MAP FILTERING IN THE DIVERSITY-ENHANCED WAVELET DOMAIN APPLIED TO ECG SIGNALS DENOISING

Marius Oltean, Jean-Marc Boucher**, Alexandru Isar**

*Universitatea "Politehnica" Timisoara, ROMÂNIA (moltean@etc.upt.ro,isar@etc.upt.ro),
** GET, ENST Bretagne, UMR CNRS TAMCIC CS 83818 BREST Cedex 3, FRANCE (jm.boucher@enst-bretagne.fr)

ABSTRACT

An effective denoising method for ECG signals affected by real sources of noise is proposed in this paper. The method is based on a maximum a-posteriori (MAP) filtering in the diversity-enhanced wavelet domain, under realistic a-priori assumptions regarding the statistical properties of the wavelet coefficients of the ECG signal. In order to evaluate the performance of the method, we studied the signal-to-noise ratio (SNR) improvement factor and the degree of the denoising influence on the automatic signal segmentation procedures. The method was tested in both synthetic and real noise conditions and it showed very promising results.

1. INTRODUCTION

Automatic ECG signal interpretation aiming detection and even prevention of cardiac illness is gaining a large popularity in both medical and signal processing communities. Unfortunately, the ECG signal acquisition process is subjected to various disturbing perturbations. The most common are power-line interferences, electromyogram (EMG) noise caused by muscle activity, motion artifacts and baseline drift due to the respiration mechanism. All these unwanted phenomena make the automatic interpretation of the signal a difficult and sometimes even impossible task.

Recently, new techniques based on wavelet transform became popular in relation with the signal denoising. The architecture of a wavelet-based denoising system relies on the wavelet transform ability to concentrate the useful signal energy into a small number of wavelet coefficients. The steps of such a denoising procedure can be readily outlined. First, a wavelet transform is applied to the original noisy signal. The wavelet coefficients are next filtered and the remaining coefficients are finally back-converted into the time domain to form the “clean” denoised signal. In [1] the authors proposed an empirical Wiener-filtering in the wavelet domain for the signal estimation in noise conditions and the method was implemented in [2] for the particular case of ECG signals. The empirical filtering in [1] seems to be particularly suitable for the ECG signal denoising, helping to the rigorous preservation of the useful waveform. Thus, using a wavelet basis function with short temporal support in the pilot estimation stage allows a good preservation of the areas around the QRS complex. On the other hand, the use of wavelets with good frequency localization in the second stage of the algorithm refines the shapes of P and T waves [2].

Note that the Wiener filter could be regarded as a particular case of a MAP filter [3]. There are two key aspects that ensure the success of such a filtering technique: realistic a-priori assumptions regarding the statistical properties of both signal and noise components and a good estimation of the parameters that describe these properties. In this paper, we propose an improved denoising method, based on high accuracy estimation of the statistical parameters of the wavelet coefficients. This improved estimation relies on the diversity enhancement of the signal to be processed. Furthermore, realistic a-priori assumptions regarding the statistical properties of the wavelet coefficients are made, well adapted to the characteristic shape of the ECG signal. The method especially considers the suppression of wide band EMG noise, but good practical results are provided for the power-line interference too.

In the particular field of ECG signals, the evaluation of the denoising quality is not a trivial task. Exact preservation of the diagnostically essential waveforms is critical. Although the classical parameters, such as SNR or mean-square error (MSE) improvement remain important measures, they do not provide a complete image of the denoising quality in this case. Thus, in this particular field, slightly different measures are also used: the correlation between before and after-denoising waveforms, the degree of the signal smoothness after denoising (highlighting the quality of P and T waveforms preservation) [4], the visual analysis of the ST-T artifacts (ripples caused sometimes by wavelet based treatments) [2,4], the RS distortion [2,5], and even the clinical evaluation of the denoised signal [6]. Unfortunately, all these approaches view the denoising as an independent process and consequently do not try to evaluate its performance in correlation with the following steps of the processing chain. Indeed, the ECG automatic processing implies signal denoising, elementary waveforms segmentation, parameters extraction and signal classification. In this context, the performance evaluation for our algorithm will also take into account the denoising effect on the automatic signal segmentation stage.

