

# A Wireless Body Sensor Network For Activity Monitoring With Low Transmission Overhead

Rubén Braojos<sup>‡</sup>, Ivan Beretta<sup>‡</sup>, Jeremy Constantin<sup>§</sup>, Andreas Burg<sup>§</sup> and David Atienza<sup>‡</sup>

<sup>‡</sup> Embedded Systems Laboratory; <sup>§</sup> Telecommunications Circuits Lab

École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

E-mail: {ruben.braojoslopez, ivan.beretta, jeremy.constantin, andreas.burg, david.atienza}@epfl.ch

**Abstract**—Activity recognition has been a research field of high interest over the last years, and it finds application in the medical domain, as well as personal healthcare monitoring during daily home- and sports-activities. With the aim of producing minimum discomfort while performing supervision of subjects, miniaturized networks of low-power wireless nodes are typically deployed on the body to gather and transmit physiological data, thus forming a *Wireless Body Sensor Network* (WBSN). In this work, we propose a WBSN for online activity monitoring, which combines the sensing capabilities of wearable nodes and the high computational resources of modern smartphones. The proposed solution provides different tradeoffs between classification accuracy and energy consumption, thanks to different workloads assigned to the nodes and to the mobile phone in different network configurations. In particular, our WBSN is able to achieve very high activity recognition accuracies (up to 97.2%) on multiple subjects, while significantly reducing the sampling frequency and the volume of transmitted data with respect to other state-of-the-art solutions.

## I. INTRODUCTION

Over the last years, Wireless Body Sensor Networks (WBSNs) have emerged as a leading technology that is poised to drastically change healthcare delivery and the everyday life of subjects. Thanks to a combination of wearable low-power sensor nodes that communicate through a wireless channel, these networks allow continuous and unobtrusive monitoring of physiological signals and activities, both for personal and medical purposes. The functionalities and the ease of use of these systems have also received a significant boost thanks to the spread of handheld devices (such as smartphones) [1], which represent the ideal high-performance complement to wearable nodes. In fact, smartphones can provide advanced features such as data logging, transmission to a remote location, and user interface, without affecting the nomadic nature of WBSNs.

In the context of personal monitoring, activity recognition attracts a high interest from researchers [2], [3]. Activity monitoring finds application mainly in the healthcare domain, such as in the supervision of patients affected by Parkinson’s disease [4], but it is also employed in sports and home monitoring [2]. The activity of a subject can be determined from kinetic data, such as acceleration [5] and orientation, collected by a set of nodes located on parts of the body that convey most of the information, such as limbs and joints. In order to minimize the subject’s discomfort while wearing a WBSN for a prolonged time, a minimum number of nodes

has to be deployed, and their size has to be miniaturized. The latter requirement poses significant limitations on the size of the batteries used by the wearable nodes, thus forcing them to limit their computational effort and the volume of data they transmit [6]. A common strategy to increase the lifetime of a WBSN [2] is the preliminary extraction of a set of features from the sensed data: this operation, performed directly on the wearable node, massively reduces the amount of transmitted data. These features are then gathered and used to feed a classification algorithm, which may run either offline on a workstation [7], or online on one of the nodes [8]. The latter option is necessary to ensure a high degree of mobility of the system, but also affects the accuracy of the classification due to the limited computation capabilities of the wearable nodes [9].

In this work, we propose a WBSN for activity monitoring that combines wireless sensor nodes and a smartphone, in order to provide different tradeoffs between classification accuracy and transmission volume (which, in turn, has a major impact on energy consumption). In an accuracy-oriented configuration, the nodes of the proposed network send their extracted features to the smartphone, which performs a precise and complex classification thanks to its high computational capabilities. Conversely, in a transmission-aware configuration, one of the wearable nodes is in charge of performing the classification and sending the result to the smartphone, which acts only as a gateway toward the user.

Specifically, with respect to the existing work in this field, the proposed WBSN includes the following contributions:

- It enables different accuracy/transmission tradeoffs, thanks to the different work assignments that can be performed among the nodes and the smartphone;
- It provides two different classification algorithms that fit the computational capabilities of the WBSN nodes and the smartphone. In particular, a floating point implementation of a neuro-fuzzy classifier as the one in [9] is proposed to perform classification on the smartphone, while a lightweight decision tree classifier is employed to perform the same task on a wearable node. Both approaches show competitive accuracies with respect to the state of the art reaching a sensitivity of 97.2% in the case of the smartphone-based classification;
- It requires a lower sampling frequency (50 Hz) with respect to most of the existing activity monitoring systems;

- It achieves similar accuracies across multiple subjects, hence it does not require an *ad hoc* training phase;
- It can be easily extended to monitor additional signals, thanks to the small bandwidth required for activity monitoring purposes. For example, the proposed network can optionally sense the electrocardiography (ECG) from the chest, in order to provide a smart holter system that associates the heartbeat rate to the current activity.

The remainder of this paper is structured as follows. Section II presents an overview of the most relevant WBSNs for activity monitoring that can be found in the literature. In Section III, we describe and motivate the design choices for the proposed network, and we then discuss the main implementation details in Section IV. Section V provides an experimental validation of the WBSN, and Section VI concludes the paper.

