JOURNAL OF COMPUTATIONAL PHYSICS 122, 185–190 (1995)

Performance of the Matched-Phase Noise Filter with Estimated Noise Spectra

B. E. McDonald and Gregory J. Orris

Naval Research Laboratory, Washington, DC 20375

Received May 24, 1994; revised March 27, 1995

We demonstrate properties of a new type of noise filter, the *matched-phase filter,* for low signal-to-noise data. It requires that the shape of the noise amplitude spectrum be approximately known. The method assumes no *a priori* information about the phases of the noise spectrum, nor anything about the signal, except that its spectrum be uncorrelated with that of the noise. We demonstrate algorithm performance in cases where one's knowledge of the noise spectrum is and is not precise. We also demonstrate that when the shape of the noise spectrum is known exactly, the matched-phase filter can reveal signals at a signal-to-noise ratio (SNR) of less than −100 dB and increase the SNR to order 0 dB. © 1995 Academic Press, Inc.

1. INTRODUCTION

A number of experiments face the task of extracting weak signals from noisy backgrounds. A contemporary example is the proposal [1] over a period of years to inject acoustic signals of order 200 W power into the ocean off California and measure their flight times across the entire Pacific. Systematic drifts in flight times could reveal global trends in deep ocean temperature. A number of considerations limit the strength of the test signal; it is therefore necessary to increase signal gain on the receiving end by a number of signal processing techniques. In cases where the background noise is reasonably steady, its spectrum can be measured before and/or after the transmission of the test signal as an estimate of the noise spectrum at the time the signal is received.

We present a new type of noise reduction algorithm (the matched-phase filter) that has recently been investigated for removal of noise from acoustic measurements; the method applies to time series in general. We had previously [2, 3] applied the method to test cases involving synthetically generated noise and found that its performance exceeded expectations. We were able to extract signals at signal-to-noise ratio (SNR) values of order −100 dB when only the shape of the noise spectrum (no phase information) was known precisely. The SNR in the algorithm output was of order 0 dB. The current paper presents the algorithm, demonstrates via perturbation expansion why its performance is as stated, and explains salient features of the algorithm's output.

2. MATCHED-PHASE FILTERING

Consider the problem of extracting a weak signal from a noisy time series. We assume that the noise component of the series may be recognized by the shape of its frequency spectrum, for which some estimates are available. We have shown elsewhere [2, 3] that information about the shape of the noise spectrum $|\tilde{N}(\omega)|$ with no phase knowledge can be used to construct a noise rejection algorithm which we call the matched-phase filter. In this paper we demonstrate that to first order in the SNR, the matched-phase filter is capable of improving an arbitrarily low SNR ($\ll 1$) to order unity when the spectrum of the noise is known exactly. We also show examples illustrating the performance of the algorithm when the noise spectrum is known only approximately.

Let us assume we are given a noise-dominated time series

$$p(t) = N(t) + \varepsilon s(t), \tag{1}$$

where N and εs are unknown noise and signal components.

If we normalize s so that $\int s^2\,dt = \int N^2\,dt$, then ε is the signal-to-noise ratio. The decibel measure of the SNR is 20 log ε. We assume no knowledge of the signal, except that it is small and that its spectrum is decorrelated from that of the noise. In attempting to remove as much noise as possible from $p(t)$, we use unknown spectral phases φ of the noise field as free parameters to construct an optimal realization of the noise field

$$N_{\text{opt}}(t) = \alpha \int d\omega |\tilde{N}(\omega)| \exp(i\varphi(\omega) - i\omega t) \tag{2}$$

which, when subtracted from the data, best quietens the time series.

Represent the given noise-dominated time series as

$$p(t) = \int d\omega \tilde{p}(\omega) e^{-i\omega t} = \int d\omega |\tilde{p}(\omega)| e^{i\phi - i\omega t}, \tag{3}$$

$$\phi(\omega) \equiv \arg(\tilde{p}(\omega))$$

0021-9991/95 $12.00

and assume knowledge of the spectral amplitude $|\tilde{N}(\omega)|$ of the noise.

