A Hardware-Assisted Energy-Efficient Processing Model for Activity Recognition using Wearables

HASSAN GHASEMZADEDH, Washington State University
RAMIN FALLAHZADEH, Washington State University
ROOZBEH JAFARI, Texas A&M University

Wearables are being widely utilized in health and wellness applications, primarily due to the recent advances in the sensor and wireless communication, which enhance the promise of wearable systems in providing continuous and real-time monitoring and interventions. Wearables are generally composed of hardware/software components for collection, processing, and communication of physiological data. Practical implementation of wearable monitoring in real-life applications is currently limited due to notable obstacles. The wearability and form factor are dominated by the amount of energy needed for sensing, processing and communication. In this paper, we propose an ultra low-power granular decision making architecture, also called screening classifier, which can be viewed as a tiered wake up circuitry, consuming three orders of magnitude less power than the state-of-the-art low-power microcontrollers. This processing model operates based on computationally simple template matching modules, based on coarse to fine grained analysis of the signals with on-demand and gradually increasing of the processing power consumption. Initial template matching rejects signals that are clearly not of interest from the signal processing chain keeping the rest of processing blocks idle. If the signal is likely of interest, the sensitivity and the power of the template matching modules are gradually increased and ultimately the main processing unit is activated. We pose optimization techniques to efficiently split a full template into smaller bins, called mini-templates, and activate only a subset of bins during each classification decision. Our experimental results on real data show that this signal screening model reduces power consumption of the processing architecture by a factor of 70% while the sensitivity of detection remains at least 80%.

Categories and Subject Descriptors: C.3 [**Computer Systems Organization**]: Special Purpose and Application-Based Systems—*Real-time and embedded systems*; J.3 [**Computer Applications**]: Life and Medical Science—*Health*; H.1.2 [**Information Systems**]: Models and Principles—*User/Machine Systems Human information processing; Human factors.*

General Terms: Design, Algorithms, Performance

Additional Key Words and Phrases: Medical Embedded Systems, Body Sensor Networks, Signal Processing, Wearable Monitoring Systems, Power Optimization.

ACM Reference Format:
Hassan Ghasemzadeh, Ramin Fallahzadeh, and Roozbeh Jafari, 2015. A Hardware-Assisted Energy-

---

This article extends an earlier version that appeared at the ACM/IEEE Second International Conference on Cyber-Physical Systems (ICCPS 2011).
This work was supported in part by the National Science Foundation, under grants CNS-1150079 and ECCS-1509063, and the TerraSwarm, one of six centers of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding organizations.
Author's addresses: H. Ghasemzadeh and R. Fallahzadeh, School of Electrical Engineering and Computer Science, Washington State University, Pullman, WA 99164-2752; email: {hassan, rfallahz}@eecs.wsu.edu; R. Jafari, Center for Remote Health Care Technologies and Systems, Departments of Biomedical Engineering, Computer Science and Engineering, and Electrical and Computer Engineering, Texas A&M University, College Station, TX 77843-3120; email: rjafari@tamu.edu.

DOI: 0000001.0000001

Efficient Processing Model for Activity Recognition using Wearables. *ACM Trans. Des. Autom. Electron. Syst.* 01, 01, Article 01 (January 2016), 28 pages.
DOI: 0000001.0000001

1. INTRODUCTION

Mobile wearable computers, that form Body Area Networks (BANs), bring to fruition many opportunities to continuously monitor individuals with sensors placed on the body, or implanted in the body. These platforms promise to revolutionize many application domains including healthcare and wellness monitoring. Examples of such applications include rehabilitation [Tamura et al. 2013], sports medicine [Jones 2007], geriatric care [Tamrat et al. 2012], gait analysis [Salarian et al. 2013; Ma et al. 2015], physical activity monitoring [Saeedi et al. 2014a; Alinia et al. 2015; Saeedi et al. 2014b], diagnosis of obesity and depression [Sigmund et al. 2014], and detection of neurodegenerative disorders such as Alzheimer's [Abbate et al. 2014], Parkinson's [Cancela et al. 2014], Huntington's [Emmerik and Wegen 2002] diseases, nutrition monitoring [Hezarjaribi et al. 2016], and monitoring and treatment of chronic diseases [Maciuca et al. 2013; Fallahzadeh et al. 2015]. In the past few years, new wearable applications have evolved and proved to be effective. Yet, one of the major obstacles is the size and weight of the sensor units. Smaller wearable units can enhance comfort and compliance. Smaller implantable units can enable many new applications. Battery size has been the dominating factor in the size of the sensors. Battery-less units operating on piezo, or units that require significantly smaller batteries, are not currently possible. The proposed technique in this paper aims to significantly reduce the power consumption of wearable units, and specifically the processing architecture.

Wearable computers, which form the core components of a BAN, are often composed of several sensors, a processing unit (e.g., a microcontroller), a communication module and a battery. Our current focus is on wearable motion sensors that are used for detection of human actions such as 'Sit to Stand' or 'Lie to Sit'. We propose an architecture, equipped with a *granular decision making module* (GDMM) which monitors incoming signals/actions. The granular decision making module attempts to detect actions that are not of interest as early as possible while consuming the least amount of energy. If the incoming action is likely of interest, the module will turn on the main signal processing unit (e.g., the microcontroller) for further processing. The granular decision making is constructed in a sequence of coarse to fine grained operations. At the beginning, the *screening* or preliminary signal processing may not exhibit high accuracy for classifying the incoming actions, but operates at an ultra low-power. The objective of the initial screening is to identify *incoming actions* that are '*obvious rejects*' or 'accepts'. As the module begins to observe the incoming actions that are *likely* of interest, more accurate decision making and screening processes are activated. Intuitively, screening at the beginning is done by a classification module with tunable parameters adjusted to consume the least amount of energy (e.g., by observing fewer samples with lower bit resolution). The tunable parameters are adjusted to enhance the accuracy of signal processing and classification as the incoming signal or incoming action travels through the screening blocks in the GDMM. The tunable parameters include time duration of actions, number and location of samples within each action, and bit resolution of sampled data. Collectively, screening blocks can select any combination of these transformations to adjust processing (or power) vs. accuracy. The decision making is performed in this fashion because often the incoming action is so dissimilar to the *action of interest* (also called target action) that it can be rejected even with a coarse-grained analysis of the signal. For incoming actions that the screening block cannot reject with high confidence levels, the main signal processing unit will be activated. The main advantage of this method is the power saving due to removing actions that

are not of interest from the signal processing chain as early as possible, deactivating the remaining modules in the signal processing chain.

Applications of wearable healthcare monitoring have unique properties motivating our proposed technique: *Events of interest* often occur with a low duty cycle (e.g., $< 1\% - 5\%$) and the randomness to the incoming signals, even in cases where the signals are not of interest, is not significant. This assumption holds for many wearable applications where the objective is to detect sparse events such as walking using motion sensors, [Dobkin 2013; Hagler et al. 2010; Guenterberg et al. 2009b], cardiac arrest [Dhulipala and Kanagachidambaresan 2014; Cong et al. 2009a] and seizures [Shahrokhi et al. 2010; Liew et al. 2009] using implantable sensors. We utilize these unique properties of the applications in order to reduce the power consumption of the wearable unit by orders of magnitude in the signal processing chain. Although in our approach every effort will be taken to ensure that granular decision making module provides acceptable precision in signal processing, in the events where it generates *false positives*, the sole cost would be the energy consumed to wake up the main signal processing unit for improved precision. Finally, the events are captured with a low sampling rate (e.g., $\leq 100Hz - 1kHz$) which implies that the processing can also be done at a slow speed.

2. RELATED WORK

Several power-aware wearables consuming hundreds of μW have been presented in the literature [Varel et al. 2014; Liew et al. 2009; Sasaki et al. 2006]. An implantable wireless monitoring device for cataract surgery is proposed in [Varel et al. 2014]. It is composed of an on-chip micro mechanical pressure sensor array, a temperature sensor, a microcontroller-based digital control unit, and an RF. The low-power interface system for implantable neural recording, presented in [Liew et al. 2009], uses power distribution techniques. The study in [Sasaki et al. 2006] presents an accelerometer with a 3D loop antenna using radio waves for power feeding. However, a microcontroller is used as control unit of the battery-free accelerometer. These systems, are either not programmable, or a microcontroller is entirely in charge of the programmability. In contrast, our signal screening module is reprogrammable and can be configured to identify various movements and templates.

Several other application-specific sensing devices have been proposed in recent years. Examples include a patient-specific seizure detector in [Shoeb et al. 2009], an implantable battery-less blood pressure monitor [Cong et al. 2009a; 2009b], an ECG sensing micro-system with adaptive RF powering [Chaimanonart and Young 2009], an implantable blood pressure monitor for in vivo measurement [Cleven et al. 2014], an implantable battery-free application specific system with telemetric and EMG recording circuitry [Parramon et al. 1997], a battery-less wireless biosensor for cardiovascular applications [Najafi and Ludomirsky 2004], an ultra low-power sensing device for measuring pulse oximetry [Huiku and Virtanen 2013], and a battery-less sensing device with radio frequency energy-harvesting [Mandal et al. 2010; Zhang et al. 2013]. All these studies, however, focus solely on power consumption of the sensing module and still require a microcontroller for processing of the data. Our proposed processing model takes into account unique properties of physical movement monitoring and BAN applications to significantly reduce the power consumption of the processing architecture.

There have been efforts towards creating ultra low-power semiconductor components and devices. Multi-threshold CMOS (MTCMOS) circuits are an example [Di and Smith 2014]. A wireless system with MTCMOS/SOI circuit technology is suggested which lowers the supply voltage of the LSIs 0.5 V and reduces the power dissipation to 1 mW [Douseki et al. 2003]. The power rating of 1 mW, however, is still larger than a

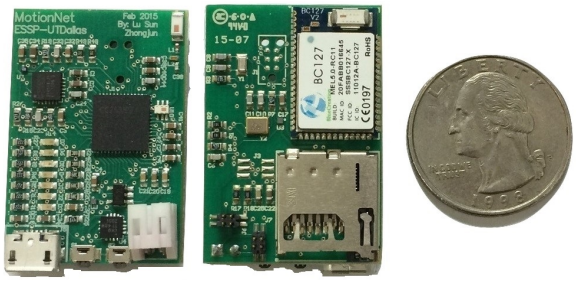

Fig. 1: The custom-designed wearable sensor node used in our data collection

typical power budget that could be supplied by energy harvesting circuits. The power budget of energy harvesting circuits is often in the range of tens of μWs. For example, a battery-less vibration-based energy harvesting system was proposed for ultra low-power ubiquitous applications which can generated 36.79 μW [Chao et al. 2007]. Although these techniques focus on circuit and transistor-level low-power design, our approach to devising ultra low-power BANs is a system-level design and optimization aiming at substantially reducing power consumption of the processing unit by keeping the main processing unit in the loop only when the event of interest is detected but these events do not happen often; hence, the amount of energy saving would be significant.

3. SENSING AND PROCESSING ARCHITECTURES

The main units of a typical BAN system are explained in the following section. We are particularly interested in movement monitoring applications that use inertial information to examine human motions for the purpose of patient monitoring, diagnosis and treatment. However, the proposed methodology may be applicable to other monitoring domains (e.g., a pacemaker that is required to detect abnormal ECG signals) because such applications also intend to detect sparse events.

3.1. Sensor Nodes

A BAN is composed of several sensor nodes attached on the patient's body, embedded with the clothing, or implanted in the human body. Motion sensor nodes are typically used in activity monitoring applications. An example of motion sensor nodes with embedded accelerometer and gyroscope is shown in Fig. 1, which is also used for our experimental data collection and validation of the proposed algorithms. Each node includes a microcontroller (i.e., TI MSP430) for signal processing, and a custom-designed sensor board including a 3-axis accelerometer and a 2-axis gyroscope for inertial data collection. The sensor node has also a radio module for communication with other sensor nodes in the network or with a gateway such as a cell phone.

