



HAL
open science

A survey of energy-efficient context recognition systems using wearable sensors for healthcare applications

Tifenn Rault, Abdelmadjid Bouabdallah, Yacine Challal, Frédéric Marin

► **To cite this version:**

Tifenn Rault, Abdelmadjid Bouabdallah, Yacine Challal, Frédéric Marin. A survey of energy-efficient context recognition systems using wearable sensors for healthcare applications. *Pervasive and Mobile Computing*, 2017, 37, pp.23-44. 10.1016/j.pmcj.2016.08.003 . hal-01466848

HAL Id: hal-01466848

<https://hal.science/hal-01466848>

Submitted on 13 Feb 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A survey of energy-efficient context recognition systems using wearable sensors for healthcare applications

Tifenn Rault^a, Abdelmadjid Bouabdallah^a, Yacine Challal^{a,c,d}, Frédéric Marin^b

^aUniversité de Technologie de Compiègne, Heudiasyc UMR CNRS 7253, 60205 Compiègne, France

^bUniversité de Technologie de Compiègne, BMBI UMR CNRS 7338, 60205 Compiègne, France

^cLaboratoire de Méthodes de Conception de Systèmes, Ecole nationale Supérieure d'Informatique, Algiers, Algeria

^dCentre de Recherche sur l'Information Scientifique et Technique, Algiers, Algeria

Abstract

Human context recognition (HCR) from on-body sensor networks is an important and challenging task for many healthcare applications because it offers continuous monitoring capability of both personal and environmental parameters. However, these systems still face a major energy issue that prevent their wide adoption. Indeed, in healthcare applications, sensors are used to capture data during daily life or extended stays in hospital. Thus, continuous sampling and communication tasks quickly deplete sensors' battery reserves, and frequent battery replacement are not convenient. Therefore, there is a need to develop energy-efficient solutions for long-term monitoring applications in order to foster the acceptance of these technologies by the patients. In this paper, we survey existing energy-efficient approaches designed for HCR based on wearable sensor networks. We propose a new classification of the energy-efficient mechanisms for health-related human context recognition applications and we review the related works in details. Moreover, we provide a qualitative comparison of these solutions in terms of energy-consumption, recognition accuracy and latency. Finally, we discuss open research issue and give directions for future works.

Keywords: State-of-the-art, Energy efficient wearable sensor networks, Human context recognition

1. Introduction

The aging of the population is a chief concern of many countries since longer life expectancy and declining fertility rates are increasing the proportion of the elderlies worldwide. It is expected that approximately 20% of the world population will be age 60 or older by 2050 [1]. This situation poses new challenges to existing healthcare systems, such as an increase in age-related diseases, dependency, and shortage of caregivers. It is hoped that technology will help in addressing these challenges, for instance by helping elderly people to live more independent lives and thus reducing the burden of caregivers.

It is against this background that the recognition of human context has received a significant attention in the medical field, because it is a task of interest for many applications, including assisted living of elderlies, physical assessment, and supervision of patient with chronic diseases such as diabetes, obesity, cognitive disorder or arrhythmia [2]. In a typical scenario, wearable sensors like accelerometers or ECG are attached to the patient to measure biomechanical, physiological and environmental data [3]. These data are further analyzed to provide a feedback to caregivers who can assess the efficiency of a new treatment, adjust medication, better study a disease or supervise a patient behavior [4]. In practice, such systems face a major challenge that limit their widespread adoption: the significant energy consumption of battery-operated devices for continuous sensing and communications. Indeed, radio transmissions, data processing

Email addresses: tifenn.rault@utc.fr (Tifenn Rault), abdelmadjid.bouabdallah@utc.fr (Abdelmadjid Bouabdallah), y_challal@esi.dz (Yacine Challal), frederic.marin@utc.fr (Frédéric Marin)

and sensing tasks actively stress the nodes reserves, whereas they are expected to operate during the course of the day without maintenance.