Let us represent Eq. (1) in the frequency domain as

$$\begin{aligned}\tilde{p}(\omega) &= \tilde{N}(\omega) + \varepsilon\tilde{s}(\omega)\\ &\equiv |\tilde{N}(\omega)|e^{i\psi} + \varepsilon|\tilde{s}(\omega)|e^{i\eta}\\ &\equiv |\tilde{p}(\omega)|e^{i\phi},\end{aligned} \tag{4}$$

where ψ, η, and ϕ are the phases of the noise, signal, and their complex sum (all implicitly functions of ω). Phases ψ and η are unknown.

We now construct a "quietened" time series $p_q(t)$

$$p_q(t) = p(t) - \alpha\int d\omega|\tilde{N}(\omega)|e^{i\varphi}e^{-i\omega t}, \tag{5}$$

by minimizing $\int |p_q(t)|^2\, dt = \int |\tilde{p}_q(\omega)|^2\, d\omega$ with respect to φ and α. The integral is minimized by the following variational procedure in the frequency domain:

$$\begin{aligned}0 &= \delta\int d\omega\big|\,|\tilde{p}(\omega)|e^{i\phi} - \alpha|\tilde{N}(\omega)|e^{i\varphi}\big|^2\\ &= \delta\int d\omega(|\tilde{p}(\omega)|^2 - 2\alpha\cos(\phi - \varphi)|\tilde{p}(\omega)|\,|\tilde{N}(\omega)|\\ &\quad + \alpha^2|\tilde{N}(\omega)|^2).\end{aligned} \tag{6}$$

It is straightforward to show that minimization with respect to φ requires

$$\varphi(\omega) = \phi(\omega).$$

Minimization with respect to α gives

$$\alpha = \frac{\int d\omega|\tilde{p}(\omega)|\,|\tilde{N}(\omega)|}{\int d\omega|\tilde{N}(\omega)|^2}. \tag{7}$$

The quietened time series p_q is then

$$p_q(t) = \int d\omega(|\tilde{p}(\omega)| - \alpha|\tilde{N}(\omega)|)e^{i\phi - i\omega t}. \tag{8}$$

The spectral phases of the quietened time series (8) are identical to those of the noisy time series (3); thus we have the name "matched-phase filter" for the algorithm. The matched-phase filter (8) simply orthogonalizes the quietened spectrum $\tilde{p}_q(\omega)$ to the optimum realization $\tilde{N}_{\text{opt}}(\omega)$ of the noise spectrum.

An expansion of (8) to first order in the SNR ε reveals that the residual SNR in p_q is of order unity. To establish this, we give the following first-order results for the phase and amplitude of $\tilde{p}(\omega)$ which are easily obtained from a complex phasor diagram of Eq. (4) (see Fig. 1):

$$\phi = \psi + \varepsilon\frac{|\tilde{s}|}{|\tilde{N}|}\sin(\eta - \psi) + O(\varepsilon^2) \tag{9}$$

$$|\tilde{p}| = |\tilde{N}| + \varepsilon|\tilde{s}|\cos(\eta - \psi) + O(\varepsilon^2). \tag{10}$$

Equations (10) and (7) give

$$\alpha = 1 + \varepsilon\frac{\int |\tilde{s}|\,|\tilde{N}|\cos(\eta + \psi)\,d\omega}{\int |\tilde{N}|^2\,d\omega} + O(\varepsilon^2). \tag{11}$$

From (8)–(11), the Fourier transform of the filtered time series is

$$\begin{aligned}\tilde{p}_q(\omega) &= -\varepsilon\left\{|\tilde{N}|e^{i\phi}\frac{\int |\tilde{s}|\,|\tilde{N}|\cos(\eta - \psi)\,d\omega}{\int |\tilde{N}|^2\,d\omega}\right\}\\ &\quad + \varepsilon\tilde{s} - i\varepsilon|\tilde{s}|e^{i\phi}\sin(\eta - \psi) + O(\varepsilon^2).\end{aligned} \tag{12}$$