3.2. Per-node Signal Processing

Each sensor node has a microcontroller which can sample motion sensors at a certain rate. The acquired signals need to undergo specific embedded signal processing tasks in order to make higher level interpretations of human movements. The goal of main signal processing chain (MSPC) is to extract useful information from sensor data. Frequently, this data is a high-level observation, such as "Is the subject running?" or "What is the stride length when the subject is walking?". In other words, the purpose of main signal processing is to provide a 'fully' software programmable

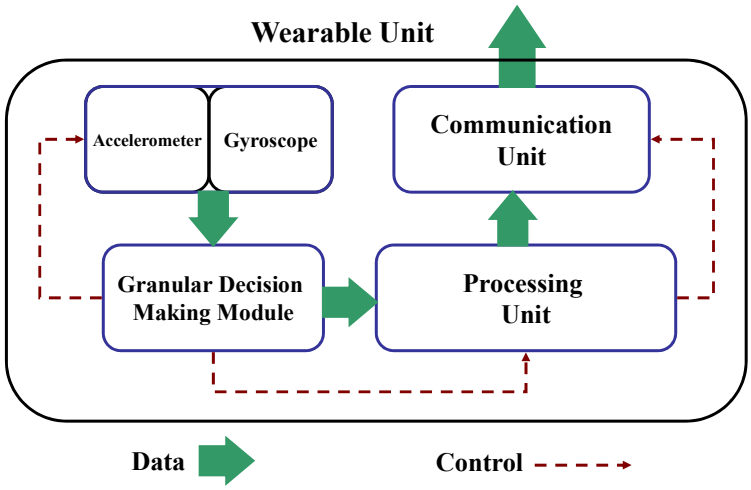

Fig. 2: Overall system architecture illustrating granular decision making module (GDMM) in connection with communication unit, sensor units, and other computing modules.

environment for development of 'highly' reliable signal processing technique for action detection/verification and extracting details from the signals (e.g., balance during 'sit to stand' when it occurs). Typically, signal processing tasks are imposed by the application of the BAN. However, a basic requirement of movement monitoring applications is to detect actions first, and perform additional processing next. This application is usually referred to as *action recognition* [Ghasemzadeh et al. 2009; Ghasemzadeh et al. 2010]. A typical signal processing chain for action recognition includes filtering, segmentation, feature extraction, feature conditioning, and classification. The sampled data are first filtered to improve signal to noise ratio. A segmentation module [Lin and Davis 2010; Hamm et al. 2013; Guenterberg et al. 2009a] then separates portions of the signal that correspond to activities from those associated with rest (non-activity). The set of statistical features extracted from individual segments is reduced in size using feature conditioning techniques to speed up and enhance the classification task. At the end of the main signal processing chain (MSPC), a classifier (e.g., k-Nearest Neighbor [Amato and Falchi 2013]) is utilized to identify the action performed by the subject.

4. ULTRA LOW-POWER MODEL

In many BAN applications, only a very small set of human actions is of interest. For example, gait analysis only is concerned with walking, fall detection with falls, Parkinson's disease monitoring with certain movements such as tremors, and sleep apnea with restless leg movements [Lanza et al. 2014]. These target actions may occur infrequently. Considerable energy is wasted processing non-target actions. Efficiently rejecting non-target actions with a screening classifier could lead to a significant increase in system lifetime.

4.1. Granular Decision Making

An overall architecture of the proposed screening approach is illustrated in Fig. 2. The granular decision making module (GDMM), which is composed of several coarse to fine grained screening classifiers, is responsible for screening sensor readings and

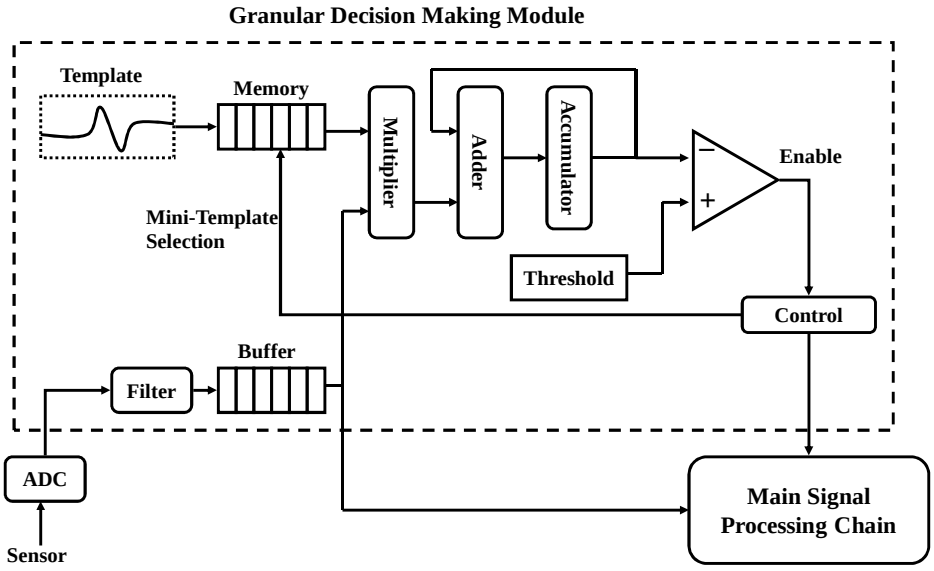

Fig. 3: Block diagram of the granular decision making module (GDMM) and the main signal processing chain (MSPC).

activating main processing unit upon arrival of an event (e.g., action/movement) of interest.

Fig. 3 shows the block diagram of the granular decision making module (GDMM) and main signal processing chain (MSPC). The main component in the diagram is the granular decision making module, communicating with main signal processing block. The main signal processing is implemented on the main processor (e.g., a microcontroller). The granular decision making module is an ultra low power screening classifier aiming to reject actions that are not of interest. This functionality is created by a multiplier-accumulator structure that implements a template matching function.

4.2. Template Matching

The screening classifier and the main signal processing form a rejecting chain of two classifiers. While the main signal processing uses classical pattern recognition techniques to classify actions, the screening classifier employs simple template matching techniques to estimate the likelihood of occurrence of a target action. An unknown action is processed by the template matching block first. If the template matching block does not reject the action, it is evaluated using the main signal processing block (i.e., the microcontroller). A template matching block functions as a binary classifier based on the cross correlation [Chen et al. 2013].Cross correlation takes the advantage of programmability and low computational complexity. The incoming signal is compared to a predefined template of the target action. The comparison assigns a score value representing similarity between the current action and the template (target-action). The cross correlation score is then compared against a threshold and the action is either accepted or rejected. Only in case of acceptance the main signal processing is activated. The cross correlation measure was chosen because it can be implemented in hardware by a series of multiplications and additions. The size of the data storage units including memory and buffer depends on the size of the mini-templates.

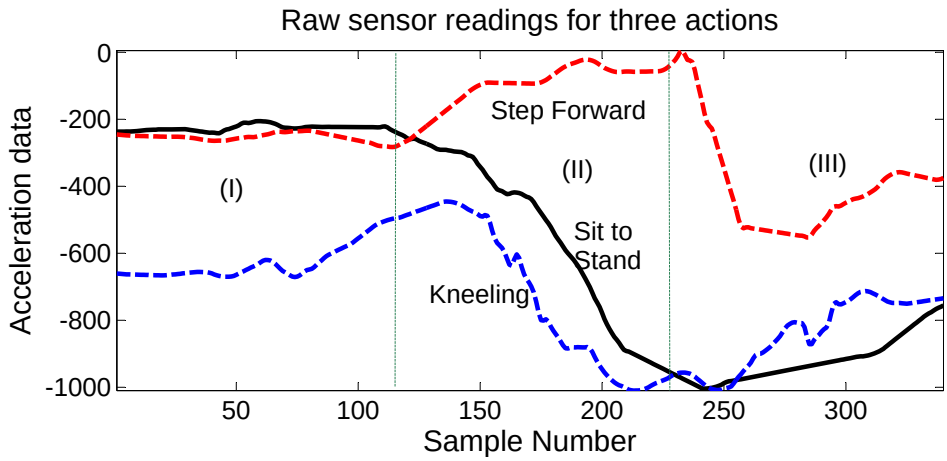

Fig. 4: An example of three templates each divided into three mini-templates.

4.3. Optimization Strategy

The template matching block described previously can be optimized for further energy saving by adjusting several tuning parameters. These parameters include time duration of actions considered for cross correlation calculation, number and location of samples [Ghasemzadeh and Jafari 2011b], and bit resolution of the sampled data [Ghasemzadeh and Jafari 2013]. This allows us to use a sequence of template matching blocks each contributing to the classification of events only to certain level while providing the benefit of low power consumption. The focus of this paper is on minimizing the number of samples used for calculation of the cross correlation function. Motivation behind this optimization is that even with a fixed bit resolution and action duration, only small portions of the template need to be considered when measuring similarity of an input signal with the template, hence offering the opportunity to further save on the computations and energy consumption. We address this optimization problem by dividing a full template into several bins, each forming a *mini-template*. Mini-template approach will further reduce power consumption of the system allowing for realization of significantly less power-hungry wearable units that can eventually enable battery-less technologies for monitoring platforms. Furthermore, mini-templates highlight prominent patterns in the signal and eliminate irrelevant portions of the signal, and therefore, improve performance of signal processing and sensitivity of the classification system.

4.4. Motivational Example

Fig. 4 illustrates motivation behind using mini-templates. This figure shows real data collected with our wearable sensors where only three actions are used for visualization. The graphs show raw sensor readings from Z-axis accelerometer of a node placed on the 'Waist' of the subject. Assume 'Sit to Stand' (bold black plot) is the action of interest and the other two actions, 'Kneeling' (dashed blue plot) and 'Step Forward' (dashed red plot) may occur as non-target. Clearly, if the entire template is considered, the target action can be distinguished from the two others based on the cross correlation measure. Assume each template is divided into three bins as indicated by (I), (II), and (III). None of the bins can solely achieve small false positive rates. For instance, if only bin (I) is used, a 'Step Forward' action may be classified as 'Sit to Stand' lead-

ing to high misclassification rate. Similarly, in bins (II) and (III) 'Kneeling' is strongly correlated with the target action ('Sit to Stand') resulting in a large number of false positives. However, assume the case where only bins (I) and (II) are activated for template matching. The only action that can be accepted by both bins is 'Sit to Stand'. If an action is accepted by bin (I), it can be confidently considered as either 'Sit to Stand' or 'Step Forward'. If the action is further accepted by bin (II), the choice of 'Step Forward' is ignored leaving 'Sit to Stand' as the final classification decision. By activating only two bins rather than the entire template, one-third of the multiply-add operations are discarded from the template matching resulting in 33% savings. Therefore, our objective is to find a minimum subset of template bins that can confidently activate the main signal processing block while maintaining low false positive rates. We note that the ordering of processing mini-templates is also important because a suboptimal ordering can result in a larger number of bins being processed.

5. MINI-TEMPLATE OPTIMIZATION

As discussed in the previous section, the template matching block in Fig. 3 can be divided into several lower cost blocks associated with a set of predefined mini-templates. We pose an optimization problem to find the minimum subset of template bins and their ordering required for detection of a target action subject to a given lower bound on sensitivity rates, called true positive rates, of the screening block. We call this optimization problem *Minimum Size Mini-Template Set* (MSMTS). Throughout this section, we use the notations in Table I to formulate this problem.

5.1. Overview of Decision Path Construction

We overview different steps needed to find an optimal decision path using the proposed GDMM. In describing these steps, we assume that a set of training examples associated with human actions are given and one action is specified as target. Furthermore, the system is aimed to minimize the power consumption such that a minimum desirable sensitivity to the classification of the target action is guaranteed. Moreover, a set of screening blocks associated with different template bins and their corresponding power consumption are given.

