# Methods for Reducing Audible Artifacts in a Wavelet-Based Broad-Band Denoising System*

**PAVAN K. RAMARAPU AND ROBERT C. MAHER,** *AES Member*

*Department of Electrical Engineering, University of Nebraska–Lincoln, Lincoln, NE 68588-0511, USA*

A wavelet-based method for denoising broad-band noisy audio signals is proposed and implemented. The method is a modified algorithm based on the method of Coifman and Wickerhauser (1992). The modified algorithm shows noticeable improvement in quality and computational overhead compared to the original method when denoising digitized music signals with different levels of additive broad-band noise. Several implementation issues are described.

## 0 INTRODUCTION

There are wide variety of applications where it is desirable to enhance degraded audio signals. The degradation of audio may be due to errors or data loss, poor analog recording techniques or storage media, the presence of high levels of background noise in the recording itself, or several of these causes. The typical objectives of enhancement are to improve perceived quality, increase intelligibility, reduce listener fatigue, and so on. Depending on the specific application, the enhancement system may be directed at only one of these objectives or at several [1]. In this engineering report the main objective is to enhance the quality of degraded audio signals, music (broad-band audio) in particular, and thereby to reduce listener fatigue.

The distinct types of degradation common in audio recordings fall into two general classes:

1) *Local degradation*: Clicks, scratches, and so on, where the signal is only disrupted at certain positions and restoration is only required at these places

2) *Global degradation*: Hiss, wow, hum, and so on, where the signal is disrupted to some extent at all times.

There are several notable methods for audio denoising. The spectral subtraction (SS) method has become a standard procedure in speech enhancement [2]. SS gained popularity because it is easy to understand and implement, that is, it only requires a discrete Fourier transform (DFT) of each frame of the noisy signal, application of a gain function, and calculation of the inverse DFT. Furthermore, the SS method makes minimal assumptions about the signal and noise, and when carefully implemented, it results in reasonably clear enhanced signals. Moreover, this approach has been heuristically developed and experimentally optimized. The main drawback of the SS method is that it can create noticeable residual noise with annoying tonal characteristics, often referred to as "musical noise." This musical noise is a common side effect of most broad-band denoising techniques [3]–[5].

Vaseghi and Frayling-Cork [3] used a frequency interpolation scheme to restore archival gramophone recordings. Their method was effective, but had the drawback that more than one copy of the recording was required. Ephraim and Van Trees [6] used a nonparametric signal subspace approach and showed that the SS method is a special case of this approach. Coifman and Wickerhauser [7] have used a wavelet-based approach to reduce broad-band noise. Other approaches to enhance noisy speech and other audio signals can be found in [1].

In all of these approaches the audio (time-domain) signal is first transformed to the frequency or time–frequency domain, and the processing of the noisy signal is accomplished in the transform domain. After processing, the altered signal is transformed back into the time domain to reconstruct the noise-reduced audio signal.

In this engineering report, signal quality enhancement of broad-band noisy audio signals using a wavelet packet basis is explored. In particular, the algorithm proposed by Coifman and Wickerhauser [7] for denoising music

* Manuscript received 1996 November 25; revised 1997 September 2.

signals is modified in several ways in order to obtain better perceptual quality.

In the next section the original algorithm of Coifman and Wickerhauser is reviewed. The modified algorithm and the simulations concerned with broad-band noise removal are described in Section 3. Finally, the report concludes with a summary and some suggestions for further research in this area.

## 1 ORIGINAL DENOISING ALGORITHM

The approach for signal quality enhancement explored in this engineering report is guided by the so-called best-basis paradigm [7]–[9]. This paradigm consists of three main steps:

1) Select a "best" basis (or coordinate system) for the problem at hand from a library of bases (a fixed yet flexible set of bases consisting of wavelets and wavelet packets).

2) Sort the coordinates (features) by "importance" for the problem at hand and discard "unimportant" features (noisy ones).

3) Use the retained coordinates to solve the problem at hand (audio signal quality enhancement) by projecting the signal onto retained bases.

The method extracts from the noisy signal a coherent part which is well represented by the given wavelets and a noisy or incoherent part which cannot be "well compressed."

This algorithm adaptively uses libraries of orthonormal waveforms, such as wavelet packets and local trigonometric waveform libraries, for separating the coherent signal from the presumably incoherent unwanted noise.

### 1.1 Denoising Algorithm

The best-basis algorithm of Coifman and Wickerhauser starts with a signal segment $x(t)$ of length $K$ as well as a library of orthonormal bases. The library could have, for example, the Fourier basis, the Haar–Walsh wavelet packets, various QMF Daubechies filters defining smoother wavelet packets, local trigonometric waveforms, and perhaps other orthonormal bases [10]–[12]. Trigonometric bases and wavelet packets were used for this purpose by Berger et al. [8]. The following sequence is used for the analysis procedure in [8].

1) The signal $x(t)$ is first expanded in each basis, that is $x(t) = \Sigma_{i=1}^{N} \alpha_i \omega_i(t)$, where $\omega_i$ are the orthogonal basis waveforms [10], $\alpha_i$ are the coefficients, and $N$ is the order of the expansion. A cost function is assigned to each expansion. Berger et al. use the Shannon second-order entropy $-\Sigma_{i=1}^{N} \alpha_i^2 \log_2 \alpha_i^2$.

2) The basis giving the least cost is chosen from the collection of bases.

3) The best-basis coefficients are sorted in the order of decreasing magnitude, that is, upon renaming, $|\alpha_1| \geq |\alpha_2| \geq \cdots$.

4) A number of leading (large-magnitude) terms are kept as the coherent part, based on a predetermined threshold cost of the remaining terms. The cost is based on the vectors of successive residual tails given by $V_1 = x(t)$, $V_2 = \Sigma_{i=2}^{N} \alpha_i \omega_i(t)$, $V_3 = \Sigma_{i=3}^{N} \alpha_i \omega_i(t)$, and so on.

5) For each $V_j$, starting at $j = 1$, the normalized entropy $H(V_j)$ is calculated. When $H(V_{j0})$ is greater than or equal to a specified threshold $\tau$, this step is stopped. Then $\Sigma_{j0}^{N} \alpha_i \omega_i(t)$ is the coherent part and the rest are the residual terms.

6) The residual terms by definition constitute the noisy part of the signal and can be treated as a new signal which in turn can be expanded and separated into its coherent and noisy components. Thus at each iteration a coherent part is extracted from the signal currently considered, and the leftover noise is treated in turn. The coherent parts, one from each iteration, are added together to produce the total coherent portion. Berger et al. [8] report that approximately seven iterations were used to process typical signals.

This published algorithm solves the problem of audio enhancement by attacking it in the manner of audio compression. The algorithm relies heavily on selecting an appropriate value of the prespecified threshold, and needs several iterations to denoise a broad-band noisy signal. In the next section this original algorithm is modified to determine whether the enhancement can be accomplished in a more efficient way.

## 2 PROBLEM FORMULATION

Consider a discrete degradation model for the broad-band noise corrupted audio signal

$$y = x + \epsilon \tag{1}$$

where $y, x, \epsilon \in \Re^n$ and $n = 2^{n_0}$. The vector $y$ represents the broad-band noisy audio signal and $x$ is the unknown true (noise-free) audio signal to be recovered. The vector $\epsilon$ is assumed to be white Gaussian noise (WGN), that is, $\epsilon \sim N(0, \sigma^2 I)$, and $\sigma^2$ is not known a priori.

Now consider an algorithm, based on the method outlined in Section 1, to estimate the desired audio signal $x$ from the available broad-band noise corrupted signal $y$. First, the library of orthonormal bases $L = \{B_1, B_1, \ldots, B_M\}$ is prepared with the Daubechies wavelets (orders 2–16) and coiflets (orders 2–6) [9], [10]. Let $B_m$ represent the best basis selected from a dictionary of orthonormal wavelet packet bases. The number of bases $M$ in this library typically ranges from 5 to 20 or more, depending on the available resources and any a priori knowledge of the signal $x$.

