

A Quality-Scalable and Energy-Efficient Approach for Spectral Analysis of Heart Rate Variability

Georgios Karakonstantis, Aviinaash Sankaranarayanan, Mohamed M. Sabry, David Atienza, Andreas Burg

Ecole Polytechnique Federale de Lausanne (EPFL), Switzerland

{georgios.karakonstantis, aviinaash.sankaranarayanan, mohamed.sabry, david.atienza, andreas.burg}@epfl.ch

Abstract—Today there is a growing interest in the integration of health monitoring applications in portable devices necessitating the development of methods that improve the energy efficiency of such systems. In this paper, we present a systematic approach that enables energy-quality trade-offs in spectral analysis systems for bio-signals, which are useful in monitoring various health conditions as those associated with the heart-rate. To enable such trade-offs, the processed signals are expressed initially in a basis in which *significant* components that carry most of the relevant information can be easily distinguished from the parts that influence the output to a lesser extent. Such a classification allows the pruning of operations associated with the *less significant* signal components leading to power savings with minor quality loss since only less useful parts are pruned under the given requirements. To exploit the attributes of the modified spectral analysis system, thresholding rules are determined and adopted at design- and run-time, allowing the static or dynamic pruning of less-useful operations based on the accuracy and energy requirements. The proposed algorithm is implemented on a typical sensor node simulator and results show up-to 82% energy savings when static pruning is combined with voltage and frequency scaling, compared to the conventional algorithm in which such trade-offs were not available. In addition, experiments with numerous cardiac samples of various patients show that such energy savings come with a 4.9% average accuracy loss, which does not affect the system detection capability of sinus-arrhythmia which was used as a test case.

I. INTRODUCTION

In the last years, the prevalence of unhealthy lifestyles have worsened personal health and increased the number of people living with chronic cardiovascular and brain disorders [1]. In an attempt to handle the need and costs for continuous, portable monitoring of such disorders, Wireless Body Sensor Networks (WBSNs) have emerged in the last decade [2, 3]. A typical WBSN comprises a set of wearable low-power nodes that record, monitor and wirelessly sent vital signals to a central network coordinator which could be another portable device, capable of analyzing the recorded signals. The latest technological progress may allow to implement portable devices for realizing such WBSNs and even to perform complex algorithms in real-time, however the challenge of energy efficient operation still remains as one of the most critical in this domain. To this end, new low-power sensor nodes [4, 5] as well as more efficient compression algorithms [3] have been proposed recently that try to reduce the amount of signals needed to be communicated and thus limit the use of power-hungry radio. In a further attempt to pre-process the signals and in a way limit the communicated data, efficient

electrocardiogram (ECG) delineation as well as heart-rate analysis algorithms were recently proposed [6, 7]. Interestingly, such algorithms allow early detection of not only heart related disorders but also brain related such as epileptic seizures [7].

Another, very useful algorithm that could enhance the WBSNs capabilities and help in the early detection of various heart and brain disorders [8, 9] is the power spectral analysis (PSA) of heart rate variability (HRV). Although in recent years it has been recognized as a powerful tool for evaluating various autonomous nervous system activities, unfortunately its high complexity is hindering its use in portable devices, especially in the WBSN nodes. Almost all existing works on PSA of HRV [10] focused mainly on the algorithmic side, trying to overcome the limitations of traditional PSA approaches, which were not suitable for unevenly sampled data, as in the case of the intervals between heart beats that are the input to such an algorithm. Even very recent works that dealt with mapping PSA systems on hardware [11] have not tried to improve their energy efficiency. It is characteristic that existing PSA systems of HRV try to provide always the maximum possible accuracy even in cases that it is not required, thus not exploiting any room for energy savings. Recent works for instance have accomplished to reduce the complexity of compression algorithms by exploiting the characteristics of ECG signals [3] and by utilizing the fact that a certain disorder can be detected with less information. Such properties remain unexploited in PSA systems and thus new approaches are required for enabling trade-offs between accuracy and energy in such systems required for allowing their integration in WBSNs.

Contributions and Outline: In this paper, we present a systematic approach for improving the energy efficiency of PSA systems with minor accuracy loss by exploiting properties of the processed signals. The proposed approach is applied on the Welch-Lomb algorithm that allows simultaneous time-frequency analysis of HRV. The contributions are summarized as follows:

- 1) Analyze the energy and performance profiles of the Welch-Lomb method on a node simulator and identify its limitations.
- 2) Exploit the approximately-sparse representation of heart beats in the Wavelet domain and systematically develop a scalable PSA system based on Discrete Wavelet Transform (DWT). An alternative representation of the power hungry block of the Welch-Lomb method based on DWT is presented.
- 3) Optimize the complexity of the modified algorithm determining the parameters that allow significant complexity reduction with minor accuracy loss.
- 4) Classify signals and associated operations based on the contribution to quality into *significant* and *less significant*. This is enabled due to the concentration of signal energy into few elements achieved due to the modified representation.
- 5) Map the developed system on a node simulator and analyze the energy-quality trade-offs achieved by pruning less-significant elements and associated operations after applying

thresholds at i) design-time and ii) run-time.