Step 1: Using a set of training instances or the target action and non-target actions, determine the sensor that best distinguishes between target action and non-target actions. The process of sensor selection is described in Section 5.2.

Step 2: Generate a template associated with the target action using the training examples. The template is generated based on the approach that will be described in Section 5.3.

Step 3: Set the threshold for each screening block such that the block meets the minimum desirable sensitivity. In order to maintain this requirement, the threshold for a given screen block is set according to discussion in Section 5.4. Intuitively, the threshold is set to the largest value that satisfies the sensitivity requirement. This is mainly due to the fact that larger values of the threshold result in lower sensitivity rates. Thus, we start with a small value (e.g., close to '0' which would result in 100% sensitivity) and increase this value as long as the sensitivity rate is still above the desirable value (λ).

Step 4: For each screening block, compute the weakly correlated set WCS_k based on Definition 5.3.

Step 5: Construct a decision path. The process of decision path construction will be defined by Problem 1 and will be suggested by Algorithm 1.

Table I: Notations

Term	Description
$\hat{a}$	target action
A	set of n not-target actions
a_i	i-th non-target action
a_i^l	l-th training trial of action a_i
T_i	template generated for action a_i
K	number of template bins
B	set of template bins due to template partitioning
b_k	k-th template bin due to template partitioning
MT_{ik}	k-th mini-template of a_i associated with b_k
$\gamma(T_i, T_j)$	similarity score between templates T_i and T_j
O	optimal subset of bins used for classification
R	size of optimal set O found by MSMTS problem
λ	lower bound on sensitivity of the screening blocks
tp_k	sensitivity or true positive rate for template bin k
thr_k	threshold value for template bin k

5.2. Prominent Sensor Selection

Given a target action $\hat{a}$ and $A = \{a_1, a_2, \ldots, a_n\}$ a set of n non-target actions, we first collect data with all these actions and from a number of wearable motion sensors. Let $S = \{s_1, \ldots, s_m\}$ denote the set of m sensors. In this article, we hypothesize that we can develop the granular decision making module for detection of the target action by examining only one sensor stream. In Section 6, we will test this hypothesis using real data and will demonstrate that it is possible to use only one sensor axis for detection of an individual target action. Our sensor selection algorithm is based on the similarity score defined below.

Definition 5.1 (*Similarity Score*). Given two time series signals f and g of length N, the similarity score $\gamma(f,g)$ between the two signals is defined based on their normalized cross correlation by

$$\gamma(f,g) = \frac{\sum_{t=1}^{N}[f(t)-\bar{f}][g(t)-\bar{g}]}{\sqrt{\sum_{t=1}^{N}[f(t)-\bar{f}]^2 \sum_{t=1}^{N}[g-\bar{g}]^2}} \tag{1}$$

where $\bar{f}$ and $\bar{g}$ denote mean values of f and g.

Let x_{ij}^l denote the sensor signal associated with the l-th trial/instance of non-target action a_i from sensor s_j. Furthermore, let $\hat{x}_j^l$ represent l-th trial of the target action $\hat{a}$ captured using sensor s_j. Without loss of generality, we assume that there exist L signal trials of each action including target and non-target. For each sensor s_j, we compute μ_j as follows.

$$\mu_j = \sum_{i=1}^{n}\sum_{r=1}^{L}\sum_{l=1}^{L}\gamma(\hat{x}_j^l, x_{ij}^r) \tag{2}$$

The value computed by μ_j represents the total amount of similarity between the training instances of the target action $\hat{a}$ and non-target actions in A. The prominent sensor used for our target action screening is then the sensor whose instances represent minimum similarity between target and non-target actions. Thus, the prominent sensor, $\hat{s}$, is given by

$$\hat{s} = \underset{j}{\operatorname{argmin}}\, \mu_j \tag{3}$$

5.3. Template Generation

Given a target action $\hat{a}$ and $A = \{a_1, a_2, \ldots, a_n\}$ a set of n non-target actions, we generate $\hat{T}$, template for target action, from the set of training trials. Templates are generated as shown in Definition 5.2 according to the similarity score in Definition 5.1.

Definition 5.2 (*Template*). Given a target action $\hat{a}$ with L training trials $\hat{X} = \{\hat{x}_j^1, \ldots, \hat{x}_j^L\}$ obtained from prominent sensor s_j and non-target actions $a_i \in A$ with training trials $X = \{X_{1j}, \ldots, X_{nj}\}$ obtained from prominent sensor s_j such that $X_{ij} = \{x_{ij}^1, \ldots, x_{ij}^l, \ldots, x_{ij}^L\}$ has L training trials, a template $\hat{T}$ for $\hat{a}$ is the best representative trial with respect to the similarity score γ between all pairs of the training trials. The trial that is most similar to the trials from the target action and most dissimilar to the trials from non-target actions is finally chosen as candidate for target action template. Thus, a template $\hat{T}_j$ using prominent sensor s_j is given by Equation 4.

$$\hat{T}_j = \underset{l}{\operatorname{argmax}}\, \theta_j^l \tag{4}$$

where θ_j^l is given by

$$\theta_j^l = \frac{\sum_{r:r\neq l} \gamma(\hat{x}_j^l, \hat{x}_j^r)}{\sum_{i=1}^{n}\sum_{r=1}^{L} \gamma(\hat{x}_j^l, x_{ij}^r)} \tag{5}$$

In Equation 5, the numerator calculates the summation of similarity scores across all training instances that are within the dataset for the target action $\hat{a}$. The denominator, in contrast, represents the overall similarity between the training instance l from target action class $\hat{a}$ and all other non-target actions a_i.

Each template is evenly divided into K bins $B = \{b_1, b_2, \ldots, b_K\}$. Each bin b_k represents a set of mini-templates associated with target action and different non-target actions. We investigate how each one of the bins contributes to detection of a target action and choose the best sequence of template bins to be examined during template matching.

5.4. Minimum Size Mini-Template Set Problem

In this section, we formally define MSMTS problem. Each template bin, b_k, is assigned a threshold value, thr_k, for the cross correlation analysis. This threshold determines how sensitive that bin is to detection of the target action $\hat{a}$. Thus, we set the threshold such that the given sensitivity (i.e., true positive) requirement of the application is met. Intuitively, the threshold is set to the largest value that satisfies the sensitivity requirement. We note that the result of the cross validation is a value ranging between 0 and 1 indicating completely dissimilar and completely similar signals respectively. Therefore, a larger thr_k value results in lower sensitivity rate because such a value requires higher similarity of a performed action to the target action template during system training. Thus, we start with a small threshold value (e.g., close to '0' which would result in 100% sensitivity) and increase this value as long as the sensitivity rate is still above the desirable value (λ). Therefore, the choice of the threshold would satisfy this requirement: $tp_k \geq \lambda$, and is given by

$$thr_k = \left\{ \underset{l}{\operatorname{argmax}}\, \gamma(\hat{MT}_k, \hat{a}_k^l) | tp_k \geq \lambda \right\} \tag{6}$$

where $\hat{MT}_k$ denotes the k-th mini-template associated with target action $\hat{a}$ and $\hat{a}_k^l$ represents the k-th segment of the target action drawn from training trial l. Therefore, the threshold thr_k is computed after the template of the target action is generated, and based on the amount of similarity of the training trials of the target action and the template.

By setting the threshold based on Equation 6, our system aims to maintain a minimum sensitivity rate of λ, to target action, on all screening blocks while attempting to minimize the false positive rates due to accepting non-target actions. We note that the choice of the threshold thr_k will impact the false positive rates due to activating a screening block or including a screening block on the decision path. In fact, lower values of thr_k translate into higher likelihood of classifying non-target actions as target, resulting in high false positive rates. As a result, we set the threshold as high as possible such that the sensitivity requirement is met. Yet, it is possible that some non-target actions look similar to the target from the point of view of some screening blocks. Thus, we need a method for identifying how well each screening block is at rejecting non-target actions. In order to quantify capability of individual screening blocks in rejecting non-target actions, we define a *Weekly Correlated Set* of actions for each bin b_k.

Definition 5.3 (*Weakly Correlated*). Let a_{ik} denote the k-th signal segment for action a_i. Within each bin b_k, the signal segment a_{ik} is referred to as weakly correlated with the target action if $\gamma(\hat{MT}_k, a_{ik}) < thr_k = 1 - \epsilon_k$, where $\hat{MT}_k$ denotes the k-th mini-template associated with target action $\hat{a}$. Similarly, for each bin b_k, a set WCS_k, *Weakly Correlated Set*, is defined as the set of actions a_i whose signal segment a_{ik} is weakly correlated.

Intuitively, an incoming signal that is weakly correlated in b_k will be rejected. The signal, however, will be further processed by subsequent bins if it is accepted by a bin b_k on the decision path. Clearly, in order to accept an event, it needs to be weakly correlated with all non-target actions. Therefore, the weakly correlated set WCS_k associated with a bin b_k indicates how many actions are likely to be rejected by b_k. As soon as WCS_k are computed for all screening blocks (i.e., bins), we need to find an optimal ordering of the blocks such that the number of such blocks is minimized and the collection of the blocks can reject all non-target actions.

In order to formally define our optimization problem, we first define *complete ordering*. Informally, a complete ordering of the template bins is a sequence of the bins, or equivalently screening blocks, which is capable of rejecting all non-target actions in A based on the concept of weakly correlated action defined previously. We then define *ordering cost* for a complete ordering. The ordering cost takes into account the number of bins that need to be examined on the decision path (i.e, bin ordering) in order to reject a non-target action. The overall ordering cost is then the summation of the costs over all non-target actions.

Definition 5.4 (*Complete Ordering*). An ordering O=$\{b_1, b_2, \ldots, b_R\}$ is complete if the following condition holds.

$$\bigcup_{k=1}^{R} WCS_k = A \tag{7}$$

Definition 5.5 (*Ordering Cost*). Let O=$\{b_1, b_2, \ldots, b_R\}$ be a complete ordering of bins and $f(a_i)$ a function that returns the index of the first bin in which the following condition holds:

$$\{a_i \mid a_i \in A\} \quad \subset \quad \bigcup_{k=1}^{f(a_i)} WCS_k \tag{8}$$

That is, $f(a_i)$ is the number of bins, in a sequence from the complete ordering O, that need to be examined to reject a_i. We note that this examination process involves choosing bins from the complete ordering in a sequence, starting with b_1 in the complete ordering set O=$\{b_1, b_2, \ldots, b_R\}$. Then the total cost of the ordering is given by:

$$Z = \sum_{a_i \in A} f(a_i) \tag{9}$$

Problem 1 (MIN SIZE MINI-TEMPLATE SET). *Given a finite set A, of non-target actions, and WCS=$\{WCS_1, WCS_2, \ldots, WCS_K\}$, a collection of subsets of A such that the union of all WCS_i forms A, MSMTS is the problem of finding a complete linear complete ordering, O, such that the cost of the ordering is minimized.*

5.5. Problem Complexity

Through the following theorem, we prove that the MSMTS problem is NP-hard.

Theorem 1. *The Min Size Mini-Template Set problem is NP-hard.*

Proof: 1. *It is straightforward to see that Min Sum Set Cover (MSSC) problem can be reduced to our MSMTS problem. The known MSSC problem is described as follows. Let U be a finite set of elements and S=$\{S_1, S_2, \ldots, S_m\}$ a collection of subsets of U such that their union forms U. A linear ordering of S is a bijection f from S to $\{1, 2, \ldots, m\}$. For each element $e \in U$ and linear ordering f, we define $f(e)$ as the minimum of $f(S)$ over all $\{S_i : e \in S_i\}$. The goal is to find a linear ordering that minimizes $\sum_e f(e)$. It is easy to see that by replacing elements of U with those of A, and also replacing subsets S_i with WCS_i we obtain the same problem as MSSC. Therefore, MSMTS is an NP-hard problem.*

Theorem 2. *There exists no polynomial-time approximation algorithm for MSMTS with an approximation ratio less than* 4.