In section 2, the architecture of the proposed system is presented. Experimental results are reported in section 3, while the final section is dedicated to concluding remarks.

2. PROPOSED DENOISING SYSTEM

The architecture of the proposed denoising system is described in figure 1. To the input we get the useful signal (s) additively perturbed by a Gaussian colored noise (p):

$$x = s + p \tag{1}$$

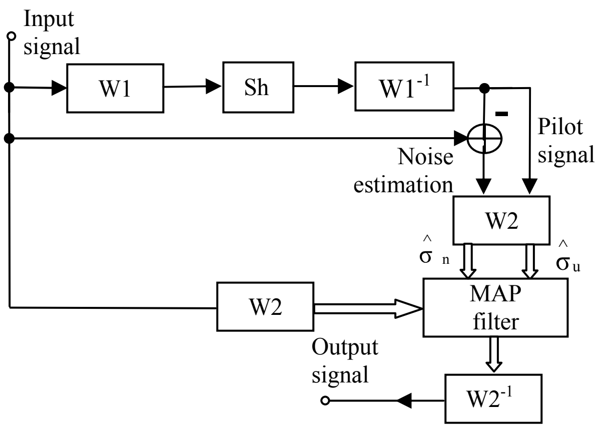

Fig.1 : Architecture of the denoising system.

The system in figure 1 implements a classical MAP filtering in the wavelet domain W2. The estimation of the statistical parameters of the wavelet coefficients is made using a "pilot" signal. The pilot is obtained by hard-thresholding the Haar wavelet coefficients of the noisy signal (the "Sh" operator, figure 1). For W2 transform we have compared two redundant wavelet transforms, providing an enhanced diversity of the signal to be processed. Their performance is presented in section 3. The sources of diversity are the type of wavelet mother used in the computation of the discrete wavelet transform (DWT) [7] and the circular translation of the signal samples respectively [8]. In the first case we consider L_1 different wavelet mothers. In the second one, L_2 circular translations of the signal samples are used, but only one wavelet mother. The two transforms are known as diversity-enhanced DWT (DEDWT) [7] and translation invariant DWT [TIDWT] [8]. In either of cases, we obtain to the output of the wavelet transform L (with $L=L_1$, or $L=L_2$) sequences of discrete wavelet coefficients:

$$ {}^l w = {}^l u + {}^l n, \quad l = 1,...,L \tag{2} $$

${}^l u$ and ${}^l n$ denoting the useful and the noise coefficients respectively, for the l-th set of wavelet coefficients. The MAP estimation of ${}^l u$ is:

$$ {}^l\hat{u}\,({}^l w) = \arg\max_{{}^l u}\left(\log(p_{w/u}({}^l w/{}^l u)\cdot p_u({}^l u))\right) = \\ = \arg\max_{{}^l u}\left(\log(p_n({}^l w - {}^l u)) + \log(p_u({}^l u))\right) \tag{3} $$

In the following, we will consider a Gaussian distribution for the noise coefficients (p_n) and a Laplacian distribution for the useful signal coefficients (p_u). In fact, the wavelet transform of an ECG signal consists into a small number of high value wavelet coefficients (especially marking the limits of the electrical activity zones) and a large number of small value coefficients (for the slow-evolution portions of the ECG). A heavy-tailed distribution for these coefficients seems therefore far more realistic than a Gaussian-one, and the particular case of a Laplacian probability density function (pdf) becomes attractive by its computational tractability. Consequently, we take:

$$ p_u(u) = \frac{1}{\sqrt{2}\sigma_u}\exp\left(-\frac{\sqrt{2}|u|}{\sigma_u}\right) \tag{4} $$