- 6) Utilize the throughput gains obtained by the pruned operations for applying voltage and frequency scaling (VFS) to obtain further energy savings.

The rest of the paper is organized as follows. Section II discusses the target system and analyzes the conventional PSA methods. Section III discusses the proposed design flow while Section IV applies the initial steps of the approach to the PSA system. Section V describes the steps for reducing the complexity of the modified PSA. Section VI presents the achievable energy, quality trade-offs. Finally, conclusions are drawn in Section VII.

II. STATE OF THE ART AND LIMITATIONS

The input to a PSA system of HRV is a fixed size window of time intervals between successive heart beats (RR intervals) that are extracted from a continuous ECG of a person using a delineation algorithm (Fig. 1(a)). PSA systems could be implemented either on a portable central coordinator of WBSN or on each sensor node since latest delineation algorithms (the output of which is needed) are already included on modern WBSN nodes and thus could provide the required RR-intervals to PSA [6]. Before describing our approach for enabling the portable integration of PSA we briefly analyze the conventional methods.

A. State of the art

Various methods can be used for the estimation of the periodogram which is the output of a PSA system but the non-periodic nature of the processed RR-intervals indicated the Lomb method as one of the most suitable one [8]. By using least squares fitting to estimate the amplitude of a given sinusoid with angular frequency ω_j over non-uniformly sampled data, the Lomb method avoids the use of interpolation and re-sampling needed in traditional PSA approaches that may alter the frequency content [8]. The Lomb periodogram of a non-periodic signal x_j is:

$$P_N(\omega) = \frac{1}{2\sigma^2} \left\{ \frac{[\sum(x_j - \mu)\cos\omega(t_j - \tau)]^2}{\sum \cos^2\omega(t_j - \tau)^2} + \frac{[\sum(x_j - \mu)\sin\omega(t_j - \tau)]^2}{\sum \sin^2\omega(t_j - \tau)^2} \right\} \quad (1)$$

where μ , σ are the mean and variance of the signal, respectively, τ is a constant offset for each angular frequency ω that makes the periodogram invariant to time-shifts, t_j is the time interval of the RR sample and the sums are taken over the corresponding window size [10]. The complexity of the Lomb method, soon led to a Fast-Lomb algorithm which uses two complex Fast-Fourier-Transform (FFT) for reducing the sums over trigonometric functions in (1) to four less-complex sums.

The overall PSA system is depicted on Fig. 1(a). Based on the varied number of RR peaks in each window the data are extrapolated (i.e., redistributed to the needed order [10]) to size N in order to meet the fixed size N (e.g. 512) of the FFT. The FFTs then calculate the four sums in (1) for each window and the Lomb calculator combines the data in order to provide the real time power-spectrum information. Interestingly, by monitoring several time intervals during a day through the application of a window $w(t)$ to the data, the time-frequency distribution can be obtained using the so called Welch-Lomb method [11]. Such a method by using a de-normalizing factor $2\sigma/N$ allows to average the variance of normalized segments obtained from the applied sliding window and to better track the time-varying components of the heart rate. Note that the Fast-Lomb method can still be used for the calculation of the periodogram within each window.

B. Profiling and Limitations of a Conventional PSA System

As an initial step for better understanding the application

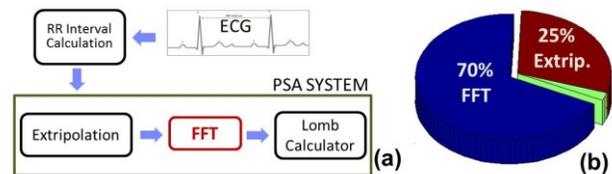


Fig. 1. (a) Overall PSA system for heart rate variability, (b) Energy profiling of a split-radix-FFT based PSA system

characteristics, we implemented a Lomb-Welch PSA system as shown in Fig. 1(a). For the implementation of the 512 sized FFT, the split-radix method was utilized, which is one of the fastest known FFT realizations. In our sliding window implementation we applied a window of 2 minutes of sampled RR-intervals with 50% overlap between the windows. The system was mapped on a simulator configured with typical sensor node parameters [13, 14] and the profiling results are depicted in Fig. 1(b). As it can be seen, the FFT block consumes most of the overall system power, which also accounts for the majority of the total computational cycles. Therefore, one can assume that by reducing the complexity of the FFT, potentially by approximating some operations, the overall system power could be reduced.

Usually, pruning methods are applied for the reduction of operations in FFT. Such methods take advantage of the potential sparsity of input or output signals in the Fourier domain and through intelligent dataflow pruning they drop redundant/useless operations. However, such methods often require extensive signal sparsity (in many cases more than 90% [15]) and they do not work well when the non-zero inputs/outputs are randomly located as in the case of the extracted RR-intervals (for each sampled window). However, most importantly, the fact that renders the straightforward pruning of the FFT in PSA of HRV systems difficult is that the RR intervals are not necessarily sparse, neither in their original representation nor in the Fourier domain. Moreover, in the conventional FFT, all the twiddle factors in the butterfly operations are unit magnitude complex numbers [16] and thus all parts of the structure are of equal importance making any approximation risky since it may lead to large accuracy loss.

