over which the statistics are stationary is assumed, thus allowing estimates of the power spectrum and total signal power at different moments in time to be made. These estimates are used to correct the distortions that are introduced by the digitization process.

The digitization process is described in § 2 along with a method for setting the “input thresholds” and the “output levels.” In § 3, the power spectrum of the digitized signal is calculated. Using this power spectrum, the magnitude of two important artifacts, which arise when removing effects of interstellar medium (ISM) dispersion from a digitized pulsar signal, are calculated in § 4. Techniques for minimizing these artifacts are also described. The fractional loss in the signal-to-noise ratio (SNR) due to the digitization process is calculated in § 5. Also presented are the optimum values for the digiti-

zation thresholds that simultaneously minimize both the SNR loss and a particular class of digitization artifacts. In § 6, the general techniques described in this paper are applied to the specific case of a 2-bit, 4-level, digitizing scheme.

2. OPTIMUM DIGITIZATION

The process of digitizing a signal involves sorting the analog input signal, $v(t)$, into a finite number of ranges, N . A digitizing scheme is defined by specifying the $N + 1$ end points or thresholds, x_k , of the N input ranges and the corresponding output levels, y_k , for each input range. Hence, when $x_k \leq v(t) < x_{k+1}$ then $\hat{v}(t) = y_k$, where $\hat{v}(t)$ is the digitized signal. In this paper, two types of level-setting schemes are discussed: a fixed output level scheme and a dynamic output level scheme. The fixed output level scheme is the standard approach to the digitization problem and it works well with statistically stationary input signals. When digitizing nonstationary signals, this method may introduce unwanted artifacts when processing the digital data. Section 4 describes such artifacts for the particular case of pulsar signal dispersion removal.

In the fixed output level scheme, input thresholds and output levels are chosen so that the "distortion" is minimized. Given the output signal $\hat{v}(t)$, the distortion, χ^2 , is defined as

$$\chi^2 \equiv \langle [f(v) - f(\hat{v})]^2 \rangle, \tag{1}$$

where angle brackets represent an ensemble average. Hence,

$$\chi^2 = \sum_{k=1}^N \int_{x_k}^{x_{k+1}} [f(v) - f(y_k)]^2 P(v) dv, \tag{2}$$

where $x_{N+1} = \infty, x_1 = -\infty, P(v)$ is the probability density function for the values of the analog signal $v(t)$, and $f(v)$ is a continuous function of v . When $f(v) = v$, the corresponding levels are called "voltage optimized" levels since the expected value of the digitized signal, $\langle \hat{v} \rangle$, equals the expected value of the undigitized signal, $\langle v \rangle$. A more important set of levels are obtained when $f(v) = v^2$. These levels are called "power optimized" levels since $\langle \hat{v}^2 \rangle = \langle v^2 \rangle$. In general, one can show from equation (6) below that $\langle f(\hat{v}) \rangle = \langle f(v) \rangle$. The thresholds and levels (i.e., x_k and y_k) are set so that χ^2 is minimized for a given $f(v)$. The necessary conditions for a minimum in χ^2 are

$$\begin{aligned} \partial_{x_k} \chi^2 &= [f(x_k) - f(y_{k-1})]^2 P(x_k) \\ &\quad - [f(x_k) - f(y_k)]^2 P(x_k) = 0 \quad (2 \leq k \leq N), \end{aligned} \tag{3}$$

$$\begin{aligned} \partial_{y_k} \chi^2 &= - \int_{x_k}^{x_{k+1}} 2 [f(v) - f(y_k)] f'(y_k) \\ &\quad \times P(v) dv = 0 \quad (1 \leq k \leq N), \end{aligned} \tag{4}$$

where $f'(v) \equiv df(v)/dv$. These conditions yield the following set of equations:

$$f(x_k) = \frac{f(y_k) + f(y_{k-1})}{2} \quad (2 \leq k \leq N), \tag{5}$$

$$f(y_k) = \frac{\int_{x_k}^{x_{k+1}} f(x) P(x) dx}{\int_{x_k}^{x_{k+1}} P(x) dx} \quad (1 \leq k \leq N). \tag{6}$$

For stationary signals, these coupled equations are solved simultaneously to obtain the optimum values for x_k and y_k . Based on a similar approach, J. Max (1976) describes a numerical algorithm for calculating these critical values and presents the results for the case of $f(v) = v$. From now on, this paper will consider only power optimized levels [$f(v) = v^2$] since the final result will be a time series of power measurements and the simplest way to obtain unbiased power measurements is to use power optimized levels.

3. THE POWER SPECTRUM OF THE DIGITIZED SIGNAL

The process of digitizing a signal distorts the signal's power spectrum. Using the fact that the power spectrum and the autocorrelation function are Fourier transform pairs, the autocorrelation function is calculated first since it is analytically easier to compute. The discrete time autocorrelation function of the original signal may be expressed as

$$\langle v(i)v(i+l) \rangle = \sigma^2 \rho(l), \tag{7}$$

where i and l are integers, $v(i) \equiv v(i\Delta t)$, $\rho(l)$ is the normalized autocorrelation function, and σ^2 is the average total power (i.e., $\langle v^2 \rangle$). Given a set of input thresholds, x_k , and output levels, y_k , σ^2 and ρ may be related to their digitized counterparts $\hat{\sigma}^2$ and $\hat{\rho}$. The digitized total power is given by

$$\hat{\sigma}^2 = \sum_{k=1}^N y_k^2 \int_{x_k}^{x_{k+1}} P(v, \sigma^2) dv, \tag{8}$$

where $P(v, \sigma^2)$ is the probability density for the values of v explicitly parameterized by the undigitized total power σ^2 . The above equation yields a functional relationship between $\hat{\sigma}^2$ and σ^2 . For example, it was shown in § 2 for the case of power optimized levels that $\hat{\sigma}^2 = \sigma^2$.

