INSTITUTE OF PHYSICS PUBLISHING PHYSIOLOGICAL MEASUREMENT

Physiol. Meas. **24** (2003) 1–15 PII: S0967-3334(03)58178-3

Time–frequency modelling and discrimination of noise in the electrocardiogram

Piotr Augustyniak

Institute of Automatics, University of Science and Technology,
30 Mickiewicza Ave. 30-059 Krakow, Poland

Received 2 January 2003, in final form 9 June 2003
Published DD MMM 2003
Online at stacks.iop.org/PM/24/1

Abstract

In widely spread home care applications of ECG recorders, the traditional approach to the problem of noise immunity is no longer sufficient. This paper presents a new ECG-dedicated noise removal technique based on a time–frequency noise model computed in a quasi-continuous way. Our algorithm makes use of the local bandwidth variability of cardiac electrical representation and splits the discrete time sequence into two sub-planes. The background activities of any origin (muscle, power line interference, etc) are measured in the regions of the time–frequency plane, situated above the local bandwidth of the signal. The noise estimate on each particular scale is non-uniformly sampled and needs to be extrapolated to the regions where the components of cardiac representation are normally expected. On the lower scales, the noise contribution is computed with the use of square polynomial extrapolation. The time–frequency representation of noise, partially measured and partially calculated, is arithmetically subtracted from the noisy signal, and the inverse time–frequency transform yields a noise-free cardiac representation. The algorithm was tested with the use of CSE database records with the addition of MIT-BIH database noise patterns. The static and dynamic performance of the algorithm is sufficient to ameliorate the signal-to-noise ratio by more than 11 dB.

Keywords: electrocardiography, noise removal techniques, time–frequency domain

1. Introduction

Noise removal is very often addressed in biomedical signal recording techniques. The primary reason is the unknown and unstable recording environment (unwanted signals, poor electrodes,

electromagnetic pollution, etc). The second reason is the high importance of these signals in providing the appropriate treatment, particularly in life-critical circumstances. **See endnote 1**

Usually, two principal sources of ECG noise can be distinguished: the 'technical' caused by the physical parameters of the recording equipment and the 'physiological' representing the bioelectrical activity of living cells not belonging to the area of diagnostic interest (also called background activity) (Zareba *et al* 2001). Both sources produce noise of random occurrence, overlapping the ECG signal in both time and frequency domains. Therefore, many noise removal techniques were recently developed including signal averaging (Moss and Stern 1996), adaptive noise cancelling (Akay 1994) or wavelet-based noise reduction (Akay 1998, Krishnan and Rangayyan 2000, Nikolaev *et al* 2001, Nikolaev and Gotchev 1998, Paul *et al*). These techniques, however, could only be applied in particular conditions and are often not suitable for home care recordings when the broadband noise contribution varies in energy. The need for an improved noise discrimination method follows the spread of wearable devices for pervasive cardiac monitoring. The expected solution should yield a signal suitable for automated interpretation, even if an untrained user operates the recorder in life-critical conditions.

Because of the unknown origin, noise has to be considered as a random sequence, which any component may occur at any time (Proakis 1989). The electrical cardiac representation, however, is predictable to a considerable extent and the local bandwidth varies with signal content (Augustyniak and Tadeusiewicz 1999). The discrete representation contains both components, and fortunately in some time periods (e.g. baseline) and above a given frequency, cardiac activity is not expected and thus the noise level can be reliably measured. The remaining part of the noise has to be extrapolated from the values measured. **See endnote 2**

The temporal distribution of diagnostic information in the discrete heartbeat representation is correlated with start- and endpoints of the P, QRS and T waves. This relationship may be deduced from the physiological background of the electrocardiography, but has also been demonstrated with the use of statistical tools (Augustyniak 2000, 2002). The practical advantage of relying on these points is that their positions are computed by standard diagnostic software with the acceptable accuracy. Among the other regions, the baseline (i.e. PQ section) is a first-choice example of a documented lack of cardiac electrical activity. During this period the stimulus is slowly conducted through the atrio-ventricular node, situated close to the geometrical centre of the heart, and therefore the baseline is commonly accepted as a reference for other measurements.

Using the baseline as an information-free region of the record, we follow the example of audio recording denoising technique (Haines 2002) and assume that the noise pattern can be measured there. This approach, unfortunately, has several limitations in a real application of ECG recordings:

- The baseline is short; typically it lasts for 60 ms, and thus the measured noise spectrum starts at ca 17 Hz.
- The baseline occurs once per heartbeat, thus the noise pattern could be updated in long irregular intervals only.
- The baseline may not be present in rare cases such as atrial flutter, fibrillation or premature R on T beats.

These limitations may be overcome using a quasi-continuous noise model, maximizing the number of noise measurement samples. This model is computed for each consecutive heartbeat with the use of individually detected wave borders and general knowledge about the locally expected bandwidth of cardiac components. This idea is the key point of our novel ECG-dedicated, denoising algorithm.