## II. RELATED WORK

An extensive survey of the literature related to activity monitoring is available in [2], which also provides a quantitative comparison of the accuracy of the existing solutions. Multiple conclusions can be derived from this analysis, regarding both architectural and methodological aspects. First, a large variety of node placements has been proposed over the years, including wrist, upper arm, chest, hip, thigh, crus, ankle, and several combinations of them. However, not all the existing solutions aim at minimizing the patient’s discomfort, often purposely increasing the number of sensors to collect a larger amount of data. Second, common patterns can be identified to increase the system lifetime: on-node feature extraction – which heavily reduces data transmission– is applied in most circumstances, and efforts to reduce the sampling frequency can also be found. However, most of the existing works still employ sampling rates of more than 100Hz because of the dynamic nature of activity recognition, and this has a negative impact on the system power consumption due to the sensing circuit. Finally, most of the existing classification techniques fail to reach accuracies above 90%, and the best results are achieved by sensors on the chest and on highly-dynamic joints (such as wrists and ankles), the latter creating major discomfort during everyday activities. In this work, we aim at minimizing the invasiveness of the network by including 3 sensors: one on the chest, which can be worn using comfortable cardio belts, and two on the legs (crus and thigh), which help to discriminate static and dynamic activities and can be integrated into smart clothes. In terms of energy consumption, we combine well-known feature extraction techniques with a classification scheme that, with the selected node placement, leads to accuracies of over 97% even with a sample frequency of only 50Hz.

In the remainder of this section, we provide a more detailed analysis of the most accurate [2] and energy-efficient solutions in the literature. In particular, the work in [7] reaches the highest classification accuracy (95%) by employing an artificial neural network that processes acceleration data coming from the wrist. Although the authors successfully manage to employ only accelerometer data, the classification is performed offline, and no indication regarding how the proposed technique would

behave on portable devices with limited resources is provided. In this work, we propose a similar classification strategy, which is however targeted for online detection on a smartphone, and is able to improve the accuracy thanks to the additional information coming from the legs.

In [8], the authors propose an online monitoring technique using a watch-like sensor on the wrist, which acquires data from accelerometers, a microphone and light sensors. In order to cope with the large amount of incoming data, the system employs feature extraction both in the time and in the frequency domain and a nearest neighbors classifier, thus requiring a relatively high computation effort that is compensated by the energy harvesting capabilities of the device. The proposed technique achieves an accuracy of approximately 91% on arm-related movements. Our approach, on the other hand, employs more sensors in order to achieve higher accuracies on activities that involve the lower limbs.

The work in [3] proposes a network to monitor general-case activities (similar to the ones targeted in our work) using a hidden Markov model, which acts on data coming from sensors on shoulder, chest and wrist. The authors focus on the importance of cross-subject accuracy and, in particular, of the ability to work with sensors that are slightly displaced when worn by different persons. The proposed classification achieves accuracies of up to 90%, but the confusion matrices show pronounced difficulties in discriminating activities where the upper body is static (e.g., sitting and standing). In our solution, we overcome the limitations of [3] by exploiting the information conveyed by the sensors on the legs, while maintaining the strong emphasis on cross-subject accuracy.

In [5], the authors propose a system for online monitoring focused on activity changes. The network is able to perform most of the computations on a node located on the chest, and to this end it employs a simple classifier based on a decision tree. The proposed methodology, combined with a set of custom features, leads to an overall accuracy of 90.8%. However, although the system achieves peaks of accuracy when detecting transitions, the performance while monitoring ongoing activity shows high classification degradations. In our WBSN, we employ a similar low-complexity decision tree, which effectively takes advantage of the additional information from the legs to make the overall classification more robust.

The most recent work in the field of activity monitoring is discussed in [4], where a system for monitoring patients affected by Parkinson’s disease is discussed. The authors propose a power saving technique that reduces sampling frequency during static activities, such as lying and sitting. The architecture can reduce the sampling rate down to 30Hz during sleep periods, but it requires high sampling frequency (up to 200Hz) during most of the other activities. Although this policy proves to be effective in the target scenario, where the patients are spending the majority of their time on static activities, our methodology proves to be a better option in the general case, as it enables a sampling frequency of only 50Hz under all circumstances. In addition, the classification accuracy of the proposed WBSN outperforms the one of [4], with the

additional benefit of working in a subject-independent fashion without requiring customized training.

### III. SYSTEM DEFINITION

The proposed WBSN for activity monitoring is illustrated in Figure 1 and it aims at detecting 7 activities: walking, sitting, standing, laying, running, and walking upstairs and downstairs. The network is composed of three wearable devices deployed on the chest, left crus, and right thigh of the subject, in order to sense data related to acceleration, orientation, and optionally other biosignals of general interest, such as the ECG. In addition, a smartphone is also incorporated in the network to provide higher computation capabilities that can be exploited either during the activity detection, or for enhanced high-level functionalities. The workload balance among the devices has an impact on the quality of the classification and on the amount of transmitted data, which in turn plays a major role for the overall energy consumption. Two tradeoffs between accuracy and transmission are identified for the proposed system, and they are discussed in depth in the following.