III. PROPOSED APPROACH

The analysis performed in the previous section indicates that in order to improve the energy efficiency of PSA systems, alternative representations, which could enable energy-accuracy trade-offs need to be investigated. This is the essential idea behind our approach depicted in Fig. 2. Initially, we find a domain (i.e. Wavelet) in which the input signal \mathbf{x} is approximately sparse and we continue by expressing it in this domain using a

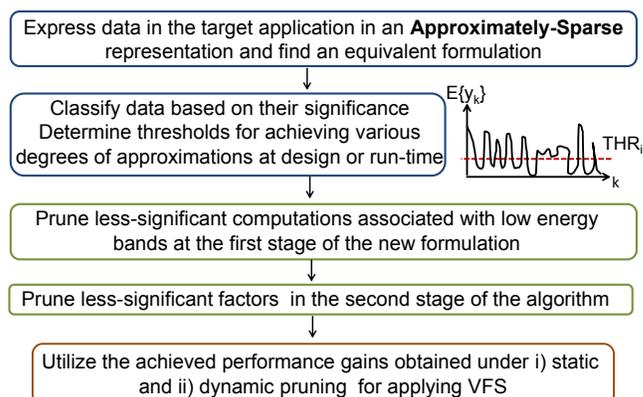


Fig. 2: Proposed design flow

suitable transformation W (i.e. DWT). Note that by approximately-sparse we mean that few signal components carry most of its energy differing significantly in magnitude by other small/close to zero components. The few large signal components are also most significant/relevant for the overall signal quality, while other signal components and the associated computations can potentially be skipped/pruned for large complexity reduction.

The applied transformation alters the signal representation and thus it requires reformulation of the initial FFT representation (F) into an equivalent one (G) for yielding the same original result y . The procedure of the first stage of our approach can be written as:

$$\mathbf{y} = \mathbf{F}_N \mathbf{x} \Leftrightarrow \mathbf{y} = \mathbf{G}_N (\mathbf{W}_N \mathbf{x}) \quad (2)$$

where $\mathbf{F}_N, \mathbf{W}_N$ and \mathbf{G}_N are the FFT, Wavelet, and equivalent transformation matrices of size N , respectively.

In the second stage of our approach the complexity of the modified algorithm is reduced by exploiting the approximately-sparse signal representation. Note that by approximately-sparse we mean that in the new signal only few components carry most of the energy. Therefore these few components are also most important for the overall signal quality, while other signal components and the associated computations play a less-significant role and can be skipped. So in this stage we identify the less significant components and prune the associated operations in the two transformations W and G that affect the output quality/accuracy only to a small extent. To achieve this, the statistics at the output of the first transform are obtained and threshold values THR_i are determined for distinguishing the elements of the output vector into significant and less-significant, i.e., according to:

$$z_k: \begin{cases} \text{significant,} & \text{if } E\{|z_k|\} \geq THR_i \\ \text{less-significant,} & \text{if } E\{|z_k|\} < THR_i \end{cases}, k = 0, \dots, N \quad (3)$$

where z_k is the k^{th} element of the output vector z of the first (W) transformation. By applying such thresholds the operations associated with the less-significant elements of the output vector are pruned. In order to ensure that any pruning under the specific threshold, results to maximum energy savings and minor quality loss we determine suitable thresholds by performing several experiments with numerous cardiac samples. Note that pruning in the second stage of the algorithm is applied based on the twiddle factors of the butterflies, which are distinguished into significant and less-significant based on their magnitude. Thresholds are also determined in this case for applying different degrees of pruning at the butterfly stages required for obtaining energy-quality trade-offs. The overall system is then mapped to a target platform and the improvement in number of computational cycles and energy based on i) *static pruning*, at design-time, through application of thresholds to the expectation (across numerous cardiac samples) of intermediate results and ii) *dynamic pruning*, at run-time, through thresholding of factors in the second stage are being obtained. Note that dynamic thresholding is expected to provide fine-grained approximations in a sample by sample case and thus to lead to less quality loss, but with an overhead due to the extra required comparisons. Finally, the reduced cycles are exploited for applying VFS for larger energy savings.

IV. EXPOSING SIGNAL SPARSITY

In this section, we start by applying the initial steps of our approach for finding an alternative sparse signal representation of the core FFT block within the PSA system.