Proof: 2. *Reducing MSSC problem to MSMTS preserves approximation of any corresponding solutions. Therefore, any lower bound for MSSC also holds for MSMTS. In [Feige and Tetali 2004], it is shown that for every $\varepsilon > 0$, it is NP-hard to approximate MSSC within a ratio of $4 - \varepsilon$. Therefore, 4 is also a lower bound on the approximation ratio of MSMTS.*

5.6. Greedy Solution

The greedy algorithm for MSMTS is adapted from the greedy algorithm for MSSC and is shown in Algorithm 1. At each step, it searches for the bin b_k that can reject largest number of remaining non-target events (by searching through the WCS_k). It then adds such a bin to the solution space O and removes the actions it can reject from further consideration. Such actions are represented by the maximum cardinality set, WCS_k, as shown in Algorithm 1. Because the algorithm will not need to examine such actions anymore, it will remove those actions from all weakly correlated sets WCS_j (j=$\{1,\ldots,K\}$) including the set with maximum cardinality (i.e., WCS_k). The algorithm terminates when all non-target actions are rejected. The approximation ratio is 4 as previously discussed.

Algorithm 1 Greedy solution for MSMTS

Calculate set WCS_k for every bin b_k
$O = \phi$
while $\left(\bigcup_{b_k \in O} WCS_k \neq A\right)$ **do**
  Select bin b_k such that WCS_k is maximum cardinality
  $O = O \bigcup b_k$
  for all $a_i \in WCS_k$ **do**
    remove a_i from all WCS_j (j=$\{1,\ldots,K\}$)
  end for
end while

6. VALIDATION

6.1. Experimental Setting

We carried out a number of experiments to collect real-data using wearable sensors and to demonstrate successfulness of our hardware-assisted proposed processing model in efficiently reducing power consumption of the processing tasks. Our experiments involved three subjects performing 14 actions, each ten times, while wearing a number of wearable sensor nodes with embedded 3-axis accelerometer and 2-axis gyroscope sensors. The wearable node was a custom-designed sensor board attached to a TelosB mote, as described in Section 3.1 and shown in Fig. 1. The experimental actions were as follows: (1) *Stand to Sit*; (2) *Sit to Stand*; (3) *Sit to Lie*; (4) *Lie to Sit*; (5) *Bend and grasp*; (6) *Rise from bending*; (7) *Kneel*; (8) *Rise from kneeling*; (9) *Look back*; (10) *Return from look back*; (11) *Turn clockwise*; (12) *Step forward*; (13) *Step backward*; and (14) *Jump*. Each subject wore seven sensor nodes secured to the upper body, lower body, and waist. The nodes were programmed to sample five sensors including x, y, z accelerometer and x, y gyroscope at 50 Hz. Our prior research [Ghasemzadeh and Jafari 2011a] shows that this sampling rate is sufficient to capture details of human daily living activities. The data were collected using a custom-designed MATLAB tool for further processing. The experiments resulted in capturing over $210,000$ samples of motion sensor data per sensor node, equivalent to a total of $1,470,000$ samples for the entire wearable network. We developed a segmentation tool for labeling the movements in order to gather ground truth labeled training data. This tool allowed us to examine the collected data and label them based on the type of activity performed by the subject. The tool used the video recording of the experiments which allowed us to observe when a movement starts and when it ends. We used 50% of the data for training, selected at random, in the template generation process as well as in finding optimal decision path. The remaining 50% were used as test data to determine the accuracy of the activity recognition and the amount of power savings.

6.2. Calculating Power Numbers

Our objective was to measure energy savings when a particular action is considered as target action (i.e., to be identified & accepted) and the rest of the actions are considered as non-target. For each action, we generated a unique template as described in Section 5.3. As discussed in Section 4.3, the power consumption of each screening block depends on several tunable parameters. In particular, the number of samples used for template matching affects the power consumption significantly.

To estimate power consumption of the template matching approach, the screening blocks were implemented using 12-bit Multiplier-ACcumulator (MAC) units. The MAC units were designed using Verilog. The cross-correlation algorithm was developed by a series of MAC steps depending on the number of incoming samples. At each clock

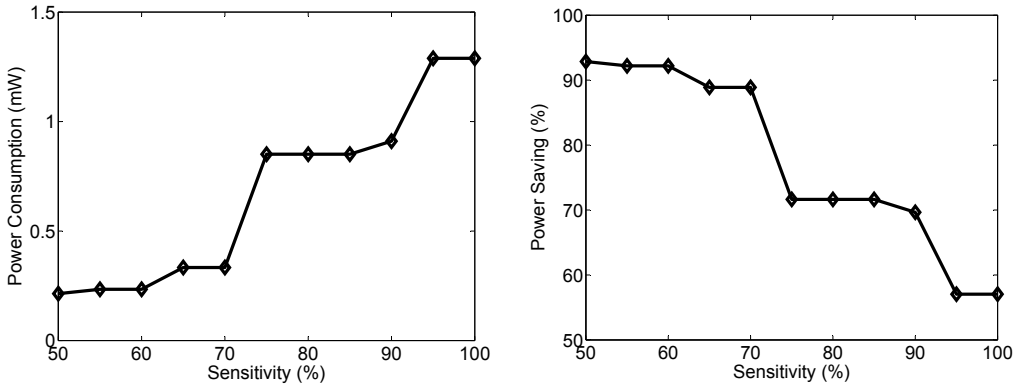

(a) Power Consumption (b) Power Saving

Fig. 5: Power performance of the proposed decision making system while screening 'Sit to Stand' movements. (a) Power consumption of the entire system; (b) System power savings achieve by the low power screening approach.

instant, the digitized template data and the incoming signal data were multiplied and added to the previous MAC value. This continued, depending on the number of samples for each incoming event, until the cross correlation value was computed for the entire action signal.

The design was synthesized using Synopsys with the 45 nm standard Cell library. The simulations of the Verilog RTL were completed using ModelSim PE. In order to extract transistor level netlist, the gate-level Verilog was imported to Cadance using Verilog-in tools. Finally, the power numbers were estimated by simulating net-list in Synopsys Hspice. The source of the power consumption can be categorized into two groups including dynamic power and leakage power. Dynamic power consumption is the amount of power dissipation due to charge and discharge of the load capacitance during operation. The load capacitance passes small current (i.e., leakage current) even when all transistors are in idle mode. There are several reasons for leakage current, including weak inversion, drain-induced barrier lowering, gate-induced drain leakage, and gate oxide tunneling [Roy et al. 2003]. We present dynamic current and leakage current together. Since the idle time increases by applying our mini-template technique, the leakage power, which is one of the most important factors in nano-meter CMOS technologies, increases. The reduction on dynamic power, however, is dominant and finally leads to achieving an overall power saving. In order to calculate dynamic power, we used the notion of average current, which presents the area under the dynamic current divided by the time duration representing the application deadline. Application deadline specifies how quickly the output must be computed. The application deadline used for all the mini-templates in our work is assumed to be $20ms$. This is due to the fact that the sampling sensor sampling frequency is $50Hz$, and it was assumed that the output must be computed before the next sample arrives. Furthermore, the operating voltage V_{DD} used in our synopsis simulation was $1.1V$ and $F_{clk} = 5MHz$ defined the circuit operating frequency. The operating frequency of $5MHz$ is commonly used in the literature for similar applications [Zhang et al. 2010].

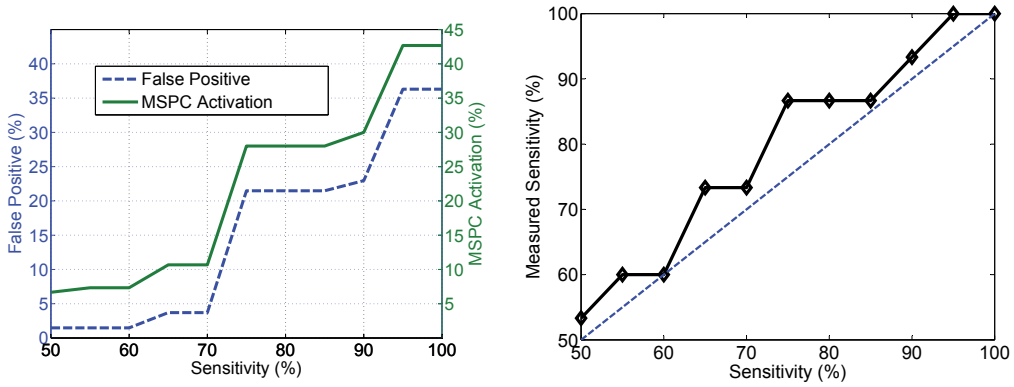

(a) FP and Micro Activation

(b) Measured Sensitivity

Fig. 6: Accuracy performance of the proposed decision making system while screening 'Sit to Stand' movements. (a) False positive rate and MSPC activation rate; (b) Measured sensitivity versus the given lower bound on sensitivity of the entire recognition system.

6.3. Full Template Scenario

In the first step, we considered 'Sit to Stand' as target action and all other actions as non-target. The main goal was to study different aspects of the proposed architecture in details by looking at one specific movement as target action. Later, in Section 6.5, we will present results for other actions as well. The main reason for choosing 'Sit to Stand' for the first set of analysis was its clinical significance and applications. 'Sit to Stand' actions are typically used as a means to assess motor function in clinical populations [Volpato et al. 2008; Kim et al. 2011; Rolland et al. 2006; Zheng et al. 2014; MAller et al. 2012; Regterschot et al. 2014; Cheng et al. 2014].

The power consumption of a screening block with full size template was computed as discussed in Section 6.2. The power consumption of the granular decision making module alone was $8.44\mu W$ which is significantly smaller than the power consumption of a typical signal processing chain (e.g., power consumption of processing unit of a Telos mote is 3 mW in active mode).

We first analyzed the effectiveness of the decision making module while a full template is used for screening. Fig. 5 and Fig. 6 illustrate the power and accuracy performance of the system obtained through this experiment. For this analysis, the given lower bound on the sensitivity of the system is assumed to range from 50% to 100%. Fig. 5(a) shows the power consumption of the system including the GDMM and microcontroller as a function of sensitivity. As expected, the power consumption increases as the sensitivity of the screening module grows. The power consumption of the system ranges from 0.21 mW for 50% sensitivity to 1.29 mW for 100% sensitivity, resulting in an average power consumption of 0.67 mW. We observe a sudden increase in the power consumption when sensitivity changes from 70% to 75% and another jump when it grows from 90% to 95%. This is in fact due to the increase in false positives as is illustrated in Fig. 6(a). This observation perhaps can be taken into consideration while exploring accuracy/power tradeoffs of the system. The graph is clearly divided into three distinct areas with low (50% to 70%), moderate (75% to 90%) and high (95% to 100%) sensitivity rates. Clearly, this provides the designer of the system with the flexibility to choose higher accuracy rates at the cost of extra power consumption and less

power savings. Fig. 5(b) shows the amount of power savings achieved using the proposed decision making module. The power saving numbers range from 57% for 100% sensitivity to 92% for the case of 50% sensitivity, with an average savings of 78%.

Fig. 6(a) shows false positive rates of the screening module as well as percentage of the times that the main processing will be activated. The false positive numbers range from 1.5% to 36.3% with an average of 15.6%. These numbers correlated with the MSPC activation which ranges from 6.7% to 42.7% resulting in an average of 22.0% chances of activating the main processor for performing more complex signal processing tasks. We note that activation of the main processor occurs on any action that is classified as positive. That includes both false positives (i.e., non-target action) and true positives (i.e., target action).

Fig. 6(b) shows actual sensitivity of the decision making module versus the desirable ones. As expected, the measured sensitivity would always exceed the lower bound sensitivity which is given as a design parameter. This is mainly due to the threshold setting mechanism which is discussed previously. The threshold for each screening block is set to guarantee the given sensitivity. This is further confirmed by all the data point in Fig. 6(b) falling into the area above the dashed line.