If signal and noise spectra are uncorrelated, the term in curly brackets will be much smaller in average magnitude than $|\tilde{s}|$ because of the oscillating cosine term inside the integral. The unknown noise phases ψ in (12) may be replaced by ϕ, since by (9) their difference is $O(\varepsilon)$. Then neglecting the curly bracket in (12), substituting $\tilde{s} = |\tilde{s}|e^{i\eta}$, $\psi = \phi + O(\varepsilon)$, and representing the sine in (12) as the difference of complex exponentials gives

$$\tilde{p}_q(\omega) \simeq \frac{\varepsilon}{2}(\tilde{s} + \tilde{s}^* e^{2i\phi}) \tag{13a}$$

$$= \varepsilon|\tilde{s}|e^{i\phi}\cos(\eta - \phi). \tag{13b}$$

To zero order, ϕ is uncorrelated to the signal, so the $\tilde{s}^* e^{2i\phi}$ term in (13a) represents residual noise. Thus (13a) implies that when the noise spectrum is known exactly, the SNR in the quietened

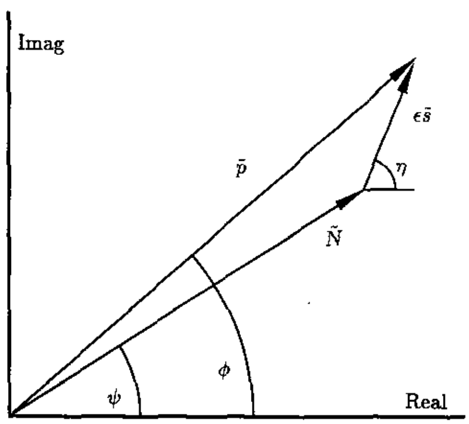

FIG. 1. Phasor diagram for $|\tilde{p}(\omega)|e^{i\phi} \equiv |\tilde{N}(\omega)|e^{i\psi} + \varepsilon|\tilde{s}(\omega)|e^{i\eta}$.

time series $p_q(t)$ emerging from the matched-phase filter is approximately unity.

Equation (13b) says that the quietened amplitude spectrum $|\tilde{p}_q(\omega)|$ oscillates below an envelope equal to the signal spectrum $\varepsilon|\tilde{s}(\omega)|$. One sees this behavior of the matched-phase filter output verified in results to be given in the next section.

3. NUMERICAL RESULTS

We will give some results of the matched-phase noise filter applied to computer generated time series with spectra represented by an analytic model [4] for Kennedy's [5, 6] sea surface noise data,

$$|\tilde{N}_{\rm fit}(\omega)|^2 = |\tilde{N}|^2_{\rm max}\,\frac{2.5(\omega/\omega_p)^{1.5}}{1 + 1.5(\omega/\omega_p)^{2.5}}, \tag{14}$$