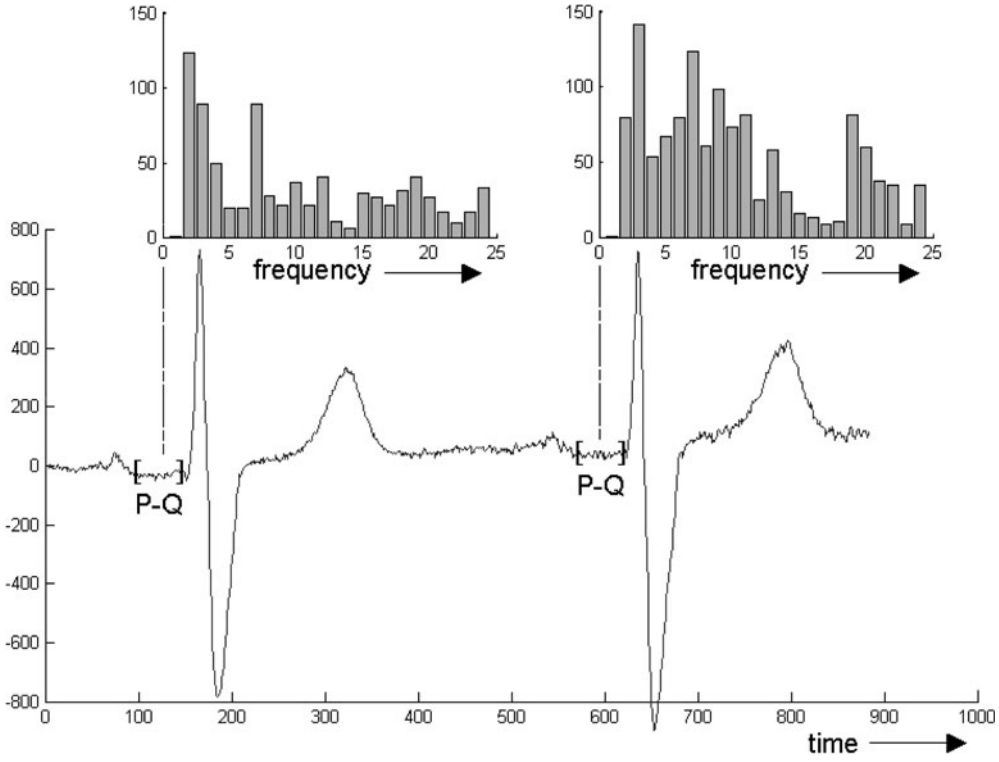

Figure 1. Measurement of the electrocardiogram noise in the baseline sections—the noiseprint spectrum is limited and unstable.

2. Materials and methods

2.1. Sparse measurement of background activity

The idea of denoising the ECG signal is based on two alternately applied procedures: updating the noise model and the subtraction of noise from the raw signal. The assumption of a lack of heart electrical activity at the baseline allows the opportunity to measure noise in two steps. First, each heartbeat is detected and the position of the PQ segment is determined with the use of a standard procedure designed for diagnostic purposes (Morlet 1986). Then the noise spectrum is computed and stored for the denoising algorithm applied immediately thereafter for the signal section representing this heartbeat (figure 1). This rather classical denoising scenario was the starting point for our research on the noise discrimination technique, based on the adaptive quasi-continuous noise model.

The background activity of muscular origin may randomly appear on the electrocardiogram. Since the time between two adjacent baseline segments is relatively long, the assumption of temporal stability of the noise source is only partially fulfilled. In such a case the use of a locally time-invariant noise pattern would produce discontinuities at the section ends falling somewhere in the PQ section. The denoised signal would be locally contaminated by noise spikes that are easily recognizable to the expert, but cause many errors during automatic signal interpretation.

Our first approach, aimed at filling the information gap between two noise measurement points, was the linear interpolation of the noise spectrum (Augustyniak 2000). This technique helped avoid abrupt changes in noise patterns, but in general performed poorly in the presence of variable-energy, wide-spectrum, background signals such as muscular activity. The particularly troublesome drawback of this method was the anticipation of noise cancelling before the noise really emerges or fades. In the case of noise changes in the P-wave surroundings, the ST segment and the T wave of the precedent heartbeat were influenced by a linear change in the noise pattern.

Another, more elaborated noise variability tracking technique was the polynomial interpolation of time–frequency coefficients in the three upper scales (Izworski and Augustyniak 2002). The noise was acquired at the baseline only and was represented on each scale by the average value of coefficients from the PQ section. The polynomial interpolation better solved the problem of rapid changes of background activity for higher polynomial orders. Unfortunately, the use of high order polynomials increased the unwanted oscillations, and the noise pattern stabilized for a long time after the noise step. Moreover, the theoretical response slope was lost at higher frequencies due to the use of averaged values. This approach did not correctly solve the problem of heart rate variability, since the series of noise measurements at consecutive heartbeats was in fact irregularly spaced.

2.2. Quasi-continuous modelling of ECG noise

Due to its simplicity, the direct measurement of the noise spectrum at the baseline provides only a rough estimate of the extra-cardiac activity representation within the signal. This was found insufficient to accurately separate the noise from the signal and was identified as a reason for poor performance or noticeable distortions. Therefore, the focus point of our new algorithm is the adaptation of the noise model to the variations of electrical background activity.

During the study of various noise characteristics recorded in the ambulatory ECG, the power line interference was found to be the principal, extra-cardiac, signal source from an energetic aspect. However, owing to the short term stability and very narrow bandwidth, it could be efficiently eliminated through technical solutions such as differential measurement (Normann 1988) or the adaptive line enhancer (Akay 1994) and, therefore, will not be discussed further in our research. The secondary, but much more problematic source of the noise was the electrical activity of muscles, because of temporal variability and wide spectrum overlapping the ECG signal. For the unsupervised recordings in the condition of a patient's everyday life, the software post-acquisition discrimination of muscular noise is the only means to improve the signal reliability.

The direct measurement of the noise spectrum at the PQ segment is too sparse to successfully meet the requirements of fast adaptation to noise variability. Moreover, it shows a high redundancy in the upper frequency bands favouring their influence to the final noise estimate, while the main energy is concentrated in the middle bands. Therefore, the processing of the time–frequency signal representation increases the temporal precision of noise measurement at the cost of losing the spectral precision.