6.4. Mini-Template Case

In the next step, we divided the entire template into several bins and used the Min Size Mini-Template Set (MSMT) problem to find minimum number of bins that are required for detecting 'Sit to Stand'. A template on Z-axis accelerometer is a vector of 340 samples that corresponds to 6.8 seconds of sensor readings. A choice of K=10 (for example) generates ten bins, each having a length of 34 samples. Using the template generation formula (i.e., equation (4)) described in Section 5.3, we computed the template for all 14 actions. Fig. 7 illustrates templates of all actions using a single-axis node worn on the 'Waist' of the user. In this case, each signal template is divided into 10 equal-length bins that are highlighted as b_1 to b_{10} on the x-axis of each graph. The vertical axis (i.e., y-axis) shows the amount of acceleration in mm/s^2. The acceleration numbers shown in the figure refer to the acceleration of the 'Waist' in sagittal plane. This figure further emphasizes on our two key hypotheses: (1) there are tangible differences between different pairs of actions (2) a relatively small portion of the template would suffice to distinguish between the target and non-target actions.

We solved our optimization problem (see Problem 1) using the greedy algorithm described in Algorithm 1. Ideally, only a small subset of the bins would suffice for reliable identification of the target action. Intuitively, the amount of power consumption would decrease as the number of bins increases. The amount of improvement, however, becomes negligible as soon as enough resolution of the signal segment is obtained due to consideration of new mini-templates. We set the number of bins to 2 at the beginning and increased this number until no significant reduction in the power consumption was observed. Fig. 8(a) shows the power consumption of the screening blocks as the number of bins grows. The power consumption of the module, which is denoted by dashed curve, ranges from 8.44 μW (for 2 bins) to 4.07 μW (for 68 bins) as the number of bins increases. As illustrated in the figure, the amount of improvement in power consumption is minimal beyond 20 bins. In fact, the power of GDMM is enhanced only 5.8% when the number of bins grows from 20 to 34. The bold curve in Fig. 8(a) shows the fraction of the template that is chosen to be active for action classification. The dashed curve in Fig. 8(b) shows the number of active bins versus the number of bins (K). The number of active bin remains almost linear after $K = 10$. The bold curve in this figure is the length of each bin with respect to K.

Fig. 9 shows the amount of power savings as a function of the number of bins used for the optimization and analysis. The amount of power savings is computed by com-

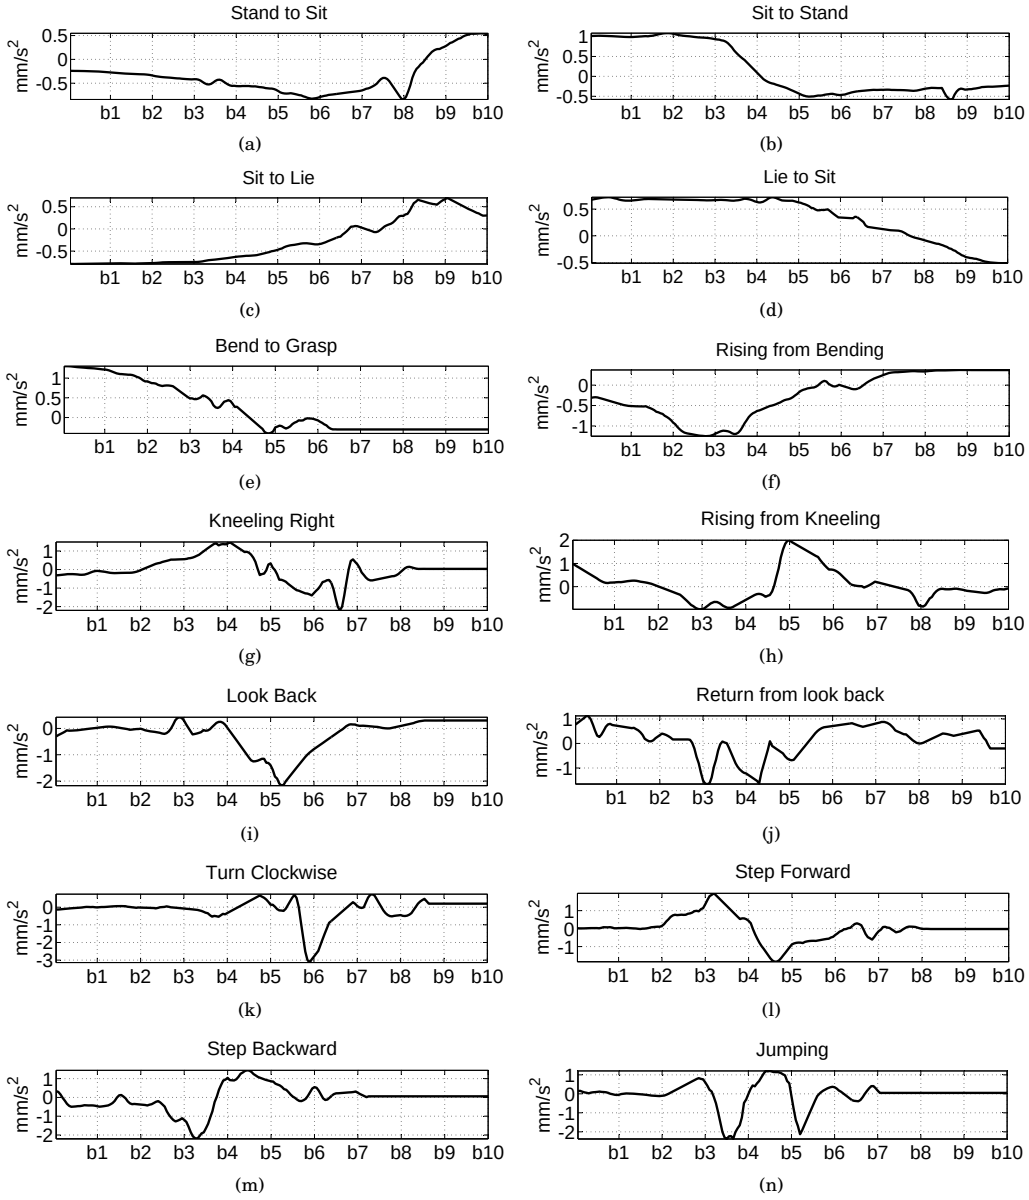

Fig. 7: Templates of various actions using 'Waist' sensor node. Each signal template is divided into 10 bins as shown on x-axis. On each graph, y-axis shows the amount of acceleration in sagittal plane (i.e., orthogonal to lateral plane of the body)

paring the total power consumption of GDMM when the mini-template approach is applied versus the full-template case. The power reduction is a result of optimizing mini-template set which leads to lower computational load in GDMM. Furthermore, Table III and Table II shows more details on how these performance metrics change as

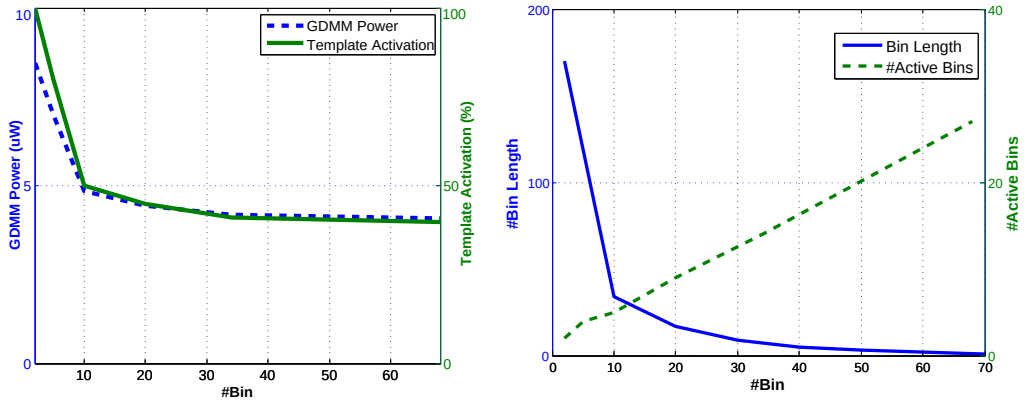

(a) GDMM Power and Template Activation (b) Bins Statistics

Fig. 8: Power performance of the mini-template. (a) Power consumption of the GDMM and template activation percentage versus number of original template bins used for power optimization. (b) Number of active bins and bin length as a function of number of bins.

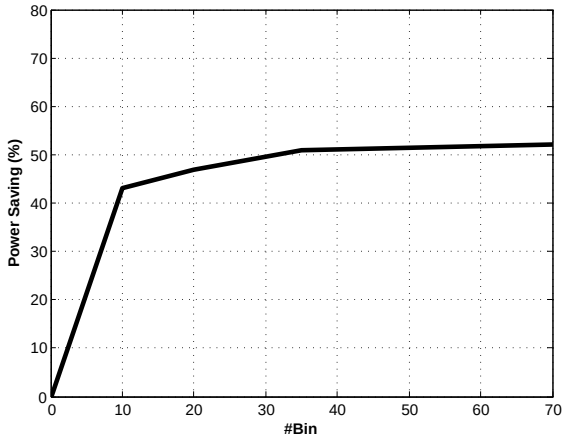

Fig. 9: Percentage of power savings as the number of bins increases.

a function of number of bins. Table II compares the power consumption of the decision paths (i.e., multiple screening blocks with mini-templates) with that of a full template to highlight the amount of additional power reduction achieved due to the construction of the mini-templates. By increasing the number of bins, the fraction of the full template which needs to be fed into GDMM decreases which results in lower dynamic power. Table II shows the power dissipation numbers in detail. The average amount of dynamic and leakage currents are shown in third and fourth column, respectively. The total power P_{Tot} is characterized by the summation of the dynamic P_{Dyn} and leakage P_{Leak} powers in the last column. The power saving values were then estimated in a

Table II: Power analysis of 12bit GDMM due to increasing number of bins for $V_{DD} = 1.1V$, $F_{clk} = 5MHz$.

No.	#Bins	$I_{Dyn}(\mu$A)	$I_{Leak}(\mu$A)	$P_{Dyn}(\mu$W)	$P_{Leak}(\mu$W)	$P_{Tot}(\mu$W)
1	2	6.54	1.13	7.20	1.24	8.44
2	5	5.23	1.13	5.75	1.24	7.00
3	10	3.27	1.13	3.60	1.24	4.85
4	20	2.90	1.13	3.18	1.24	4.43
5	34	2.66	1.13	2.92	1.24	4.17
6	68	2.57	1.13	2.83	1.24	4.07

Table III: Improvements due to increasing number of bins

No.	#Bins	Step Improv. (%)	Template Activation (%)	Overall Improv. (%)
1	2	NA	100	0
2	5	20	80	20
3	10	37	50	50
4	20	10	45	55
5	34	8	41	58
6	68	3	39	60

similar fashion for templates and mini-templates of different lengths on the incoming data.

The power consumption of the entire system is 0.85 mW on average, which results in an average power savings of 71.6%. Furthermore, the system exhibits 6.7% false positive rates on average, and the main signal processing chain is activated 28% of the times. We should note that power saving results presented in this article are based on the assumption that all actions occur with equal probabilities. In reality, however, target actions may occur less frequently, resulting in higher power savings due to infrequent activation of the main processor.

In Table III, step improvements and overall improvements are listed versus the number of bins. Step improvement refers to the amount of improvement achieved by increasing the number of bins (K). For instance, for $K = 5$, the template activation is 80% (4 bins out of 5 are activated) and for $K = 10$, this number is 50%. Therefore, by increasing the number of bins from 5 to 10, we achieve 37% step improvement ($(80-50)/80 = 37\%$). Overall improvement shows percentage of improvement in power optimization achieved by our optimization compared to the baseline (i.e., full size template), as proposed in Section 5. We note that if the full template is divided into a small number of bins (e.g., $K = 2$) most of the bins might be on the decision path, which results in a small or no improvement. Therefore, it is important to divide the template into a sufficiently large number of bins (K) and find only a small number of bins (R) for screening as suggested by the Min Size Mini-Template Set problem. In our experiments, a ratio of 10% to 15% between length of mini-templates and length of full-template (e.g., 34 and 340 for 'Sit to Stand') leaves sufficient information within each mini-template for classification, according to our experiments. Therefore, the number of bins can be set to satisfy this requirement. As suggested in Table III, we obtained an average improvement of 40.7% with mini-templates for detecting 'Sit to Stand'.

