A Methodology for Embedded Classification of Heartbeats Using Random Projections

Rubén Braojos, Giovanni Ansaloni *and* David Atienza
Embedded Systems Laboratory
École Polytechnique Fédérale de Lausanne, Switzerland
Email: {ruben.braojoslopez},{giovanni.ansaloni}, {david.atienza}@epfl.ch

***Abstract*—Smart Wireless Body Sensor Nodes (WBSNs) are a novel class of unobtrusive, battery-powered devices allowing the continuous monitoring and real-time interpretation of a subject's bio-signals. One of its most relevant applications is the acquisition and analysis of Electrocardiograms (ECGs). These low-power WBSN designs, while able to perform advanced signal processing to extract information on hearth conditions of subjects, are usually constrained in terms of computational power and transmission bandwidth. It is therefore beneficial to identify in the early stages of analysis which parts of an ECG acquisition are critical and activate only in these cases detailed (and computationally intensive) diagnosis algorithms. In this paper, we introduce and study the performance of a real-time optimized neuro-fuzzy classifier based on random projections, which is able to discern normal and pathological heartbeats on an embedded WBSN. Moreover, it exposes high confidence and low computational and memory requirements. Indeed, by focusing on abnormal heartbeats morphologies, we proved that a WBSN system can effectively enhance its efficiency, obtaining energy savings of as much as 63% in the signal processing stage and 68% in the subsequent wireless transmission when the proposed classifier is employed.**

I. Introduction

Wireless Body Sensor Nodes (WBSNs) are miniaturized, wearable systems able to measure and wirelessly transmit biological signals of patients. A major field of application of WBSNs is the ambulatory acquisition of electrocardiograms (ECGs), which evaluates the electrical activity of the heart. In this context, these devices allow long time monitoring of subjects producing little discomfort and requiring minimal medical supervision.

A recent trend in WBSNs, driven by the progress in semiconductor technologies, has been the emergence of "smart" wireless nodes [1]. These smart WBSNs (Figure 1) can, in addition to acquiring and transmitting data wirelessly, perform advanced digital signal processing filtering noise typically corrupting the signals and/or executing an automated diagnosis.

In the field of ambulatorial electrocardiography, an important early diagnosis step is to separate normal and pathological heartbeats. The benefit of this separation is two-fold: first, it can provide helpful information for speeding up the visual inspections of lengthy recordings by medical staff; second, it can be used to activate a more detailed analysis of only those beats presenting pathological characteristics.

The second scenario, investigated in this paper, can lead to two non-obvious, yet substantial, benefits. On the one hand, if a detailed diagnosis is performed off-node, it can be desirable to transmit or store only pathological beats on the WBSN, greatly reducing either the energy employed for wireless transmission or the data storage requirements, respectively. On the other hand, even if the analysis of beats is executed on the wireless node, computation effort can be reduced by activating it only when abnormal beats are detected, increasing energy efficiency.

While off-line algorithms have been proposed to classify the different heartbeat morphologies [2][3], their real-time implementation poses a considerable challenge on WBSN platforms, due to their high computational requirements. A neuro-fuzzy classifier (NFC) approach [4] is instead potentially a good candidate for this application. Its simple feed-forward structure makes it eligible to be optimized for, and executed on, the constrained resources typically present on WBSNs.

An important factor to consider is the high dimensionality of the heartbeat representation problem as, for every beat, tens of samples before and after the peak have to be considered to perform a reliable classification. To effectively address this problem, we propose and evaluate a methodology based on random projections (RPs) [5] to reduce the input size of the classifier. The approximation error introduced by random projections is theoretically bounded, nonetheless empirical evidence shows that certain projections perform better than others. Our experiments show that even a rather simple optimization, such as the one performed by a genetic algorithm [6] in few generations, can find a proper projection to obtain optimal classification results.