Calculation of the normalized autocorrelation function, $\hat{\rho}(l)$, for the digitized signal requires a knowledge of $P(v(i), v(i+l), \rho(l))$, the joint probability distribution for $v(i)$ and $v(i+l)$, which is parameterized by the undigitized autocorrelation function $\rho(l)$. By defining $v \equiv v(i)$ and

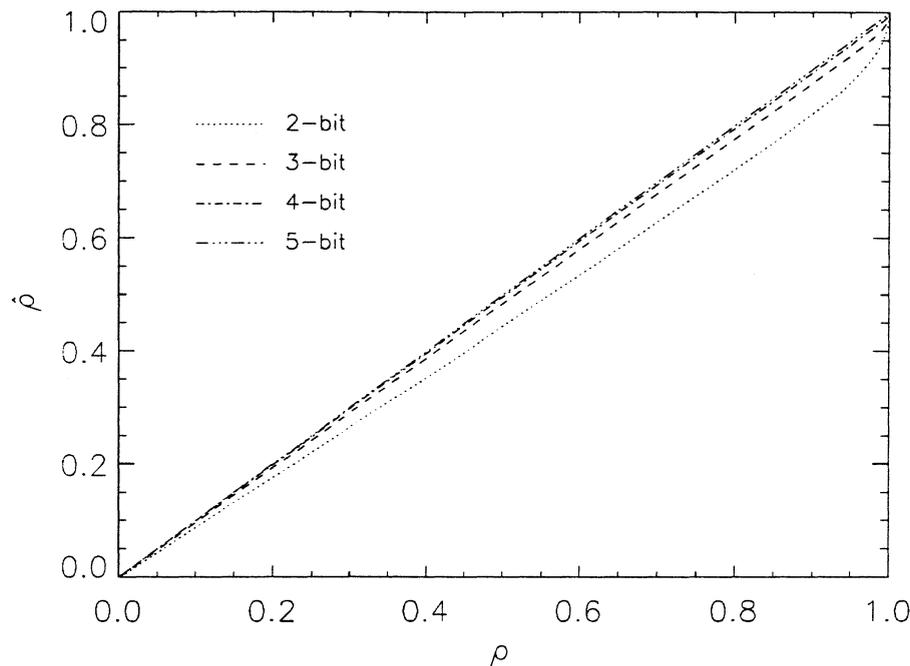


FIG. 1.—Plot of the autocorrelation function of a digitized signal, $\hat{\rho}$, vs. the signal's undigitized autocorrelation function, ρ , for 2-, 3-, 4-, and 5-bit systems. Voltage optimized levels were used along with a Gaussian-distributed input signal. The corresponding plot for power optimized levels is indistinguishable from this one.

$v' \equiv v(i + l)$, it follows that

$$\hat{\rho}(l) = \frac{\langle \hat{v}(i)\hat{v}(i + l) \rangle}{\langle \hat{v}^2 \rangle} \tag{9}$$

$$= \frac{1}{\hat{\sigma}^2} \sum_{k=1}^N \sum_{j=1}^N \int_{x_k}^{x_{k+1}} \int_{x_j}^{x_{j+1}} y_k y_j P(v, v', \rho(l)) dv dv'. \tag{10}$$

Given $\rho(l)$, equation (10) produces the corresponding $\hat{\rho}(l)$. From this, one obtains $\hat{\rho}(\rho)$. Alternatively, one can use simulated random data to numerically calculate $\hat{\rho}(\rho)$ for a given joint probability distribution function.

Consider a stationary Gaussian noise process with a normalized autocorrelation function, $\rho(l)$. The joint probability distribution function is given by

$$P[v(i), v(i + l), \rho(l)] \tag{11}$$

$$= \exp \left(- \frac{v(i)^2 + v(i + l)^2 - 2\rho(l)v(i)v(i + l)}{2[1 - \rho(l)^2]\sigma^2} \right)$$

$$\times [2\pi\sigma^2\sqrt{1 - \rho(l)^2}]^{-1}.$$

Using equations (10) and (11), Cooper (1970) worked out $\hat{\rho}(\rho)$ analytically for a 4-level (2-bit) digitizer. Since the analytical method becomes cumbersome for a large number of levels, a Monte Carlo technique is used to analyze 8-, 16-, and 32-level systems (3-, 4-, and 5-bit systems, respectively). Figure 1 shows $\hat{\rho}$ versus ρ for each such system. In all cases $\hat{\rho}(\rho)$ is extremely linear until $\rho \approx 0.9$. At this point, the curves quickly rise toward $\hat{\rho}(1) = 1$. This property of $\hat{\rho}(\rho)$ is utilized in order to calculate the power spectrum of the digitized signal.