An innovative solution applied in our new algorithm is the heuristic bandwidth variability function, expressing the local bandwidth of a typical electrocardiogram (Augustyniak and Tadeusiewicz 1999). This general function is adapted to the shape of each consecutive heartbeat and defines the local width of frequency content of cardiac components. This approach relates the position of the noise measurement sections to the distribution of diagnostic data in the signal and prevents the misinterpretation of low energy cardiac nuances

as noise components. The expected bandwidth at the time point n is expressed by a discrete function $f(n)$:

$$f : \forall n \in \{0, 1, \ldots, N\} \rightarrow f(n) \in [0; 0.5) \tag{1}$$

representing the local relative cut-off frequency.

The adaptation of the general bandwidth variability function $f(n)$ for each particular heartbeat cancels the beat-to-beat differences in the positioning of wave borders. Let $k_1, \ldots, k_5 \in \{0, 1, \ldots, N\}$ be irregularly spaced samples of the function f and represent the standard positions of wave borders. Their order is $k_1 < k_2 < L < k_5$ and k_1 corresponds to the P-wave onset and k_5 to the T-wave end, respectively. Let $h_1, \ldots, h_5 \in \{0, 1, \ldots, M\}$ be selected by the wave border computation procedure in the current heart cycle represented by M samples. The points $h_1, \ldots, h_5$ are also irregularly spaced, but their distribution differs from those of $k_1, \ldots, k_5$. To compute the adapted bandwidth function $f'(m)$, the general function f is piecewise projected to the new irregular space as so $f'(h_i) = f(k_i)$ for each point $i = 1, \ldots, 5$:

See endnote 3

$$\forall n \in [k_i, k_{i+1}] \qquad \forall m \in [h_i, h_{i+1}] \quad f'(m) = P^{S_i}(f(n)) \tag{2}$$

and the projection scale S_i varies from section to section,

$$S_i = \frac{h_{i+1} - h_i}{k_{i+1} - k_i}. \tag{3}$$

Figure 2 displays the general bandwidth variability function $f(n)$ and its projection to two real heartbeats differing in length and wave distribution.

The adapted bandwidth function distinguishes the noise measurement region $N(j, m)$ in the time–frequency representation of the current heartbeat $R(j, n)$, where j is the scale number, with reference to the start- and endpoints of the P, QRS and T waves. The relative frequency range of each scale may be expressed as $[2^{-j-1}, 2^{-j})$. The instantaneous cardiac representation $C(j, m)$ is present only below the adapted bandwidth function $f'(m)$ and thus $C(j, m) = R(j, m)$ if $f'(m) > 2^{-j-1}$, and 0 otherwise. The complementary region $N(j, m)$, where the noise level measurement is performed, contains coefficients j, m for which $f'(m) \leqslant 2^{-j-1}$ (figure 3).

All time–frequency atoms not expected to contain the representation of cardiac activity are now used as measurement points for the frequent updating of the noise model. The percentage of the length of noise measurement sections in the total signal duration is very high on the first scale (85% on average) and decreases to 65% and 23%, respectively, on the next two scales. Consequently, the noise measurement points become sparse and irregularly spaced. To keep the noise model up to date and ready for compensating the noise variability, our algorithm uses the interpolation technique in the time domain for each scale, independently.

Below a certain frequency, the reliable measurement of noise is never possible as, according to the general bandwidth function, the scales are entirely occupied by the representation of cardiac activity. For the typical ECG signal sampled at 500 Hz, the noise measurement is performed on scales $j = 1, \ldots, 3$ (32–250 Hz) and the results are then extrapolated to frequencies below 32 Hz ($j > 3$). In fact, the low-edge frequency of the fourth scale corresponds to an interval length of 64 ms, which cannot always be entirely marked off in the reference PQ section.

2.3. Modelling noise from irregular samples

Since the local bandwidth of the cardiac representation is computed continuously, the position of noise measurement sections varies from beat to beat on each scale. Therefore, the