To tackle this problem, many solutions have been proposed to save energy in context recognition applications. Among them we can cite sensor set selection [5], deactivation of power-hungry sensors [6], adaptive sampling rate [7], communications reduction [8]. All these solutions leverage energy consumption for latency and recognition accuracy. This is because health-related HCR applications have specific requirements, and energy-efficiency is usually understood to mean a satisfactory trade-off between three optimization criteria: low energy consumption, high recognition accuracy and low latency. Although numerous methods have been proposed to address the energy problem of HCR based on body sensors, there does not exist a comprehensive survey on the techniques available in the literature. In this paper, we survey the state-of-the-art of energy-efficient HCR solutions making use of wearable sensors. Our aim is to provide designers of HCR systems with a survey that offers a holistic view of energy-saving solutions while taking into consideration the specific requirements of healthcare applications.

This survey is organized as follows. In the next section, we introduce health-related human context recognition applications using wearable sensors. Then, in Section 3, we discuss the general components of context recognition systems and their main requirements. Afterwards, in section 4, we introduce our classification of the existing energy-efficient mechanisms developed for HCR applications. From Section 5 to 8, we present solutions related to respectively power-on time reduction, communication reduction, computation reduction, and battery charging. Subsequently, in Section 9, we provide a qualitative comparison of the reviewed solutions, and in Section 10 we identify open research issues.

2. Human context recognition

In this section, we briefly introduce human context recognition (HCR) applications making use of wearable sensor networks.

2.1. Applications

A human context recognition system using wearable sensors exploits sensors signals that allow monitoring a user and its environment in order to infer ongoing tasks and living conditions. Human context recognition encompasses both personal context recognition (e.g activity, health status, emotion) and environmental context recognition (e.g. location, dust-level, social interaction), as described in Table 1. The comprehension of the personal and environmental contexts is of great importance to improve the patient follow up for medical and wellbeing applications.

Human context	Context type	Sensor type
Personal context	activity, emotion, health status	accelerometers, gyroscope, blood pressure, ECG, EMG, temperature
Environmental context	location, place, noise-level, pollution, social interaction	GPS, WiFi traces, CO_2 sensor, luminosity, microphone, Bluetooth scans

Table 1: Human context recognition.

In a typical scenario, patients wear sensors that supervise their vital parameters in order to identify emergency situations and allow caregivers to respond effectively. This kind of application includes remote vital sign monitoring [9], and sudden fall or epilepsy seizure detection [10]. Moreover, body sensors can be used to gather clinically relevant information for rehabilitation supervision, elderly monitoring [11] or to provide support to a person with cognitive disorder [12]. For example, in [12], Lin et al. use a GPS-enabled wearable device to detect potential disorientation behaviors of cognitively-impaired elders. Zhan et al. [11] infer activities of daily living using accelerometer and video data embedded in conventional reading glasses

in order to remotely supervise elder and disabled patients. Another type of health-related system aims to use human context information to promote a more active and healthy lifestyle. For example, the BeWell+ application monitors user's everyday activities using the accelerometer and microphone sensors on phone [13]. It computes and displays wellbeing scores (sleep score, physical activity score and social interaction score) in order to encourage healthy behaviors. In [3], the authors designed a system to promote physical activity among elderlies: the platform uses wearable sensors to provide a real-time feedback to the user regarding his movement. It is worth mentioning that human context recognition can be of great interest in other domains of application such as gaming, industry and social interaction. For instance, Lee et al. [14] propose SocioPhone, a mobile platform for face-to-face interaction monitoring.

However, health-related applications have specific requirements in terms of energy-consumption, latency and recognition accuracy. These specifications justify the fact that dedicated energy-efficient solutions have been developed in the literature, although they may be transferable to other applications.

2.2. History

Early researches in context recognition were conducted using vision-based systems [15], but these systems were found expensive, intrusive and limited to well-equipped laboratories because of their cost and complexity. With the development of hardware technologies, sensing devices - lightweight enough to be worn - were developed. They provided a low-cost and effective alternative for human context recognition [16]. As a consequence, the first studies on human context recognition using wearable sensors emerged around the 1990s. They mostly relied on inertial sensors such as accelerometers and gyroscopes, and focused on recognizing a special subset of everyday activities (e.g. ambulation, walking, sitting, typing) [17]. At that time, a portable data logger was used to gather the data from body mounted sensors. Although these research prototypes were still bulky and did not propose real time feedback, they offered promising perspectives for ubiquitous and personal applications.