6.5. Extension to Other Actions

In order to measure the power consumption of our system for screening individual actions, we consider each action as target, and find minimum number of mini-templates

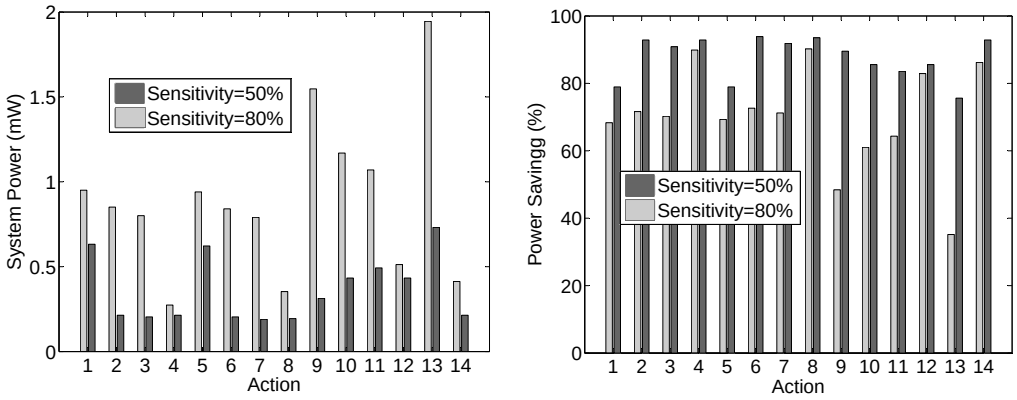

(a) Power Consumption (b) Power Saving

Fig. 10: Power performance of the system for two cases with 80% and 50% sensitivity rate. (a) Power consumption of the system including screening blocks and microcontroller; (b) Power saving due to the preliminary signal processing and early rejection of irrelevant movements.

needed for screening that particular action. The value of acceptance/rejection threshold (see Definition 5.3) was set to guarantee a minimum sensitivity (also known as true positive rate) which is given by the user. This problem was repeated for two cases with sensitivity rates of 50% and 80%. The test was done setting number of bins (K) to 20.

Fig. 10(a) shows the total power consumption of the system (including screening block and microcontroller). For the case of 80% sensitivity, power consumption values range from 0.27 mW to 1.95 mW with an average of 0.89 mW over all the experiments. For the 50% sensitivity, power consumption of the system ranges from 0.19 mW to 0.73 mW, and the average value is 0.36 mW. As it can be observed from this analysis, the power consumption of the system decreases as the sensitivity rates decrease. This is mainly because with a smaller sensitivity rate, less actions will be accepted as target action due to the low precision of the screening blocks. On the other hand, higher sensitivity rates will result in more actions being processed by the main processor (as they are accepted by the screening blocks) which results in overall higher power consumption of the entire system due to significantly higher power consumed by the microcontroller.

The overall power savings are achieved for screening different actions are shown in Fig. 10(b). The 80% sensitivity results in power savings ranging from 35.2% to 90.2% with an average of 70.1%. The amount of saving in the overall power consumption ranges from 75.6% to 93.9% with an average of 87.6%.

We also calculated the actual sensitivity rates that are obtained as a result of threshold setting for template matching. Fig. 11(a) shows the measured sensitivity rates for both cases of 50% and 80% desirable performance. When the given sensitivity is set to 80%, the measured values are between 80% and 86.7%. These number give an average measured sensitivity of 81.9%. Similarly, measured sensitivity ranges from 50.3% to 66.7% when the given lower bound is set to 50%.

Fig. 11(b) illustrates the false positive rates for the analysis cases and for each one of the 14 studied actions. The false positive rate ranges from 1.0% to 62.9% for the case

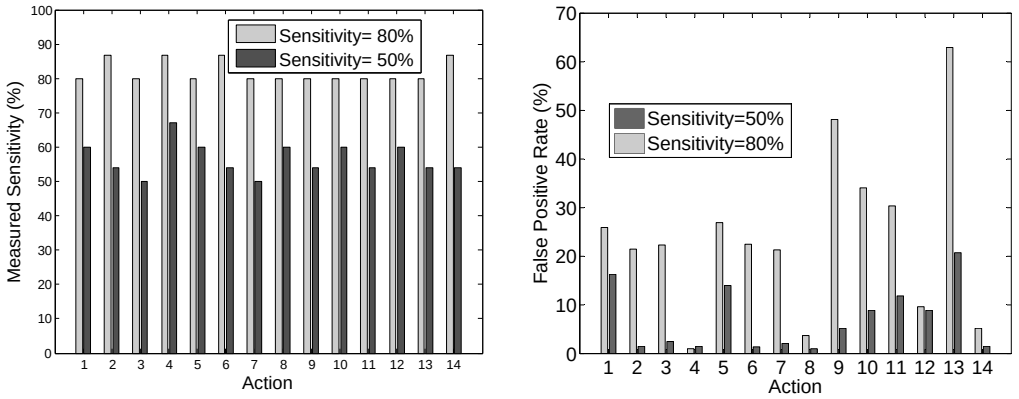

(a) Measured Sensitivity (b) False Positive Rate

Fig. 11: Accuracy performance of the system for two cases with 80% and 50% senility rate. (a) Sensitivity of the screening blocks; (b) False positive rate of the screening blocks (the entire granular decision making module).

of 80% desirable sensitivity. The average false positive rate across all the actions is 23.9%. In the case of 50% sensitivity, the obtained false positive rate is between 1.0% and 20.7% with an average of 6.9%. We note that the false positive rate decreases as the sensitivity goes down from 80% to 50%. This can be explained as follows. A lower desirable sensitivity rate (e.g., 50% compared to 80%) would results in the thresholding algorithm to choose a larger threshold for the template matching. The large threshold would diminish the precision of the template matching algorithms, and therefore, less actions will be accepted by the screening blocks and passed to the next processing level (i.e., microcontroller).

Finally, Fig. 12 shows the percentage of the time that the main processor is activated for each action as target and for two scenarios under analysis. The activation of the microcontroller differs for different actions and ranges from 8.7% to 64.7% when guaranteeing 80% sensitivity, and ranges from 6.0% to 24.0% for 50% sensitivity assurance.

6.6. Robustness of Template Generation

In Section 5.3, we described the process of template generation. Intuitively, we find a trial in the set of training trials associated with the target action. The target action trial is chosen such that it best represents the target action but is also most dissimilar to non-target actions. However, the non-target action set could be potentially large. Therefore we investigate the possible impact of non-target actions on the outcome of our template generation method. In this paper, we used a set of 13 non-target actions. In order to demonstrate the robustness of the template generation process to changes of the non-target action set, we performed an analysis to show how the output of our template generation algorithm changes as new non-target actions are added to the system. We start by randomly choosing one non-target action at a time and computing the template trial each time a new action is added to the non-target set. We repeated this process for 5 scenarios each representing a sequence of randomly selected non-target actions.

Fig. 13(a) shows 5 scenarios where each scenario represents a random sequence of 13 non-target actions. For this analysis, we considered 'Sit to stand' as our target action.

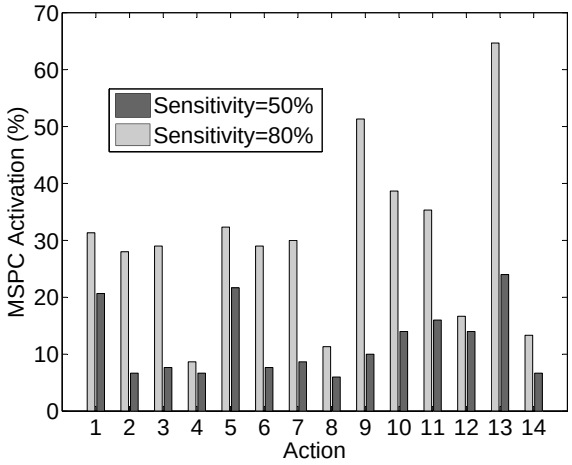

Fig. 12: Activation of the main signal processing chain for different movements considered as target.

In Fig. 13(b), we show the trial number of the target action which has been identified, by our template generation algorithm, as the template for the target action. For example, for scenario1, the non-target action set initially contains only one action, 'Sit to lie'. The system selects trial '12' of the target action as template. When the second non-target action, 'Lie to sit', is added, the template trial changes to '25'. By adding the third action, 'Step forward', the trial number for scenario1 changes to '30'. We, however, note that for the rest of the sequence (4th to 13th non-target actions for scenario1), the same trial (i.e., trial '11') is consistently chosen as template. Thus, for scenario1, the effect of non-target action on template generation disappears after including only 4 non-target actions. As it can be observed from Fig. 13(b), for scenario2, scenario3, scenario4, and scenario5, this effect disappear after adding 9, 3, 8, and 2 non-target actions respectively. On average, the set of non-target actions included only 5.2 actions in order for the template generation to stabilize.

Fig. 13(c) shows the amount of Θ in Equation 5 for the generated template for each scenario and number of non-target actions. The value of Θ depends on the target trials as well as non-target trials. This is the reason behind the fluctuation of Θ by addition of each non-target trial. However, we must take into the consideration that although the chosen template depends on the set of other actions' trials as well, it is, regardless of other actions, chosen within the target action training set. This is the reason why for example in scenario1, after adding 4 actions the chosen trial does not change regardless of the minor changes in Θ. This analysis demonstrates that our template generation scheme is robust to potential changes in the non-target action set and tends to quickly converge to a stable outcome.

7. DISCUSSION AND FUTURE WORK

We used cross-correlation scores to perform preliminary low power signal processing by quantifying similarity between incoming signals and target action. This approach is promising and allows for significant power saving while achieving acceptable accuracy performance. Higher power savings can be obtained in the expense of increase in false positive rates. In order to maintain smaller false positive rates, more complicated computing blocks (as alternatives for cross-correlation) can be used. Clearly, there are

Scenario1	Scenario2	Scenario3	Scenario4	Scenario5
Sit to lie	Step backward	Rise from bending	Step forward	Rise from bending
Lie to sit	Step forward	Sit to lie	Bend & grasp	Jump
Step forward	Kneel	Rise from kneeling	Jump	Step forward
Look back	Return from look back	Step forward	Rise from kneeling	Rise from kneeling
Step backward	Stand to sit	Jump	Rise from bending	Step backward
Kneel	Jump	Return from look back	Step backward	Kneel
Bend & grasp	Bend & grasp	Kneel	Lie to sit	Return from look back
Turn clockwise	Turn clockwise	Bend & grasp	Return from look back	Sit to lie
Jump	Lie to sit	Turn clockwise	Stand to sit	Step forward
Rise from kneeling	Step forward	Lie to sit	Sit to lie	Lie to sit
Rise from bending	Rise from kneeling	Step backward	Step forward	Bend & grasp
Return from look back	Sit to lie	Stand to sit	Kneel	Turn clockwise
Stand to sit	Rise from bending	Step forward	Turn clockwise	Stand to sit

(a) Scenarios showing five sequences of randomly selected non-target actions. For this example, 'Sit to Stand' is considered as target action. The table shows how the 13 non-target actions are added to the system for each scenario.