In order to calculate the digitized signal power spectrum, $\hat{\rho}(\rho)$ is approximated by the following form:

$$\hat{\rho}(l) = A\rho(l) + (1 - A)\delta(1 - \rho(l)). \tag{12}$$

Here

$$\delta(x) = \begin{cases} 1, & x = 0, \\ 0, & x \neq 0, \end{cases} \tag{13}$$

and

$$A = \left. \frac{d\hat{\rho}}{d\rho} \right|_{\rho=0}. \tag{14}$$

Figure 2 plots both the exact $\hat{\rho}(\rho)$ and approximate form of $\hat{\rho}(\rho)$ for a 4-level system with Gaussian statistics. For systems with more levels, the agreement is even better. In practice, the

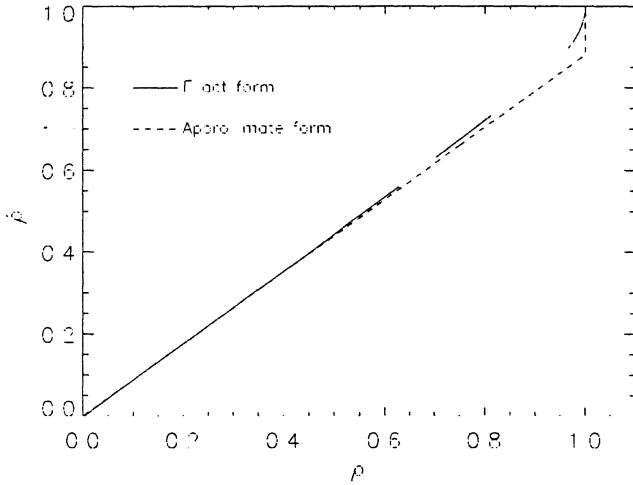


FIG. 2.—Comparison of the exact form of $\hat{\rho}(\rho)$ with its approximate form (see eq. [12]) for a 2-bit voltage optimized scheme with a Gaussian input signal. The corresponding plot for power optimized levels is indistinguishable from this one.

normalized autocorrelation function of the undigitized data rarely goes above 0.5 except at $l = 0$, where it is 1 by definition. Hence, this first-order approximation is sufficient for almost any conceivable real data.

Using the approximate form of $\hat{\rho}(\rho)$, the power spectrum of the digitized signal is given by

$$\hat{P}(\omega_k) = \hat{\sigma}^2 \frac{1}{M} \sum_{l=0}^{M-1} e^{i\omega_k l} \hat{\rho}(l) \quad (15)$$

$$= \hat{\sigma}^2 \frac{1}{M} \sum_{l=0}^{M-1} e^{i\omega_k l} \{A\rho(l) + (1-A)\delta(1-\rho(l))\} \quad (16)$$

$$= \hat{\sigma}^2 A \left[\frac{1}{M} \sum_{l=0}^{M-1} e^{i\omega_k l} \rho(l) \right] + (1-A) \frac{\hat{\sigma}^2}{M} \quad (17)$$

$$= A \frac{\hat{\sigma}^2}{\sigma^2} P(\omega_k) + (1-A) \frac{\hat{\sigma}^2}{M}, \quad (18)$$

where equation (12) is substituted into equation (15) to obtain equation (16) and it is assumed that $\rho(l) = 1$ only at $l = 0$ to arrive at equation (17). Here $\omega_k = 2\pi k/M$, M is the number of time samples in the data segment, and $i \equiv \sqrt{-1}$. In order to keep the mathematics as clear and concise as possible, equation (18) is derived for the case of real sampled data. One can show that equation (18) also holds for the case of ideal complex sampled data.

4. PULSAR SIGNAL DISPERSION REMOVAL

High time resolution observations of radio pulsars have revealed that a pulsar signal may be modeled as amplitude modulated noise (Rickett 1975; Cordes 1975; Jenet et al. 1998). Hence, a pulsar signal is a good example of a nonstationary noise process. As the pulsar signal travels through the ISM,

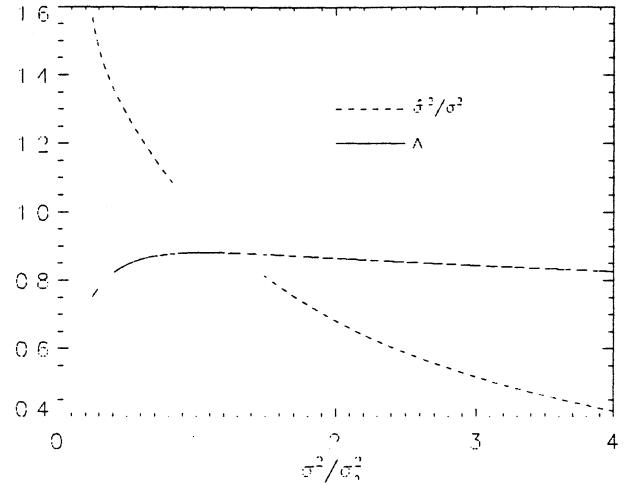


FIG. 3.—(Solid line) Plot of A vs. the normalized total power $\hat{\sigma}^2/\sigma^2$, where σ^2 is the background noise level. (Dashed line) Plot of the digitized total power normalized to the total power level vs. the normalized total power (i.e., $\hat{\sigma}^2/\sigma^2$ vs. σ^2/σ_n^2). Four level (2-bit) power optimized input thresholds and output levels were used to calculate these graphs. Notice that when the power level changes from σ_n^2 to $2\sigma_n^2$, A changes by only 3%, while $\hat{\sigma}^2/\sigma^2$ changes by 37%. Thus, A is relatively insensitive to total power fluctuations.

higher frequencies propagate faster than lower frequencies. Hence, the radio signal “sweeps” through the observed bandwidth, arriving at higher frequencies first. In order to obtain a signal that best represents the pulsar signal before propagation through the ISM, this dispersive effect must be removed by filtering the data.