The audio signals treated in this work are all discrete: they are vectors of a finite number of real-valued samples. The audio signals are monophonic (single channel), 5 to 10 seconds in duration, and sampled at a rate of 44 100 samples per second. The signals are divided into frames of 2048 samples (about 46 ms) in order to follow the local characteristics of the audio. The best-basis algorithm is applied separately on each frame.

In the next section the modified algorithm is discussed. The reason to modify the published algorithm of Coifman et al. is threefold: 1) to select adaptively the threshold level for each frame while processing the noise corrupted signal, 2) to further reduce the clicks and whistles, which are subjectively more annoying, in the processed output signal, and 3) to optimize the algorithm for the audio enhancement case since the existing method is geared toward solving the problem of data compression.

## 3 MODIFIED DENOISING ALGORITHM

It is assumed that the noise-free signal frames $x$ can be completely represented by basis $B_m$,

$$x = W_m \alpha_m \tag{2}$$

where $W_m \in \Re^{n \times n}$ is an orthogonal matrix whose column vectors are the basis elements of $B_m$ and $\alpha_m \in \Re^n$ is the vector of expansion coefficients of $x$. The basis $B_m$ is the best basis calculated for the frame $x$ using the minimum-entropy criterion [7].

In order to reduce the sensitivity of the algorithm to $\tau$ (threshold level), the wavelet coefficients obtained from the best-basis decomposition are subjected to signal-dependent thresholding. The rationale for this method is that the signal can be well decomposed into a wavelet basis whereas noise cannot be represented by any of the bases. For example, if the frame under consideration has only noise, then it cannot be represented by any of the bases. So the threshold level should be high. On the other hand, when no noise is present in the frame under consideration, the signal is well represented by the best wavelet basis and so the threshold level should be near zero. Hence a method is needed to adjust the threshold according to the nature of the frame under treatment.

As can be seen by the contour map of the wavelet coefficients in Fig. 1, the broad-band noise cannot be represented well by the basis [no visible structure in Fig. 1(b)], whereas the noiseless audio signal (music alone) can be effectively represented by a few components. A different viewpoint is shown in Fig. 2 with the wavelet scales (parameter $j$) and the time (parameter $i$) shown along the $x$ axis itself. This gives a better representation for sorting the coefficients and removing the ones below the threshold level, as shown in Fig. 3.

The signal-dependent thresholding chosen here ex-

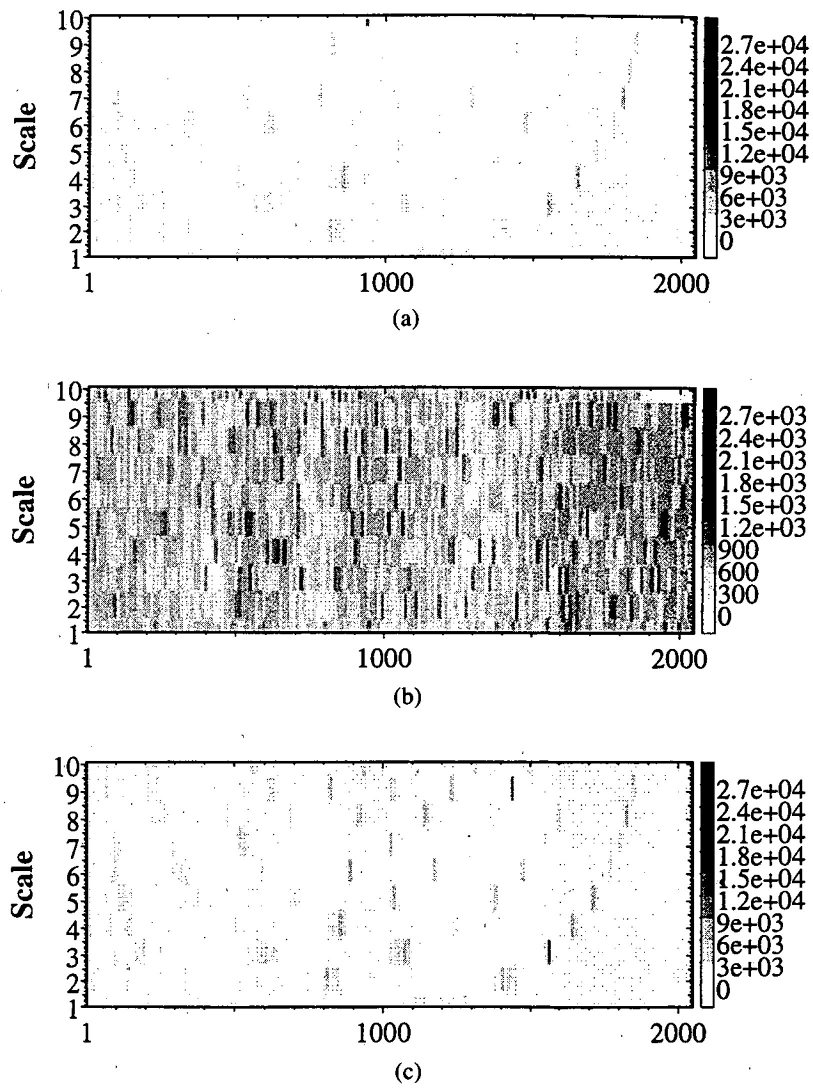


Fig. 1. Contour plots of wavelet transforms. Note that different bases are used for each plot. (a) Noisefree audio signal. (b) Broad-band noise. (c) Broad-band noise corrupted audio signal. $x$-axis—time index; 3rd axis (darkness of pixel location)—magnitude of wavelet coefficient.

ploits the second-order statistics of the signal and adaptively selects the threshold level. The threshold level for soft thresholding is selected with the algorithm shown in Fig. 4. The characteristic function to select the threshold is derived from experimentation with sample audio data. The threshold function depends on the maximum value of the decomposed wavelet coefficients, the root mean square (rms) value, and the average (mean) value of the wavelet coefficients for the signal under consideration.

So-called birdy noise occurs as components are switched on and off near the threshold. Although the chirp effect is not as annoying as the broad-band noise, the birdy noise is often irritating if it is not masked by the signal itself.

One way to alleviate the birdy noise problem is to use a soft threshold instead of a hard one. After selecting the nominal threshold level using the scheme shown in Fig. 4, soft thresholding is done to the coefficients below the threshold level. The governing equation for the soft threshold is given by

$$\alpha_m(i,j) \leftarrow \alpha_m(i,j)\, e^{[\alpha_m(i,j)/\tau - 1]/\gamma}, \qquad \alpha_m(i,j) \leqslant \tau \tag{3}$$

where $i$ is the shift parameter and $j$ the dilation parameter, $\gamma$ is the decay constant, and $\tau$ the threshold level. A small value of $\gamma$ gives results that are similar to hard thresholding, and a higher value implies that thresholding is hardly done. The suggested value of $\gamma$, chosen after some experimentation, is 0.2.

An exponential function was chosen as the kernel for the soft thresholding since all functions in general can be approximated and represented by exponential functions. Moreover, the exponential function is a smooth function and decays asymptotically to zero.