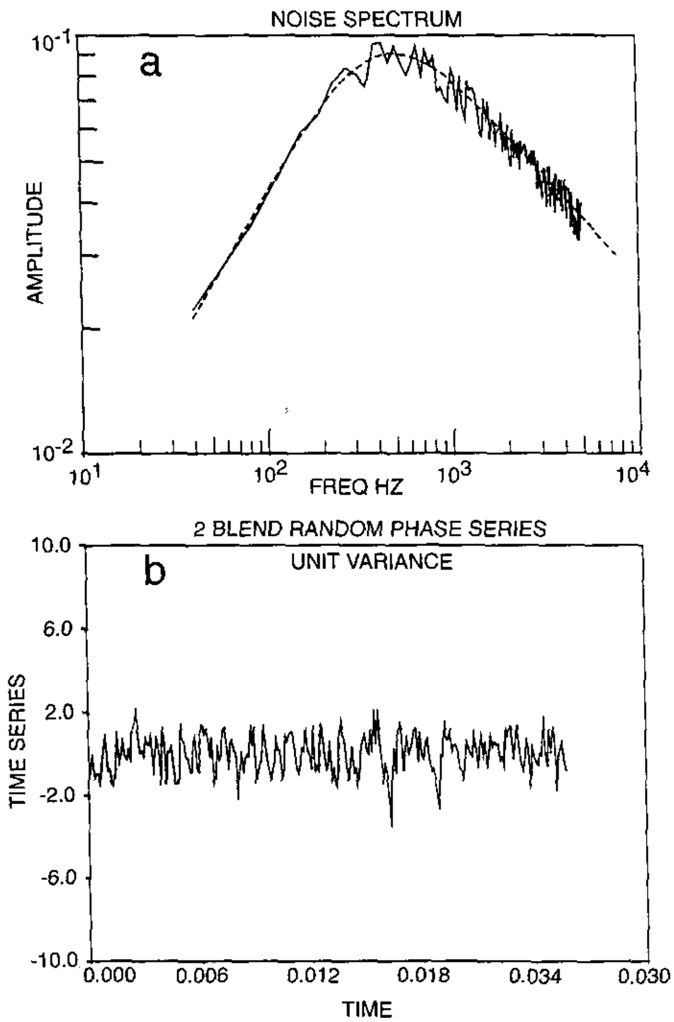

FIG. 2. (a) Dashed line: Eq. (14) for f_p = 500 Hz and n = 1.5. Solid line: Spectrum for linear combination of two random phase realizations $N_1 + 0.1N_2$. (b) Time series for $N_1 + 0.1N_2$.

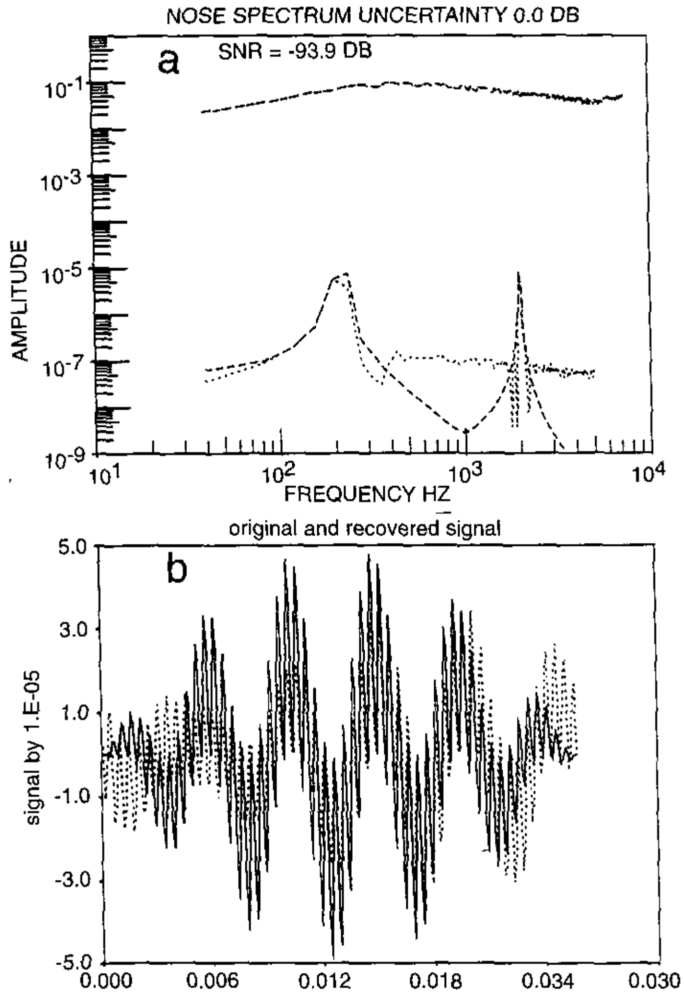

FIG. 3. (a) Upper dashed curve: random phase noise plus bitonal signal. Lower dashed curve: spectrum of the signal; dotted curve: spectrum of the quietened time series p_q recovered in single precision. (b) Solid line: true signal time series $\varepsilon s(t)$; dotted line: $p_q(t)$ recovered from the matched-phase filter.

where ω_p is the radian frequency at spectral peak. Figure 2a (dashed line) illustrates an arbitrarily normalized spectrum of the form (14) with f_p = 500 Hz. Spectra that might be measured in an ocean experiment would not be so smooth or simple, so we add some structure to the spectrum as follows. We generate two different noise time series $N_1(t)$ and $N_2(t)$ by assigning two different sets of random phase factors to Eq. (14). Then we take $N(t) = N_1(t) + 0.1N_2(t)$ and normalize $N(t)$ to unit variance. Time series $N(t)$ is shown in Fig. 2b.