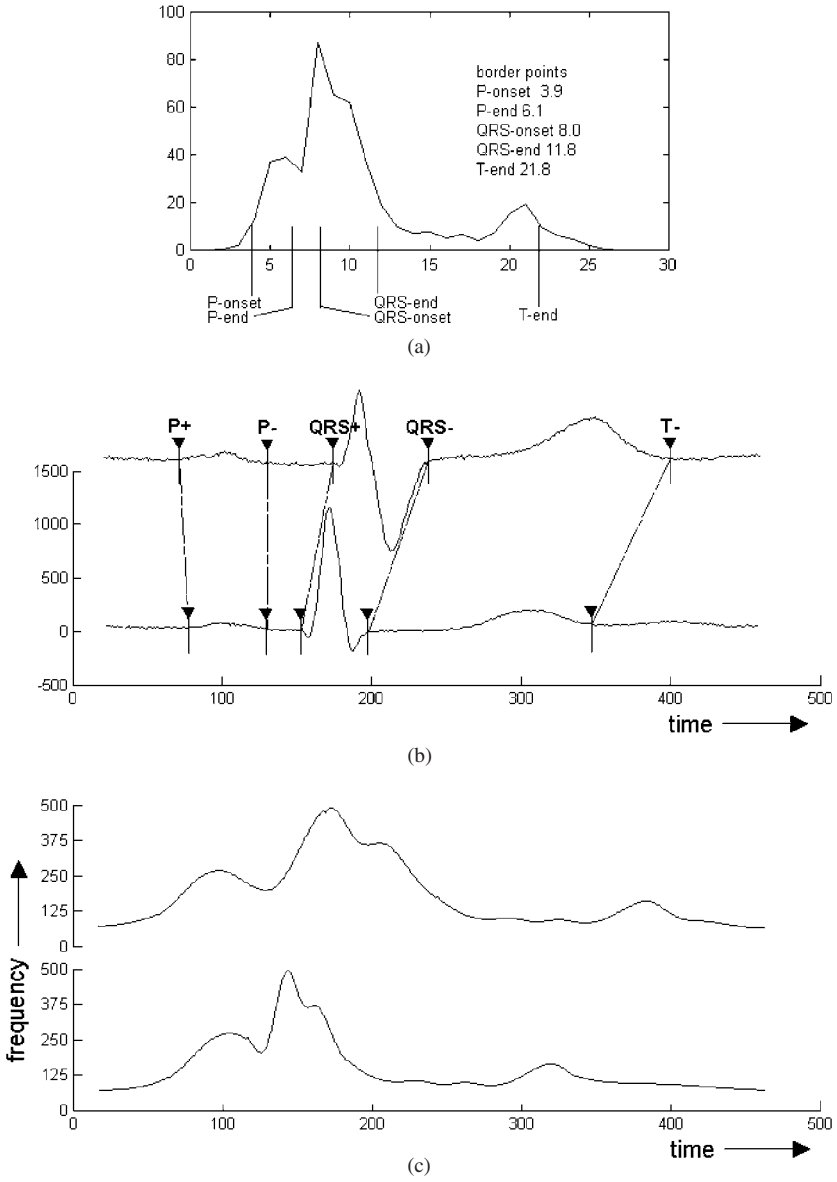

Figure 2. (a) General bandwidth variability function with standard positions of wave borders, (b) two examples of heartbeats with different segmentation point positions and (c) adapted bandwidth variability function corresponding to the heartbeats in (b).

measurement algorithm has to pick up coefficients of each scale j as a separate non-uniformly sampled time series $\{\phi_n\}_{n\in\Gamma}$ satisfying the frame condition (Aldroubi and Feichtinger 1998) for $A > 0, B > 0, f \in \mathrm{H}$:

$$A\|f\|^2 \leqslant \sum_{n\in\Gamma} |\langle f, \phi_n\rangle|^2 \leqslant B\|f\|^2. \tag{4}$$

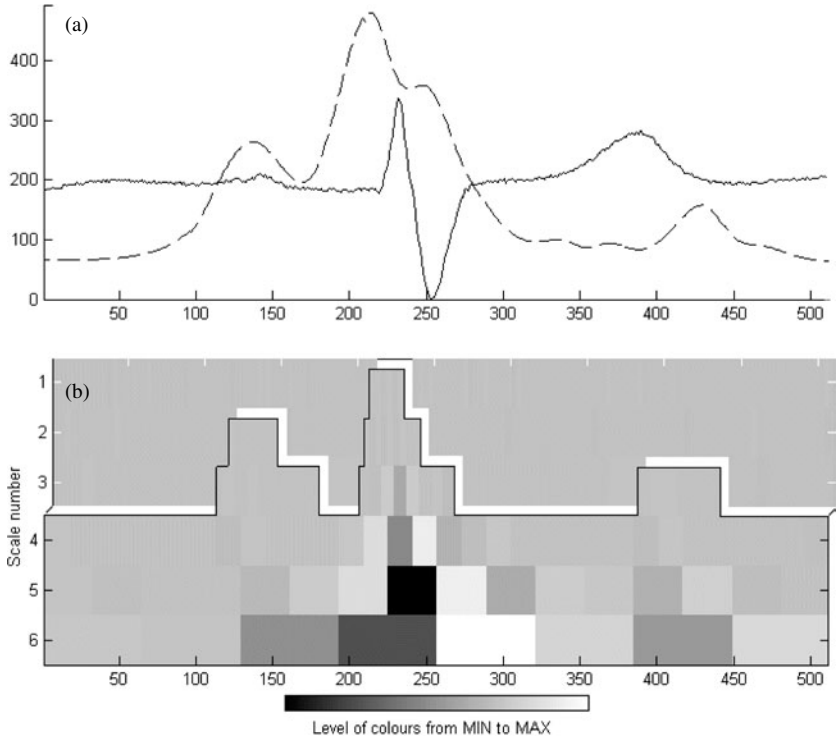

Figure 3. (a) The example heartbeat (solid line) and the adapted bandwidth variability function (dashed line), and (b) the split of the time–frequency signal representation in the noise measurement region and the cardiac representation region.

There are three sources of this irregularity:

- variability of the R–R interval;
- restriction on noise measurement in the sections containing cardiac representation;
- changes in the position of heartbeat components, which cause poor beat-to-beat reproducibility of the adapted bandwidth function.

In the sections containing the atoms of cardiac activity, or where such activity is expected, the contribution of noise has to be estimated from the previous and subsequent measured values (figure 4). This operation can be done by the projection of each non-uniformly sampled time series $N_j(\{n, v(n)\})$ to the regular space (Aldroubi and Feichtinger 1998). For this purpose, in our algorithm, we applied the cubic spline interpolation given by a continuous function:

$$S_i(x) = a_i + b_i(x - x_i) + c_i(x - x_i)^2 + d_i(x - x_i)^3 \tag{5}$$

$x \in [x_i, x_{i+1}]$, $i \in \{0, 1, \ldots, n-1\}$ best fitted to the time series N_j. The uniform representation of the noise is then obtained by sampling the $S_i(x)$ at the time points m:

$$N'_j(m) = \sum_m S_i(x)\,\delta(x - mT). \tag{6}$$

Taking into account all three processed scales N'_1, N'_2 and N'_3, the partially measured and partially computed atoms complete the time–frequency plane of noise N', from which the time-domain noise pattern can be recovered with the precision of the original signal's

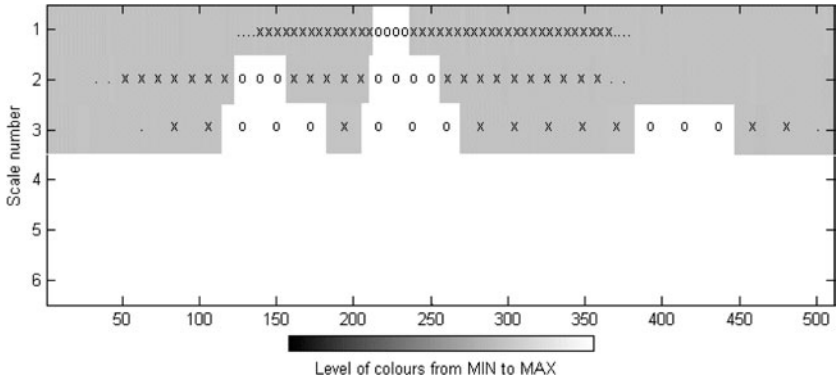

Figure 4. The distribution of noise measurement and interpolation samples on each scale. The missing values 'o' are estimated from the previous and subsequent measured values 'x'.

sampling interval. This provides justification for the name of quasi-continuous, background activity model.