Reconstruction (via the inverse wavelet transform) is performed using the soft thresholded coefficients. The resulting signal is expected to contain a reduced amount

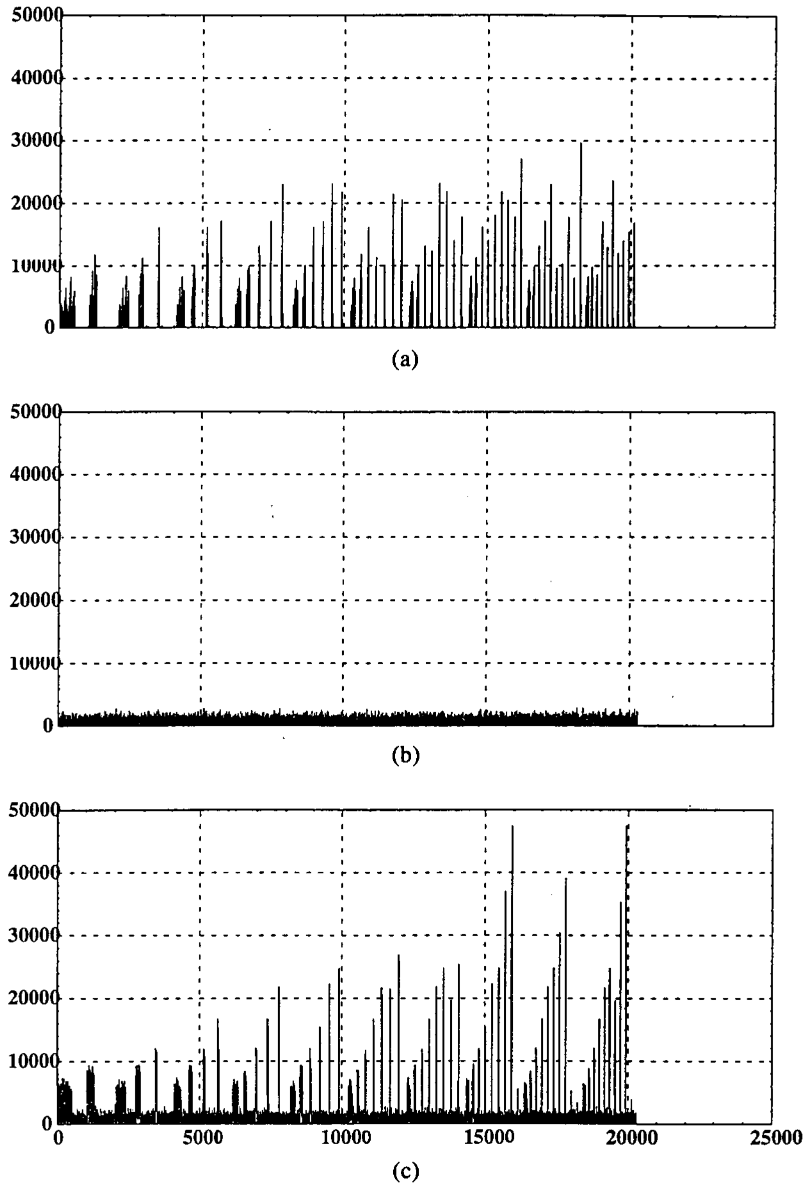


Fig. 2. Wavelet coefficients. (a) Noisefree audio signal. (b) Broad-band noise. (c) Broad-band noise corrupted audio signal. $x$ axis—wavelet index, $y$ axis—magnitude of wavelet coefficient.

of broad-band noise.

The block diagram of the modified algorithm is shown in Fig. 5.

### 3.1 Overlapping Windows

A naive way to segment a signal into frames is to divide the signal into nonoverlapping segments, which is equivalent to multiplying the signal by an abrupt segmentation function. This type of windowing can give rise to sharp, regularly spaced discontinuities in the reconstructed signal. On playback, the background noise is noticeably reduced in the enhanced audio file, but the periodic clicks caused by the abrupt window segmentation degrade the subjective quality of the reconstructed signal.

A better approach to alleviate the problems caused by abrupt window transitions is to employ overlapping windows. In this overlapping window approach the input signal is divided into overlapping frames. Each frame is decomposed using the best basis, followed by soft thresholding and then reconstruction of the time-domain signal from the thresholded wavelet coefficients. A cross-fade function derived from the raised-cosine (Hann) window technique is used to reconstruct the sample values in the overlapping regions. This scheme is shown in Fig. 6. The cross-fade technique merges the reconstructed sample values of the previous frame with the reconstructed sample values of the current frame at the overlap region of the two frames. An overlap length of 128 samples was used with satisfactory results. The frame "hop" or spacing is 2048 − 128 samples long.

### 3.2 Additional Reduction of Birdy Noise

Although the soft threshold reduced the audibility of the birdy noise, additional reduction was still desirable. Since the wavelet coefficients usually do not vary much

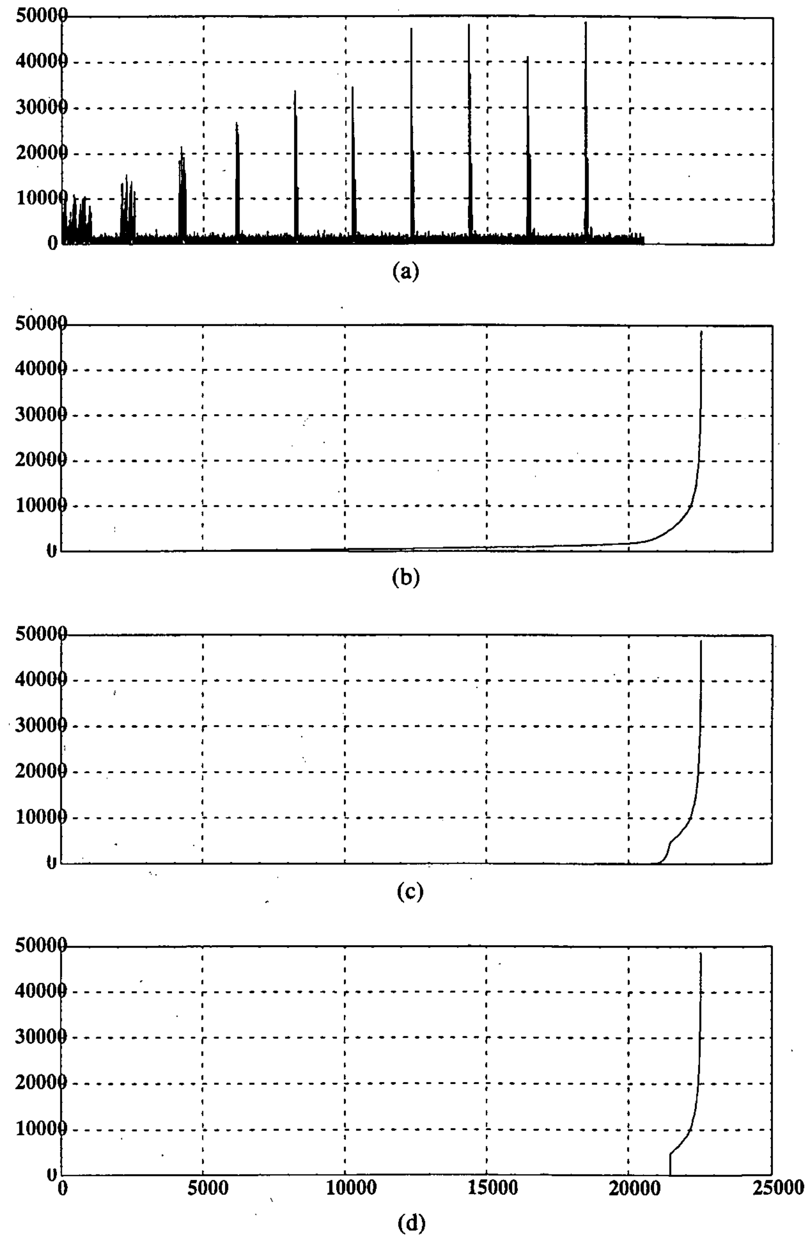


Fig. 3. Noisy audio signal. (a) Unsorted wavelet coefficients. (b) Sorted wavelet coefficients. (c) Soft thresholded coefficients. (d) Hard thresholded coefficients. $x$ axis—wavelet index; $y$ axis—magnitude of wavelet coefficient.