A. Seeking an Approximately-Sparse Basis for RR-intervals

Interestingly, ECG signals are known to be approximately

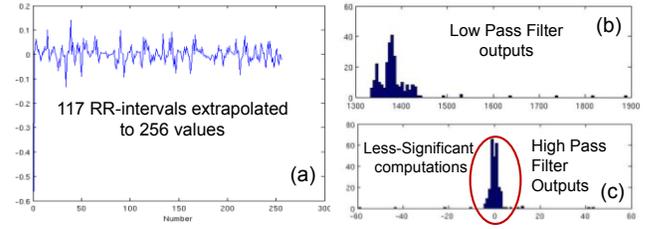


Fig. 3: Example of (a) extrapolated RR-intervals, (b) lowpass and (c) highpass wavelet filter outputs for a Haar based DWT.

sparse in the wavelet domain which was exploited for compressing them [12, 18], motivating us to investigate if such an attribute applies also to the RR-intervals (inputs to PSA).

Representation in the wavelet domain occurs through sub-band decomposition. In particular, the original signal passes through a pair of filters and is then down sampled by a factor of 2. Given the low pass-filter (LF) and high-pass filter (HF) that satisfy the Wavelet constraints, Wavelet decomposition can be compactly expressed as a linear transformation matrix \mathbf{W}_N constructed from LF and HF with N denoting the size of the matrix. The decomposition can then be expressed as:

$$\mathbf{W}_N \mathbf{x} = \begin{bmatrix} \mathbf{a}_{N/2} \\ \mathbf{c}_{N/2} \end{bmatrix} \quad (4)$$

where $\mathbf{a}_{N/2}$ and $\mathbf{c}_{N/2}$ is the decomposed low-pass and high-pass signal, respectively [16]. The DWT consists of one or more stages depending on the degree of the desired resolution and each of them contains a HF and a LF that compute the so-called approximation and detail coefficients, respectively. Note that the filter order depends on the basis of the mother Wavelet used, i.e., Haar, Db2, Db4 etc. Interestingly, after processing extrapolated RR-intervals of numerous heart samples from the MIT-BIH arrhythmia database [17] (Fig. 3(a)) with DWT of various known bases (i.e. Haar, Db2 and Db4) we observed that the two output bands of high and low frequencies differ significantly in terms of their magnitude. Specifically, we observed that the HPF outputs were distributed around zero (Fig. 3(c)) indicating that such parts could be pruned for complexity reduction with potentially minimum impact on quality. In other words, RR-intervals are approximately-sparse in the Wavelet domain, which can be utilized for finding an approximately sparse representation.

B. Finding an Equivalent Formulation in the Sparse Basis

To begin with, an N^{th} order FFT can be written as:

$$\mathbf{F}_N = \mathbf{F}_N \mathbf{S}'_N \mathbf{S}_N = \begin{bmatrix} \mathbf{I}_{N/2} & \mathbf{T}_{N/2} \\ \mathbf{I}_{N/2} & -\mathbf{T}_{N/2} \end{bmatrix} \begin{bmatrix} \mathbf{F}_{N/2} & \mathbf{0} \\ \mathbf{0} & \mathbf{F}_{N/2} \end{bmatrix} \mathbf{S}_N \quad (5)$$

where $\mathbf{T}_{N/2}$ is the diagonal matrix with twiddle factors on the diagonal and \mathbf{S}_N is an $N \times N$ even-odd separation matrix. The first part of the new transform, the DWT \mathbf{W}_N , obeys $\mathbf{W}'_N \mathbf{W}_N = \mathbf{I}_N$ since it is an orthogonal linear transformation. Based on this property the Fourier transform can be written as $\mathbf{F}_N = \mathbf{F}_N \mathbf{W}'_N \mathbf{W}_N$. Considering also (5), the following factorization can be written:

$$\mathbf{F}_N = \mathbf{F}_N \mathbf{W}'_N \mathbf{W}_N = \begin{bmatrix} \mathbf{A}_{N/2} & \mathbf{B}_{N/2} \\ \mathbf{C}_{N/2} & \mathbf{D}_{N/2} \end{bmatrix} \begin{bmatrix} \mathbf{F}_{N/2} & \mathbf{0} \\ \mathbf{0} & \mathbf{F}_{N/2} \end{bmatrix} \mathbf{W}_N \quad (6)$$

where $\mathbf{A}_{N/2}, \mathbf{B}_{N/2}, \mathbf{C}_{N/2}, \mathbf{D}_{N/2}$ are all diagonal matrices, which was also mathematically proven in [16] but its use was limited up to now. The values on the diagonal of $\mathbf{A}_{N/2}$ and $\mathbf{C}_{N/2}$ correspond to the length- N FFT of the lowpass filter of the wavelet transform, whereas the values on the diagonal of $\mathbf{B}_{N/2}$ and $\mathbf{D}_{N/2}$ are the length- N FFT of the highpass filter of the wavelet transform. The factorization shown in (6) suggests a DWT based FFT algorithm, whose block diagram for an order $N = 8$ is depicted in Fig. 4. The