2.4. Extrapolating noise to lower scales

Frequency domain noise extrapolation uses the second-order polynomial generated by coefficients of first three scales $i = 1, \ldots, 3$. The accuracy of this extrapolation is not critical with respect to muscle noise removal, since the spectrum of muscular activity representation falls rapidly at lower frequencies and the signal-to-noise ratio is usually acceptable.

Due to the dyadic structure of the time–frequency decomposition, this extrapolation includes all atoms of embedded trees originating from the considered coefficient (figure 5). Consequently, an estimation of the noise level at a given time point k on the scale j is based on three average values $M_j(k, i)$ of all corresponding atoms $s(n, i)$ on each of the first three scales:

$$M_j(k,i) = \frac{1}{2^{j-i}} \sum_{n=2^{j-i}k}^{2^{j-i}(k+1)-1} s(n,i) \tag{7}$$

For a better correlation with the skew spectrum of muscle representation, the logarithmic frequency scale (octave number instead of the absolute frequency value) is maintained during extrapolation,

$$N'(k,\, j) = a_{k,j}\, j^2 + b_{k,j}\, j + c_{k,j} \tag{8}$$

and since $i = 1, \ldots, 3$:

$$\begin{aligned} b_{k,j} &= 4(M_j(k,2) - M_j(k,1)) - \tfrac{3}{2}(M_j(k,3) - M_j(k,1)) \\ a_{k,j} &= \tfrac{1}{3}((M_j(k,2) - M_j(k,1)) - b_{k,j}) \\ c_{k,j} &= M_j(k,1) - a_{k,j} - b_{k,j}. \end{aligned} \tag{9}$$

In the lower frequency bands, the noise is more stable than in the first three scales, because the time–frequency atoms last for a longer time according to the uncertainty principle. Since the lower scales ($j > 3$), where noise measurement cannot be performed, contain 12.5% of the total number of time–frequency atoms, the frequency domain extrapolation of background activity has only little influence on the total computation costs.

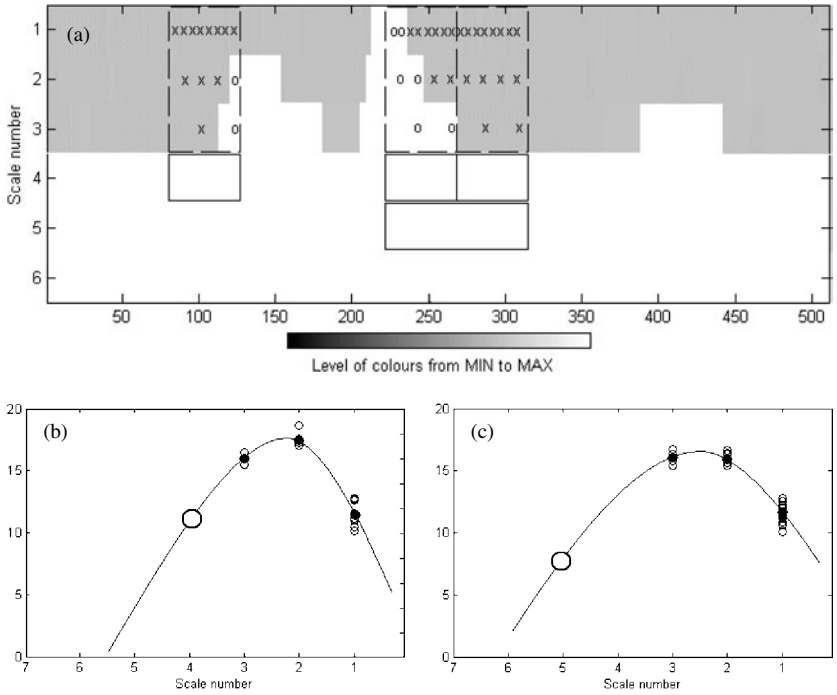

Figure 5. The extrapolation of noise values to low frequency bands: (a) averaging the noiseprint in the time domain, (b) example of the polynomial extrapolation of the fourth scale value (marked by a large empty circle) and (c) example of polynomial extrapolation of the fifth scale value.

A degree of polynomial extrapolation was chosen as a trade-off between the computational complexity and the stability of the noise model. Higher order polynomials leave excessive degrees of freedom when fitted to three octaves and for the baseline length, limiting the lower boundary of the reference spectrum the only solution would be the use of oversampled signals (1000 Hz or more). The use of lower order polynomials is equivalent to linear extrapolation and, due to the slope raising monotonically, poorly fits the average spectrum of muscular activity.