from one frame to the next, coefficients that have not been thresholded in the adjacent frames are left without thresholding, even if they lie below the current adaptive threshold level. This knowledge of preceding and succeeding frames is used to reduce the birdy noise level in the processed signals in the manner of hysteresis. For example, if $\alpha^{(k-1)}(i,j)$, $\alpha^{(k)}(i,j)$, $\alpha^{(k+1)}(i,j)$ are wavelet coefficients for the frames $k-1, k, k+1$, respectively, and $\tau^{(k-1)}$, $\tau^{(k)}$, $\tau^{(k+1)}$ are the threshold levels for the wavelet coefficients for the frames $k-1$, $k$, $k+1$,

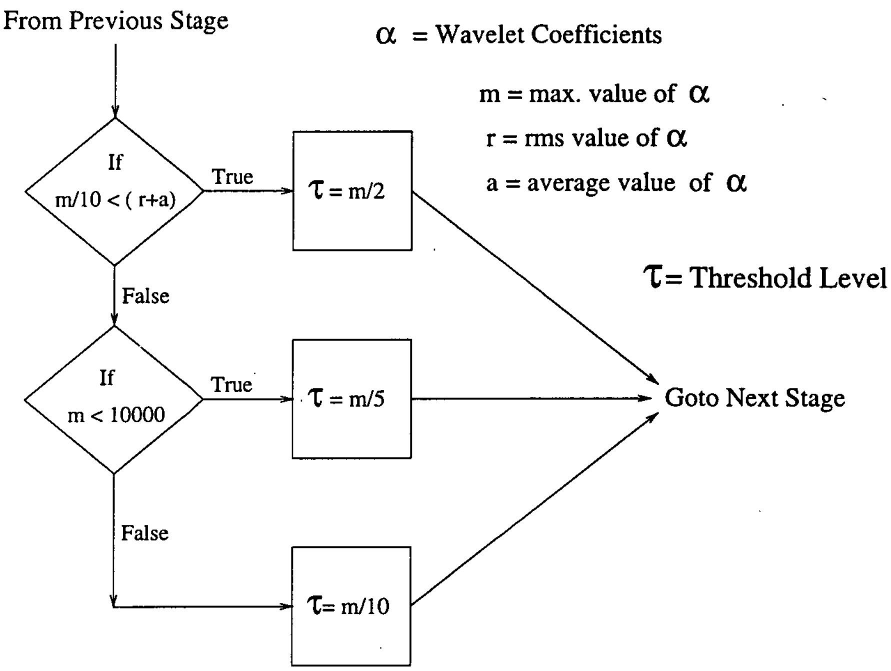


Fig. 4. Scheme for threshold level selection.

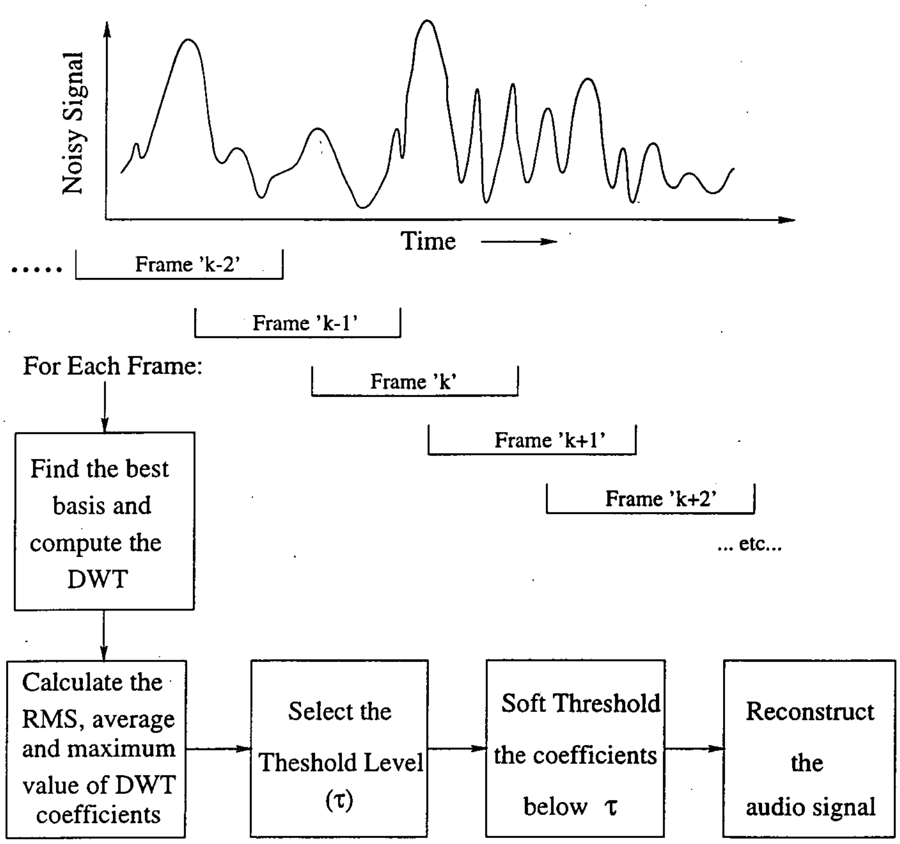


Fig. 5. Block diagram of modified algorithm.

respectively, then

$$\text{if } \alpha^{(k-1)}(i,j) \leq \tau^{(k-1)} \quad \text{and} \quad \alpha^{(k+1)}(i,j) \leq \tau^{(k+1)} \tag{4}$$

the threshold *is* applied to coefficient $\alpha^{(k)}(i, j)$ if $\alpha^{(k)}(i,j) \leq \tau^{(k)}$. However, if $\alpha^{(k)}(i,j)$ is above the threshold level, then *no* thresholding is done. With this approach the birdy noise was reduced appreciably, although it still could be heard faintly on certain examples. In any case, this approach seems to be more effective than averaging in the time or the frequency domain.

## 4 SIMULATIONS

To test the techniques described in Section 3 various audio signals with 5- to 10-second duration were considered. All of the examples were 16-bit mono signals sampled at 44.1 kHz. These signals were artificially corrupted with −10-dB, −20-dB, and −30-dB levels of additive broad-band (white) noise, and then processed with the modified algorithm proposed earlier.

The enhanced signals were then compared with the undegraded (original) audio signals. Informal listening tests performed on these signals show that the proposed method gave a noticeable reduction of the additive noise. The quality of the enhanced signal was judged quite close to that of the original signals, except for the introduction of occasional clicks and background warble or birdy noise due to the processing of the signal. In the informal tests, 11 out of 14 subjects preferred the enhanced signal from the modified algorithm when compared to the enhanced signal from the original method.

Fig. 7 shows the results of processing a few time frames of the *xan.aif* (synthesized music) audio signal corrupted with −10 dB of broad-band noise, and Fig. 8 gives the corresponding Fourier transform.

Fig. 9 shows the results of processing a few time frames of the *vibraphone.aif* audio signal corrupted with −10 dB of broad-band noise, and Fig. 10 gives the corresponding Fourier transform. The enhanced waveforms in both cases can be seen to be largely free of broad-band noise.

Figs. 11 and 12 present the wavelet coefficients of the enhanced frame and the error between the wavelet coefficients of the corrupted and the enhanced frames.