In early 2000s, medical sensors have incorporated wireless connections, such as Bluetooth and ZigBee, to communicate wirelessly to nearby computers, personal digital assistants, or mobile phones. Meanwhile, advances in miniaturization permitted sensors to be embedded within wrist band, belts, bracelets, rings [18]. Since then, mobile and real-time systems incorporating multi-modal sensors that could be used to recognize physical activities, social interactions or location have emerged [19]. Still, these applications were constrained by the sensors resources limitations. On the one hand, low memory and computation capabilities of sensors required the designer to develop and/or use low-resource-demanding algorithms in order to take into consideration these specific new constraints. On the other hand, the sensors' energy consumption became a key challenge for the successful development of pervasive systems which necessitate to be worn during an extended period of time for in-wild experiments and long-term monitoring applications.

2.3. Sensors

The most commonly types of sensors used for personal context recognition are probably accelerometers, because they are small, cheap and unobtrusive. Other types of sensors used for health-related data collection include pressure sensors, body temperature sensors, electromyography, oximetry sensors, and electrocardiographs. More kinds of sensors have been employed to improve the context-awareness of the application. For example, GPS or WiFi traces are used to determine the user localization [20], microphone records and Bluetooth traces are analyzed to detect ongoing interaction [21], CO_2 and dust sensors are employed to estimate the pollution level [22]. With the advances in miniaturization, increasingly pervasive sensing devices have been developed: sensors are now integrated into clothes, smart glasses and wearable rings [11, 18, 23]. A step further, some applications now instrument objects in the user's environment with RFID tags, and use data from a wearable RFID reader to infer household activities (such as making coffee, doing laundry, washing dishes) [8]. However, the use of RFID tags restricts these approaches to closed instrumented environments.

Depending on the application objective, a recognition system can use the same type of sensors at multiple body locations or combine complementary sensors that can help in recognizing personal and environmental contexts, e.g motion and audio, motion and location [24, 20]. Traditionally, information sources were deployed for a specific goal, but new applications try to dynamically adapt to available sensor data [25].

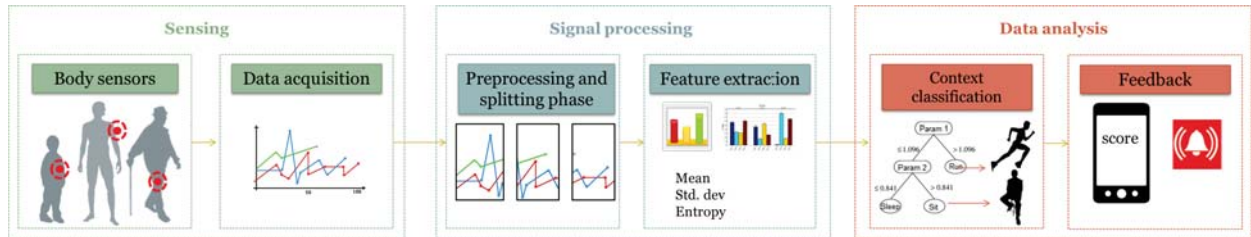


Figure 1: The main stages of human context recognition systems.

Indeed, the available source of information now varies depending on the user co-location with other body area networks [26], and on the user location when there exist space-embedded sensors [27].

2.4. Context information

Having the viewpoint of human context sensing provides a great deal of advantages. First, the recognition of human environment helps to efficiently monitor the current human activity, and vice versa. For instance, in [8], Wang et al. employ location information to infer household activities. Nath propose ACE [28], a system that can learn associative rules between contexts such as 'if the user is not walking and is not driving, he/she is indoor'. In CenceMe [29], Miluzzo et al. show that activity classification accuracy varies when the phone is carried in the pocket or in the hand. This is why the system uses phone sensing context to drive the classification model by switching in the right technique to enhance the application performances. Moreover, the result of a context recognition application can be used a clue to reduce the energy consumption of the recognition system itself. For example, a human activity recognition system can reduce its energy consumption when the user is not moving or performs activities with low intensity [7, 30]. A place detection application can monitor the user mobility and location in order to take advantage of the existence of external devices in the environment to save energy [24, 27]. Finally, personal context recognition can improve the energy efficiency of an ambient context recognition system, while the environment recognition can certainly enhance the energy-efficiency of personal context recognition application. For example, in SensLoc [20], when the movement detector detects that the user is stationary, it switches off the path recognition application. Inversely, one can imagine that if a location recognition system detects that the user is at home, a vision-based system in charge of monitoring the user's personal context may be activated, thus relieving the wearable sensors.