4.1. Digitization Artifacts

When one digitizes the signal before removing the dispersive effects, unwanted systematic artifacts are introduced into the data. Figure 4a displays an average pulse profile of the Vela pulsar (PSR B0833-45). The data were complex sampled and digitized with 2-bits using thresholds of $(-\sigma, 0, \sigma)$ and power optimized fixed output levels. Approximately 200 individual pulses were averaged together. Pulse-to-pulse variations in amplitude and temporal structure were large for this segment of data. The total power level in the dispersed data remained below approximately twice the background noise power level for the entire stretch of data. The “negative” dips to either side of the average pulse profile are due to the first term in equation (18): $A\hat{\sigma}^2 P(\omega_k)/\sigma^2$. Both A and $\hat{\sigma}^2$ are nonlinear functions of the total undigitized power σ^2 . As the pulsar signal enters the band of interest, σ^2 increases. If the thresholds and output levels are kept constant, $A\hat{\sigma}^2/\sigma^2$ decreases. This decrease is mainly due to $\hat{\sigma}^2$ since A is relatively insensitive to total power fluctuations (see Fig. 3). Hence, the digitized signal underestimates the total power in the undigitized signal. When $A\hat{\sigma}^2/\sigma^2$ decreases, it can be seen from equation (18) that the power at all frequencies decreases. Hence, the decrease will occur even at frequencies where the pulsar signal is not present. Figure 4b shows a gray-scale plot of the frequency structure of the dispersed signal that

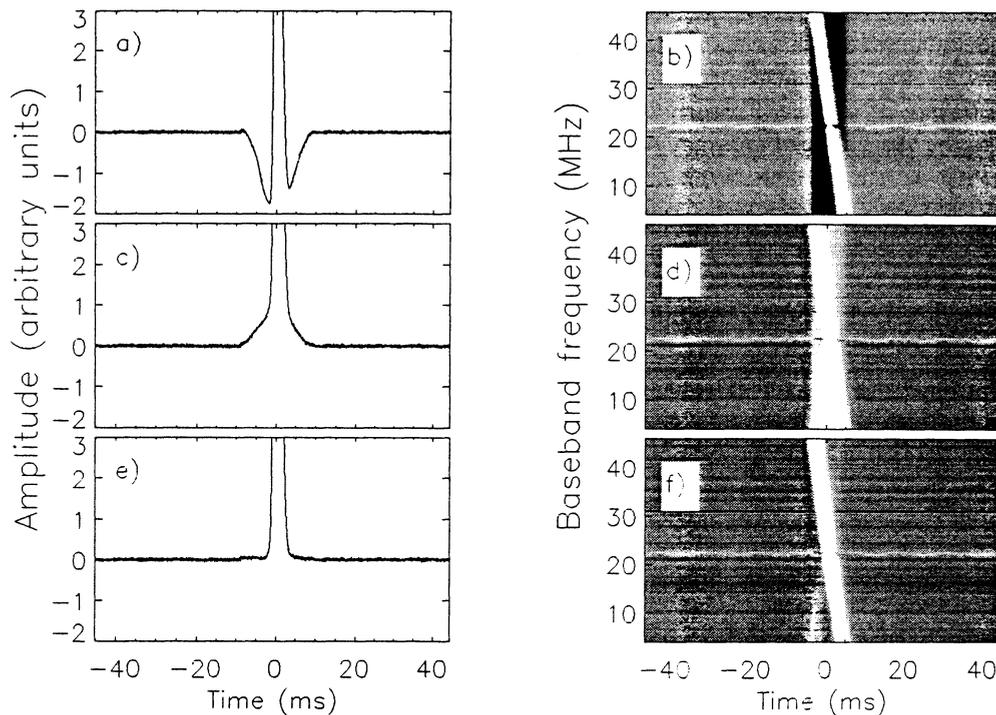


FIG. 4.—Average pulse profiles of the Vela pulsar (PSR B0833–45) and the corresponding gray-scale images of the average pulse frequency structure. The pulse profiles are magnified to show the digitization artifacts. The gray-scale images are all plotted using the same color stretch. The figures were calculated using power optimized output levels and (a, b) no dynamic level setting or scattered power correction, (c, d) dynamic level setting only, or (e, f) dynamic level setting and the scattered power correction. The data were complex sampled and digitized with 2-bits using thresholds of $(-\sigma, 0, \sigma)$. The narrowband artifacts at 22 MHz in the gray-scale images are due to external interference. The remaining artifacts in (e, f) are due to nonideal image rejection in the complex sampling process prior to digitization.

was used to calculate the folded profile shown in Figure 4a. The background noise level is seen at $t < -4$ ms and $t > 5$ ms. At $t = -4$ ms, the pulse enters the band, hence reducing the factor $A\hat{\sigma}^2/\sigma^2$. Notice that the power at frequencies where the pulse is not present is reduced below the off-pulse noise level. When these channels are added together so as to account for dispersion, the power just before and after the pulse will be underestimated, thus creating the dips seen in Figure 4a.

Since A is insensitive to variations in σ^2 , the problem of power underestimation may be resolved by finding a level-setting scheme such that $\hat{\sigma}^2 = \sigma^2$. In the next section, a “dynamic” level-setting scheme in which $\hat{\sigma}^2 \approx \sigma^2$ is described. This will effectively remove the off-pulse dips. Once these dips are removed, the next-order artifact that becomes apparent in the average pulse profile is a shallow increase in power on either side of the pulse profile (see Fig. 4c). This artifact is caused by what is known as the “quantization noise” power. As the pulsar signal enters the frequency band of interest, a fraction of this increased power is “scattered” uniformly across the entire band as can be seen in Figure 4d. The magnitude of the scattered power is given by the second term in equation (18): $(1 - A)\hat{\sigma}^2/M$.