## 5 DISCUSSION

The modified algorithm presented in this engineering report removes broad-band noise to a great extent from the music signals, but at the same time introduces faint

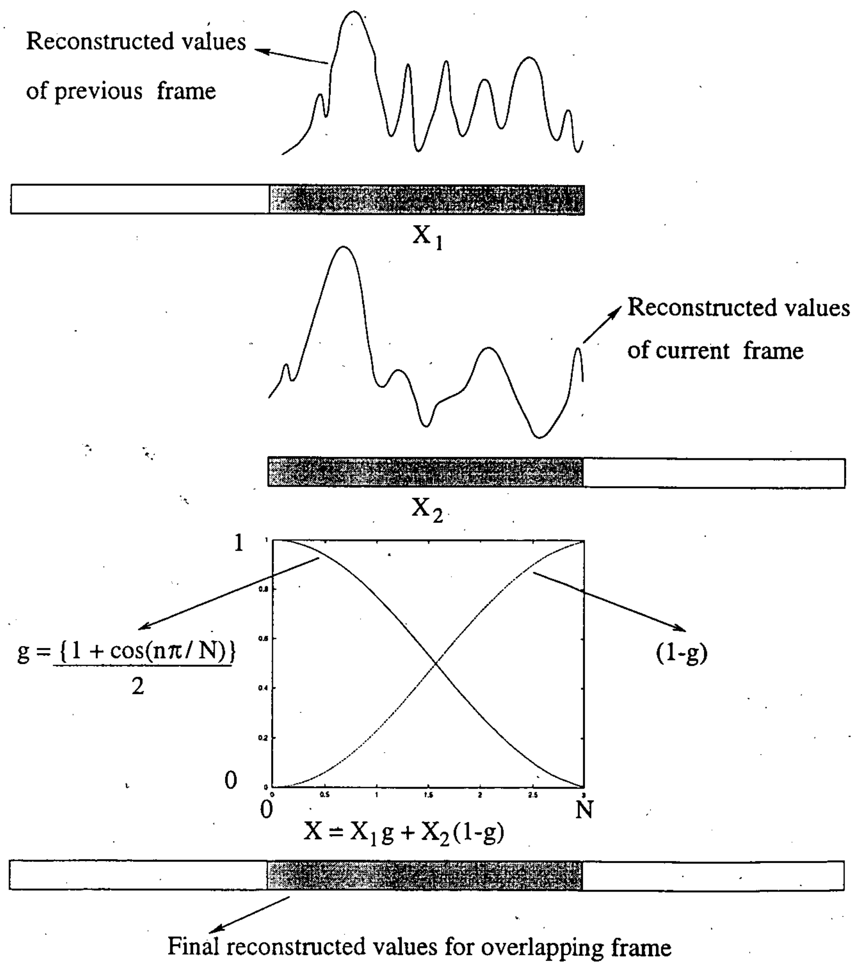


Fig. 6. Scheme for cross-fade technique.

birdy noise in the processed output signals. The algorithm has some difficulty with high percussion instruments, presumably because of the similarity between the inharmonic components of the percussion signals and the corrupting broad-band noise. Apart from degrading the high percussion instruments, the algorithm works well on a wide range of music signals.

One problem with the denoising algorithm is that it seems to work unevenly in the regions of high and low activity of the audio signal. The high-energy portions of the music are cleaned and made more vibrant, whereas denoising the low-energy portions of music sounds unnatural because too much harmonic content is removed.

The incremental modifications made to the method of Berger et al. [8] provide the following benefits:

- The enhanced audio signal retains more of the desired signal structure and contains fewer clicks and whistles compared to the original method.
- The original method is geared more toward the problem of data compression, whereas the proposed method is intended for the problem of audio enhancement only.
- The threshold value in the original method must be specified in advance, whereas the proposed method automatically selects the threshold adaptively for

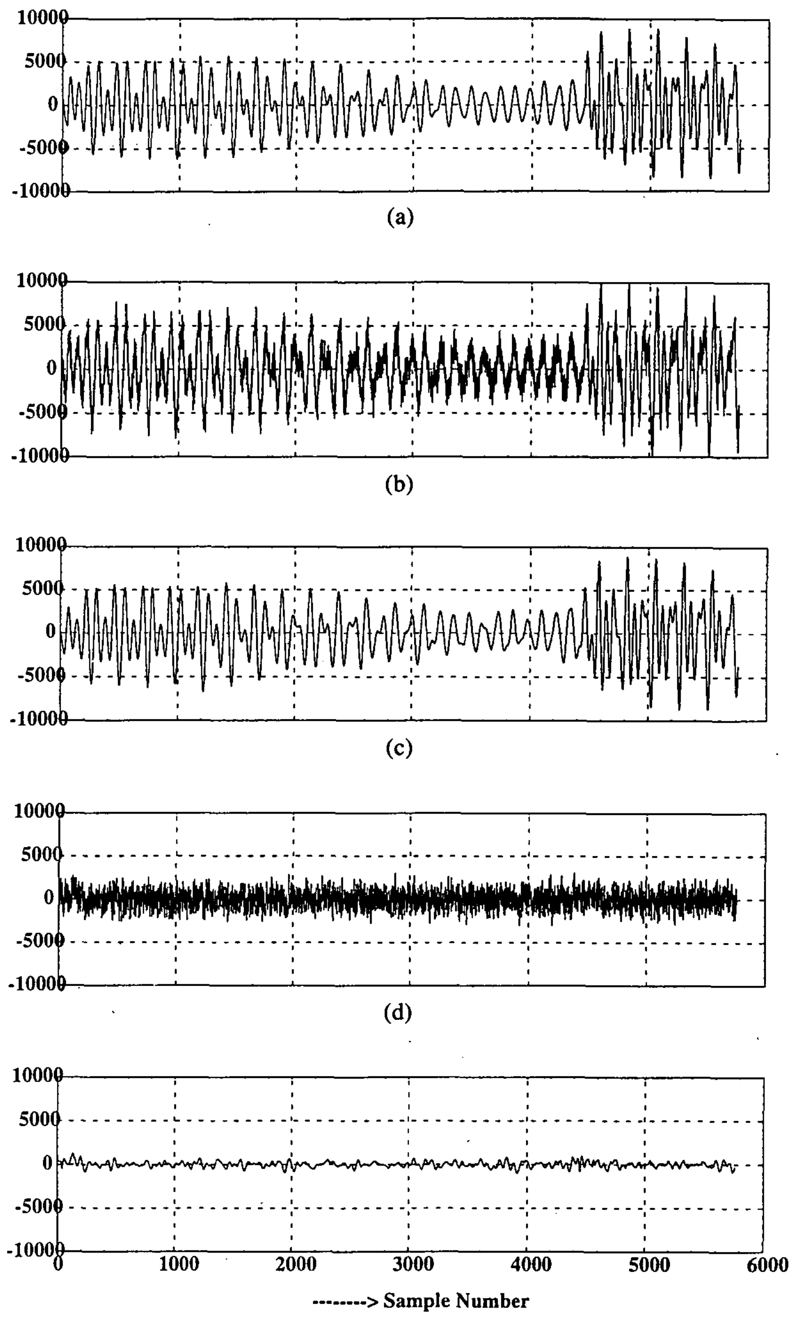


Fig. 7. (a) Time frames of *xan.aif* signal. (b) Signal corrupted with broad-band noise $-10$ dB below signal level. (c) Enhanced signal using modified algorithm. (d) Difference between (b) and (c). (e). Residual error between (a) and (c). *x* axis—wavelet index; *y* axis—magnitude of wavelet coefficient.

each frame.

- There is less computational overhead in the proposed noniterative algorithm when compared to the original method.

## 6 CONCLUSION

A denoising algorithm based on the method of Berger et al. [8] has been suggested and successfully implemented to reduce broad-band noise in recorded music signals. The algorithm is based on the assumption that the broad-band noise cannot be efficiently represented by any of the bases in a library of wavelets. The algorithm has been implemented and evaluated with the goal that the enhanced music should be as clean as possible without removing a noticeable amount of the musical content. The enhanced audio signals processed by this scheme are found to be largely free from the broad-band noise, although some minor problems, such as faint birdy noise and warble arising from the processing method, still remain for future investigation.