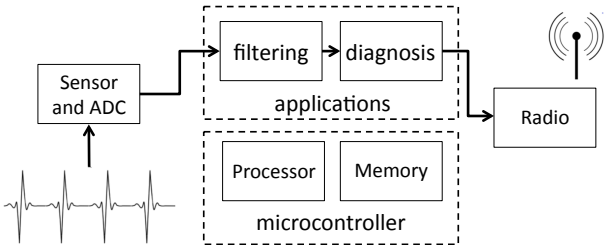

Fig. 1. Block scheme of a WBSN platform.

The research illustrated in this paper has been supported by the IcyHeart European project (FP7, capacities. Proj. num. 286130.)

To evaluate real-time performance of the trained and optimized classifier on a resource-constrained WBSN, we executed it on the IcyHeart System-on-Chip (SoC) [7]. This platform integrates a wireless transmitter, a multi-channel ADC converter and a low-power microprocessor (featuring a clock frequency of 6 MHz and an embedded RAM of 96 KBs), on a single die.

The resulting embedded software includes the RP-classifier, a filtering stage and a peak detector, used to isolate beats. Three-lead delineation has been instead adopted as an example of detailed analysis, and is activated by the classifier only for beats identified as pathological.

Test heartbeats were retrieved from the MIT-BIH Arrhythmia database [8], considering all beats presenting three different morphologies: normal sinus rhythms, left bundle branch blocks and premature ventricular contractions. Experimental results show that the proposed methodology can identify more than 97% of abnormal beats, while using a small fraction of the available SoC memory and computing resources.

Main contributions of this paper are the following:

- We explore the effectiveness of random projections to reduce the heartbeat representation dimensionality and therefore the classification problem complexity.
- We propose a complete design methodology to derive a real-time, lightweight heartbeat classifier based on a neuro-fuzzy structure. Moreover, we detail the required optimizations to efficiently execute it on a WBSN.
- We evaluate the accuracy and run-time performance of the resulting application targeting the state-of-the-art IcyHeart platform.

The paper proceeds as follows: Section II acknowledges related efforts in the field, Section III describes the implementation of the training and test phases of the NFC and the RP, detailing the optimizations performed to embed the trained application on a WBSN. Classification performance is reported in Section IV, while Section V concludes the paper.

II. Related work

Neuro-fuzzy classifiers (NFCs) [9] have been extensively studied in the literature. Their ability to explicitly express uncertainty in classification, given by the employed *fuzzy values*, makes them particularly well-suited to the problem of heartbeat classification, as proposed in [10].

NFCs can be effectively trained using established methods, the most common being the gradient descent algorithm described in [9] and the scale conjugate gradient introduced in [11] and [12], which is employed in this work. The approach is both computationally simpler and presenting lower memory requirements than comparable methods, like the ones based on support vector machines [2] and linear discriminants [13]. Therefore, it is more suitable for execution in embedded WBSNs.

Several state-of-the-art strategies for neuro-fuzzy classification of ECGs can be distinguished based on the methodology employed to extract the features of individual heartbeats generating the classifier input. Solutions to the features extraction

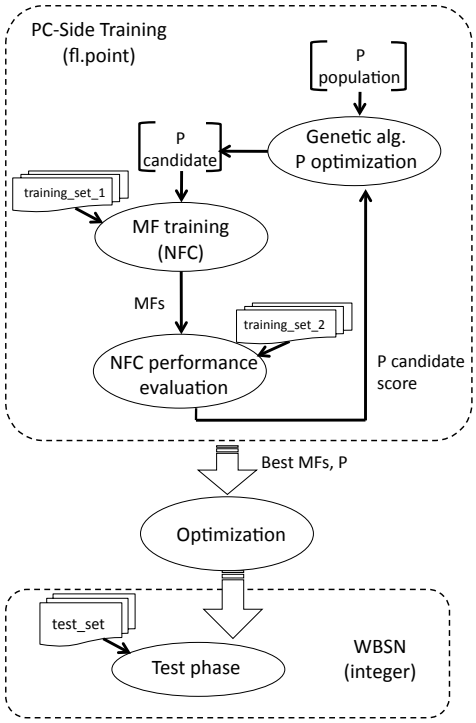

Fig. 2. RP-based fuzzy classification framework: PC-based training (top) and WBSN execution (bottom).

problem include the ones based on Independent or Principal component analysis (ICA [14] / PCA, [3]), Discrete Wavelets Transform (DWT [15]) and Discrete Cosine Transform (DCT [4]). As opposed to our work, these methods demand a computation effort not compatible with WBSN resources. A different approach, based on detection of morphological features, is presented in [16]; it nevertheless requires a detailed analysis of heartbeats before classification, conversely to our goal of early identification of pathological beats.