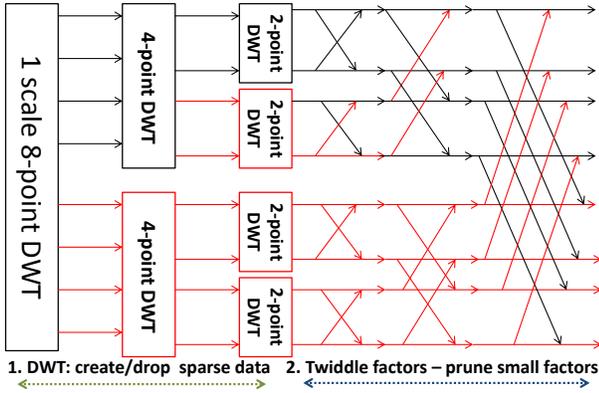


Fig. 4: DWT-based FFT consists of two stages i) DWT, ii) Twiddle-factors. Less-significant data approximated/pruned in red.

algorithm consists of 2 main stages; the highpass and the lowpass DWT outputs go through separate length-4 DFTs, and then they are combined with butterfly operations. The same scheme can be applied iteratively to shorter length DFTs to get the full DWT based algorithm. The full system is equivalent to a binary tree wavelet packet followed by modified FFT butterfly operations, where the twiddle factors are the frequency response of the wavelet filters. The overall process represents a new transformation with the distinguished processes W (DWT) and G that we looked for in (2), as part of the first step of our approach.

At this point the question that arises is how the complexity of the new algorithm compares to the original FFT. To this end, we have evaluated the complexity of the algorithm with $N=512$ using various wavelet bases (i.e. Haar, Db2, Db4) and compared it with a split-radix FFT, one of the fastest known FFT implementations. Results show that (without pruning or exploitation of the sparsity of the signal) the wavelet-based FFT comes with 36%, 49%, and 76% increased number of computations compared to the split-radix FFT in case of Haar, Db2, and Db4 DWT bases, respectively. Therefore, there is a need to reduce the complexity, which is the goal of second part of our approach.

V. COMPLEXITY REDUCTION OF NEW FORMULATION

The second part of our approach is distinguished into two parts as we discussed in Section III; discovering less-significant signal components and pruning the associated operations within the first stage of the algorithm (DWT) and then in the second stage.

A. Pruning Operations in DWT

As we discussed in Section IV, the first stage of the new formulation is the DWT, which after processing the RR-intervals distinguishes the signal into two groups/bands; the high energy (LF outputs) and the low energy (HF outputs) bands. Based on such differences the highpass-detail computations associated with the less-significant signal components can be pruned, eliminating the corresponding half band of the DWT as highlighted in red in Fig. 4. By doing so the multiplications and additions associated with the dropped signals in the second stage of the wavelet-based FFT can also be pruned. After pruning the operations, the new sparse wavelet-based FFT process ($F_{N,SP}$) can be written as:

$$F_{N,SP} = F_N W'_N W_N = \begin{bmatrix} A_{N/2} & \mathbf{0} \\ C_{N/2} & \mathbf{0} \end{bmatrix} \begin{bmatrix} F_{N/2} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} W_N \quad (7)$$

where $B_{N/2}$ and $D_{N/2}$ as well as the lower right $F_{N/2}$ in (6) are set to zero due to their less-significant content.

We now repeat the complexity comparison and we observe that the number of computations is reduced by 28%, 21%, and 8%

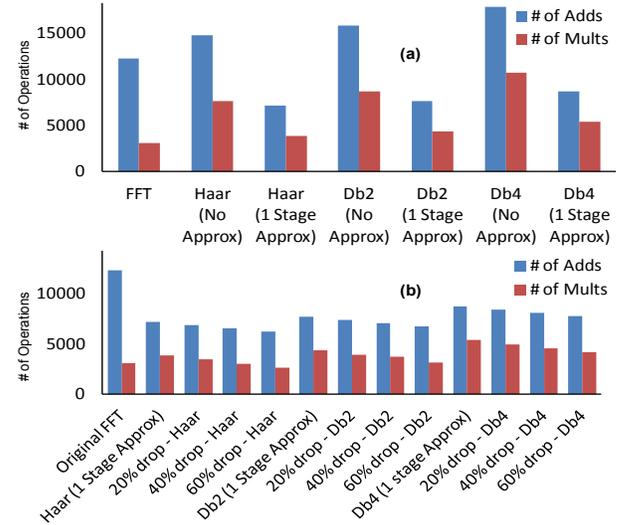


Fig. 5: Complexity comparison of proposed approach with various Wavelet basis and approximations applied (a) in 1st stage and (b) in 2nd stage. (Reference 512-FFT split radix)

compared to the split radix FFT if the Haar, the Db2, or the Db4 are used as DWT basis, respectively (Fig. 5(a)). In the above approximation the highpass band is eliminated at design-time by applying an appropriate threshold, however based on the specific samples we could also apply such a threshold at run-time for fine-grained pruning.