The magnitude of the dedispersed profile distortions can be calculated using equation (18). The total digitized power far

from the pulse is given by

$$\hat{P}_{\text{off}} = \sum \hat{P}(\omega_k) = \hat{\sigma}_n^2, \quad (19)$$

where $\hat{\sigma}_n^2$ is the digitized off-pulse noise power. After correcting for dispersion, the total power just before and after each individual pulse is approximately given by

$$\hat{P}_{\text{dip}} = \sum \hat{P}(\omega_k) = A(\sigma^2) \frac{\hat{\sigma}^2}{\sigma^2} \sigma_n^2 + [1 - A(\sigma^2)] \hat{\sigma}^2, \quad (20)$$

where $\hat{\sigma}^2$ and σ^2 are the digitized and undigitized dispersed on-pulse power, respectively. The following dimensionless quantity, D , is defined as a measure of the dedispersed pulse distortion for a given level-setting scheme:

$$D \equiv \frac{\hat{P}_{\text{dip}} - \hat{P}_{\text{off}}}{\hat{P}_{\text{off}}} \quad (21)$$

$$= A(\sigma^2) \left(\frac{\hat{\sigma}^2 \sigma_n^2}{\sigma^2 \hat{\sigma}_n^2} - 1 \right) + [1 - A(\sigma^2)] \left(\frac{\hat{\sigma}^2}{\hat{\sigma}_n^2} - 1 \right). \quad (22)$$

For an ideal digitizing scheme, $D = 0$. It can be seen from

equation (22) that if $\hat{\sigma}^2$ varies linearly with σ^2 , then the first term is zero. The remaining distortion is due to the scattered power since it vanishes for $A = 1$. When using a constant level-setting scheme, $D \neq 0$ since $\hat{\sigma}^2$ varies nonlinearly with σ^2 and $A \neq 1$. This results in the artifacts seen in Figure 4. The next section discusses both a dynamic level-setting scheme that keeps $\hat{\sigma}^2 \approx \sigma^2$ and a technique that reduces the effects of the scattered power on the dedispersed pulse profile.

4.2. Artifact Minimization Techniques: Dynamic Level Setting and the Scattered Power Correction

For the purposes of this paper, dynamic level setting is the process of adjusting the output levels in order to maintain a near linear relationship between the undigitized power and the corresponding digitized power. This correction is performed after the data has been recorded with fixed input thresholds. Using a model for the statistics of the input signal and the values of the fixed input thresholds, the undigitized power level may be estimated from the data. Using this estimate, the output levels may be chosen in order to maintain a linear power response. There are many possible choices. In this paper, equation (6) with $f(v) = v^2$ is used to obtain the output levels. This minimizes χ^2 for the data segment with the constraint of fixed input thresholds. For illustrative purposes, Appendix B compares this with an alternative prescription.

Once an appropriate output level-setting prescription is adopted, the dynamic level-setting technique will generate a time series such that $\hat{\sigma}^2 \approx \sigma^2$. Dynamic level setting does not attempt to assign optimum values for the input thresholds. Section 5 describes how to obtain a set of optimum input thresholds by maximizing the signal-to-noise ratio of the data.

As described above, dynamic level setting keeps $\hat{\sigma}^2 \approx \sigma^2$ by estimating the undigitized power and setting the output levels according to a chosen prescription. In order to implement this procedure, the digitized data is divided up into several small overlapping or nonoverlapping segments each containing L points. The total undigitized power, σ^2 , within each segment is estimated by inverting the following relationship:

$$\Phi = \frac{1}{\sqrt{2\pi\sigma}} \int_{x_l}^{x_h} e^{-v^2/2\sigma^2} dv, \quad (23)$$

where Φ is the observed fraction of samples that fall within the range x_l to x_h and Gaussian statistics are assumed. Given this estimate for σ^2 , the values of the output levels are calculated using equation (6) with $f(v) = v^2$ and

$$P(v) = \frac{1}{\sqrt{2\pi\sigma}} e^{-v^2/2\sigma^2}. \quad (24)$$

Using these output levels, appropriate numerical values are

assigned to each sample in the current data segment. This procedure is repeated for each set of L points. Once the data has been completely “unpacked,” the digitized total power, $\hat{\sigma}^2$, may be calculated from the resulting time series. In Appendix A, it is shown that

$$\hat{\sigma}^2 = \sigma^2 + O\left(\frac{1}{L}\right) \quad (25)$$

for dynamic level setting. Hence, this method asymptotically approaches the ideal response, $\hat{\sigma}^2 = \sigma^2$, as $L \rightarrow \infty$. Figures 4c and 4d were produced using this scheme with nonoverlapping segments, and, as these figures indicate, the dips have been removed. The remaining digitization artifacts are due to the scattered power.

The scattered power may be removed when dedispersing the data with an incoherent filter bank. An incoherent filter bank takes a band-limited (bandwidth B) time series and divides it into M time series, or channels, wherein each channel is a measure of the total power of the signal within a small band $\Delta f = B/M$ (Jenet et al. 1997). After calculating the power in M frequency channels, the scattered power, $(1 - A)\hat{\sigma}^2/M$, may be subtracted from each channel before the channels are time shifted by the dispersion delay and added as usual. Note that $M\Delta t$ should be as large as possible to decrease the effects of the $O(1/M)$ term while remaining less than the smallest time-scale of interest (e.g., the dispersion smearing time), where Δt is the sampling time. Figures 4e and 4f show the pulse profile and the frequency structure of PSR B0833–45 calculated using both dynamic level setting and the scattered power correction, respectively. For these data, A was held constant at its value for the background noise, $A = A(\sigma_n^2)$. In general, one could dynamically set the value of A according to the local power level but this was unnecessary for these data since the total power never rose above approximately twice the system noise power level (see Fig. 3).