#### A. Device Taxonomy

The proposed WBSN includes two classes of devices, each one characterized by its different computation capabilities, size and energy budget. In particular, the network is composed of:

- A set of **wearable sensor nodes**, which are in charge of sampling the signals of interest by incorporating accelerometers, gyroscopes, and optionally other sensors. These devices are battery-powered and are required to be small to minimize the subject's discomfort: the combination of these requirements leads to a limited energy budget, which translates into low-power systems with limited computation capabilities. In fact, wearable nodes typically feature very basic microcontroller architectures with integer arithmetic modules and modest memory resources, thus imposing significant limits on the complexity of the algorithms they can execute. Sensor nodes usually include a radio for data transmission, which typically represents the most energy-hungry component of the system [6].
- A **smartphone**, which provides higher computational capabilities due to more complex processing cores, and larger batteries that are generally recharged on a daily basis. These devices typically include hardware support for floating point operations, larger volatile and non-volatile memories, and multiple wireless interfaces. As a consequence, these devices are ideal for executing complex and accurate classification algorithms, as well as high-level operations, e.g., interfacing with the user, logging the recorded data, and possibly communicating this data to a remote service over the internet.

The two kinds of devices have different roles in the WBSN. In particular, the classification can be performed either on a selected wearable node, which collects the data from the other sensors and then executes the classification algorithm, or by the smartphone, which receives the data from all the

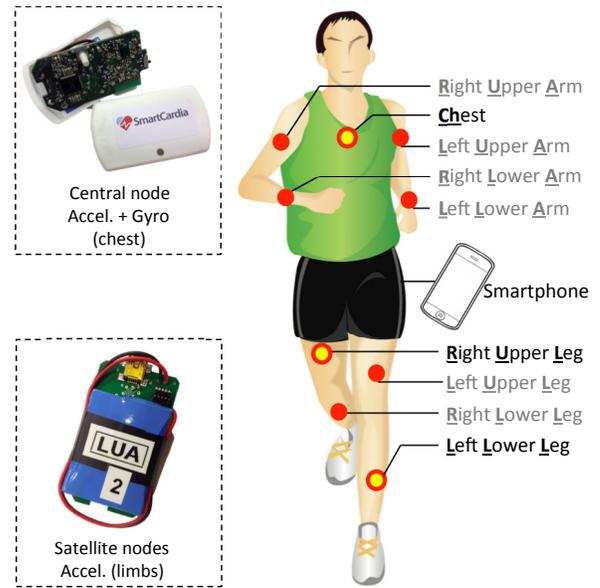


Fig. 1. Overview of the proposed WBSN. The three nodes in yellow are the ones belonging to the final configuration, whereas the remaining nodes have only been used for data collection purposes

sensor nodes. These two strategies generate different traffic volumes through the WBSN and, while the former approach is only applicable when the classification algorithm is relatively simple, the latter provides a more general solution that allows to implement complex classifiers with higher precision.

#### B. Network Structure and Task Assignment

In the proposed WBSN, three wearable nodes are placed on the body, as shown in Figure 1: two nodes including accelerometers are located on the left crus and on the right thigh, and a more complex node that also includes gyroscopes is located on the chest. The node on the chest can be optionally used to sense and analyze the ECG signal [10], thus extending the functionality of the network. The position of the nodes has been selected after a detailed analysis of multiple possible combinations, which is reported in Section V, aiming at deploying a minimum number of nodes while achieving a high classification accuracy.

The structure of the communication among the nodes, as well as the workload distribution throughout the network, is dictated by a tradeoff between the volume of transmitted data and the final classification accuracy, and can possibly be changed dynamically. A network with localized workload would send a lower amount of data through the wireless channel, but most of the computation would be performed with lower precision on microcontrollers with limited resources. In order to tackle this tradeoff, we propose two alternative methods, which are illustrated in Figure 2. The two strategies are defined as follows:

- **Smartphone-based classification** (Figure 2a): in this configuration, each wearable node performs a local feature extraction on the data, and then sends the resulting information to the smartphone, which acts as the center

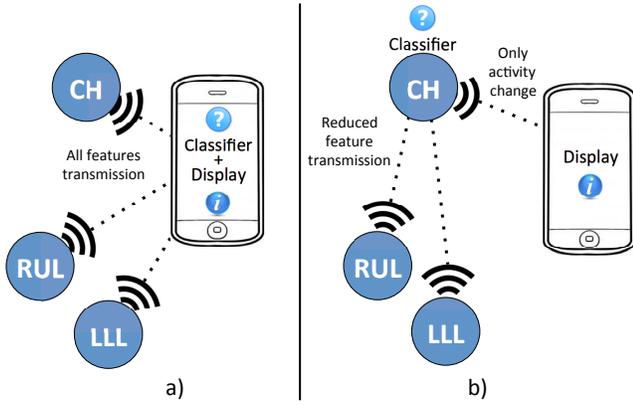


Fig. 2. The two proposed classification strategies: smartphone-based (a), and on-node (b). In both scenarios, the phone acts as network coordinator.

of a star topology. Then, the smartphone performs a complex real-time classification on its CPU, enabled by the floating point support and the larger amount of resources available on the device. This configuration allows a more accurate classification, but also requires all the nodes to transmit their data, thus generating a high traffic volume;

- **On-node classification** (Figure 2b): in this configuration, a selected wearable node is in charge of collecting data from the other devices, performing a local classification, and then sending the result to the smartphone only when a change of activity is detected (periodic packets are also exchanged between the node and the smartphone to probe the connectivity). However, the on-node classification is less accurate than the classification performed on the smartphone, because of the limited computational resources. In the proposed network, the node on the chest is selected as the center of the network, as it can be assumed to have a clearer transmission path to the smartphone with respect to the devices on the leg. Moreover this node is responsible of sensing a higher number of channels (i.e. accelerometers and gyroscope) and selecting it as the center of classification reduces the amount of data to be transmitted within the network. Finally, as the node on the chest is worn over a cardio belt, its size and its battery can be enlarged to ensure a sufficient lifetime in spite of the higher workload.