B. Significance-Driven Twiddle-Factor Pruning

In the second stage of the algorithm, the DWT outputs are multiplied with twiddle factors that are the frequency response of the filter coefficients of the chosen wavelet basis (Haar, Db2, Db4 etc.). Such factors carry the unique property that they do not lie on the unit circle but they differ in their magnitude substantially as opposed to the FFT twiddle factors. Specifically, we observe that the twiddle factors, elements of the diagonal sub-matrix $A_{N/2}$ in (7) [$A_{11}, A_{22}, \dots, A_{N/2N/2}$] decrease in magnitude ($A_{11} > A_{22} > \dots > A_{N/2N/2}$), whereas the factors within the diagonal sub-matrix $C_{N/2}$ in (7) increase in magnitude. For instance, in case of $N=8$ we have in C_4 : $C_{51} < C_{62} < \dots < C_{84}$. In particular, usually elements $A_{4N/2N/2}$ and $C_{(N/2+1)1}, C_{(N/2+2)2}$ etc. have a small magnitude (close to zero). Fig. 7 shows the values of the elements of the matrices $A_{N/2}$ and $C_{N/2}$ in case of $N=512$ where we can observe that many factors have a very small value. This indicates that the operations associated with these small factors might also influence the output result to a smaller extent. To determine the significance of each factor and of the associated operations we performed a sensitivity analysis by pruning various small twiddle

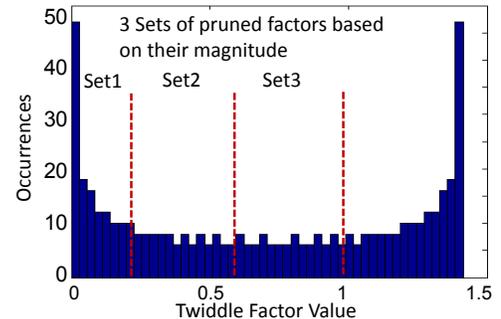


Fig. 6: Values of twiddle factors in $A_{N/2}$ and $C_{N/2}$ matrices, $N=512$

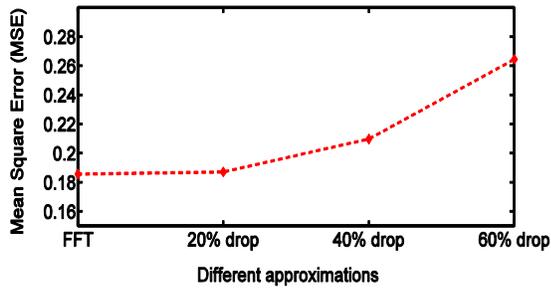


Fig. 7: Mean Square Error (MSE) for various approx. in 2nd stage

factors and observed the impact on output quality, quantified in terms of the mean-square-error (MSE) between the original output signal and the approximated one. Three sets of factors were finally distinguished based on their magnitudes and the number of pruned computations that can achieve (i.e., Set1 corresponds to 20% pruned factors/operations, Set2 to 40%, and Set3 to 60%). Our results, after applying the above degrees of approximations to the modified PSA system and using various cardiac samples [17] show that the output quality in terms of MSE deteriorates slightly compared to the outputs of the conventional algorithm. Such distortions as we will show in next section do not affect the monitoring capability of the system.

Note also that as the wavelet basis and thus DWT filter sizes increases (e.g., in case of Db4), the number of small-valued/zero twiddle-factors in the second stage of the algorithm also increases. However, at the same time the number of computations in the first DWT stage is also increasing. Therefore, there is a clear trade-off between the approximations applied in the 2nd stage and the number of operations in the 1st stage (DWT). Fig. 5(b) compares the proposed approach to the conventional FFT approach in case of using various wavelet bases and for the different degrees of approximations (Mode1-20%, Mode2-40%, and Mode3-60%) of less significant computations in the 2nd stage. It is evident that the choice of Haar wavelet basis has the lowest complexity when compared to all other choices. Therefore, Haar was chosen as the wavelet basis for the implementation of the PSA system since it can lead to low-complexity. Overall, we observed that the proposed approximations can reduce by 52% the number of additions and 17% the multiplications compared to a conventional split-radix FFT algorithm. Note that the savings increase with the order (i.e. in case of $N=1024$ then we obtain we obtain further 12% fewer multiplications and 8% fewer additions) due to the logarithmic complexity growth of the original FFT with the order. Of course, the proposed approximations lead to quality degradation which we investigate next along with the achieved energy savings.