Random Projections (RPs) [17] offer an efficient solution for representing heartbeats with few coefficients. In particular, Achlioptas projections [5], employing matrices consisting only of the elements 0, 1 and -1, can be adopted to project ECG samples, while guaranteeing an upper bound on the approximation error and lowering computational complexity. An early work exploring RPs [18] for heartbeat classification gives a hint of their effectiveness. This results served as inspiration to investigate the methodology presented in the current manuscript.

III. Classification Methodology

The training and test phases of the classification framework, which are illustrated in Figure 2, have multiple and different constraints. On the one side, the training phase is performed off-line, on a host (PC) platform and using high-precision (floating-point) data representation to obtain the most accurate framework set-up. On the other side, the test phase is executed on an embedded WBSN, being therefore tightly constrained in memory footprint and run-time, admitting only integer arithmetic and avoiding exponential operations.

It is therefore mandatory to transform the classifier, after training and before execution, to lower its computational requirements according to the embedded platform capabilities. Our proposed methodology to realize this step is detailed in Section III-B, while Section III-A describes the algorithm used for off-line training of the RP matrix and the NFC. Finally, the result of our complete framework is a classifier which is both accurate and efficient.

A. Training phase

Random Projection. The first step of our proposed training phase is the generation of a $k \times d$ Achlioptas matrix (**P**), where d is the number of digital samples acquired for each heartbeat and k is the number of desired coefficients in the random projection, with $k \ll d$. **P** elements are defined [5] as:

$$\mathbf{P}_{k,d} = \begin{cases} +1 & \text{with probability } \frac{1}{6} \\ -1 & \text{with probability } \frac{1}{6} \\ 0 & \text{with probability } \frac{2}{3} \end{cases}$$

Each line of **P** indicates which elements of an input vector v have to be added (possibly negated) to derive the corresponding vector u of randomly-projected coefficients: $u = \mathbf{P}v$.

Neuro-fuzzy classifier. The coefficients are the inputs to the multi-layered NFC (Figure 3). Its first *membership* layer employs Membership Functions (MFs) to compute, for each coefficient k, a membership grade μ for each of the three classes l: normal beats (N), left branch block (L) and premature ventricular contraction (V). During the training phase, MFs are gaussian curves, defined by their center c and variance σ:

$$\mu_{k,l}(u_k) = \exp\left(\frac{-(u_k - c_{k,l})^2}{2\sigma_{k,l}^2}\right)$$

where $l \in \{N, V, L\}$.

In the subsequent *fuzzification* layer, the membership grades of all coefficients for each class are multiplied: $f_l = \prod_k \mu_{k,l}$. The resulting fuzzy value expresses how strongly an examined heartbeat belongs to that specific class: the larger the value (with respect to the values of the other classes), the higher the confidence in the correctness of the assignment.

The third *defuzzification* layer of the NFC marks each beat as either normal or pathological, by considering the first and second maximum of the fuzzy values ($M1_f$, $M2_f$) and their sum $S = \sum_l f_l$. If $(M1_f - M2_f) \geq \alpha_{train} S$ (with $\alpha_{train} \in [0, 1]$), the beat is assigned to the class with the maximum fuzzy value (N, V, L). Otherwise, the beat is marked as Unknown (U). V, L, U beats are considered as possibly pathological, while N beats are marked as normal.