### C. On-Node Feature Extraction

It is worth noting that, in both strategies defined in the previous section, the wearable nodes are in charge of performing feature extraction on the sensed signal. This is a common technique [2] to avoid the transmission of the entire signal, which would not be feasible for small battery-powered devices. After acquiring the samples from the sensor, each wearable node extracts relevant information from the raw data stream using the embedded microcontroller. This strategy also has a positive effect on the classification, as the classifier can act on data with a low degree of redundancy.

Multiple features can be computed on a wearable node, each

one conveying different information, and requiring different computational efforts to be extracted. In the literature [2], features are typically classified into four groups:

- *Time-domain features*, which include mean, standard deviation, median, percentiles, derivatives, zero crossings, and many others. They are extracted directly from the signal, and provide relevant information about its waveform and its statistical behavior. Their computation complexity is linear with respect to the signal size, thus making them ideal for low-power on-node extraction;
- *Frequency-domain features*, which can be derived from the fast Fourier transform (FFT) of the signal, and include the spectral energy and entropy, the principal frequency energy, and a selection of the first  $n$  coefficients of the FFT. Their complexity is superlinear;
- *Wavelet features*, which include the coefficients of the wavelet transform of the signal, and can detect variations in the frequency components over time;
- *Custom features*, which are derived *ad hoc* from domain-specific considerations about the target problem. In the context of activity monitoring, these features include signal magnitude area [5], inter-axis correlation [11], and time-domain gait detection [12].

A quantitative comparison of different feature sets is provided in [13], though the reference classifier is different from the ones employed in the proposed WBSN. Results show that, while frequency-domain features lead to the best classification accuracy, time-domain features provide the most intriguing trade-off between accuracy and complexity. Conversely, wavelet features provide good results when detecting transitions among different activities [2], but they do not perform well overall [13]. In the proposed WBSN, we extract two time-domain features for each axis of the accelerometers and the gyroscope: mean and standard deviation. These features, along with the aforementioned low computational effort, lead to the best accuracy when combined with the proposed classification techniques and node placements.

In order to derive the features from the stream of samples, a segmentation based on the common sliding window technique [14] is included on each node. The selected time-domain features are then extracted over a set of  $L$  consecutive samples, which form a window, before moving to the next set of samples, which may partially overlap with the previous window. The overlapping between two consecutive windows determines the time that elapses between the production of two features, and hence the responsiveness of the WBSN.

Finally, it is important to point out that no filtering is applied on the sensed data. Even though the samples are affected by high-frequency noise, the hardware-level low-pass filtering inherently applied by the sensor has been shown to be sufficient for classification purposes [15], without explicit data filtering via digital signal processing on the microcontroller.

### D. Classification Algorithms For Activity Monitoring

The classification algorithm is in charge of estimating the current activity, starting from the features transmitted by the

nodes. As mentioned in Section III-A, the complexity and the accuracy of the algorithm is influenced by the execution environment, namely, by the availability of hardware support for floating point arithmetic and by the memory size of the target device.

A large range of classification techniques for activity monitoring has been proposed in the literature [13] [2]. Among them, several techniques show low complexity – including *Naïve Bayes* classifiers–, while more complex approaches employ algorithms with floating point operations, such as *nearest neighbor* and *support vector machines* (SVM). Other approaches based on simple feed-forward artificial neural networks [7] have also shown competitive accuracies on activity monitoring applications. Adaptations of such floating point algorithms for integer arithmetic [9] typically lead to a classification accuracy loss, thus making the porting for wearable devices challenging.

In the proposed WBSN, we advocate two separate classifications techniques, depending on the device that performs the classification. In the context of smartphone-based classification, a *neuro-fuzzy classifier* (NFC) is employed. Thanks to a set of well-established training techniques [9], an NFC can handle the number of classes required for activity monitoring applications with high accuracy, even when subject-independent training is employed (as shown in Section V). In addition, an NFC has a simpler structure with respect to most of the mentioned classifiers, which translates into lower computational complexity and memory requirements, making it suitable for mobile applications.

Similarly to many floating point classifiers, NFCs also show a perceivable loss of accuracy when integer approximations are used [9]. For this reason, in the case of on-node classification, we introduce a simplified algorithm based on a *decision tree*. The decision tree is a set of thresholding rules that discriminate activities based on the value of one feature at time. If the tree is properly balanced, the complexity of this technique is logarithmic in the number of classes, it has limited memory requirements, and it is suitable for an integer microcontroller. In addition, the rules that we derive for activity monitoring are sufficiently general to ensure that the same decision tree can maintain its performance across different subjects.

The implementation details of the proposed classification algorithms will be discussed in depth in the following section.

#### IV. SYSTEM IMPLEMENTATION

In this section, we discuss the most relevant implementation aspects of the proposed WBSN. We start by describing the hardware characteristics of the nodes in the network, and we then detail the two classification techniques. In order to discuss the training of the software algorithms, we also illustrate the WBSN that has been used for data collection.