2.5. Use of the time–frequency model in noise discrimination

The time–frequency, ECG background activity model can easily be converted to the time domain, and we often performed such conversions for verification purposes of noise statistical properties. However, with respect to noise discrimination, it is very interesting to continue the processing in the time–frequency domain. The values of time–frequency atoms in the noise model $N'(j, m)$ are subtracted from the values of the corresponding atoms in the representation of the raw signal $R(j, m)$:

$$D(j, m) = R(j, m) - N'(j, m). \tag{10}$$

This operation yields a modified time–frequency plane representing the distilled cardiac signal $D(j, m)$. This plane is then fed to the inverse wavelet transform, which produces the time-domain ECG signal with discriminated noise. The noise removal by subtraction of the noise model from the raw signal representation in the time–frequency domain is possible

Table 1. The average difference of denoised and original database signals for most frequent patterns of continuous noise.

	PRD (%)				SNR (dB)			
Noise pattern	50	20	10	5	−3	−7	−10	−13
Poor electrode contact	47	13	4.7	2.7	−3.1	−8.8	−13.2	−15.7
Electromagnetic interference	17	4.3	1.3	0.95	−7.7	−13.6	−18.9	−20.2
Muscle fibrillation	14	2.3	1.0	0.85	−8.5	−16.3	−20.0	−20.7

due to the linearity of the discrete wavelet transform and by making the assumption of no correlation between the 'cardiac' and 'muscular' sources that is usually well fulfilled.

Applying further restrictions to the transformation, we aimed at increasing the computational efficiency of our algorithm. The use of pure real-valued wavelets simplifies the computations very much and helps avoid problems with cubic splines for complex samples. The use of pure integer transformation (Calderbank *et al* 1996) has an advantage of yielding the distilled signal in the integer domain.

3. Testing

Testing the noise removal procedure is a difficult task because of the lack of noise-free real signals. At the same time, the simulated electrocardiogram, although noise free, differs in diagnostically meaningful details from the real one. For this reason we performed two continuous noise tests using real and artificial records. For verification of dynamic features of the adaptive noise model, we performed two additional tests using artificial records and noise modulated by a sine and by square waves. The noise patterns in use were representative of the most frequent noise sources:

- poor electrode contact (mainly low frequency and abrupt baseline changes);
- electromagnetic interference (power line sinus wave, 60 Hz);
- muscle fibrillation (a clipping-free section was selected).

All patterns were taken from the MIT-BIH database (12 bits, 360 Hz) (Moody 1997), resampled and mixed with the signal at the following four levels of total energy: 50%, 20%, 10% and 5% (corresponding to –3 dB, –7 dB, –10 dB and –13 dB signal-to-noise ratio). The real ECG records were taken from the CSE multilead database (12 bits, 500 Hz) (Willems 1988) and the artificial records were acquired using an ECG recorder (12 bits, 500 Hz) from the test generator.

3.1. Real signals from the CSE database

The real ECG recordings already contain the noise of various origins. The original time–frequency content of the recordings is close to the expected bandwidth, but does not match it perfectly. The recorded noise has characteristics very close to those of the noise artificially added from the database. Therefore, the global noise estimate includes both noise sources in a similar way, and the discrimination procedure is not expected to yield a distilled signal close to the original. The divergence of the distilled and original signals is quite important and may easily be confused with ECG distortion. For this reason, the correlation of the original and distilled records (table 1) is not an adequate estimate for the performance of noise discrimination for the original recordings.

Table 2. The average difference of denoised and original synthesized signals for most frequent patterns of continuous noise.

	PRD (%)				SNR (dB)			
Noise pattern	50	20	10	5	−3	−7	−10	−13
Poor electrode contact	46	11	4.3	2.1	−3.4	−9.6	−13.7	−16.8
Electromagnetic interference	17	4.3	1.3	0.95	−7.7	−13.6	−18.9	−20.2
Muscle fibrillation	10	1.4	0.71	0.33	−10.0	−18.5	−21.5	−24.8

Table 3. The average difference of denoised and original synthesized signals for most frequent patterns of sine-modulated noise.

	PRD (%)				SNR (dB)			
Noise pattern	50	20	10	5	−3	−7	−10	−13
Poor electrode contact	47	13	4.5	2.4	−3.1	−8.8	−13.5	−16.2
Electromagnetic interference	17	4.4	1.4	1.1	−7.7	−13.6	−18.5	−19.6
Muscle fibrillation	11	1.6	0.78	0.37	−10.0	−17.9	−21.1	−24.3

3.2. Artificial signals with added continuous noise

The artificial ECG signals are synthesized mathematically in the generator's hardware with no random component. The time–frequency representation of the record, before mixing the database noise, perfectly matches the expected local bandwidth variability. In this case the distilled signal is expected to be identical to the original and any difference is interpreted as the remaining noise. The results given in table 2 represent the performance of our noise discrimination algorithm for most frequent noise patterns and four noise mixing levels. The best performance is observed in muscle fibrillation noise removal. The contribution of noise added at the level of half of the signal energy (−3 dB) was reduced to 10% (−10 dB).

3.3. Artificial signals with added sinus-modulated noise

See endnote 4

The test using static noise only estimates the global noise removal performance, while the advantage of our algorithm is its fast adaptability to the muscular noise changes. To reveal the dynamic properties of the proposed method, the changes of muscular activity were simulated by the modulation of noise energy with a sine wave of frequency sweeping from 1 to 10 Hz. In order to maintain the average noise level comparable with previous tests, the maximum amplitude of the mixed noise was twice as high as in the case of continuous noise. The result of this test is better represented by the time series of the remaining noise (figure 6) than by a global quantitative estimate as given in table 3. However, it is noteworthy that the quantitative results are only slightly worse than in the case of continuous noise (cf table 2).