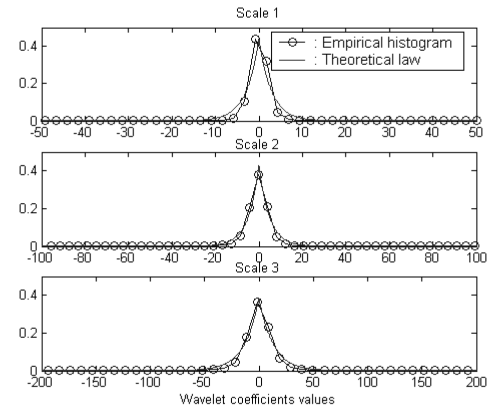

Fig. 2: Empirical histogram fitted on the theoretical Laplacian pdf.

In figure 2, the empirical histogram of the detail wavelet coefficients for three decomposition scales is fitted on the theoretical Laplacian pdf, in order to support our work frame.
Under the considered hypothesis, the equation (3) becomes:

$$ \frac{{}^l w - {}^l\hat{u}}{\sigma_n^2} - \frac{\sqrt{2}}{\sigma_u}\operatorname{sgn}\,{}^l\hat{u} = 0 \tag{5} $$

Finally, the solution of (5) can be expressed as:

$$ {}^l\hat{u} = \operatorname{sgn}({}^l w)\left(\left|{}^l w\right| - \sqrt{2}\,{}^l\sigma_n^2 / {}^l\sigma_u\right)_+ \tag{6} $$

where $(X)_+=X$ for $X>0$ and 0 otherwise. In the equation (6), we denoted by ${}^l\sigma_n^2$ the noise variance and by ${}^l\sigma_u$ the standard deviation of the l-th set of useful signal coefficients. In practice, these parameters are not known and therefore they must be estimated. The relation (6) reduces to a soft-thresholding of the noisy coefficients [3].
Note that the wavelet transform of the ECG signal will contain zones of high-amplitude coefficients (marking the ruptures in the signal) that alternate with areas of small-amplitude coefficients. Therefore, $\hat{\sigma}_u$ must be locally adapted, in order to accurately track these variations. Using the wavelet coefficients of the pilot signal, this parameter is separately estimated for each coefficient, using a sliding window:

$$ {}^l\hat{\sigma}_u(j,k) = \frac{\sqrt{\sum_i |\xi(j,i)|^2}}{\upsilon}, \quad i = k - \frac{\upsilon - 1}{2},...,k + \frac{\upsilon - 1}{2} \tag{7} $$

where $\xi(j,i)$ represents the wavelet coefficient of the pilot signal, j standing for the decomposition scale and i for the position within the scale. ν is the length of the sliding window. To the limit, even the choice of $\nu=1$ provides satisfactory results, since experimental work showed that increasing the window length does not lead to a significant gain. For the noise variance estimation, we use the wavelet coefficients of a "purely noise" signal, obtained as the difference between the pilot estimation and the initial signal. The variance is then calculated at each decomposition level as being simply the variance of the purely noise wavelet coefficients at that level. Thus, the MAP filtering becomes locally adapted. Indeed, for each noisy wavelet coefficient ${}^{l}w(j,k)$, the threshold value for the soft-thresholding operator defined by (6) is different.

In order to obtain the denoised signal, the inverse transform is applied to the L sets of wavelet coefficients. In the case of DEDWT, this consists in applying the L_1 correspondent inverse discrete wavelet transforms (IDWT). For TIDWT, each set of wavelet coefficients is converted into the time domain using the same IDWT and then each version of the signal is correspondingly un-shifted [8]. In either case, we have L different versions of a signal estimated in an "optimal" manner denoted by ${}^{l}\hat{s}(m), \quad l = 1,...,L$. The final result is obtained by averaging the correspondent samples of these versions:

$$\hat{s}(m) = \frac{1}{L}\sum_{l=1}^{L} {}^{l}\hat{s}(m) \qquad (8)$$

Note that the averaging operation will still eliminate part of the residual noise that remained after the MAP filtering of the wavelet coefficients. Without loss of generality, we can assume that the residual noise is a zero-mean stationary Gaussian process, whose L different realizations are the noise samples that perturb the underlying "original" signal. Therefore, any sample-by-sample averaging operation will tend to cancel-out these noise artifacts and to preserve the useful signal. The effectiveness of DEDWT was illustrated in [7]. Furthermore, in the case of TIDWT, averaging over shifts significantly reduces the artifacts near the discontinuities, which are inherently connected to the denoising with the decimated version of the DWT [8].

In some cases, a second iteration of the method could be performed. Thus, the denoised signal obtained after the first iteration is reused as an improved version of the pilot signal. The MAP filtering procedure is next performed once more on the noised signal.

3. RESULTS

We tested our algorithm on a large number of ECG signals from the CHU Brest database. The sampling frequency for these signals is of 1000 Hz, with a 16 bits/sample resolution. The method was tested in both artificially generated and real noise conditions. For the generation of the synthetic noise, a second-order AR-process was used, resulting in a colored Gaussian noise. This simulates the physical EMG noise, which is a wide-band colored signal, whose dominant energy spans in the 50 – 150 Hz range.

In order to obtain the pilot estimation, we shrinked the Haar coefficients of the noisy signal, with the threshold value $T(j) = s(j)\sqrt{2\log M}$, where s(j) represents the standard deviation of the noisy wavelet coefficients at the decomposition level j and M is the length of the data block [9], namely M=4096 samples. For the second stage of the algorithm, we have chosen for DEDWT implementation L_1=10 different wavelet mothers with good frequency localization, from the Daubechies, Coiflet and Symmlet families. In the case of TIDWT, we used the "fully" TIDWT [8], which averages over all circulant shifts of the signal. That is, in this case, we get L_2=4096. The wavelet mother used was Daubechies-8, chosen due to its similarity to the ECG trace.

In order to evaluate the performance of the proposed method, we used five "clean" ECG records of 60 second each. Artificially generated noise was added to the useful signal, resulting in SNR ratios between 10 and 20 dB. The output SNR is calculated for the entire ECG signal as well as for the fragments delimiting the P wave, the most sensitive to noise (this last measure is denoted by PwSNR). For each input SNR the experience was repeated 10 times and the results were averaged. The number of iterations of the method was optimized in order to maximize the output SNR.

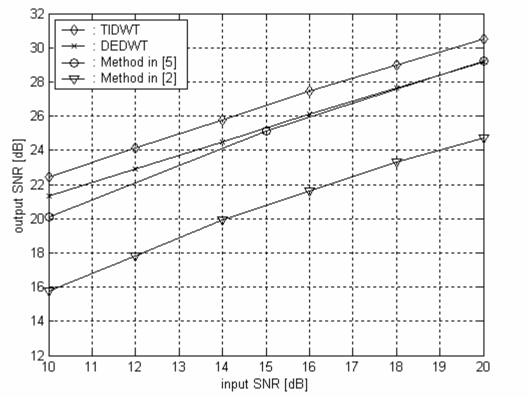

Fig.3: SNR denoising performance in colored Gaussian noise.

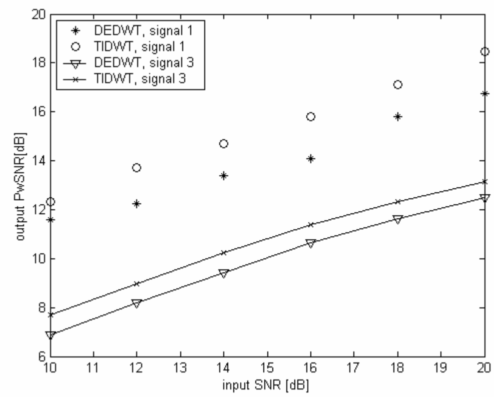

Fig.4: PwSNR denoising performance: two illustrative examples.