##### A. Node Specifications

The proposed WBSN is implemented using two different kinds of custom nodes: one targeting the chest, and one targeting the limbs. All the nodes in the network embed a *Jenic*

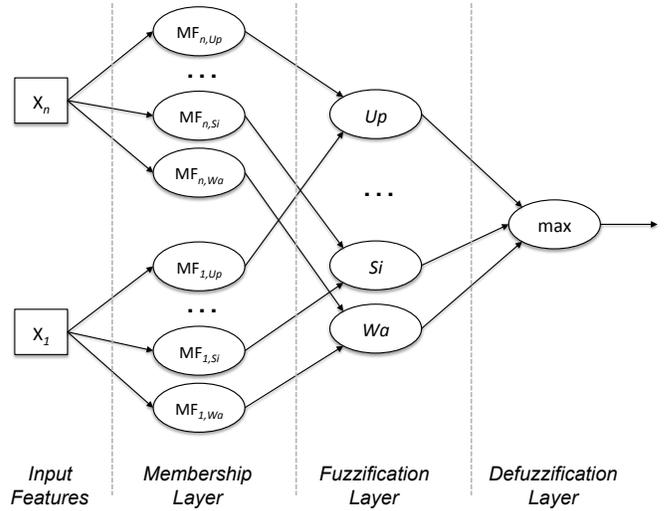


Fig. 3. Structure of the proposed neuro-fuzzy classifier

*JN5148* microcontroller [16] specifically designed for *ZigBee* applications, and comprising the IEEE 802.14.4 transceiver [17], 128 kB of ROM, and 128 kB of RAM. The node on the chest additionally includes an *LSM330DLC* inertial module by STMicroelectronics [18], which features a triaxial accelerometer and gyroscope. The nodes targeting the limbs, on the other hand, include an *LIS3LV02DQ* inertial sensor by STMicroelectronics [19], which only provides a triaxial accelerometer. Finally, the mobile phone is a Sony Xperia U running Android 4.1.2, which interfaces towards the *ZigBee* network by means of a custom dongle, which also embeds a *Jenic* chip and a serial converter by FTDI.

The accelerometer and gyroscope sensors are configured to sample at a frequency of 50 Hz for any activity being monitored, which is a major improvement with respect to the state of the art [4]. In order to perform segmentation with the sliding window technique, the nodes use part of the available RAM to store windows of 4 s, with an overlap of 2 s. The window length has been selected to store a large number of samples, thus avoiding outliers and inconsistencies, as the target feature is expected to have a smooth behavior over such a long time interval. The overlap determines the responsiveness of the network: in the proposed implementation, a new estimation of the activity is produced every 2 s, which fits the slow dynamics of human movements.

##### B. Smartphone-Based Classification

In the proposed smartphone-based classification, the handheld device is in charge of receiving the features from all the wearable nodes (24 features in total), and assign the current activity to one of the 7 classes: walking (*Wa*), sitting (*Si*), standing (*St*), laying (*La*), running (*Ru*), and walking upstairs (*Up*) and downstairs (*Dw*). The decision is taken using a neuro-fuzzy classifier, as discussed in Section III-D, which follows the feed-forward structure shown in Figure 3.

The NFC structure is composed of three layers. The first *membership layer* takes as input the features extracted by the

nodes, and for each feature computes a membership grade for each of the activities. In particular, the grade  $\gamma_{i,k}$  for the  $i$ th feature and class  $k$  is defined according to a membership function  $MF_{i,k}$ , i.e., a gaussian distribution with a mean value  $\mu_{i,k}$  and a standard deviation  $\sigma_{i,k}$ . These statistical parameters are determined during the training phase, and at run-time these values are employed to compute the membership grade  $\gamma_{i,k}$

$$\gamma_{i,k}(x_i) = MF_{i,k}(x_i) = \exp\left(-\frac{(x_i - \mu_{i,k})^2}{2\sigma_{i,k}^2}\right), \quad (1)$$

where  $x_i$  denotes the value of the  $i$ th feature. In the subsequent *fuzzification layer*, the membership grades of the same class are multiplied as  $\prod_i \gamma_{i,k}$ . The resulting value represents how likely the activity belongs to class  $k$ . In the third *defuzzification layer*, the maximum among the different values of  $f_k$  is selected, and the activity is classified accordingly.

In order to perform the offline training of the NFC, a traditional scaled conjugate gradient method has been employed [20]. The training has been performed using data collected from three different subjects. This strategy ensures a high accuracy without requiring a personalized training phase, and makes the WBSN robust when operating in conditions that differ from the training setup.

### C. On-Node Classification

In the proposed on-node classification configuration, a simplified decision tree is used to classify the current activity on the microcontroller of the node located on the chest. The node then sends a notification to the mobile phone only when the activity changes, thus greatly reducing data transmission.

The proposed decision tree is based on an analysis of the parameters that are expected to discriminate different activities, and its structure is summarized in Figure 4. For example, the variance of the accelerometer data coming from the  $x$ -axis on the crus can discriminate static activities (sitting, standing and laying) from dynamic ones: a high value of the standard deviation indicates a movement on that axis, thus indicating a dynamic activity. It can be observed that, excluding the features that are directly sensed on the chest, the nodes on the legs are only required to transmit three features (the standard deviation of the  $x$ -axis of the crus, the mean of the  $z$ -axis of the crus, and the mean of the  $x$ -axis of the thigh), thus greatly reducing the transmission workload.

The thresholds have been determined by analyzing the data of three different subjects. Because of the general nature of the rules included in the tree, even the on-node classification proves to be suitable for cross-subject usage.