3.4. Artificial signals with added square-modulated noise

The noise modulated with a square wave of frequency ranging from 1 to 10 Hz was used to simulate the sudden occurrence of noise in virtually any correlation with the cardiac signal. This test was performed to determine the adaptability limits of the proposed algorithm. The real interference of muscular biopotentials does not start or end abruptly, because the muscular

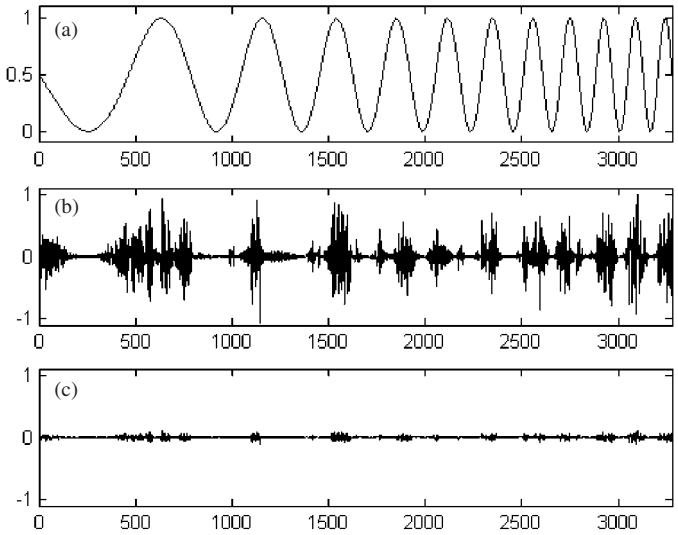

Figure 6. Result of the denoising test with the use of sinus-modulated noise: (a) modulating function, (b) added noise pattern and (c) remaining noise (amplitude scale the same as (b)).

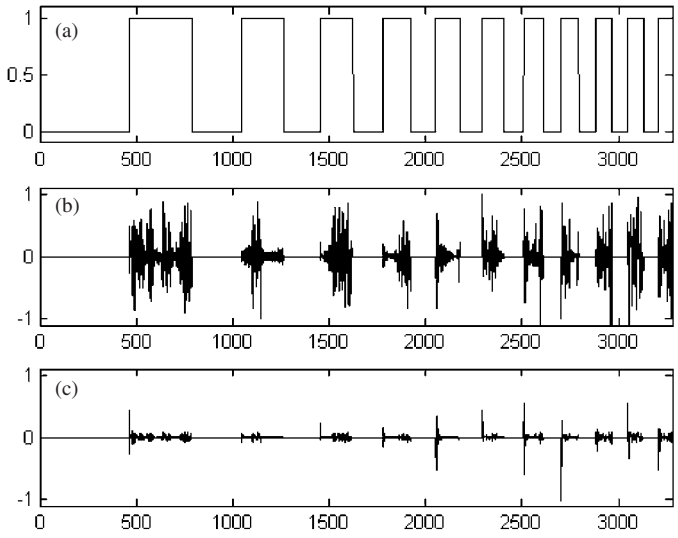

Figure 7. Result of the denoising test with the use of square-modulated noise: (a) modulating function, (b) added noise pattern and (c) remaining noise (amplitude scale the same as (b)).

fibre's triggers are not perfectly synchronized. Here again, the result is represented by the time series of the remaining noise (figure 7), where the most interesting are regions close to the noise changes. Table 4 gives a quantitative general estimate of noise removal performance for comparative purposes with those of previous tests. The results, in particular for muscle fibrillation noise, are considerably worse than in the case of static (cf table 2) or sine-modulated (cf table 3) noise, and this is due to the errors in the modelling of noise steps.

Table 4. The average difference of denoised and original synthesized signals for most frequent patterns of square-modulated noise.

	PRD (%)				SNR (dB)			
Noise pattern	50	20	10	5	−3	−7	−10	−13
Poor electrode contact	47	16	5.7	3.9	−3.1	−8.0	−12.4	−14.1
Electromagnetic interference	23	9.1	4.4	1.9	−6.4	−10.4	−13.6	−17.2
Muscle fibrillation	19	4.7	2.1	1.2	−7.2	−13.3	−16.8	−19.2

4. Discussion

The newly proposed algorithm for ECG noise discrimination is based on the definition of the local signal bandwidth. This approach maximizes the number of noise measurement points in the time–frequency plane. In consequence, the contribution of true measured values is favoured over the interpolated data in the quasi-continuous noise model. All properties of the denoising procedure are influenced by the local bandwidth definition, thus the results reported here correspond to the average temporal distribution of diagnostic information in the discrete heartbeat representation. The concept of the general bandwidth function gives the user an opportunity to define his or her own profile of interest. The commonly used example is the pursuit of ventricular late potentials (VLP), for which the high frequency cardiac components are expected in the terminal QRS section. The wideband approach, defining the signal bandwidth as equal to the Nyquist frequency everywhere, allows no assumption to be made on local signal properties, but does eliminate the points where the noise can be reliably measured. The opposite strategy defines the signal bandwidth to be much lower than the Nyquist frequency everywhere and allows the noise measurement at any time point, as in the case of oversampled signals. Appropriate bandwidth tracking enables the optimal measurement of noise, thanks to the use of general knowledge about signal properties adapted to the current content of the data stream. These general remarks are valid for a wide class of physiological and technical signals characterized by irregular data distribution.

The noise patterns applied during tests differ in spectra:

- the noise originating from poor electrode contact has a high energy in the low frequency range;
- the power line interference is a narrow-bandwidth noise very stable in time;
- the muscle fibrillation produces wide-bandwidth representation relatively stable in time with maximum energy in the middle of the frequency range of cardiac representation.

All these characteristics influence the performance of the noise modelling and discrimination algorithm. Since the poor electrode contact noise concentrates on lower scales, it cannot be correctly estimated using values measured from the three upper scales. Although tables 1–4 suggest a slight improvement in signal quality, the noise is cancelled only in the middle frequency range (fourth to sixth scales). The power line interference concentrating all the energy in practically one frequency influences the second and third scales (32, . . . , 128 Hz) all the time. The presence of this time-invariant component lowers the sensitivity of noise adaptability in those scales and leads to the inaccurate frequency-domain estimation of the noise in lower scales. The characteristics of muscle fibrillation best fit the initial assumptions: in the time domain signal changes are relatively slow and in the frequency domain, thanks to the presence of the maximum, the parabolic spectrum approximation is quite accurate. This is the reason for the best results of the newly proposed algorithm obtained in the tests with muscle noise.