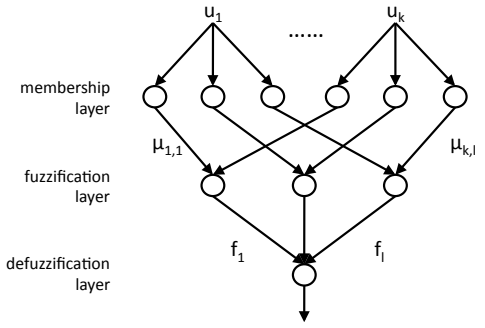

Fig. 3. Three-layers neuro-fuzzy classifier.

The choice of a proper defuzzification coefficient α_{train} gives the flexibility to unbalance the classifier training process, fixing the percentage of abnormal beats incorrectly classified as N. The performance (score) of the trained classifier is then the corresponding percentage of normal beats correctly detected.

Two-step training. The aim of the training is to search for high-performance **P** matrix and MFs. The algorithm starts from an initial population of random projection matrices $\mathbf{P}_{pop}$, optimizing the MFs of the related classifier over a first set of projected heartbeats (*training_ set_1*) using the scale conjugate method [11]. The performance of the random projection is then the score of the associated NFC over a second training set of beats (*training_set_2*).

To optimize the projection, we use a genetic algorithm that considers each $\mathbf{P} \in \mathbf{P}_{pop}$ as chromosomes and, performing crossover and mutation operations on them, derive higher-performance candidates. For the experiments presented in Section IV, we employed an initial population of 20 matrices and let the algorithm run for 30 generations.

B. Resource-constrained optimization phase

The optimized projection and the trained classifier cannot be employed "as they are" in a WBSN platform. A first consideration is that data must be represented as integers, as opposed to the floating point format used in the training phase.

Then, several other elements are particularly critical. First, gaussian MFs employed in the NFC make use of exponential operations, which are not easily (or not at all) implementable in embedded platforms. Second, the NFC fuzzification layer must be analyzed to prevent overflows when performing the product operation. Finally, the memory required to store the random projection matrix can exceed the WBSN resources.

Membership functions linearization. Given their center c and standard deviation σ, gaussian MFs are approximated to the integer range $[0, (2^{16} - 1)]$ during the optimization step using four segments:

$$MF_{lin}(x) = \begin{cases} 0 & \text{if } |c - x| \geq 4S \\ 1 & \text{if } 4S > |c - x| \geq 2S \\ lin.approx1 & \text{if } 2S > |c - x| \geq S \\ lin.approx2 & \text{if } S > |c - x| \end{cases}$$

where MF_{lin} is the linearized MF and $S = 2.35\sigma$. The linear approximation segments are graphically represented in Figure 4. This formulation has the desirable property to be positive in a large range; hence, it is rare that a fuzzy value becomes 0 after the defuzzification (product) classifier stage.

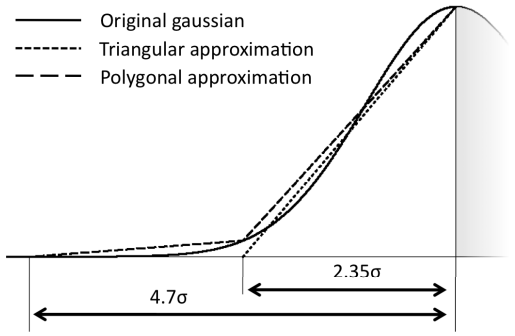

Fig. 4. Linear approximation of gaussian MFs in the range $[-4.7\sigma, 0]$, compared to a gaussian curve and a simpler triangular interpolation.

Fuzzification/defuzzification. Only the ratio between the fuzzy coefficients f_l, as opposed to their value, is relevant for the defuzzification layer. From this observation stems the opportunity to optimize the fuzzification step, retaining the maximum precision given the 32-bit representation used for the results of the defuzzification products. In our proposed implementation, the membership grades $\mu_{k,l}$ related to the two first coefficients are multiplied for each of the three classes. The three resulting numbers are left-shifted to the maximum amount so that none of them overflow and then the rightmost 16 bits are discarded. All subsequent membership grades are then processed in a similar fashion, obtaining the fuzzy values of the beats for the different classes.