### D. Data Collection System

The offline training of the classification algorithms described above requires the collection of a large amount of data from different subjects, while they perform multiple activities. This step is also needed to evaluate the goodness of the selected feature set, and the most suitable node placement on the subject's body. In order to perform these analyses, data

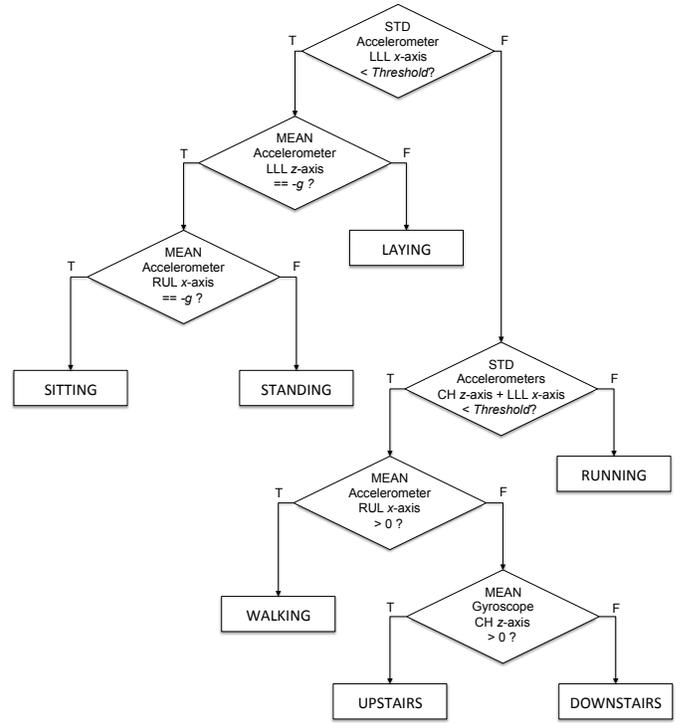


Fig. 4. Structure of the proposed tree for on-node classification

needs to be collected with no on-node manipulation, and many different node positions have to be evaluated at the same time.

A different WBSN has been designed to perform data gathering for offline analysis. The data collection network, which is more invasive than the final system, includes 9 wearable nodes: two for each arm (upper arm and forearm), two for each leg (thigh and crus), and one on the chest, as shown in Figure 1. Each node collects the data from the available sensors – i.e., accelerometers on the limbs, accelerometer and gyroscope on the chest–, and sends it to an external network coordinator connected to a workstation.

The main design challenge in the data collection network is handling the high throughput generated by the nodes. In fact, ZigBee natively supports moderate data rates of up to 250 kbps, but in practice the overall network throughput is heavily reduced by contention, interference and data framing [21]. As a consequence, in order to avoid dropped packets, each node is allowed to buffer the data up to its maximum memory capacity, and then is synchronized by the central coordinator to be the only transmitting node while flushing its buffer contents, thus emulating a collision-free burst transmission. As an additional measure to avoid bottlenecks and data losses, the data gathering has been performed in short sessions (approximately one minute), in which one activity at a time has been monitored.

The data collection phase has been iterated on multiple subjects to gather the samples used for training and cross-validation of the proposed classifiers. The corresponding results are reported in the following section.

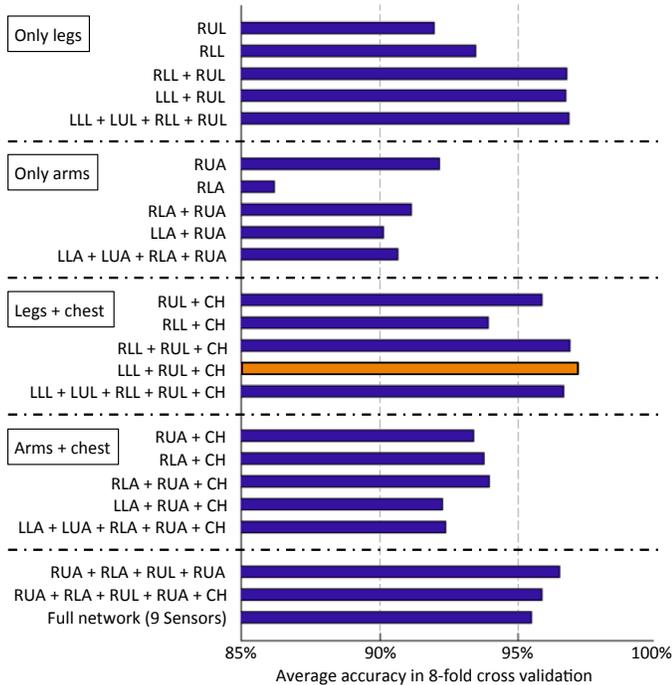


Fig. 5. Classification accuracies of the proposed NFC for smartphone-based classification, with multiple combinations of node placements

## V. EXPERIMENTAL EVALUATION

The proposed WBSN has been implemented and tested on three subjects for training and testing purposes. In this section, we show a quantitative evaluation of the WBSN, in order to show the tradeoffs between accuracy and amount of transmitted data, when executing either a smartphone-based or an on-node classification.