3. General framework and requirements

In this section, we first describe the general structure of any human context recognition system, and we propose a generic architecture able to support such applications. Then, we present the main requirements of HCR applications.

3.1. General framework

General models for context recognition systems are given by Bulling et al. [31], Hoseini-Tabatabaei et al. [19], and Lara and Labrador [2]. The main stages for human context recognition include *sensing*, *signal processing* and *data analysis*, as illustrated in Figure 1. More specifically, sensing corresponds to the data acquisition at sensor nodes which produces raw data. Then, the data stream can be preprocessed to filter out signal variability. At the splitting stage, the processed data are split into time windows that are likely to contain an information. Afterwards, features extraction is applied to each window in order to obtain relevant characteristics of the signal, e.g. standard deviation, average, discrete Fast Fourier Transform (FFT) coefficients. During the training phase, the extracted features and corresponding ground truth labels are used as input to train a classifier model. During the testing phase, the extracted features and the trained model are used to identify the ongoing activity or current environment. Finally, depending on the recognized context, the application can provide the user with a feedback (e.g display a wellbeing score, raise an alarm).

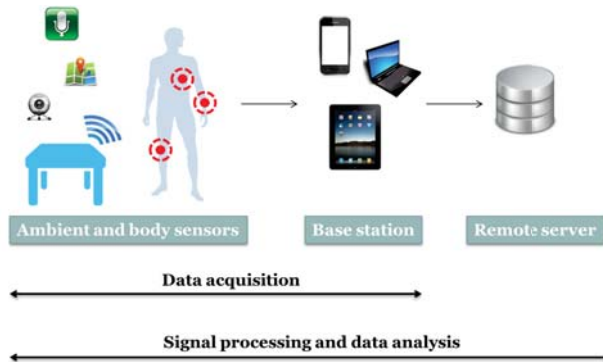


Figure 2: A generic human context recognition architecture.

To support this context recognition chain, a typical HCR architecture relies on the following devices: body sensors and possibly ambient sensors, a base station and possibly a remote server. This generic architecture for HCR is shown in Figure 2. In the first place, wearable sensors worn on the body measure attributes of interest such as ECG, motion or temperature. Ambient sensors (e.g. microphone and tagged objects) may also be used to improve the context-awareness of the application. The sensors communicate with a base station, which can be a laptop or a mobile phone, and whose main role is to run the healthcare monitoring application. In order to do that, the base station can communicate with a remote server on which the collected data are stored. Data acquisition is performed at the sensors, and at the base station when it has integrated sensors as is often the case of today’s mobile phone. The other tasks (signal processing and data analysis) can be conducted on any devices depending on the architecture. The different devices (sensors / base station / server) have various computing, memory and energy capabilities.