Defuzzification marks each beat as normal or pathological, much like in the training phase. The chosen implementation does not employ divisions, and can therefore be efficiently implemented in WBSNs. Moreover, it is possible to tune the defuzzification coefficient α_{test} independently of the α_{train} chosen during the training phase (described in Section III-A), giving the opportunity to adjust the ratio of detected normal and abnormal beats.

Random Projection matrix representation. The **P** matrix is generated in such a way that its elements only assume three values ($+1, -1$ and 0). We therefore use a compact representation where each element is coded using two bits, which requires 1/4 of the memory with respect to a corresponding matrix of 8-bits values.

Downsampling can also be applied to reduce the memory occupied by the random projection matrix. If, for example, one every four samples of the acquired signal is considered, the size of the matrix is reduced by a factor of four.

IV. Experimental Results

A. Experimental Setup

To evaluate the proposed RP-based ECG classifier, we investigated its performance when evaluating the normal (N), ventricular contraction (V) and left branch block (L) beats present in the MIT-BIH Arrhythmia Database recordings, available on the Physiobank website [8].

For each beat, classification in either the normal or pathological classes is performed. As mentioned in Section I, the goal of the classifier is to recognize abnormal beats, activating only in these cases a detailed analysis such as the three-lead delineation proposed in Figure 6. In this context, figures of merit of the classifier are the Normal Discard Rate (NDR) and Abnormal Recognition Rate (ARR). NDR assesses the rate of normal beats that are correctly identified as such and thus discarded. Complementarily, ARR reports the percentage of abnormal beats that correctly activate the delineation block.

	N	**V**	**L**	**Total**
training_set_1	150	150	150	450
training_set_2	10024	892	1084	12000
test_set	74355	6618	8039	89012

TABLE I
SIZE AND COMPOSITION OF THE TWO TRAINING SETS AND OF THE TEST SET OF ECG HEARTBEATS.

Two randomly-selected excerpts of the database where used for training the NFC and optimizing the random projection matrix, respectively. The first training set is smaller than the second because training of the NFC is more computationally complex than its evaluation; in this way, it was possible to retrieve "good" solutions in a reasonable time. The *test_set* comprises all N, V, L beats present in the database (the composition of the sets is shown in Table I). Across experiments, the defuzzification coefficient α_{train} was chosen to have a minimum ARR of at least 97% on *training_set_2*.

ECG recordings on the database are acquired at 360 Hz; we define each heartbeat as spanning 100 samples before and 100 samples after its peak. Peaks are automatically detected using a wavelet-based technique, firstly proposed in [1], which decompose the input signals in four dyadic scales at different frequencies, retrieving the peak as the zero-crossing point on the first scale in-between couples of maximum-minimum points across scales.

B. Evaluation of a suitable RP coefficients number

A first important design choice is to determine the number of random projection coefficients. A larger coefficient set impacts both the size of the random projection matrix and the complexity of the NFC, the decision being therefore a trade-off between classification accuracy and real-time performance (run-time and required memory). To explore this aspect, we trained and tested the framework with a number of coefficients ranging from 8 to 32. Table II reports the NDR corresponding to a minimum ARR of 97% on the *test_set*. For comparison, the table also presents the classification performance obtained using the off-line Principal Component Analysis (PCA) algorithm proposed in [3] to reduce the representation dimensionality.

coefficients	**8**	**16**	**32**
NDR-PC	93.74	95.16	93.05
NDR-WBSN	92.31	92.53	93.04
PCA-PC	93.66	95.78	89.75

TABLE II
NORMAL DISCARD RATE (NDR) ON *test_set* FOR A FIXED ABNORMAL RECOGNITION RATE (ARR) OF 97%, VARYING THE NUMBER OF COEFFICIENTS.