# of non-target actions	Scenario1	Scenario2	Scenario3	Scenario4	Scenario5
1	12	3	12	19	12
2	25	14	12	19	11
3	30	14	11	18	11
4	11	14	11	14	11
5	11	14	11	14	11
6	11	14	11	14	11
7	11	14	11	14	11
8	11	14	11	11	11
9	11	11	11	11	11
10	11	11	11	11	11
11	11	11	11	11	11
12	11	11	11	11	11
13	11	11	11	11	11

(b) Evolution of trial section for 'Sit to Stand' as the number of non-target actions grows. The table shows the trial of 'Sit to Stand' selected as template as the number of non-target actions changes from 1 action to 13 actions. The results are presented for 5 randomly generated sequences of non-target actions.

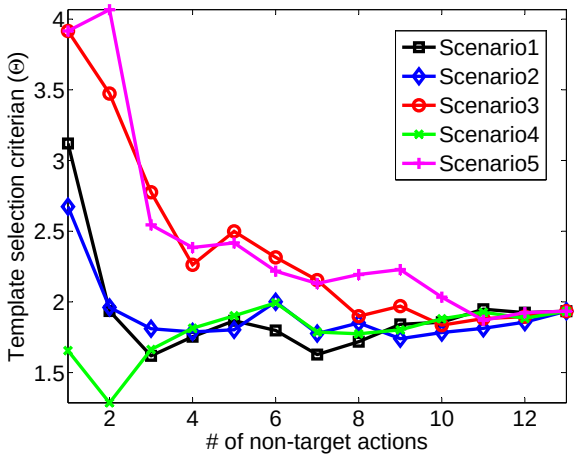

(c) The value of Θ in Equation 5 for the target action trial selected as template as the number of non-target action grows from 1 action to 13 actions. The results are presented for 5 randomly generated sequences of non-target actions, specified as 'Scenario1' to 'Scenario5'.

Fig. 13: Analysis of the robustness of the template generation approach due to addition of non-target actions.

tradeoffs between complexity of screening blocks and desired sensitivity rates. Currently, we are investigating alternative similarity measures for cross-correlation, and studying different system design tradeoffs.

The data storage units including template storage and the buffer are needed for all the cases as shown in Fig. 3. For example assuming a 50Hz sampling frequency, a 2-second movement and template, and a 12-bit Analog-to-Digital Convertor (ADC), the total data storage for a full template scenario is 2400 bits. This is because 'Memory' and 'Buffer' in Fig. 3 will require $50 \times 2 \times 12 = 1200$ bits each. As another example, if the window size is increased to 6.8 seconds, as in our experiments, the total amount of data storage with a full template will be $50 \times 6.8 \times 12 = 4080$ bits for each one of 'Memory' and 'Buffer' in Fig. 3 resulting in a total data storage of 8160 bits. This amount of storage, however, is not always necessary because, depending on the amount of sensitivity rate and number of template bins, only a portion of the template will be active. For instance, with a given sensitivity rate of 80% and number of template bins of 10, the storage size for 'Memory' in Fig. 3 will be 2040 bits $(4080/2)$ since the template activation is only 50% as suggested in Fig. 8(a). Thus, the total storage will be $2040 + 4080 = 6120$ bits (765 Bytes) in this case.

There might be some concerns about whether or not the set of non-target activities are the best representative of the possibly broad range of non-target activities. To address this concern, two facts need to be considered. First, a broad range of applications in the area of BANs are designed for controlled scenarios (such as the applications mentioned in Section 1) where a reasonable set of non-target actions can be listed. Secondly, the proposed energy-aware model assumes that the set of all activities are given. This information can be fed into our model either by using a supervised learning activity recognition approach or alternative approaches such as semi-supervised learning models [Stikic et al. 2011; Lara and Labrador 2013] which aim to tackle the common issue of supervised learning (i.e., the set of non-target actions is sometimes large). Prior studies [Logan et al. 2007; Huynh et al. 2008] which have conducted open world experiments consisted of long-term daily activity recordings of individuals in nonlaboratory conditions suggest that in real world scenarios the final set of activities ranges from 8 to 20 which leads to a reasonable memory overhead of 144×8 to 144×20 Bytes for a given/desirable specific sensitivity of 80%.

The granular decision making module presented in this article assumes a fixed-length template based on normal speed of the human movements. In reality, however, human movements can occur at different speeds. The decision making module can still work by dynamically adjusting the length of the template. Variations in the speed can be handled through a signal scaling method. The signal scaling can be done during template matching and based on the detected speed. In such scenarios, the template can be adjusted using a normalization factor by down sampling or up sampling. In fact, the system can start with a low power saving if the speed of the incoming signal is much different than that of the template. However, the main processing module can report back to the GDMM over time about the speed of the current action. Then the GDMM will adjust the template length accordingly. As part of our future work, we plan on evaluating the performance of our system with respect to temporal changes in the incoming signals.

The amount of power savings that can be achieved by our screening approach highly depends on the frequency of occurrence of the target action. For our experiments, we assumed that all actions are equally likely, and therefore, 'Sit to Stand' occurs 7.1% of the times. In reality, however, human actions are sparse occurring a lot more infrequently, which results in much higher power savings.

In this article, we focused on analyzing the accuracy performance of the granular decision making module independent of the exact types of the computing algorithms

that take place in the main signal processing chain. We assumed that the main signal processing chain will perform more complex signal processing for activity recognition and/or extraction of more detailed information from the sensor signals when an action of interest occurs. Thus, the accuracy performance metrics discussed in this paper (e.g., false positive rates and true positive rates) are reported for the granular decision making module rather than the main signal processing chain.

Our experimental validation involves an implicit assumption that actions occur continuously and at equal rates. In reality, however, human actions may be occurring at significantly lower rates. We note that infrequent occurrence of the target action will result in activating the main processor less frequently, and therefore, achieving more power savings.

Our ongoing research on devising granular decision making involves refinement of the template matching for further optimization with respect to other tunable parameters such as bit resolution of sensor readings and duration of incoming actions in both time and frequency domains.

In this article, we focused on development of the GDMM and studied its impact on reducing power consumption of the main processor. As part of our ongoing research, we are currently investigating full integration of the GDMM within the main processor and refining our power model to account for additional power sources, such as I/O power dissipation, due to such an integration.

8. CONCLUSION

In this article, we presented a new ultra low-power signal screening approach with the goal of significantly reducing energy consumption of wearable computers. We also demonstrated the efficacy of the proposed processing model for energy saving in healthcare applications. The proposed signal screening model is based on a sequence of template matching operations each associated with a fraction of a pre-specified signal template, called mini-template. The idea is to reject actions that are unlikely to be the target action of interest and initiate detailed processing by the main processor only if the incoming signal is highly correlated with the predefined template. Our experimental results demonstrate the efficacy of the proposed architecture significantly reducing the power consumption of the system. Specifically, we obtained an energy savings of over 70.1% while detecting transitional daily activities while maintaining 80% sensitivity in activity recognition.

REFERENCES

Stefano Abbate, Marco Avvenuti, and Janet Light. 2014. Usability Study of a Wireless Monitoring System among Alzheimers Disease Elderly Population. *International Journal of Telemedicine and Applications* 2014 (2014).

P. Alinia, R. Saeedi, B. Mortazavi, A. Rokni, and H. Ghasemzadeh. 2015. Impact of sensor misplacement on estimating metabolic equivalent of task with wearables. In *Wearable and Implantable Body Sensor Networks (BSN), 2015 IEEE 12th International Conference on*. 1–6. DOI:http://dx.doi.org/10.1109/BSN.2015.7299385

Giuseppe Amato and Fabrizio Falchi. 2013. On knn classification and local feature based similarity functions. In *Agents and Artificial Intelligence*. Springer, 224–239.

Jorge Cancela, Matteo Pastorino, Maria T Arredondo, Konstantina S Nikita, Federico Villagra, and Maria A Pastor. 2014. Feasibility Study of a Wearable System Based on a Wireless Body Area Network for Gait Assessment in Parkinsons Disease Patients. *Sensors* 14, 3 (2014), 4618–4633.

N. Chaimanonart and D.J. Young. 2009. A wireless batteryless in vivo EKG and body temperature sensing microsystem with adaptive RF powering for genetically engineered mice monitoring. *TRANSDUCERS 2009 - 15th International Conference on Solid-State Sensors, Actuators and Microsystems* (2009), 1473–1476.

L. Chao, C.-Y. Tsui, and W.-H. Ki. 2007. A batteryless vibration-based energy harvesting system for ultra low power ubiquitous applications. In *Proceedings - IEEE International Symposium on Circuits and Systems*. 1349–1352.

Fang Chen, Cun Ji Zhang, Bin Wen Zhao, and Jin Fei Shi. 2013. A Speed-Up Way of Template Matching Using Normalized Cross-Correlation. *Applied Mechanics and Materials* 313 (2013), 1188–1191.

Yuan-Yang Cheng, Shun-Hwa Wei, Po-Yin Chen, Mei-Wun Tsai, I Cheng, Ding-Hao Liu, Chung-Lan Kao, and others. 2014. Can sit-to-stand lower limb muscle power predict fall status? *Gait & Posture* (2014).

Nina J Cleven, Peter Isfort, Tobias Penzkofer, Anna Woitok, Benita Hermanns-Sachweh, Ulrich Steinseifer, and Thomas Schmitz-Rode. 2014. Wireless Blood Pressure Monitoring with a Novel Implantable Device: Long-Term In Vivo Results. *Cardiovascular and interventional radiology* (2014), 1–9.

P. Cong, N. Chaimanonart, W.H. Ko, and D.J. Young. 2009a. A wireless and batteryless 10-bit implantable blood pressure sensing microsystem with adaptive RF powering for real-time laboratory mice monitoring. *IEEE Journal of Solid-State Circuits* 44, 12 (2009), 3631–3644.

P. Cong, N. Chaimanonart, W.H. Ko, and D.J. Young. 2009b. A wireless and batteryless 130mg 300W 10b implantable blood-pressure-sensing microsystem for real-time genetically engineered mice monitoring. In *Digest of Technical Papers - IEEE International Solid-State Circuits Conference*.

VR Sarma Dhulipala and GR Kanagachidambaresan. 2014. Cardiac Care Assistance using Self Configured Sensor Networka Remote Patient Monitoring System. *Journal of The Institution of Engineers (India): Series B* (2014), 1–6.

Jia Di and Scott Christopher Smith. 2014. Ultra-low power multi-threshold asynchronous circuit design. (March 4 2014). US Patent 8,664,977.

Bruce H Dobkin. 2013. Wearable motion sensors to continuously measure real-world physical activities. *Current opinion in neurology* 26, 6 (2013), 602–608.

T. Douseki, T. Tsukahara, Y. Yoshida, F. Utsunomiya, and N. Hama. 2003. A batteryless wireless system with MTCMOS/SOI circuit technology. In *Proceedings of the Custom Integrated Circuits Conference*. 163–168.

Van Emmerik and Van Wegen. 2002. On the Functional Aspects of Variability in Postural Control. *Exercise and Sport Sciences Reviews* 30, 4 (2002), 177–183.

R. Fallahzadeh, M. Pedram, R. Saeedi, B. Sadeghi, M. Ong, and H. Ghasemzadeh. 2015. Smart-Cuff: A wearable bio-sensing platform with activity-sensitive information quality assessment for monitoring ankle edema. In *Pervasive Computing and Communication Workshops (PerCom Workshops), 2015 IEEE International Conference on*. 57–62. DOI:http://dx.doi.org/10.1109/PERCOMW.2015.7133994

Uriel Feige and Prasad Tetali. 2004. Approximating Min Sum Set Cover. *Algorithmica* 40, 4 (2004), 219–234. DOI:http://dx.doi.org/10.1007/s00453-004-1110-5

H. Ghasemzadeh, E. Guenterberg, and R. Jafari. 2009. Energy-Efficient Information-Driven Coverage for Physical Movement Monitoring in Body Sensor Networks. *IEEE Journal on Selected Areas in Communications* 27 (2009), 58–69.