From the generic architecture presented in Figure 2, we identify different scenarios in Table 2. Indeed, depending on where the context recognition tasks are executed, the energy/performance profile of the body sensor network varies. For example, in *scenario 1*, the sensors have the possibility to directly store the raw data, leaving all the processing for later. The drawbacks of this method is that it prevents real-time feedback to the patient and it requires a lot of memory, even if to save space, the sensor can store extracted features. Another solution, in *scenario 2*, consists in sending the raw data to the base station, but it results in a huge amount of energy consumption due to transmissions. In the meantime, using a more powerful base station allows to have access to more complex algorithms and models. Furthermore, in *scenario 3*, the sensor can send only the extracted features or the compressed signal, which decreases communication cost but increases the computation load. Indeed, computations that run on a mobile device can also consume a considerable amount of energy: Milluzzo et al. [29] noticed that sampling the phone’s microphone and running a discrete Fourier transform on the sound sample uses more power than sampling the accelerometer and classifying the accelerometer data. In *scenario 4*, the entire context recognition process is performed on the sensor. From the energy point of view it may be optimal, but sensor scarce resources limit the implementation to lightweight algorithms. Actually, some classification algorithms are very expensive in terms of memory and computation requirements, which makes them not suitable for sensor platforms [2]. Therefore, there exist a trade-off between the accuracy that can be achieved through local computation and the additional power required for transferring the data. Besides energy consideration, when the context recognition is conducted on sensors, it could prevent designers from using these sensors as sources of information in opportunistic applications since the results will be application- or user-centric. In *scenario 5*, we present a more complex hierarchical architecture, where the sensor calculates extracted features and sends them towards the user’s mobile phone. Using these samples, the phone performs personal context recognition (e.g activity recognition), and sends the recognized activity, as well as its own raw data obtained from an integrated microphone, to a remote server, which performs the social context recognition.

Inferring the context locally has been emphasized to provide a number of advantages [29]. It presents


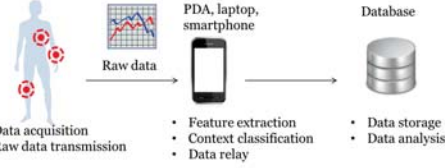
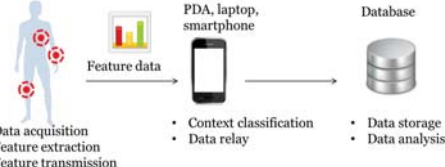
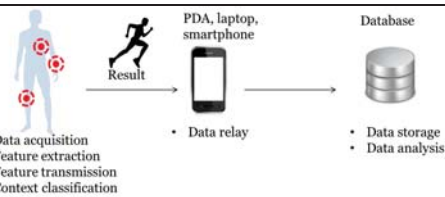
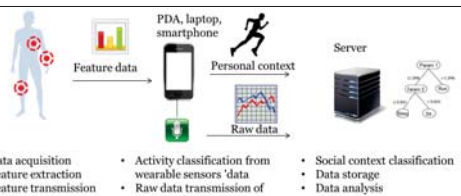
Scenario	Architecture	Advantages	Disadvantages
1	 <ul style="list-style-type: none"> • Data storage 	<ul style="list-style-type: none"> • no energy consumption for transmission and processing 	<ul style="list-style-type: none"> • memory needs • no real-time feedback
2	 <ul style="list-style-type: none"> • Data acquisition • Raw data transmission <ul style="list-style-type: none"> • Feature extraction • Context classification • Data relay <ul style="list-style-type: none"> • Data storage • Data analysis 	<ul style="list-style-type: none"> • no computation at sensor • access to more complex algorithms 	<ul style="list-style-type: none"> • increase energy consumption for transmission • increase energy consumption at the mobile phone
3	 <ul style="list-style-type: none"> • Data acquisition • Feature extraction • Feature transmission <ul style="list-style-type: none"> • Context classification • Data relay <ul style="list-style-type: none"> • Data storage • Data analysis 	<ul style="list-style-type: none"> • reduce energy for transmission at sensor 	<ul style="list-style-type: none"> • lightweight pre-processing algorithms at sensor
4	 <ul style="list-style-type: none"> • Data acquisition • Feature extraction • Feature transmission • Context classification <ul style="list-style-type: none"> • Data relay <ul style="list-style-type: none"> • Data storage • Data analysis 	<ul style="list-style-type: none"> • reduce energy for transmissions • increase privacy • increase resilience to wireless dropouts 	<ul style="list-style-type: none"> • increase energy for computation at sensor • not for opportunistic applications since results are application- or user-centric
5	 <ul style="list-style-type: none"> • Data acquisition • Feature extraction • Feature transmission <ul style="list-style-type: none"> • Activity classification from wearable sensors' data • Raw data transmission of integrated sensors <ul style="list-style-type: none"> • Social context classification • Data storage • Data analysis 	<ul style="list-style-type: none"> • possibility to merge different sources of information • opportunities for resources sharing • allow complex processing 	<ul style="list-style-type: none"> • need for security • trade-off between local and remote computing must be carefully studied • dependent on network access