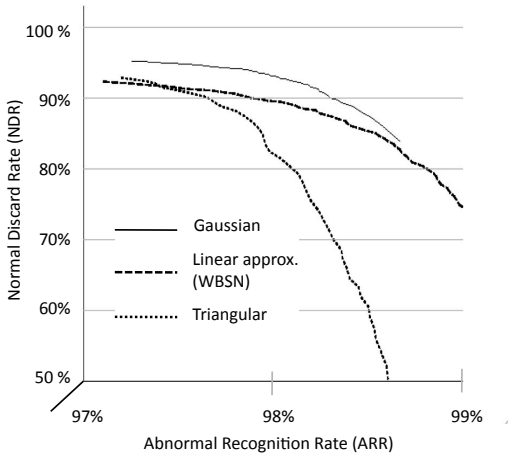

Fig. 5. Pareto fronts of the best NDR/ARR solutions, employing gaussian, linearly approximated or triangular MFs.

Values in the table are classified in three settings. The first row ($NDR-PC$) refers to a PC-based implementation were no approximation is performed: it therefore uses floating point data representation and gaussian membership functions. The second row ($NDR-WBSN$) instead shows the performance of the corresponding WBSN version, employing linearized MFs, integer data and a sampling rate of 90 Hz (a four-times downsampling of the original recordings). In the third row ($PCA-PC$), coefficients are derived from PCA instead of random projections.

Two conclusions can be drawn from these experiments. Firstly, only a small number of randomly-projected coefficients are sufficient to achieve a NDR of over 90%. Actually, increasing this value from 8 to 32 does not benefit classification in a tangible way. Secondly, the difference between $NDR-PCA$, $NDR-PC$ and $NDR-WBSN$ is rather small, being always in the range of few percentage points. This shows that the proposed methodology does not penalise accuracy, even if the computation effort required in the three cases differs greatly.

C. Impact of MF linearization

The error induced by linearizing the MFs is plotted in Figure 5. In this set of experiments, we adjusted the α_{train} parameter to have a minimum ARR of 97% in *training_set_2*, and scaled α_{test} to obtain different NDR/ARR trade-offs on *test_set*. In all cases, 50 samples acquired at 90 Hz were randomly projected on 8 coefficients. The figure compares the NDR/ARR pareto fronts obtained from gaussian MFs, from the linear approximation described in Section III-B and from the simpler triangular approximation shown in Figure 4.

Triangular MFs, even if achieving good results for lower ARRs, cannot scale well if higher recognition rates of abnormal beats are desired. Performance of the linearly-approximated NFC instead closely follows the one of the gaussian classifier, while still requiring little computational effort. In the gaussian and linear approximation cases, it is possible to correctly recognize 98.5% of abnormal beats, with a NDR of 87%. Using triangular MFs, the NDR figure for the same recognition rate drops to 62%.

	Code Size (KB)	Duty Cycle
RP-classifier	1.64	< 0.01
RP + filtering + peak detection (1)	30.29	0.12
Multi-lead delineation (2)	46.39	0.83
Proposed system (3)	76.68	0.30

TABLE III
CODE SIZE AND DUTY CYCLE OF THE SUB-SYSTEMS IDENTIFIED IN FIGURE 6 USING 8 COEFFICIENTS. TESTS PERFORMED ON THE ICYFLEX WBSN RUNNING AT 6 MHZ

D. Run-time and memory size evaluation

RP-based detection of pathological beats must not be the computation bottleneck of a WBSN system during real-time execution. As described in Section I, the role of the classifier is to activate a detailed analysis for abnormal beats, therefore, it should require considerably less effort than performing analysis over the full signal.

In this section we investigate, as an example of a complete diagnosis application, a system (system (3) in Figure 6), in which RP-classification, performed on a single lead (sub-system (1)), is used to trigger three-leads delineation of pathological heartbeats (sub-system (2)).

Figure 6 shows that, apart from the RP-classification block, two additional stages have to be incorporated in the classification sub-system (1). Firstly, a filtering stage is required to remove artifacts caused by respiration and muscle contractions usually corrupting ECG signals. Secondly, a peak detector has to be employed to identify heartbeats.

We employ state-of-the-art solutions for the filtering stages, the peak detector and the delineation block, proposed by the authors of [1]. Filtering is realized using morphological operators, a wavelet-based algorithm is used for peak detection and a delineation algorithm using multi-scale morphological derivatives (MMDs) is executed over the combination of the three filtered leads in the subsystem (2). Their implementation has been highly optimized for execution on embedded WBSNs. Eight coefficients are used in the RP-classifier stage.