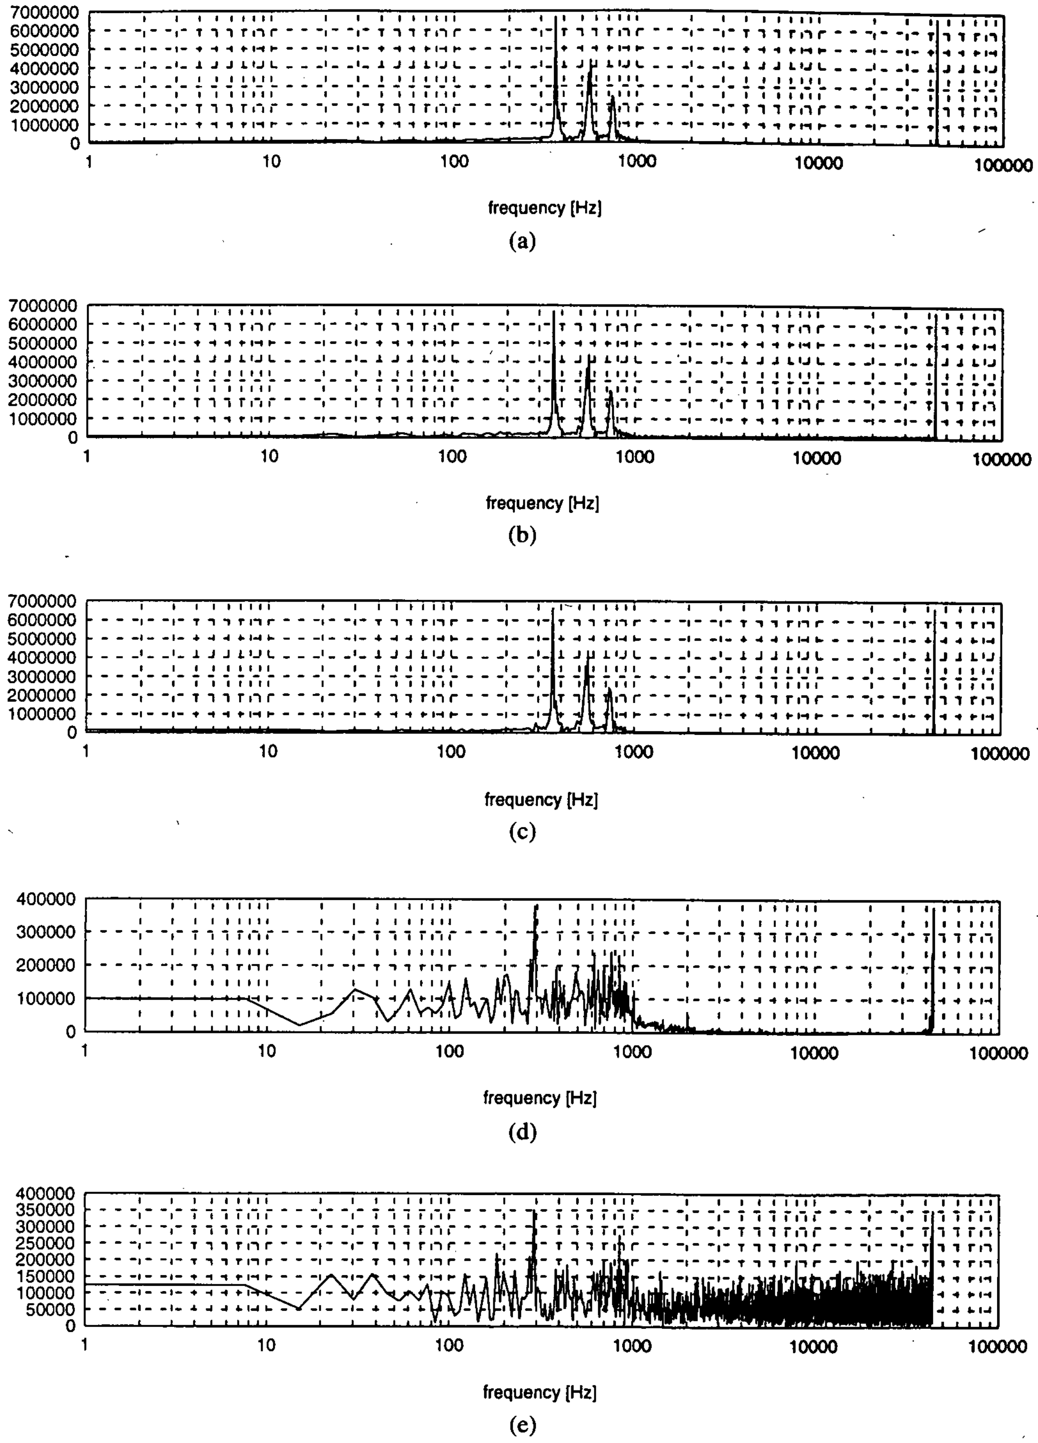


Fig. 8. Fourier transforms. (a) Time frames of *xan.aif* signal. (b) Signal corrupted with broad-band noise − 10 dB below signal level. (c) Enhanced signal using modified algorithm. (d) Difference between (b) and (c). (e) Residual error between (a) and (c). $y$ axis—magnitude of Fourier coefficient.

## 7 REFERENCES


[1] J. S. Lim, *Speech Enhancement* (Prentice-Hall, Englewood Cliffs, NJ, 1983).

[2] S. F. Boll, "Suppression of Acoustic Noise in Speech Using Spectral Subtraction," *IEEE Trans. Acoustics, Speech, Signal Process.*, vol. ASSP-27, pp. 112–120 (1979 Apr.).

[3] S. V. Vaseghi and R. Frayling-Cork, "Restoration of Old Gramophone Recordings," *J. Audio Eng. Soc.*, vol. 40, pp. 791–801 (1992 Oct.).

[4] S. V. Vaseghi and P. J. W. Rayner, "A New Application of Adaptive Filters for Restoration of Archived Gramophone Recordings," in *Proc. IEEE Int. Conf. on Acoustics, Speech, Signal Processing* (1988), pp. 2548–2551.

[5] S. V. Vaseghi, "Algorithms for Restoration of Archived Gramophone Recordings," Ph.D. thesis, University of Cambridge, Cambridge, UK (1988).

[6] Y. Ephraim and H. L. Van Trees, "A Signal Subspace Approach for Speech Enhancement," *IEEE Trans. Speech and Audio Process.*, vol. 3, pp. 251–266 (1995 July).