With both dynamic level setting and the scattered power correction in place, the remaining distortion, D , may be estimated. The off-pulse power is now given by

$$\hat{P}_{\text{off}} = A(\sigma_n^2)\sigma_n^2, \quad (26)$$

and the power just off pulse but in the distorted region is now given by

$$\hat{P}_{\text{dip}} = A(\sigma^2)\sigma_n^2. \quad (27)$$

From the definition of D (see eq. [21]),

$$D = \frac{A(\sigma^2)}{A(\sigma_n^2)} - 1. \quad (28)$$

Since A is relatively insensitive to changes in σ , D is small. In

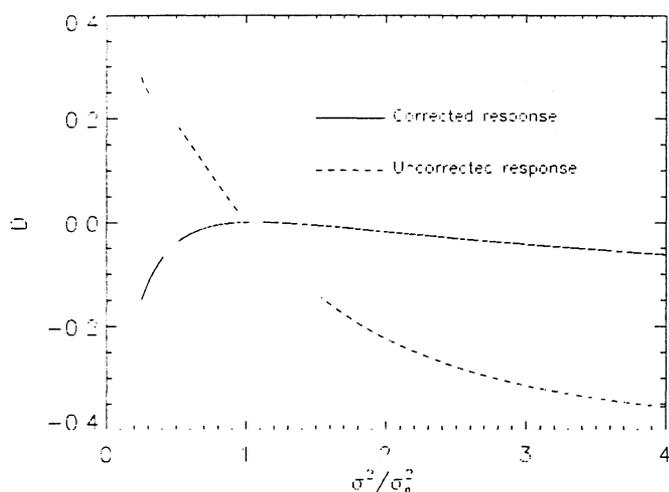


FIG. 5.—Total distortion, D , due to power underestimation and scattered power effects vs. σ^2/σ_n^2 for both the power optimized fixed level-setting scheme (uncorrected response) and the dynamic level-setting scheme with the scattered power correction (corrected response). The curves were calculated for a 2-bit system.

Figure 5, plots of D versus σ^2/σ_n^2 are shown for both the uncorrected power-optimized level-setting scheme (i.e., constant output levels with no scattered power correction) and the corrected scheme using both dynamic level setting and the scattered power correction.

5. SIGNAL-TO-NOISE LOSS AND OPTIMUM THRESHOLDS

As stated in the previous section, the dynamic level-setting scheme specifies only the values of the output levels given a set of input thresholds. Here it is shown how to set the input thresholds such that both the remaining distortion, D , and the fractional loss in the signal-to-noise ratio are minimized.

The scattered power introduced by the digitization process is a source of incoherent noise that reduces the signal-to-noise ratio. By defining a renormalized power spectrum $S(\omega_i) \equiv \sigma^2 \hat{P}(\omega_i)/A\hat{\sigma}^2$, it follows from equation (18) that

$$S(\omega_i) = P(\omega_i) + \frac{1-A}{A} \frac{\sigma^2}{M}. \quad (29)$$

Consequently, it can be seen that S is the undigitized power spectrum plus the “quantization” noise power spectrum. The total quantization noise power, N_q , is obtained by summing the last term in the above equation over all ω_i :

$$N_q = \frac{1-A}{A} \sigma^2. \quad (30)$$

The signal-to-noise ratio for the digitized signal, $\hat{\gamma}$, may be

expressed as

$$\hat{\gamma} = \gamma \frac{N}{N + N_q} \quad (31)$$

$$\approx \gamma \left(1 - \frac{N_q}{N}\right), \quad (32)$$

where γ is the undigitized SNR and N is the undigitized noise power. The fractional loss in the SNR, l , is defined as

$$l \equiv \frac{\gamma - \hat{\gamma}}{\gamma}. \quad (33)$$

Substituting equation (32) this becomes,

$$l = \frac{N_q}{N}. \quad (34)$$

From equation (30) and $N = \sigma_n^2$,

$$l = \frac{1-A}{A} \frac{\sigma^2}{\sigma_n^2}. \quad (35)$$

Taking σ^2 to be the largest value of the undigitized dispersed power, the above expression for l is an upper bound on the fractional loss in the SNR.

From the expressions for l and D (eqs. [35] and [28], respectively), it is apparent that maximizing $A(\sigma_n^2)$ simultaneously minimizes the remaining distortions in the data and the fractional loss in the SNR. Hence, the input thresholds should be set so that $A(\sigma_n^2)$ is a maximum, where σ_n^2 is the background noise power level. Assuming Gaussian statistics, one can show from the definition of A (see eq. [14]) and equation (11) that

$$A = \frac{1}{2\pi\hat{\sigma}_n^2} \left\{ \sum_{k=1}^N y_k \left[\exp\left(\frac{-x_{k+1}^2}{2\sigma_n^2}\right) - \exp\left(\frac{-x_k^2}{2\sigma_n^2}\right) \right] \right\}^2. \quad (36)$$

The output levels are set according to equation (6) with $f(v) = v^2$:

$$y_k^2 = \left[\sigma_n^2 \operatorname{erf}\left(\frac{x}{\sqrt{2}\sigma_n}\right) - \frac{2x\sigma_n}{\sqrt{2\pi}} \exp\left(-\frac{x^2}{2\sigma_n^2}\right) \right]_{x_k}^{x_{k+1}} \quad (37)$$

$$\times \left[\operatorname{erf}\left(\frac{x}{\sqrt{2}\sigma_n}\right) \right]_{x_k}^{x_{k+1}}^{-1}.$$

Table 1 presents the optimum input thresholds, x_k , for 2–8-bit digitizers with arbitrarily spaced input thresholds. These values were calculated using the multidimensional downhill simplex