We generate a windowed bitonal signal to be combined with the noise as

$$s(t) = w(t)(\sin(2\pi f_1 t) + \sin(2\pi f_2 t)) \tag{15}$$

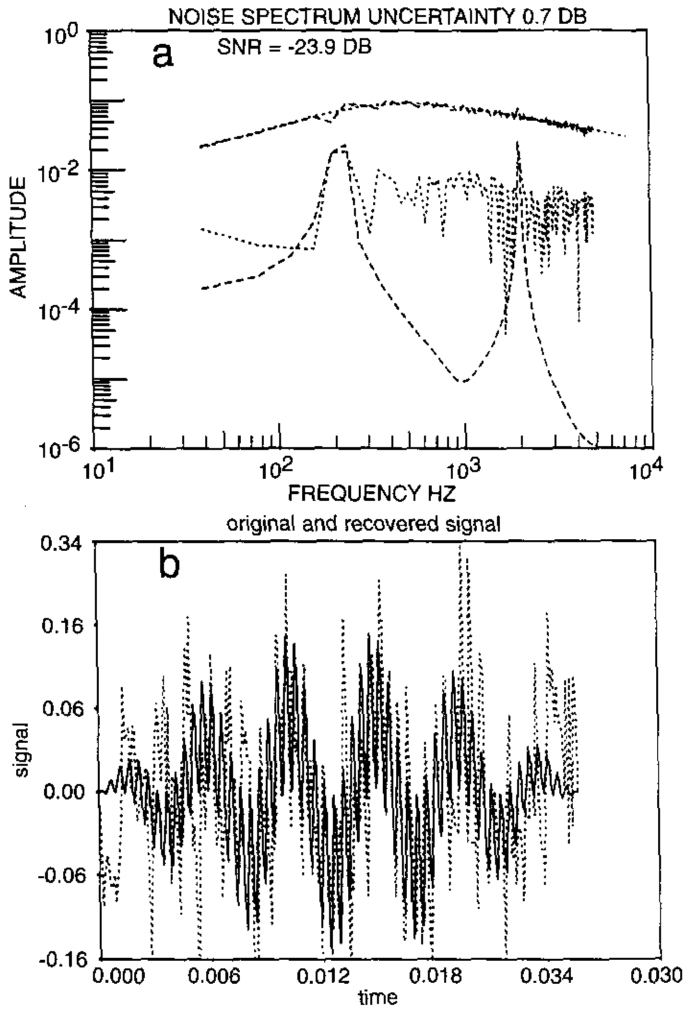

FIG. 4. See Fig. 3 caption for details. In (a) the upper dotted curve is the smoothed estimate of the noise spectrum.

where

$$w(t) = \begin{cases} 4\dfrac{t}{\Delta T}\left(1 - \dfrac{t}{\Delta T}\right), & 0 < t < \Delta T \\ 0, & \text{otherwise} \end{cases}$$

In the examples given here, f_1 and f_2 are 220 Hz and 2 kHz, respectively. We use a Fourier transform of length 256 at a sampling rate of 10 kHz, so $\Delta T = 0.0256$ s. The choice of these frequencies is arbitrary, but reflects a range of frequencies that might be used in sub-mesoscale to mesoscale acoustic oceanography.

Case 1. We take $\varepsilon = 2.56 \times 10^{-5}$ and construct $p_q(t)$ from Eqs. (7)–(8), given precise knowledge of the noise spectral amplitude $|\tilde{N}(\omega)|$. Figure 3a shows upper dashed curves corresponding to $\tilde{p}(\omega)$ which is indistinguishable on this scale with $|\tilde{N}(\omega)|$. Below in Fig. 3a is a dashed line representing the source spectrum $\varepsilon|\tilde{s}(\omega)|$ and a dotted line for the output of the matched-phase filter $|p_q(\omega)|$. The filter has clearly revealed the bitonal signal at SNR of -94 dB subject to *perfect* knowledge of the shape of the noise spectrum. For some frequencies in Fig. 3a, $|p_q(\omega)|$ exceeds $\varepsilon|\tilde{s}(\omega)|$ only because of single-precision roundoff error. Double-precision arithmetic would be sufficient in this case to bring $|p_q(\omega)|$ back below $\varepsilon|\tilde{s}(\omega)|$ as predicted by (13b). (Case 3 below illustrates improvements introduced by the use of double precision.)