Table III reports the code size and the run-time of the different parts of the considered system when executed on the IcyHeart WBSN operating at a 6 MHz. The first row reports the figures obtained for the RP-classifier alone, while the second one also considers the filtering and peak detection stages of sub-system (1).

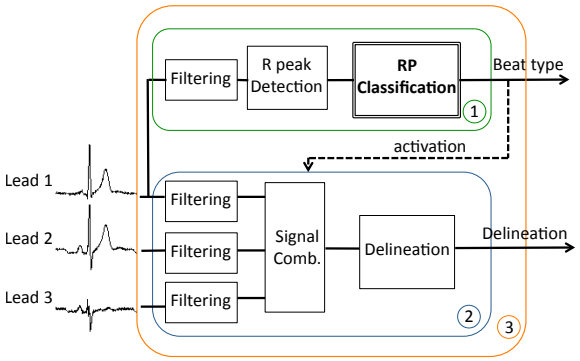

Fig. 6. Experimental scenario: the RP-based neuro-fuzzy classifier is used to activate an accurate multi-lead delineation only in case of heartbeat abnormality.

In the third row, a system performing multi-lead MMD delineation on the full input signal is investigated. This setting reflects the performance of a delineator that is always active, analyzing both normal and pathological heartbeats. Finally, in the last row of the table, the values for the complete system are given, where delineation is performed only on abnormal beats, as identified by the RP classification.

A first observation that can be derived from these experiments is that the RP classification requires little resources (less than 1% of the duty cycle and 2 KB of memory). As an assessment of its efficiency, we observe that most of the computation performed by sub-system (1) is actually needed for input data filtering and peak detection, as opposed to the RP-based NFC itself.

The second, and most important, observation is that the duty cycle of the complete application is tangibly lower than an equivalent one that performs a full delineation on all beats. Indeed, experimental evidence shows that the run-time of sub-system (3) is 63% lower than that of sub-system (2), while presenting a memory size overhead of 30 KB, thus justifying our assumptions on the benefits of early classification.

E. Improvement of energy efficiency

Detection of pathological heartbeats can be exploited to obtain considerable gains in energy efficiency. In addition to the reduction in computation effort discussed in the previous section, a further benefit derives from the optimization of the data to be transmitted from the WBSN on the power-hungry wireless link.

Assuming a scenario in which a WBSN reports only the peak of normal beats, and all fiducial points (onset, peak and end of the three characteristic waves composing the beat) for abnormal ones, usage of the wireless link can be substantially reduced, with respect to the case were all fiducial points of all beats are communicated. In fact, considering all beats in the *test_set* described in Table I as input signals, we achieve a 68% energy consumption reduction in the wireless module and 63% reduction in the energy consumption of the bio-signal analysis part. Thus, overall we achieve an estimated 23% total energy reduction, as computation and wireless communication are two major contributors to the power budget of smart WBSNs (their combined figures accounting for approximately 34% total energy in typical WBSN implementations [1]).

V. Conclusion

In this work, we have presented a real-time framework for detecting pathological heartbeats on an embedded WBSN platform using a neuro-fuzzy classifier.

The paper addressed two major issues: on the one hand, a solution based on random projections and optimized off-line using a genetic algorithm was investigated to reduce the classifier input dimensionality. On the other hand, we proposed a strategy for adapting the trained NFC to the constraints of an embedded platform. The optimized RP-classifier only uses integer arithmetic and can process heartbeats efficiently, while still being able to correctly discern normal and abnormal morphologies with high confidence. Consequently, substantial savings are achieved in run-time and energy required for computation and wireless communication.

Experiments carried on the complete set of normal, premature ventricular contraction and left branch block beats present in the MIT-BIH Arrhythmia database, showed that the proposed RP-classifier can successfully detect up to 97% of the abnormal beats, while only misinterpreting 7% of normal ones as pathological.