TABLE 3
OPTIMUM INPUT THRESHOLD SPACING FOR A
UNIFORM DIGITIZER

Bits	Spacing	A	l
2	0.9674	0.8808	0.1353
3	0.5605	0.9635	0.0379
4	0.3188	0.9890	0.0111
5	0.1789	0.9967	0.0033
6	0.09925	0.9990	0.0010
7	0.05445	0.9997	0.0003
8	0.02957	0.99992	0.00008

NOTES.—The optimum threshold spacing in units of σ and the corresponding A and l values for a digitizer with uniformly spaced levels. By symmetry, the central threshold is zero.

following form (see eq. [36]):

$$A = \frac{2}{\pi} \frac{\{y_3 [1 - \exp(-t^2/2\sigma^2)] + y_4 \exp(-t^2/2\sigma^2)\}^2}{(y_3^2 - y_4^2) \operatorname{erf}(t/\sqrt{2}\sigma) + y_4^2}. \quad (43)$$

The undigitized power, σ^2 , may be estimated from a set of digitized data samples by measuring the fraction of samples, Φ , between x_2 and x_4 ,

$$\Phi = \frac{1}{\sqrt{2\pi}\sigma} \int_{x_2}^{x_4} e^{-v^2/2\sigma^2} dv \quad (44)$$

$$= \operatorname{erf}\left(\frac{t}{\sqrt{2}\sigma}\right). \quad (45)$$

We recommend that $L > 100$ and $L\Delta t < T_{\text{dm}}$, where L is the number of points used to measure Φ , Δt is the sample time and T_{dm} is the dispersion smearing time across the narrowest

frequency band of interest (i.e., a scintillation band). This constraint allows for a reasonably accurate measurement of σ^2 over a timescale for which the signal is stationary.

7. DISCUSSION

The process of digitizing a nonstationary stochastic signal will introduce systematic distortions into the statistics of the resulting digitized data. These distortions may be investigated analytically by approximating the autocorrelation function of the digitized signal with a form that can easily be transformed into the corresponding power spectrum. Once the form of the digitized power spectrum is known (eq. [18]), the resulting artifacts that may arise from further data processing can be calculated. This technique was applied specifically to the case of dispersion removal from a recorded pulsar signal. The resulting profile distortion is given by equation (22). This distortion is large for standard fixed output level digitization schemes. A dynamic output level scheme was introduced that, to a specified accuracy, removes the power underestimation effects leaving only the effects of the scattered power. In addition, proper setting of the input thresholds will simultaneously reduce the scattered power distortion and the fractional SNR loss since they are caused by the same phenomenon. With the dynamic level-setting technique, the data may be dedispersed using coherent or incoherent dispersion techniques. In cases where the scattered power is large (i.e., for bright pulsars like Vela), a technique was presented that further reduces the scattered power when incoherently dedispersing the data.

We would like to thank Shri Kulkarni, Tom Prince, Patrick Brady, and David Nice for helpful discussions. This research was supported by the National Science Foundation under grant ASC-9318145.

APPENDIX A

THE VARIANCE OF THE DIGITIZED SIGNAL

In this Appendix, it is shown that

$$\hat{\sigma}^2 = \sigma^2 + O\left(\frac{1}{L}\right) \quad (A1)$$

for dynamic level setting where $\hat{\sigma}^2$ and σ^2 are the digitized and undigitized total power, respectively. When dynamically setting the levels for a set of L points, an estimate of the undigitized power of these points must be made. This may be done by measuring

the fraction of samples, Φ , which fall within a range of thresholds. The expected value of Φ is

$$\langle \Phi \rangle = \int_{x_i}^{x_h} P(v) dv. \quad (A2)$$

In general, $\langle \Phi \rangle = g(\sigma^2)$. Inverting this relationship to obtain an estimate of σ for a given realization of Φ yields

$$\sigma^2 = g^{-1}(\Phi). \quad (A3)$$

This value of σ^2 is used to set the levels for the L points in question. In order to calculate the corresponding digitized power, $\hat{\sigma}^2$, the function $f(\Phi)$ is defined as

$$f(\Phi) = \sum_{i=1}^{N+1} y_i^2(\Phi) \int_{x_i}^{x_{i+1}} P(v) dv. \quad (A4)$$

Here x_i and y_i are the input thresholds and output levels, respectively. Notice that the output levels are functions of Φ . The expected value of $f(\Phi)$ is the digitized power $\hat{\sigma}^2$:

$$\hat{\sigma}^2 = \langle f(\Phi) \rangle = \sum_{\Phi} \mathcal{P}(\Phi) f(\Phi), \quad (A5)$$

where $\mathcal{P}(\Phi)$ is the discrete probability distribution for Φ . This probability distribution is given by a binomial distribution:

$$\mathcal{P}(\Phi) = \frac{L!}{(\Phi L)!(L - \Phi L)!} \langle \Phi \rangle^{\Phi L} (1 - \langle \Phi \rangle)^{L(1 - \Phi)}. \quad (A6)$$

The binomial distribution is highly peaked at $\Phi = \langle \Phi \rangle$. Approximating $f(\Phi)$ by a Taylor series about $\Phi = \langle \Phi \rangle$ and performing the

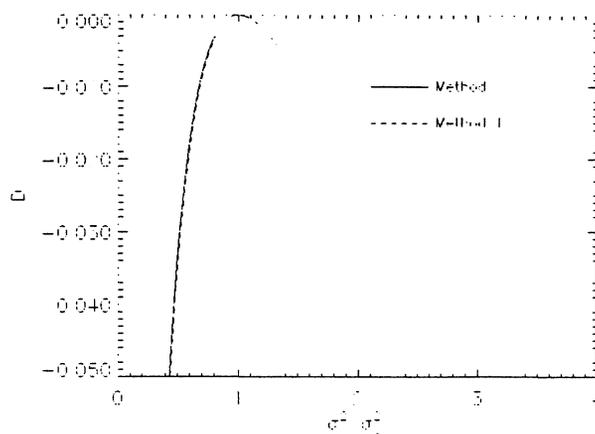


FIG. 6.—Total distortion, D , due to power underestimation and scattered power effects vs. σ^2/σ_n^2 for the two dynamic level-setting methods compared in Appendix B. These curves would be indistinguishable if plotted with the same scale as in Fig. 5.

distortion parameter, $D(\sigma^2)$, for both methods using two bits. As the signal power increases, the distortion in method II grows slightly faster with the input signal power than does method I. At $\sigma^2/\sigma_n^2 = 2$, the relative distortion in method I is 4% smaller than method II.