In Fig. 3b we show the signal $\varepsilon s(t)$ as a solid line and the recovered signal $p_q(t)$ from the matched phase filter as a dotted line. The recovered signal is quite recognizable despite moderate distortion.

Case 2. We will demonstrate a more realistic situation in which the noise spectrum is not known exactly. We will again

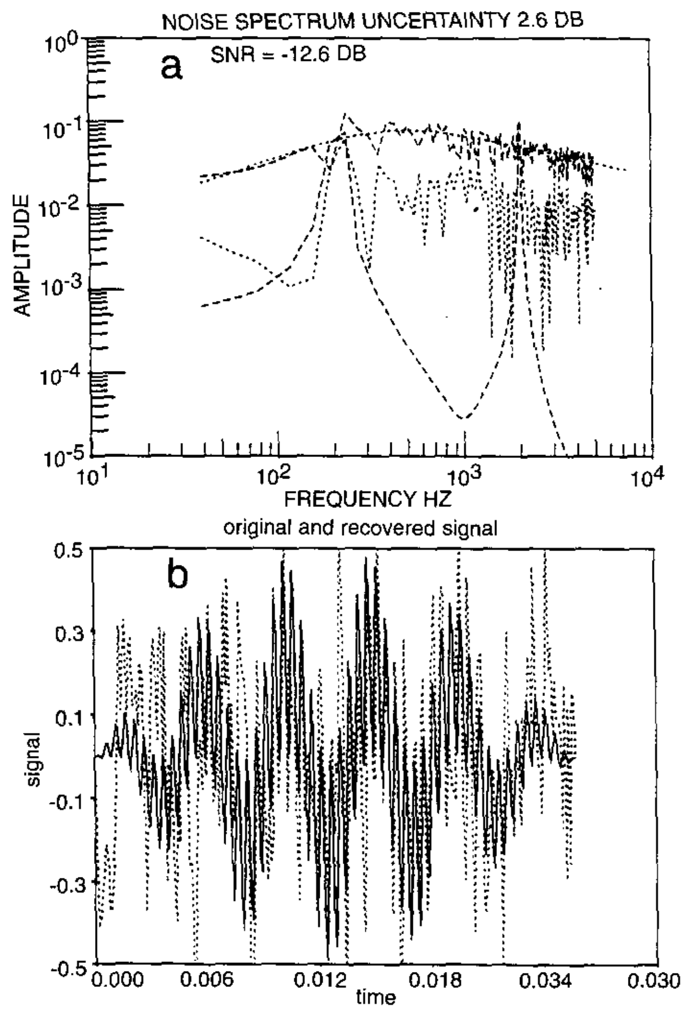

FIG. 5. See Fig. 3 caption for details. In (a) the upper dotted curve is the smoothed estimate of the noise spectrum.