VI. ENERGY QUALITY TRADE-OFFS EVALUATION

In order, to evaluate the energy savings and the distortions obtained by the proposed approximations we have mapped the conventional and modified PSA system (with $N=512$ FFT) on a single RISC processor simulator configured with typical, available sensor node characteristics [13, 14]. In particular, we are simulating a typical single-core sensor node that is accompanied with a 64KB SRAM. The available power consumption values of the processor in a low leakage 90nm technology node [14] are also utilized and the conventional and proposed systems are cross compiled for the target architecture. To determine the impact of the proposed modified PSA system

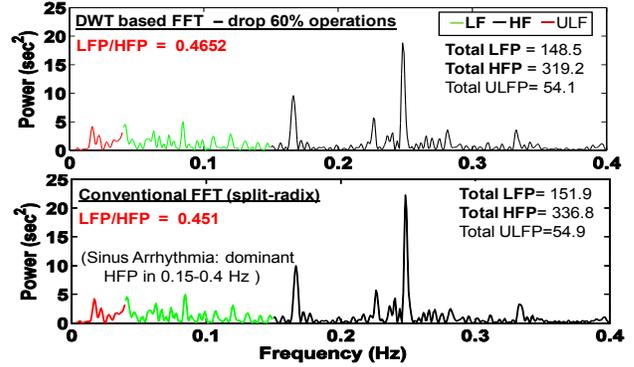


Fig. 8: PSA of conventional (split-radix FFT based) vs. proposed system (with applied approximations in the modified FFT block)

and the applied approximations we have evaluated the ratio between the overall low frequencies power (LFP) and the overall high frequencies power (HFP) of the heart rate spectra. Specifically, we extracted the RR-intervals from numerous sinus-arrhythmia and healthy samples from PhysioNet [17] and provided them as input to the proposed and conventional (Welch-Lomb based on split radix FFT) PSA system as described in Section II.B. The overall HFP within 0.15 – 0.4 Hz and LFP within 0.04-015 Hz is then calculated by integrating the total power within these bands, at the output of the systems by Lomb calculator (Fig. 1(a)). Typically, a ratio of LFP over HFP much less than 1 indicates a sinus arrhythmia condition and is an appropriate quality metric for such an application. Note that sinus-arrhythmia was used as a test case for quantifying the impact of our accuracy loss. Such a test case was also used by other recent papers on PSA [11]. In any case, our method can be utilized for the detection of any other condition requiring PSA.

A. Quality under various approximations

Table I depicts the average ratio LFP/HFP for various samples and for different modes of approximations applied statically at design time. We observe that the ratio remains close to the original value and much less than 1 even when the highpass-band in the first DWT stage and 60% of the twiddle factors in the second stage are pruned. The periodogram obtained from the proposed PSA system (with the highpass band and 60% of the twiddle factors in the modified FFT block being pruned) and the conventional system (based on split radix FFT), for a patient with sinus arrhythmia is depicted in Fig. 8. Interestingly, we can observe that even if we drop 60% of the operations we obtain only 3% difference in LFP/HFP ratio compared to conventional PSA system, thus we can still easily identify the sinus arrhythmia condition. Furthermore, by using a sliding window configuration as described in Section II, we obtained time-frequency distributions of hourly monitoring of various sinus arrhythmia patients. By obtaining the LFP over HFP ratios for the various

TABLE I
Average LFP: HFP ratio under static and dynamic pruning

LFP/HFP Ratio	Orig. FFT based PSA	PSA based on prop. FFT with 1 st stage approx. and various approx. in 2 nd stage			
		1 st stage band drop	Set1	Set2	Set3
Static Pruning	0.45	0.465	0.465	0.483	0.492
Dynamic Pruning	0.45	0.465	0.467	0.47	0.471

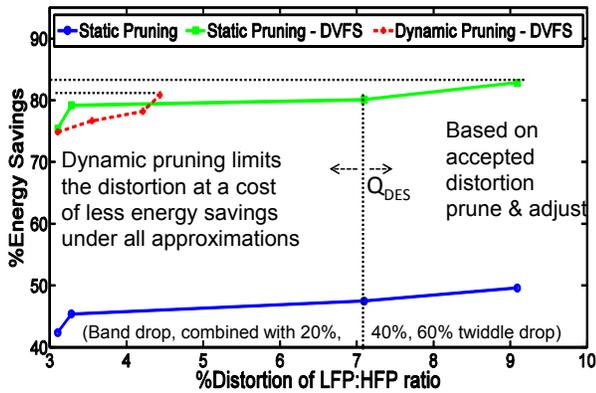


Fig. 9: Energy-Quality trade-offs of the proposed PSA system for various degree of static and dynamic pruning (savings compared to the conventional PSA system)

time intervals of such time-frequency distributions, using heart rate samples of 16 patients we find that on average our approach results in approximately 4.9% of error in such ratio and in all cases we could correctly identify the sinus-arrhythmia condition.

B. Energy Savings under static pruning and VFS

Fig. 9 depicts the energy savings and the distortion in the HFL/LFP ratio obtained after applying various degrees of pruning as mentioned in previous sections. It can be observed that after dropping the highpass band of DWT and pruning 60% of the twiddle factors in the second stage we obtain 51% energy savings with 9.2% error in the LFP/HFP ratio for the considered samples which does not affect the arrhythmia detection capability. In any case the degree of pruning could be tuned for obtaining maximum energy savings based on the acceptable distortion (Q_{DES}).