In practice, the two methods are essentially the same. As expected, method I removes the distortions slightly better than method II, while method II has a slightly better SNR. Since this paper deals with the removal of systematic digitization artifacts, method I is used as the optimum level-setting scheme.

REFERENCES

- Cooper, B. F. C. 1970, *Australian J. Phys.*, 23, 521
 Cordes, J. M. 1975, *ApJ*, 195, 193
 Jenet, F. A., Anderson, S. B., Kaspi, V. M., Prince, T. A., & Unwin, S. C. 1998, *ApJ*, 498, 365
 Jenet, F. A., Cook, W. R., Prince, T. A., & Unwin, S. C. 1997, *PASP*, 109, 707
 Max, J. 1976, in *Waveform Quantization and Coding*, ed. N. S. Jayant (New York: IEEE), 16
 Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, *Numerical Recipes: The Art of Scientific Computing* (Second ed.; Cambridge: Cambridge Univ. Press)
 Rickett, B. J. 1975, *ApJ*, 197, 185

sum in equation (A5) results in

$$\hat{\sigma}^2 = \sum_{i=0}^{\infty} \frac{1}{i!} f^{(i)}(\langle \Phi \rangle) \langle (\Phi - \langle \Phi \rangle)^i \rangle. \quad (\text{A7})$$

The various moments, $\langle (\Phi - \langle \Phi \rangle)^i \rangle$, of the binomial distribution can be calculated from the following relationship:

$$\langle \Phi \rangle^j \left(1 - \frac{1}{L}\right) \left(1 - \frac{2}{L}\right) \dots \left(1 - \frac{j-1}{L}\right) = \left\langle \Phi \left(\Phi - \frac{1}{L}\right) \left(\Phi - \frac{2}{L}\right) \dots \left[\Phi - \frac{(j-1)}{L}\right] \right\rangle. \quad (\text{A8})$$

One derives this relationship directly from equation (A6). Using this, one finds that

$$\langle \Phi^j \rangle = \langle \Phi \rangle^j + \frac{\beta_j}{L} + O\left(\frac{1}{L^2}\right), \quad (\text{A9})$$

where β_j is a constant that depends on the moment in question. From equation (A9) and the following identity

$$\langle (\Phi - \langle \Phi \rangle)^j \rangle = \sum_{i=0}^j \frac{j!}{i!(j-i)!} \langle \Phi^i \rangle \langle \Phi \rangle^{j-i} (-1)^{j-i}, \quad (\text{A10})$$

the moments take the following form:

$$\langle (\Phi - \langle \Phi \rangle)^j \rangle = \sum_{i=0}^j \frac{j!}{i!(j-i)!} \left[\langle \Phi \rangle^i + \frac{\beta_i}{L} + O\left(\frac{1}{L^2}\right) \right] \langle \Phi \rangle^{j-i} (-1)^{j-i} \quad (\text{A11})$$

$$= \frac{1}{L} \sum_{i=0}^j \frac{j!}{i!(j-i)!} \beta_i \langle \Phi \rangle^{j-i} (-1)^{j-i} + O\left(\frac{1}{L^2}\right). \quad (\text{A12})$$

From this relationship it is seen that for $j > 1$ all moments $\langle (\Phi - \langle \Phi \rangle)^j \rangle$ go as $1/L$ in leading order. Using this, equation (A7) becomes

$$\hat{\sigma}^2 = f(\langle \Phi \rangle) + \sum_{i=2}^{\infty} \frac{1}{i!} f^{(i)}(\langle \Phi \rangle) \langle (\Phi - \langle \Phi \rangle)^i \rangle \quad (\text{A13})$$

$$= f(\langle \Phi \rangle) + O\left(\frac{1}{L}\right). \quad (\text{A14})$$

Since $f(\langle \Phi \rangle) = \sigma^2$ for the adopted level-setting prescription, this is desired result.

APPENDIX B

COMPARISON OF LEVEL-SETTING TECHNIQUES

For illustrative purposes, this Appendix compares an alternative dynamic level-setting technique to the one presented in the text. The method described in the text (method I) uses equation (6) with $f(v) = v^2$ to set the output levels and constrains the remaining degrees of freedom (i.e., the input thresholds) by maximizing A . The alternative method (method II) sets the output levels by forcing $\hat{\sigma}^2 = \sigma^2$ and maximizing A is a function of the output levels with fixed thresholds. The input thresholds are set so that A is at an absolute maximum when σ^2 is equal to the background noise level.

Since method I is based on minimizing χ^2 , it should correct distortions better than method II. Method II should have a better signal-to-noise ratio since it is based almost entirely on maximizing A . Numerically, method II yields $A(\sigma_n^2) = 0.8825$, while method I yields $A(\sigma_n^2) = 0.8808$ for a 2-bit system. This corresponds to a 1.6% improvement in the SNR. Figure 6 plots the