The static performance of the noise suppression algorithm meets our expectations. The improvement in the average signal-to-noise ratio (by 11.6 dB, see table 2, last row) is slightly better than for the procedure described in (Nikolaev *et al* 2001). The dynamic performance of noise discrimination is also very good (by 11.1 dB, see table 3, last row) as long as the sinus-modulated noise is used. This modulation corresponds to the natural onset of muscular activity and thus justifies the practical applicability of our method. For the square-modulated noise the average signal improvement falls to 6.5 dB due to the inaccuracy of model adaptation particularly when a noise step occurs in a section where the full bandwidth is used by cardiac components. The measured noise data are locally missing and the interpolation cannot reconstruct the correct values of time–frequency coefficients. Nevertheless, this drawback is only of minor practical importance, because the full bandwidth is used only for the QRS section (typically 100 ms), which usually has a high amplitude of cardiac components. Consequently, the signal masks the noise estimation errors.

The only problem that may limit the application of our denoising algorithm to the interpretative workstations is its high complexity. The main steps of the processing: the ECG signal segmentation, the wavelet transform and the time–frequency arithmetic can be performed with the use of fixed-point data representation. However, the cubic spline noise interpolation in the time domain and the polynomial extrapolation of noise in the frequency domain still involve the use of floating-point operations.

5. Conclusion

A new ECG-dedicated method for noise modelling and discrimination was developed and tested. The noise model is quasi-continuous and optimizes the response to the physiological changes of background activity. The use of the bandwidth function allows the user to define his/her own profile of interest and to even adapt the method to other signals of variable information density.

References

Akay M 1994 *Biomedical Signal Processing* (San Diego: Academic)

Akay M (ed) 1998 *Wavelets in Biomedical Signal Processing* (New York: IEEE)

Aldroubi A and Feichtinger H 1998 Exact iterative reconstruction algorithm for multivariate irregularly sampled functions in spline-like spaces: the L^p *theory Proc. Am. Math. Soc.* **126** 2677–86

Augustyniak P 2000 Compression and denoising of the ECG using standard bandwidth variability function *Proc. IEE 1st Int. Conf. on Advances in Medical Signal and Information Processing* pp 48–53

Augustyniak P 2002 Pursuit of the ECG information density by data cancelling in time–frequency domain *IFMBE Proc.* vol 2 pp 152–3

Augustyniak P and Tadeusiewicz R 1999 The bandwidth variability of a typical electrocardiogram *Proc. European Medical and Biological Engineering Conf. EMBEC '99* pp 394–5

Calderbank A R, Daubechies I, Sweldens W and Yeo B-L 1996 Wavelet transforms that map integers to integers *Technical Report* Department of Mathematics, Princeton University

Haines R (The Coriolis Group) 2002 *Digital Audio* **See endnote 5**

Izworski A and Augustyniak P 2002 ECG noise modelling in time–frequency domain using the polynomial extrapolation *IFMBE Proc.* vol 3 pp 564–5

Krishnan S and Rangayyan R M 2000 Automatic denoising of knee-joint vibration signals using adaptive time–frequency representations *Med. Biol. Eng. Comput.* **38** 2–8

Moody G B 1997 *The MIT-BIH Arrhythmia Database CD-ROM* 3rd edn Harvard-MIT Division of Health Sciences and Technology

Morlet D 1986 Contribution a l'analyse automatique des electrocardiogrammes—algorithmes de localisation, classification et delimitation precise des ondes dans le systeme de Lyon (in French) *PhD These* INSA, Lyon

Moss A and Stern S 1996 *Noninvasive Electrocardiology–Clinical Aspects of Holter Monitoring* (London: Saunders)

Nikolaev N and Gotchev A 1998 Denoising of ECG signals using wavelet shrinkage with time–frequency dependent threshold *Proc. European Signal Processing Conf. EUSIPCO-98 (Island of Rhodes, Greece)* pp 2449–53 **See endnote 6**

Nikolaev N, Gotchev A, Egiazarian K and Nikolov Z 2001 Suppression of electromyogram interference on the electrocardiogram by transform domain denoising *Med. Biol. Eng. Comput.* **39** 649–55

Normann R A 1988 *Principles of Bioinstrumentation* (New York: Wiley)

Paul J, Reedy M and Kumar V A transform domain SVD filter for suppression of muscle noise artefacts in exercise ECG's *IEEE Trans. Biomed. Eng.* **47** 645–62 **See endnote 7**

Proakis J 1989 *Digital Communications* 2nd edn (New York: McGraw-Hill)

Willems J L 1988 Common standard for quantitative electrocardiography multilead atlas—measurements results data set 3 commission of the European communities Medical and Public Health Research, Leuven

Zareba W, Maison-Blanche P and Locati E H 2001 *Noninvasive Electrocardiology in Clinical Practice* (New York: Futura Publishing)

Endnotes

(1) Author: Please check the term ‘The second reason’. Is ‘The secondary reason’ more appropriate?
(2) Author: The meaning of ‘which any component may occur at any time’ is unclear. Please check.
(3) Author: The meaning of the sentence is unclear. Please check.
(4) Author: Please check the term sinus here and elsewhere. Is sine meant here?
(5) Author: Please provide complete information in Haines (2002) if possible.
(6) Author: Please confirm year of references Nikolaev and Gotchev and Augustyniak and Tadeusiewicz here and also in the text.
(7) Author: Please provide year in reference Paul *et al* and also in the text.