Table 2: Possible architectures for context recognition.

resilience to cellular or WiFi dropouts and minimizes the data transmitted to the backend server, which in turn improves the system communication load efficiency. It also protects user privacy and the data integrity by keeping the raw data on the mobile device. However, the complexity of the applications is limited by the available local computing and storage resources. Uploading data to a mobile phone or a backend server provides better opportunities for the exploitation of aggregate data from a large number of users, while allowing for the realization of more complex applications. However, it demands more serious considerations of privacy, and requires a careful study of the additional delay and power consumption needed for the remote context delivery. As we can see, context recognition architectures often leverage energy consumption for accuracy and latency performances, and vice versa.

Finally, when considering multiple sensors communicating with a mobile phone, two alternative approaches are generally envisioned: *pull* and *push* models. In push approaches, the sensors simply continuously transmit their samples (or extracted features) to the smartphone [32]. On the one hand, this technique increases the accuracy by maximizing the amount of acquired data. On the other hand, continuous transmission of high data rate streams can often impose impractical traffic loads on existing wireless technologies likely to be used between the sensors and the mobile gateway (e.g. ZigBee, Bluetooth). Although sensors can take decision on whether to transmit or not the data based on internal processing, such as accelerom-

eter classification, the data acquisition process is generally decoupled from the actual context recognition. Inversely, in pull services, the mobile phone acts as a coordinator, and carefully chooses the order and the numbers of data items to request from each individual sensor [4, 33]. These solutions spare bandwidth and allow implementing more fine-grained optimization since the device retrieve information from sensors depending on the application needs. In the same time, the latency will certainly be increased.

3.2. Requirements

Health-related application making use of wearable sensors have specific requirements, and designers of energy-efficient user context recognition systems face challenges associated with the trade-off between high recognition accuracy, low latency and low power consumption.

The energy consumption is decisive in order to ensure long term monitoring of the patients and to foster the acceptance of these technologies by the patients. Generally, most of the energy drainage is due to the communications, followed by the sensing and computation tasks. When considering the general architecture depicted in Figure 2, it is important to notice that both the sensors and the base station may be energy constrained which currently dramatically limit the use of context monitoring applications. Traditionally, powerful laptop were used as base station, but for mobility consideration most of the works now consider mobile phone as base station. Although phones are more powerful devices compared to wearable sensors, the truth is that they also face energy consumption problems [34]. Thus, we will consider studies that are interested in either the energy consumption of the sensors or the base station (e.g. mobile phone). Note that few approaches aim to tackle the energy problem of both devices [6, 26].

The latency corresponds to the system reactivity, that is the time that elapsed between the beginning of an activity and its detection by the system. It includes the time required to acquire, process and analyze the data. The latency is crucial for some critical applications such as epilepsy seizure detection and sudden fall detection. In this case, low detection latency is the chief concern so that less attention may be paid to the node battery charge.

The recognition accuracy characterizes how well the system detects and identifies current contexts from sensor data. This usually requires that the recognized context be compared to ground truth labels representing the actual environment and activity being performed. Besides correct classification, context recognition systems can miss, confuse or falsely detect events that did not occur [31]. Therefore, the *recognition accuracy* is defined as the number of correct predictions over the total number of predictions [6, 8]. There exist several other metrics to evaluate the performances of a recognition system, such as precision, recall or F-Measure [35] but most of the reviewed architectures consider the accuracy.

We used the data of SociableSense [24] to represent in Figure 3 the latency and accuracy achieved by this architecture in function of the energy consumption. This behavior illustrates the traditional energy-latency-accuracy trade-off faced by context recognition applications. As we can see, when the energy decreases, the latency increases and the classification accuracy degrades. This is because energy-efficient mechanisms spare resources to increase the devices lifetime.