[7] R. R. Coifman and M. V. Wickerhauser, "Entropy-


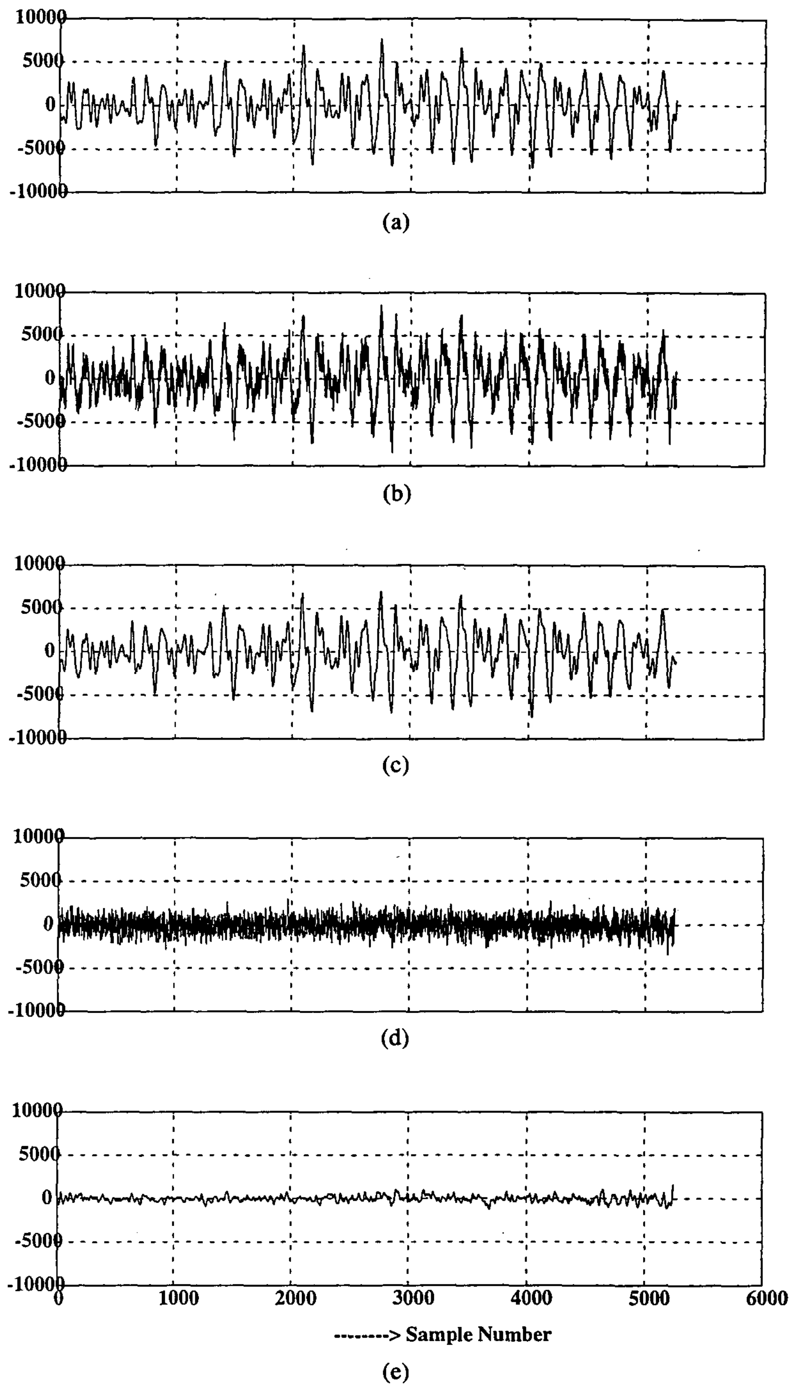


Fig. 9. (a) Time frames of *vibraphone.aif* signal. (b) Signal corrupted with broad-band noise −10 dB below signal level. (c) Enhanced signal using modified algorithm. (d) Difference between (b) and (c). (e) Residual error between (a) and (c): $x$ axis—wavelet index; $y$ axis—magnitude of wavelet coefficient.

Based Algorithms for Best Basis Selection," *IEEE Trans. Inform. Theory*, vol. 38, pp. 713–718 (1992 Mar.).

[8] J. Berger, R. R. Coifman, and M. J. Goldberg, "Removing Noise from Music Using Local Trigonometric Bases and Wavelet Packets," *J. Audio Eng. Soc.*, vol. 42, pp. 808–818 (1994 Oct.).

[9] N. Saito, "Local Feature Extraction and Its Applications Using a Library of Bases," Ph.D. dissertation, Yale University, New Haven, CT (1994 Dec.).

[10] G. Kaiser, *A Friendly Guide to Wavelets* (Birkhäuser, Boston, 1994).

[11] I. Daubechies, "Orthonormal Bases of Compactly Supported Wavelets," *Comm. Pure Appl. Math.*, vol. 41, pp. 909–996 (1988).

[12] I. Daubechies, "Ten Lectures on Wavelets," in *CBMS-NSF Regional Conference Series in Applied Mathematics*, vol. 61, pp. 909–996 (1992).

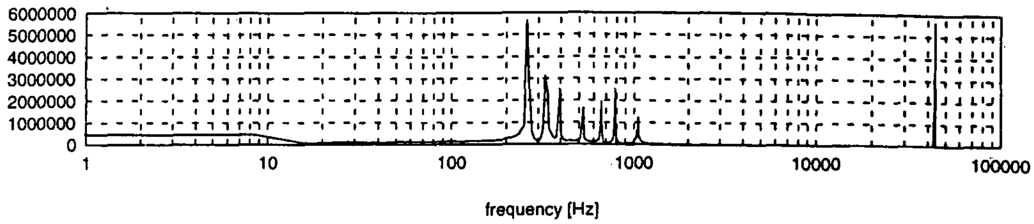


(a)

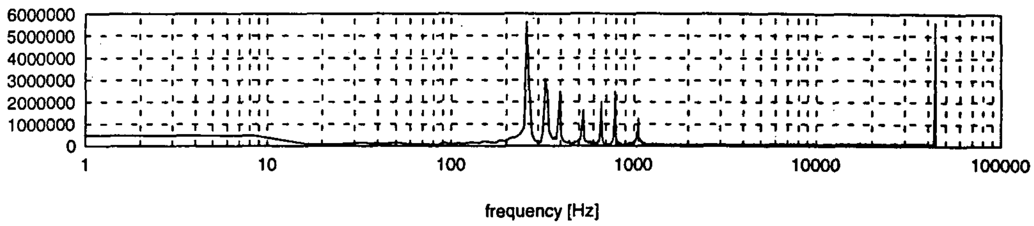


(b)

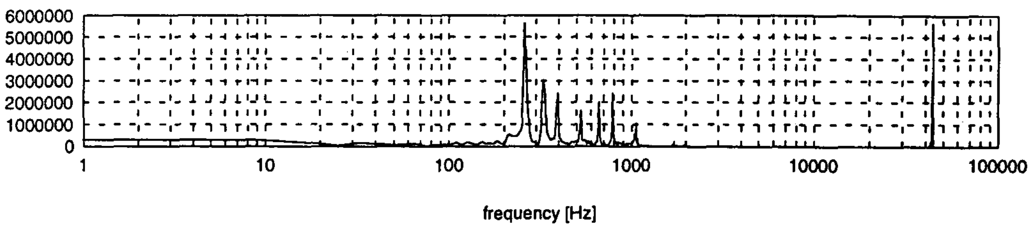


(c)

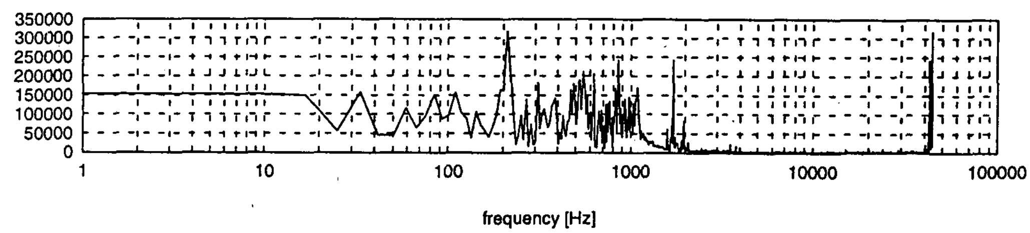


(d)

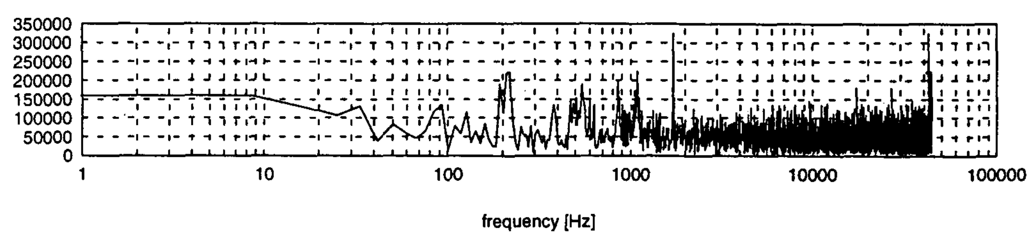


(e)

Fig. 10. Fourier transforms. (a) Time frames of *vibraphone.aif* signal. (b) Signal corrupted with broad-band noise − 10 dB below signal level. (c) Enhanced signal using modified algorithm. (d) Difference between (b) and (c). (e) Residual error between (a) and (c). *y* axis—magnitude of Fourier coefficient.