### A. Assessment of Classification Accuracy

We first investigate the accuracy of the proposed classification methodologies, and we motivate the selected node placement. The accuracy results of the proposed neuro-fuzzy classifier, implemented on the smartphone using floating point precision, are reported in Figure 5. The accuracy is determined using an 8-fold cross validation process over the data coming from three subjects, gathered using the collection network discussed in Section IV-D. Thanks to the availability of data from the limbs and the chest, we have performed an exhaustive comparison among the possible node placements. According to the data in Figure 5, the configuration that achieves the best results with a limited number of nodes is the one comprising the chest, and two sensors located on the right thigh and on the left crus. This result reflects the ability of the nodes on the legs to discriminate among static activities, and the ability of the node on the chest to distinguish among dynamic ones. Overall, the proposed WBSN effectively exploits the information coming from the three sensors to achieve an accuracy of 97.2%, thus improving the results of [7] (95%), which is based on a similar classification method and is currently the most accurate work in the literature.

TABLE I

CONFUSION MATRIX OF THE PROPOSED NEURO-FUZZY CLASSIFIER  
(THE VALUES ARE ROW-NORMALIZED)

	Dw	Up	Si	La	St	Wa	Ru
Downstairs	88.7	4.1	1.0	4.1	0.0	2.1	0.0
Upstairs	3.4	95.5	0.0	0.0	0.0	1.1	0.0
Sitting	0.0	0.0	100.0	0.0	0.0	0.0	0.0
Laying	0.0	0.0	0.0	100.0	0.0	0.0	0.0
Standing	0.0	0.0	0.0	0.0	100.0	0.0	0.0
Walking	3.3	0.0	0.0	0.0	0.0	96.7	0.0
Running	0.0	0.0	0.0	0.0	0.0	0.0	100.0

Table I shows the row-normalized confusion matrix for the NFC. As expected, most of the activities are correctly predicted. The main source of error is the distinction between walking upstairs and downstairs: this can be explained by the similar mechanics of the two activities, which are discriminated by the classifier relying only on the orientation of the chest.

On the other hand, the node-based classification (using the decision tree discussed in Section IV-C) achieves lower accuracy, which is caused by the limited computational capabilities. In particular, the proposed classifier achieves an accuracy of 88% with the three selected nodes, a result that is comparable to the existing approaches based on decision trees [2]. Although the on-node classification is 9.2% less accurate than the NFC running on the smartphone, it requires a considerably lower overhead in terms of power consumption and data transmission, as discussed in the following.

### B. Assessment of Transmission Volume Reduction

To conclude this section, we present a quantitative analysis of the bandwidth required by the on-node and the smartphone-based classification strategies. For the sake of illustration, the bandwidth estimation does not account for the packet overhead introduced by the underlying transmission protocol. In both configurations the traditional feature extraction provides a massive reduction of data transmission with respect to raw data streaming, which can be estimated to be approximately 98.4% with the selected sampling frequency (50Hz). On top of this, significant savings can be achieved by employing on-node classification, and by accepting its lower accuracy.

On the one hand, in the proposed smartphone-based classification, each of the nodes sends a new set of features to the smartphone with a rate of  $\rho$  transmissions per second (in the proposed WBSN,  $\rho = 1/2$ , i.e., a new set every 2 s). Each of the features is encoded with 2 Bytes. Hence, a total of 48 Bytes is required to encode all the 24 transmitted features, i.e., the mean and the variance of the three axes of all the accelerometers, plus the mean and variance of three axes of the gyroscope on the chest. As a consequence, the required bandwidth for the proposed network is  $48 \text{ Bytes} \cdot \rho = 24 \text{ Bytes/s}$ .

In the proposed on-node configuration, on the other hand, only the two sensors on the legs are sending features to the node on the chest, and only three of these features are required

(see Figure 4): the standard deviation of the  $x$ -axis of the crus, the mean of the  $z$ -axis of the crus, and the mean of the  $x$ -axis of the thigh, totaling to 6 Bytes. The node on the chest is in charge of transmitting the following data to the smartphone:

- One byte to communicate a new activity, every time a change is detected. This contribution to the overall bandwidth can be further divided into two parts: the first one is due to the actual number of times the subject changes activity during the observation period (let us denote this number by  $\Delta$ ). The second contribution is due to the situations when the classifier incorrectly predicts the current activity: in the worst case, this requires a transmission to communicate the erroneous activity, and a second one when the prediction is corrected. Let us denote the accuracy of the on-node classifier by  $\alpha_{node}$  (in the proposed network, it is equal to 88%). The traffic generated due to erroneous classification can thus be estimated as  $2(1 - \alpha_{node})\rho$  Bytes/s;
- A periodic probe message (1 Byte) to check whether the connection is alive. We denote this overhead traffic by  $\Omega$  (measured in Bytes/s).

Overall, the number of packets transmitted in the on-node classification scenario is equal to:

$$B_{on-node} = 6 \cdot \rho + \Delta + 2 \cdot (1 - \alpha_{node}) \cdot \rho + \Omega \text{ [Bytes/s]} \quad (2)$$

Assuming that the actual activity changes ( $\Delta$ ) are in the order of ten times per hour (which is realistic, considering that sleep periods are also included in the average), and that the overhead ( $\Omega$ ) is in the order of 1 packet per minute, the contribution of the mispredicted activities dominates the traffic between the node on the chest and the smartphone. Still, according to this worst-case estimation, the global network traffic is equal to approximately 3.14 Bytes/s, i.e., 86.9% savings with respect to the more accurate smartphone-based classification.