References

[1] F. Rincon, J. Recas, N. Khaled, and D. Atienza, "Development and evaluation of multilead wavelet-based ECG delineation algorithms for embedded wireless sensor nodes," *Information Technology in Biomedicine, IEEE Transactions on*, vol. 15, no. 6, pp. 854–863, Nov. 2011.

[2] X. Jiang, L. Zhang, Q. Zhao, and S. Albayrak, "ECG arrhythmias recognition system based on independent component analysis feature extraction," in *TENCON 2006. 2006 IEEE Region 10 Conference*, Nov. 2006, pp. 1–4.

[3] R. Ceylan and Y. Ozbay, "Comparison of FCM, PCA and WT techniques for classification ECG arrhythmias using artificial neural network," *Expert Systems with Applications*, vol. 33, no. 2, pp. 286–295, 2007.

[4] V.-E. Neagoe, I.-F. Iatan, and S. Grunwald, "A neuro-fuzzy approach to classification of ECG signals for ischemic heart disease diagnosis," in *American Medical Informatics Association Annual Symposium*, 2003, pp. 494–498.

[5] D. Achlioptas, "Database-friendly random projections: Johnson-Lindenstrauss with binary coins," *Journal of Computer and System Sciences*, vol. 66, no. 4, pp. 671– 687, June 2003.

[6] J. Yang and V. Honavar, "Feature subset selection using a genetic algorithm," *Intelligent Systems and their Applications, IEEE*, vol. 13, no. 2, pp. 44 – 49, mar/apr 1998.

[7] IcyHeart European Project. [Online]. Available: www.icyheart-project.eu/

[8] PhysioBank. [Online]. Available: www.physionet.org/physiobank/

[9] C.-T. Sun and J.-S. Jang, "A neuro-fuzzy classifier and its applications," in *Fuzzy Systems, 1993., Second IEEE International Conference on*, 1993, pp. 94–98 vol.1.

[10] R. Ceylan, Y. Ozbay, and B. Karlik, "A novel approach for classification of ECG arrhythmias: Type-2 fuzzy clustering neural network," *Expert Systems with Applications*, vol. 36, no. 3, Part 2, pp. 6721– 6726, Aug. 2009.

[11] M. F. Moller, "A scaled conjugate gradient algorithm for fast supervised learning," *Neural Networks*, vol. 6, no. 4, pp. 525–533, 1993.

[12] B. Cetisli and A. Barkana, "Speeding up the scaled conjugate gradient algorithm and its application in neuro-fuzzy classifier training," *Soft Computing*, vol. 14, no. 4, pp. 365–378, Oct. 2009.

[13] P. de Chazal, M. O'Dwyer, and R. Reilly, "Automatic classification of heartbeats using ECG morphology and heartbeat interval features," *Biomedical Engineering, IEEE Transactions on*, vol. 51, no. 7, pp. 1196 –1206, July 2004.

[14] S.-N. Yu and K.-T. Chou, "Integration of independent component analysis and neural networks for ECG beat classification," *Expert Systems with Applications*, vol. 34, no. 4, pp. 2841–2846, May 2008.

[15] I. Guler and E. D. Ubeyli, "ECG beat classifier designed by combined neural network model," *Pattern Recognition*, vol. 38, no. 2, pp. 199 – 208, Feb. 2005.

[16] C. Ye, M. Coimbra, and B. Vijaya Kumar, "Arrhythmia detection and classification using morphological and dynamic features of ECG signals," in *Engineering in Medicine and Biology Society (EMBC), 2010 Annual International Conference of the IEEE*, Sept. 2010, pp. 1918 – 1921.

[17] E. Candes and T. Tao, "Near-optimal signal recovery from random projections: Universal encoding strategies?" *Information Theory, IEEE Transactions on*, vol. 52, no. 12, pp. 5406 –5425, Dec. 2006.

[18] I. Bogdanova, F. J. Rincon, and D. Atienza, "A multi-lead ECG classification based on random projection features," in *37th IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Mar. 2012, pp. 625–628.