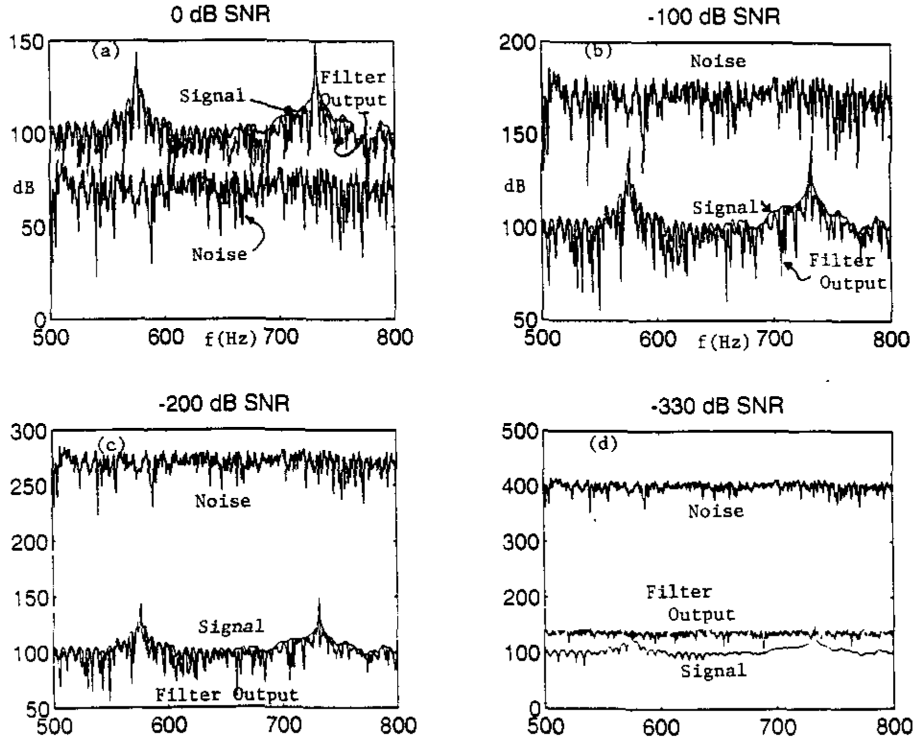

FIG. 6. 500–800 Hz detail of the spectrum of a door chime immersed in synthetic ocean noise for varying SNR values. Perfect knowledge of the noise spectrum allows the signal to be revealed until double precision roundoff error level is reached in panel (d).

use $\tilde{N}(\omega)$ (Fig. 2a, solid line) in constructing $p(t)$, but we will provide the matched-phase filter with only an appropriate spectrum, $\tilde{N}_{fit}(\omega)$, from (14) (Fig. 2a, dashed line). Results are shown in Fig. 4 for $\varepsilon = 2.56 \times 10^{-1.5}$. The noise spectrum uncertainty (i.e., the 20 log amplitude variance between $\tilde{N}(\omega)$ and $\tilde{N}_{fit}(\omega)$) is 0.7 dB. The SNR is now -24 dB. This is a 70 dB SNR increase over Case 1, but because the noise spectral shape is not precisely known, the recovered signal in Fig. 4b is quite distorted although still recognizable.

In Fig. 5 we show results of increasing the noise spectrum uncertainty to 2.6 dB by constructing $N(t) = N_1(t) + 0.3N_2(t)$, where the contribution from $N_2(t)$ has been tripled from the results of Fig. 4. With $\varepsilon = 0.256$, the SNR is -12.6 dB, or 11.3 dB higher than in Fig. 4. Because the noise spectrum is less well prescribed, however, the recovered signal in Fig. 5b has about the same level of distortion as that of Figure 4b. In fact one can see in the spectrum of the recovered signal in Fig. 5a (dotted line) a mirror image of the discrepancy between $\tilde{N}(\omega)$ and $\tilde{N}_{fit}(\omega)$.

Case 3. Figure 6 shows the result of removing noise from a sample of 8192 points, sampled at 8 kHz. The noise is generated by combining separate random phase realizations of the spectrum $\tilde{N}_{fit}(\omega)$, and the signal is a door chime consisting of two bell tones, 740 Hz followed $\frac{1}{6}$s later by 580 Hz. The door chime time series is added to noise time series in amounts determined to give desired values for SNR.

The purpose of Case 3 is to illustrate how far the matched-phase filter can be pushed when the noise spectrum is known exactly, and to demonstrate the property predicted in Eq. (13b) that the recovered spectrum oscillates below an envelope equal to the true signal spectrum. In Fig. 6 we show examples of noise removal at SNR values of 0, -100, -200, and -330 dB. Most of these values are experimentally unreasonable, but they serve to illustrate the power of the method itself. In going from the 256-point sample in Cases 1 and 2 to 8192 points we must now use double precision in order to retain accurate inner products in (7). Each panel of Fig. 6 shows the spectral interval 500 to 800 Hz. The excellent performance of the matched-phase filter is due to (a) perfect knowledge of $\tilde{N}(\omega)$ and (b) lack of correlation between the signal and noise.

Figure 6a shows recovery of the signal at SNR 0 dB. The signal spectrum happens to exceed the noise spectrum by about 25 dB in the 500–800 Hz range, but power integrated from 0 to 8 kHz is the same for signal and noise. The quietened spectrum oscillates rapidly, with the envelope of oscillations tracing out the signal spectrum as predicted by eq. (13b). Figures 6b and c show recovery of the door chime signal diminished in amplitude by succeeding factors of 10^5 below the noise to which it is added. Finally in 6d we show that a SNR lower than double precision roundoff error results in loss of all detail in the recovered spectrum.