## VI. CONCLUSIONS

In this work, we have proposed an activity monitoring system based on a Wireless Body Sensor Network, targeting minimal discomfort for the subject, high accuracy and low overhead in terms of data transmission. The final configuration includes three nodes placed on the chest, right thigh and left crus of the subject, and a ZigBee-enabled smartphone.

Two strategies have been proposed in order to achieve the desired goals. First, we have implemented a highly accurate neuro-fuzzy classifier based on a relatively large set of features extracted from the kinetic sensors, and executed on the smartphone using floating point arithmetic. Secondly, we have explored a more constrained, yet accurate, on-node classification algorithm based on a decision tree, which analyses a minimal subset of features, and communicates only the activity transitions to the mobile device.

Experimental results have shown that, while the first approach provides accurate results outperforming the state of the art and achieving a classification accuracy of 97%, the second method reduces the required network bandwidth, saving more

than 86% of air-time while still obtaining a competitive sensitivity (88%).

## ACKNOWLEDGEMENTS

The authors wish to thank SmartCardia [22] for providing the sensor nodes used in this work.

This work has been partially funded by the *BodyPowered-SenSE* (no. 20NA21 143069) and *ObeSense* (no. 20NA21 143081) RTD projects, evaluated both by the Swiss NSF and funded by Nano-Tera.ch with Swiss Confederation financing.

## REFERENCES

- [1] J. Lewandowski et al., "Logic-centred architecture for ubiquitous health monitoring," *Journal of Biomedical and Health Informatics*, vol. PP, no. 99, pp. 1–1, 2014.
- [2] A. Avci et al., "Activity recognition using inertial sensing for healthcare, wellbeing and sports applications: A survey," in *Proc. of Int. Conf. on Architecture of Computing Systems (ARCS)*, 2010, pp. 1–10.
- [3] J. Lester et al., "A practical approach to recognizing physical activities," in *Proc. of Int. Conf. on Pervasive Computing*, 2006, pp. 1–16.
- [4] F. Casamassima et al., "Context aware power management for motion-sensing body area network nodes," in *Proc. of Design, Automation and Test in Europe (DATE)*, 2014, pp. 1–6.
- [5] D. Karantonis et al., "Implementation of a real-time human movement classifier using a triaxial accelerometer for ambulatory monitoring," *IEEE Trans. on Information Technology in Biomedicine*, vol. 10, no. 1, pp. 156–167, 2006.
- [6] I. Beretta et al., "Design exploration of energy-performance trade-offs for wireless sensor networks," in *Proc. of Design Automation Conference (DAC)*, 2012, pp. 1043–1048.
- [7] J.-Y. Yang et al., "Using acceleration measurements for activity recognition: An effective learning algorithm for constructing neural classifiers," *Pattern Recognition Letters*, vol. 29, no. 16, pp. 2213–2220, 2008.
- [8] C. Lombriser et al., "On-body activity recognition in a dynamic sensor network," in *Proc. of Int. Conf. on Body Area Networks (BodyNets)*, 2007, pp. 17:1–17:6.
- [9] R. Braojos et al., "A methodology for embedded classification of heartbeats using random projections," in *Proc. of Design, Automation and Test in Europe (DATE)*, 2013, pp. 899–904.
- [10] F. Rincon et al., "Multi-lead wavelet-based ECG delineation on a wearable embedded sensor platform," in *Computers in Cardiology*, 2009, pp. 289–292.
- [11] L. Bao and S. S. Intille, "Activity recognition from user-annotated acceleration data," Springer, 2004, pp. 1–17.
- [12] G. Bieber and C. Peter, "Using physical activity for user behavior analysis," in *Proc. of Int. Conf. on Pervasive Technologies Related to Assistive Environments (PETRA)*, 2008, pp. 94:1–94:6.
- [13] S. Preece et al., "A comparison of feature extraction methods for the classification of dynamic activities from accelerometer data," *IEEE Trans. on Biomedical Engineering*, vol. 56, no. 3, pp. 871–879, 2009.
- [14] E. Keogh et al., "An online algorithm for segmenting time series," in *Proc. of Int. Conf. on Data Mining (ICDM)*, 2001, pp. 289–296.
- [15] N. Ravi et al., "Activity recognition from accelerometer data," in *Proc. of Int. Conf. on Innovative Applications of Artificial Intelligence (IAAI)*, 2005, pp. 1541–1546.
- [16] NXP, *JN5148-001 IEEE802.15.4 Wireless Microcontroller (JN-DS-JN5148-001 1v9) [Online]*, 2013.
- [17] IEEE, "IEEE Std 802.15.4-2006 (revision of IEEE Std 802.15.4-2003)," 2006.
- [18] STMicroelectronics, *LSM330DLC iNEMO inertial module: 3D accelerometer and 3D gyroscope [Online]*, 2012.
- [19] STMicroelectronics, *LIS3LV02DQ MEMS Inertial Sensor [Online]*, 2005.
- [20] M. F. Moller, "A scaled conjugate gradient algorithm for fast supervised learning," *Neural Networks*, vol. 6, no. 4, pp. 525 – 533, 1993.
- [21] T. R. Burchfield et al., "Maximizing throughput in ZigBee wireless networks through analysis, simulations and implementations," in *Proc. of Int. Workshop on Localized Algorithm Protocols WSNs*, 2007, pp. 15–29.
- [22] <http://www.smartcardia.com>