Figure 3 illustrates the dependence between the input and output SNRs obtained by applying our method. The results are compared with other results reported in the literature. The use of TIDWT provides the best overall performance: a gain of about 1 dB compared to the DEDWT. Between the other methods, it is the method in [5] that provides comparable results, but this method uses more complex preliminary operations as beat splitting and alignment. Yet, this comparison must be regarded with circumspection, since the work databases are different.

In the case of the P wave denoising, the important variance of the results made undesirable an averaging operation in order to obtain a single PwSNR curve, as in the case of the overall SNR. However, the TIDWT provides superior results in all cases, with a maximal gain between 0.82 dB (test signal number 3) and 1.7 dB (signal 1). The PwSNR performance for these two particular examples is illustrated in figure 4. It must be also noted that for the entire signal set and for both wavelet transforms used, the PwSNR improvement is significant: for an overall input SNR of 22 dB or less, the denoising gain for the P wave is always superior to 10 dB.

Next, the denoising influence on the automatic P wave segmentation process was studied.

The segmentation method used in this purpose was explained and implemented by the authors in [10]. The method captures the dependencies that exist between the wavelet coefficients situated at different decomposition levels in the form of a probabilistic Markov tree with hidden states. In this context, we added synthetically generated noise to five ECG signals for which the automatic segmentation procedure provides very good results. The overall considered input SNR was between 10 dB and 20 dB. In the P wave region, this corresponds to low and very low PwSNRs (< 2 dB in all the cases). The error measure is the difference between the automatically detected limits of the P wave in the case of the denoised signals, with respect to the manually annotated database. For each considered signal, a number of 20 P waves were segmented, and the errors were averaged in order to obtain the mean segmentation error, for both the onset and the end of this wave. For the whole signal subset, the test showed very promising results. The automatic detection error of the P wave was in almost all cases inferior to the tolerable error, which is of 25 ms accordingly to the cardiologist, beginning from an input SNR of 10 dB. There was a single exception (one test signal), for which the tolerable error was obtained for 14 dB. This suggests that our method can be considered for the correct segmentation of relatively low-SNR signals.

In order to provide a deeper analysis of the way that the denoising procedure can influence the segmentation, we applied our method as a pre-treatment step for 25 relatively clean signals from CHU Brest database, that were next segmented using the procedure in [10]. The segmentation results for P wave were compared with the case where another denoising procedure [11] is applied (a SURE filtering [12], followed by a Wiener filtering with the protection of the QRS coefficients) (see table 1).

	Method in [11]	DEDWT	TIDWT
Onset Error	11.16 ms	11.01 ms	10.22 ms
End Error	11.37 ms	8.87 ms	7.99 ms
Segmentation Error Rate	15.96 %	15.2 %	13.46 %

Table 1 : Denoising effects on the automatic segmentation of the P wave.

The results in table 1 show that our method does not degrade the useful P waveform by excessively shrinking the coefficients in high SNR conditions. Moreover, this pre-treatment consistently improves the segmentation results.

For the validation of our method in real conditions, we applied it on a high number of ECG signals perturbed by real noise. The signals are raw data, provided by Task Force Monitor 3040i, from CNS Systems. In figure 5, the denoising result for such an ECG signal is shown, illustrating the effectiveness of our method. In figure 5(a), we have a portion of the input noisy signal. A “fist look” over this signal could lead us to the conclusion that the noise is not white. In figure 5 (b), the denoised signal is shown. The noise is effectively removed, and the signal preserves its useful waveform characteristics.

4. CONCLUSIONS

In this paper an effective, low-complexity method for the denoising of the ECG signals has been presented. This method consists on a MAP filtering in the wavelet domain, under realistic a-priori assumptions for the ECG wavelet coefficients statistics. The filtering is made in the domain of a diversity-enhanced wavelet transform, which provides robustness and superior performance to our method. The performance evaluation is realized by measuring the SNR improvement factor and the degree of influence on the automatic segmentation of the P wave. The results are very promising, showing an excellent SNR improvement and positive influence on the automatic P wave segmentation procedure.