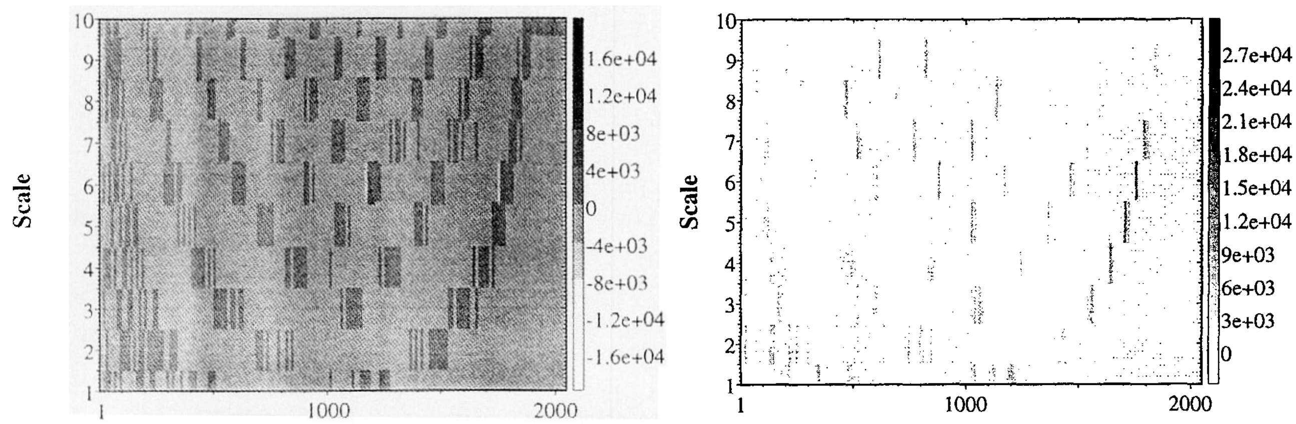


(a) (b)

Fig. 11. Contour plots. (a) Wavelet transforms for enhanced audio signal. (b) Difference in wavelet transforms between corrupted and enhanced signals. $x$ axis—time index, 3rd axis (darkness of pixel location)—magnitude of wavelet coefficient.

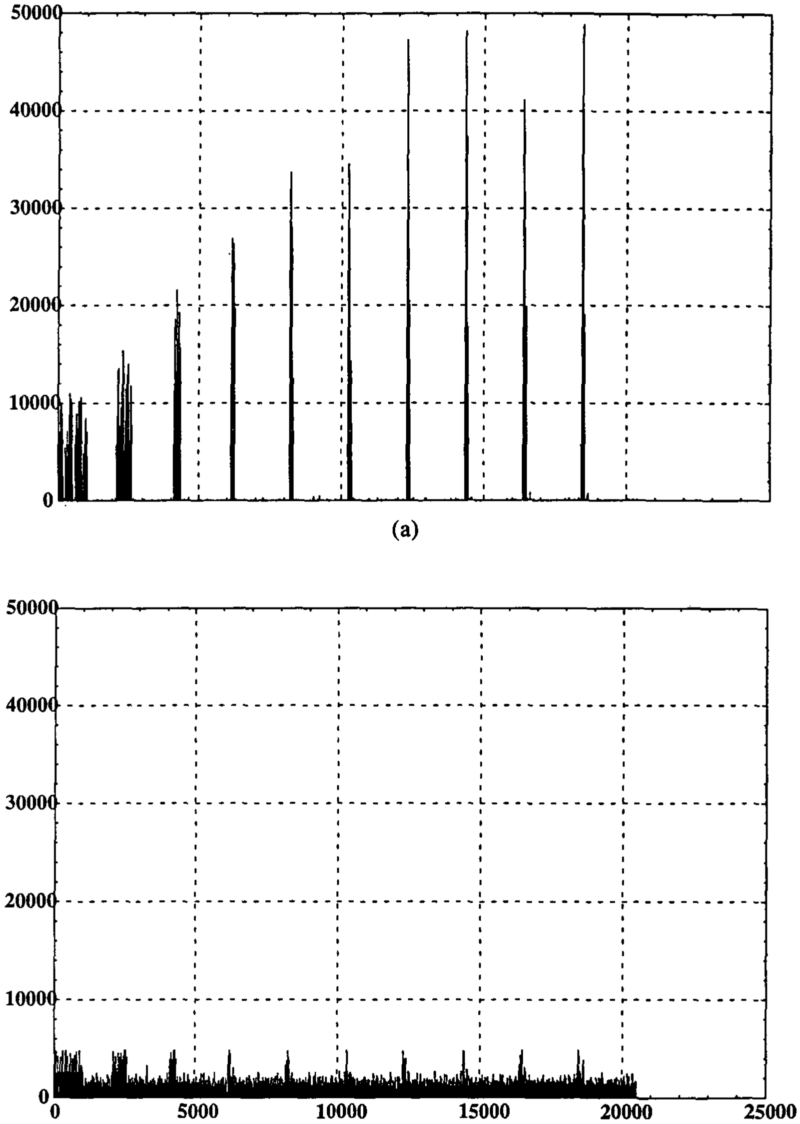


(b)

Fig. 12. Wavelet coefficients. (a) Enhanced audio signal. (b) Error between corrupted and enhanced audio signals. $x$ axis—wavelet index, $y$ axis—magnitude of wavelet coefficient.

## THE AUTHORS

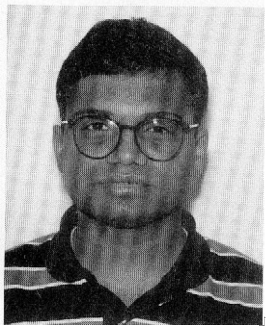

P. K. Ramarapu

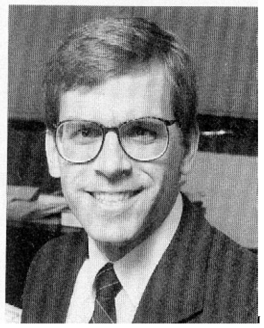

R. C. Maher

Pavan K. Ramarapu was born in Hyderabad, India, in 1973. He received a B.E. degree from Osmania University, Hyderabad, Indian, in 1994, and an M.S. degree from the University of Nebraska-Lincoln, in 1996, both in electrical engineering. From 1996 to 1997 he was a doctoral student at the University of Nebraska. In 1997 he joined Zilog, Inc., of Campbell, California, where he is involved with modem system design. His interests include digital signal processing and telecommunications.

•

Robert C. (Rob) Maher was born in Cambridge, UK, in 1962. He holds a B.S. degree from Washington University, St. Louis, MO (1984), an M.S. degree from the University of Wisconsin-Madison (1985), and a Ph.D. degree from the University of Illinois-Urbana (1989), all in electrical engineering. While a doctoral student, he worked as a graduate research assistant under James Beauchamp with the University of Illinois Computer Music Project. His doctoral dissertation dealt with the separation of voices in musical signals. Dr. Maher joined the University of Nebraska-Lincoln in 1989, serving first as an assistant professor (1989–1995) and then as a tenured associate professor (1995–96) with the Department of Electrical Engineering. His primary teaching and research interests involved applications of digital signal processing to audio engineering, music synthesis, and acoustics. He received several teaching awards, including a University of Nebraska College Distinguished Teaching Award in 1995.

Dr. Maher left academia in 1997 to become director of engineering with EuPhonics, Inc., of Boulder, Colorado, where he was recently promoted to vice president of engineering. His research and development efforts involve advanced digital music synthesis, sound localization techniques, and audio DSP architectures. He is a member of several professional organizations, including the Audio Engineering Society, IEEE, Acoustical Society of America, ASEE, ICMA, and the Tau Beta Pi, Eta Kappa Nu, Phi Kappa Phi, and Sigma Xi honor societies.