Another advantage of our approach is that any static pruning of operations results in reduced number of cycles and less execution time. This means that we can relax the frequency of operation allowing us to also reduce the supply voltage V_{dd} , which can lead to quadratic energy savings. Specifically, the execution time is given by $Exec_time = \#cycles * freq(V_{dd})$, where $freq(V_{dd})$ is the frequency of operation at a given supply-voltage. In our case for determining the degree of VFS that can be applied while maintaining the same processing time, we noted the performance improvement that we obtain with the new system. We found that the improvement ranges from 40% up to 51% depending on the degree of applied approximations. Based on the reduction in execution time we calculated the voltage scaling and the resulted energy savings that can be obtained under various modes of approximations. We observe that the proposed PSA when combined with VFS can lead up to 82% energy savings. Note that the additional energy savings come with the same distortions obtained under static pruning (without VFS).

C. Dynamic Pruning

In the above analysis we have applied a fixed number of approximations by pruning the highpass band in the 1st stage of the modified FFT and the three different sets of factors in the 2nd stage for each sample based on expected values of individual intermediate variables that we analyzed during design time. However, as we said in Section IV, it is possible to apply dynamic thresholding at run-time for achieving fine-grained approximations on sample by sample case. To do so we have altered the application software by including some extra comparison instructions during the 2nd stage. Data and twiddle

factors that are below a set of thresholds are eliminated on the fly for the various samples. By doing so we can achieve fine-grained approximations, limiting the distortions (for the same amount of approximations) as shown in Fig. 9. This is due to the elimination of only the computations that are below a threshold for the specific sample. However, the reduction in distortion comes at an approximately 10% energy overhead compared to the static case due to the extra instructions i.e., comparisons limiting also the degree of the applied VFS. All in all, the proposed approach leads to the design of an energy efficient PSA system that could adapt its energy and performance based on the allowed distortion (Q_{DES}) as opposed to recent (non-scalable) implementations [11].

VII. CONCLUSION

An approach for the design of an energy efficient PSA system by utilizing the sparsity of the processed signals in the Wavelet domain was presented. An equivalent transformation for the core kernel of the PSA system was found that exposes signal sparsity and helps in classifying the signal components into significant and less significant based on their contribution to output quality, which was not possible before. Based on the alternative formulation, less-significant operations were identified and pruned in the various stages. Results show that the proposed system can lead up-to 82% energy savings with minor distortions that do not affect the detection capability of sinus arrhythmia in all samples that we studied. The proposed system allows the widespread use of PSA systems in WBSNs enabling the enhancement of their monitoring and analysis capabilities.

REFERENCES

- [1] World Health Organization, "Cardiovascular diseases," 2009. [Online].
- [2] M-k Suh, et al. "A Remote Patient Monitoring System for Congestive Heart Failure," *JMS*, 2011.
- [3] H. Mamaghanian, et al. "Compressed Sensing for Real-Time Energy-Efficient ECG Compression on Wireless Body Sensor Nodes," *IEEE TBE*, 2011.
- [4] Shimmer Research. (2008). [Online].
- [5] R. F. Yazicioglu, et al., "Ultra-low-power wearable biopotential sensor nodes," *IEEE EMBS*, 2009
- [6] F. Rincon, et al., "Development and Evaluation of Multilead Wavelet-Based ECG Delineation Algorithms for Embedded Wireless Sensor Nodes," *IEEE TITB*, 2011.
- [7] F. Massien, et al. "Miniaturized wireless ECG monitor for real-time detection of epileptic seizures," *ACM TECS*, 2013.
- [8] AI Maistrou, "Implicit Comparison of Accuracy of Heart Rate Variability Spectral Measures Estimated via Heart Rate and Heart Period Signals," *IEEE CinC*, 2008.
- [9] J. F. Thayer, et al. "A meta-analysis of heart rate variability and neuroimaging studies: Implications for heart rate variability as a marker of stress and health," *NBR*, 2012.
- [10] W. H. Press, et al, "Fast algorithm for spectral analysis of unevenly sampled data," *Astrophysical Journal*, 1989.
- [11] S-Y Tseng, et al, "An effective heart rate variability processor design based on time-frequency analysis algorithm using windowed Lomb periodogram," *IEEE BioCAS*, 2010.
- [12] K. Kanoun, "A real-time compressed sensing-based personal electrocardiogram monitoring system," *IEEE DATE*, 2011.
- [13] MPARM. www-micrel.deis.unibo.it/sitoweb/research/mparm.html.
- [14] A. Y. Dogan, et al. "Low-power processor architecture exploration for online biomedical signal analysis", *CDS, IET*, 2012.
- [15] H. Hassanieh, et al. "Faster GPS Via the Sparse Fourier Transform," *ACM MOBICOM*, 2012.
- [16] H. Guo, et al, "Wavelet Transform Based Fast Approximate Fourier Transform," *IEEE ICASSP*, 1997.
- [17] MIT-BIH arrhythmia database. (2005). [Online].
- [18] L. Sörnmo, et al, "Bioelectrical Signal Processing in Cardiac and Neurological Applications," Elsevier, 2005.