H. Ghasemzadeh and R. Jafari. 2011a. Physical Movement Monitoring Using Body Sensor Networks: A Phonological Approach to Construct Spatial Decision Trees. *Industrial Informatics, IEEE Transactions on* 7, 1 (Feb 2011), 66–77. DOI:http://dx.doi.org/10.1109/TII.2010.2089990

H. Ghasemzadeh and R. Jafari. 2011b. An Ultra Low Power Granular Decision Making Using Cross Correlation: Minimizing Signal Segments for Template Matching. In *Cyber-Physical Systems (ICCPS), 2011 IEEE/ACM International Conference on*. 77 –86.

Hassan Ghasemzadeh and Roozbeh Jafari. 2013. Ultra Low-power Signal Processing in Wearable Monitoring Systems: A Tiered Screening Architecture with Optimal Bit Resolution. *ACM Trans. Embed. Comput. Syst.* 13, 1, Article 9 (Sept. 2013), Article 9, 23 pages. DOI:http://dx.doi.org/10.1145/2501626.2501636

H. Ghasemzadeh, V. Loseu, and R. Jafari. 2010. Structural Action Recognition in Body Sensor Networks: Distributed Classification Based on String Matching. *Information Technology in Biomedicine, IEEE Transactions on* 14, 2 (March 2010), 425 –435.

E. Guenterberg, H. Ghasemzadeh, and R. Jafari. 2009a. A Distributed Hidden Markov Model for Fine-grained Annotation in Body Sensor Networks. In *Proceedings of the 2009 Sixth International Workshop on Wearable and Implantable Body Sensor Networks-Volume 00*. IEEE Computer Society, 339–344.

E. Guenterberg, A.Y. Yang, H. Ghasemzadeh, R. Jafari, R. Bajcsy, and S.S. Sastry. 2009b. A method for extracting temporal parameters based on hidden markov models in body sensor networks with inertial sensors. *IEEE Transactions on Information Technology in Biomedicine* 13, 6 (2009), 1019–1030.

S. Hagler, D. Austin, T.L. Hayes, J. Kaye, and M. Pavel. 2010. Unobtrusive and ubiquitous in-home monitoring: A methodology for continuous assessment of gait velocity in elders. *IEEE Transactions on Biomedical Engineering* 57, 4 (2010), 813–820.

Jihun Hamm, Benjamin Stone, Mikhail Belkin, and Simon Dennis. 2013. Automatic annotation of daily activity from smartphone-based multisensory streams. In *Mobile Computing, Applications, and Services*. Springer, 328–342.

Niloofar Hezarjaribi, Ramin Fallahzadeh, and Hassan Ghasemzadeh. 2016. A Machine Learning Approach for Medication Adherence Monitoring Using Body-Worn Sensors. *IEEE/ACM Design, Automation and Test in Europe (DATE)* (2016).

Matti Huiku and Juha Virtanen. 2013. Method for reducing power consumption in pulse oximeter systems, pulse oximeter system and pulse oximeter sensor. (Oct. 29 2013). US Patent 8,571,622.

Tâm Huynh, Mario Fritz, and Bernt Schiele. 2008. Discovery of activity patterns using topic models. In *Proceedings of the 10th international conference on Ubiquitous computing*. ACM, 10–19.

W.D. Jones. 2007. Helmets Sense The Hard Knocks [News]. *Spectrum, IEEE* 44, 10 (Oct. 2007), 10–12. DOI:http://dx.doi.org/10.1109/MSPEC.2007.4337656

Min H. Kim, Chung H. Yi, Won G. Yoo, and Bo R. Choi. 2011. EMG and kinematics analysis of the trunk and lower extremity during the sit-to-stand task while wearing shoes with different heel heights in healthy young women. *Human Movement Science* 30, 3 (2011), 596 – 605.

Giuseppe Lanza, Mariagiovanna Cantone, Bartolo Lanuzza, Manuela Pennisi, Rita Bella, Giovanni Pennisi, and Raffaele Ferri. 2014. Distinctive patterns of cortical excitability to transcranial magnetic stimulation in obstructive sleep apnea syndrome, restless legs syndrome, insomnia, and sleep deprivation. *Sleep medicine reviews* (2014).

Oscar D Lara and Miguel A Labrador. 2013. A survey on human activity recognition using wearable sensors. *Communications Surveys & Tutorials, IEEE* 15, 3 (2013), 1192–1209.

W.-S. Liew, L. Zou, X. abd Yao, and Y. Lian. 2009. A 1-V 60-uW 16-channel interface chip for implantable neural recording. In *Proceedings of the Custom Integrated Circuits Conference*. 507–510.

Z. Lin and L.S. Davis. 2010. Shape-based human detection and segmentation via hierarchical part-template matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32, 4 (2010), 604–618.

Beth Logan, Jennifer Healey, Matthai Philipose, Emmanuel Munguia Tapia, and Stephen Intille. 2007. *A long-term evaluation of sensing modalities for activity recognition*. Springer.

Yuchao Ma, R. Fallahzadeh, and H. Ghasemzadeh. 2015. Toward robust and platform-agnostic gait analysis. In *Wearable and Implantable Body Sensor Networks (BSN), 2015 IEEE 12th International Conference on*. 1–6. DOI:http://dx.doi.org/10.1109/BSN.2015.7299366

Andrei Maciuca, Mircea Strutu, Dan Popescu, and Grigore Stamatescu. 2013. Cell-based sensor network for complex monitoring at home of patients with chronic diseases. In *Electrical and Electronics Engineering (ISEEE), 2013 4th International Symposium on*. IEEE, 1–6.

Andreas Buch MAller, Bo Martin Bibby, Anders Guldhammer Skjerbk, Ellen Jensen, Henrik Srensen, Egon Stenager, and Ulrik Dalgas. 2012. Validity and variability of the 5-repetition sit-to-stand test in patients with multiple sclerosis. *Disability and Rehabilitation* 0, 0 (2012), 1–8.

S. Mandal, L. Turicchia, and R. Sarpeshkar. 2010. A Low-Power, Battery-Free Tag for Body Sensor Networks. *Pervasive Computing, IEEE* 9, 1 (January-March 2010), 71 –77.

N. Najafi and A. Ludomirsky. 2004. Initial animal studies of a wireless, batteryless, MEMS implant for cardiovascular applications. *Biomedical Microdevices* 6, 1 (2004), 61–65.

J. Parramon, P. Doguet, D. Marin, M. Verleyssen, R. Munoz, L. Leija, and E. Valderrama. 1997. ASIC-based batteryless implantable telemetry microsystem for recording purposes. In *Annual International Conference of the IEEE Engineering in Medicine and Biology - Proceedings*, Vol. 5. 2225–2228.

G Ruben H Regterschot, Marjanne Folkersma, Wei Zhang, Heribert Baldus, Martin Stevens, and Wiebren Zijlstra. 2014. Sensitivity of sensor-based sit-to-stand peak power to the effects of training leg strength, leg power and balance in older adults. *Gait & posture* 39, 1 (2014), 303–307.

Yves Rolland, Valerie Lauwers-Cances, Matteo Cesari, Bruno Vellas, Marco Pahor, and Hlne Grandjean. 2006. Physical Performance Measures as Predictors of Mortality in a Cohort of Community-dwelling Older French Women. *European Journal of Epidemiology* 21 (2006), 113–122. Issue 2. http://dx.doi.org/10.1007/s10654-005-5458-x 10.1007/s10654-005-5458-x.

Kaushik Roy, Saibal Mukhopadhyay, and Hamid Mahmoodi-Meimand. 2003. Leakage current mechanisms and leakage reduction techniques in deep-submicrometer CMOS circuits. *Proc. IEEE* 91, 2 (2003), 305–327.

Ramyar Saeedi, Navid Amini, and Hassan Ghasemzadeh. 2014a. Patient-Centric On-Body Sensor Localization in Smart Health Systems. In *The Asilomar Conference on Signals, Systems, and Computers*. IEEE.

Ramyar Saeedi, Brian Schimert, and Hassan Ghasemzadeh. 2014b. Cost-sensitive feature selection for on-body sensor localization. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*. ACM, 833–842.

Arash Salarian, Pierre R Burkhard, François JG Vingerhoets, Brigitte M Jolles, and Kamiar Aminian. 2013. A novel approach to reducing number of sensing units for wearable gait analysis systems. *Biomedical Engineering, IEEE Transactions on* 60, 1 (2013), 72–77.

S. Sasaki, T. Seki, and S. Sugiyama. 2006. Batteryless accelerometer using power feeding system of RFID. In *2006 SICE-ICASE International Joint Conference*. 3567–3570.

F. Shahrokhi, K. Abdelhalim, D. Serletis, P.L. Carlen, and R. Genov. 2010. The 128-channel fully differential digital integrated neural recording and stimulation interface. *IEEE Transactions on Biomedical Circuits and Systems* 4, 3 (2010), 149–161.

A. Shoeb, D. Carlson, E. Panken, and T. Denison. 2009. A micropower support vector machine based seizure detection architecture for embedded medical devices. In *Proceedings of the 31st Annual International Conference of the IEEE Engineering in Medicine and Biology Society: Engineering the Future of Biomedicine, EMBC 2009*. 4202–4205.

Erik Sigmund, Dagmar Sigmundová, Romana Šnoblová, and Andrea Madarásová Gecková. 2014. ActiTrainer-determined segmented moderate-to-vigorous physical activity patterns among normal-weight and overweight-to-obese Czech schoolchildren. *European journal of pediatrics* 173, 3 (2014), 321–329.

Maja Stikic, Diane Larlus, Sandra Ebert, and Bernt Schiele. 2011. Weakly supervised recognition of daily life activities with wearable sensors. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 33, 12 (2011), 2521–2537.

Tigest Tamrat, Margaret Griffin, Sonia Rupcic, Stan Kachnowski, Tom Taylor, and James Barfield. 2012. Operationalizing a wireless wearable fall detection sensor for older adults. In *Pervasive Computing Technologies for Healthcare (PervasiveHealth), 2012 6th International Conference on*. IEEE, 297–302.

Toshiyo Tamura, Masaki Sekine, Hiroaki Miyoshi, Yutaka Kuwae, and Toshiro Fujimoto. 2013. Wearable inertia sensor application in the rehabilitation field. *Advances in Science and Technology* 85 (2013), 28–32.

Çağdaş Varel, Yi-Chun Shih, Brian P Otis, Tueng S Shen, and Karl F Böhringer. 2014. A wireless intraocular pressure monitoring device with a solder-filled microchannel antenna. *Journal of Micromechanics and Microengineering* 24, 4 (2014), 045012.

Stefano Volpato, Margherita Cavalieri, Gianluca Guerra, Fotini Sioulis, Monica Ranzini, Cinzia Maraldi, Renato Fellin, and Jack M. Guralnik. 2008. Performance-Based Functional Assessment in Older Hospitalized Patients: Feasibility and Clinical Correlates. *The Journals of Gerontology Series A: Biological Sciences and Medical Sciences* 63, 12 (2008), 1393–1398.

Jinyong Zhang, Zedong Nie, Jin Huang, Li Yu, and Lei Wang. 2010. Towards low frequency low noise low power body sensor network-on-chip. In *Green Circuits and Systems (ICGCS), 2010 International Conference on*. IEEE, 115–117.

Yanqing Zhang, Fan Zhang, Yousef Shakhsheer, Jason D Silver, Alicia Klinefelter, Manohar Nagaraju, James Boley, Jagdish Pandey, Aatmesh Shrivastava, Eric J Carlson, and others. 2013. A batteryless 19 W MICS/ISM-band energy harvesting body sensor node SoC for ExG applications. *Solid-State Circuits, IEEE Journal of* 48, 1 (2013), 199–213.

Enhao Zheng, Baojun Chen, Xuegang Wang, Yan Huang, and Qining Wang. 2014. On the Design of a Wearable Multi-sensor System for Recognizing Motion Modes and Sit-to-stand Transition. *Int J Adv Robot Syst* 11 (2014), 30.