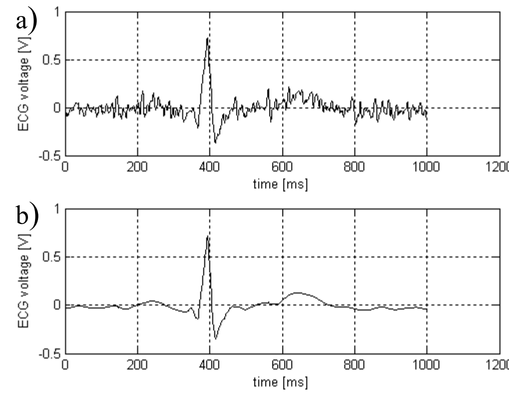

Fig. 5: Denoising method applied on ECG affected by physical sources of noise: the acquired signal (a) and the processed signal (b).

Furthermore, the method allows the correct segmentation of low SNR ECG records. Tested on real ECG signals affected by noise, the method also showed good results, effectively eliminating the noise and preserving the shape of the useful waveforms.

5. REFERENCES

[1] S. Ghael, A. Sayeed and R. Baraniuk, *Improved wavelet denoising via empirical wiener filtering*, Proceedings of SPIE 3169, San Diego, U.S.A., July 1997.

[2] N. Nikolaev and A. Gotchev, *ECG signal denoising using wavelet domain Wiener filtering,* Proceedings of the European Signal Processing Conf. EUSIPCO-2000, Tampere, Finland, pp. 51-54, September 2000.

[3] L. Sendur and I. Selesnick, *Bivariate shrinkage functions for wavelet based denoising exploiting interscale dependency,* IEEE Trans. Signal Proc., vol. 50, pp. 2744-2756, Nov. 2002.

[4] P. Romaniuk, *Non-regular distorsions in ECG signal introduced by wavelet denoising*, Medical Informatics & Technologies - MIT 2000, Ustron, Poland.

[5] A. Gotchev, N. Nikolaev and K. Egiazarian, *Improving the transform domain ECG denoising performance by applying inter-beat and intra-beat decorrelating transforms*, Proceedings of the IEEE International Symposium on Circuits and Systems, ISCAS 2001 Volume II, pp. 17-20, Sydney, Australia, May 2001.

[6] P. Romaniuk and K. Chizynski, *Clinical evaluation of ECG signal wavelet denoising method*, Medical Informatics & Technologies - MIT 2001, Ustron, Poland.

[7] A. and D. Isar, *Adaptive denoising of low SNR signals*, ARI, The Bulletin of the Istanbul Technical University, Volume 53, Number 2, pp. 31-37, September 2003.

[8] R. Coifman et D. Donoho, *Transaltion-invariant de-noising, in Wavelets and Statistics*, A. Antoniadis and G. Oppenheim Eds, Springer-Verlag, pp. 125-150, New York, 1995.

[9] I. Jonstone et B. Silverman, *Wavelet threshold estimators for data with correlated noise*, Journal of Royal Statistical Society, Volume B 59, number 2, pp. 319-351, 1997.

[10] S. Graja and J.M. Boucher, *Multiscale Hidden Markov Model Applied to ECG Segmentation,* Intelligent Signal Processing, 2003 IEEE International Symposium , pp 105-109, Sept. 2003.

[11] R. Le Page, *Detection and analysis of the P wave of an ECG signal: application to the atrial fibrillation detection,* Ph. D.Thesis, Université de Bretagne Occidentale, Brest, France, February 2003.

[12] D. Donoho and I. Johnstone, *Adapting to unknown smoothness via wavelet shrinkage*, Journal of American Statistical Assoc., vol. 90, pp. 1200-1224, December 1995.