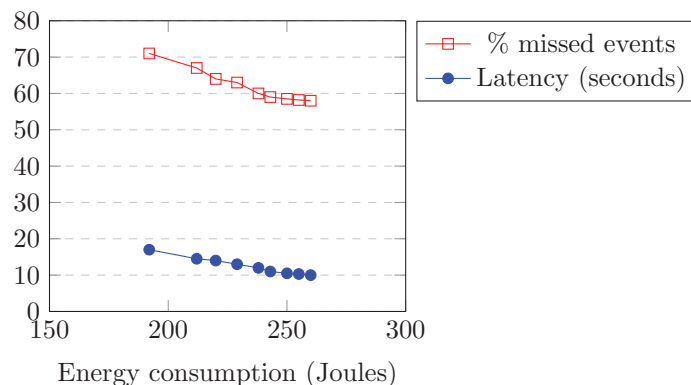


Figure 3: The energy-latency-accuracy trade-off of an activity recognition application [24].

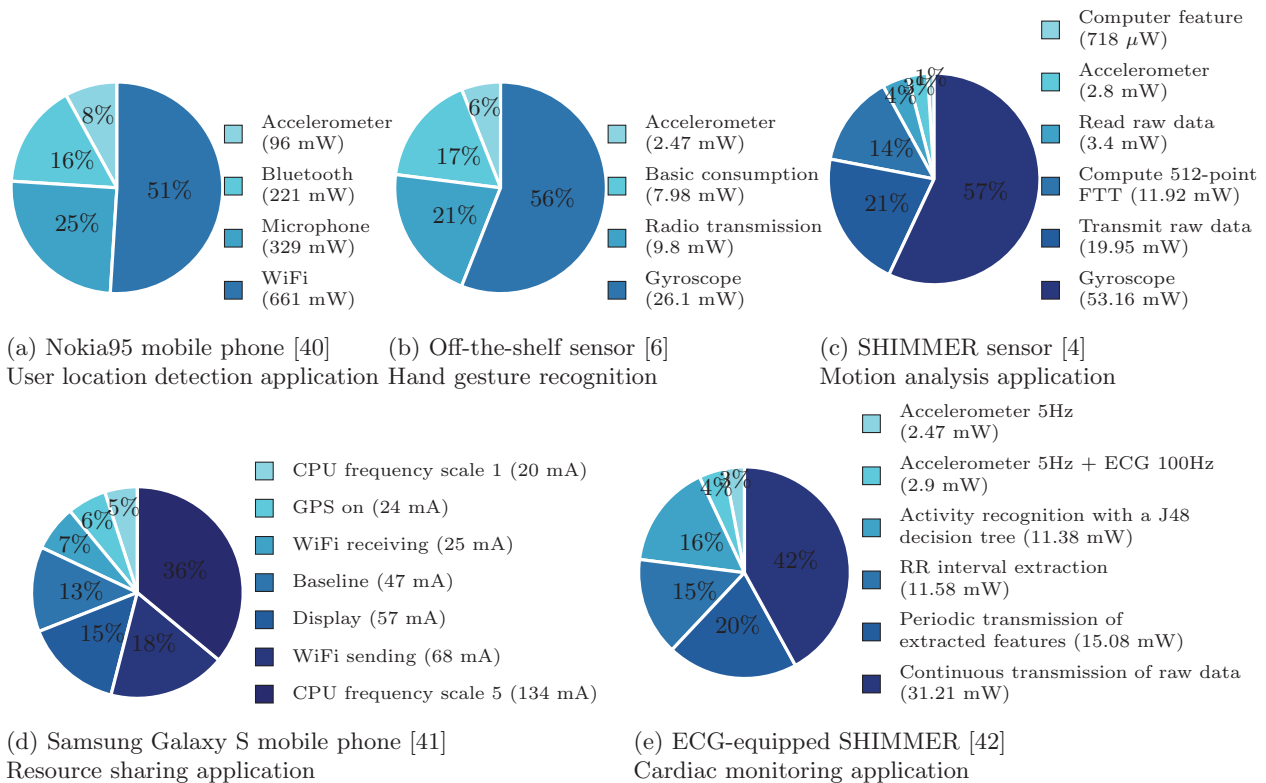


Figure 4: Energy consumption profile of different devices.

4. Classification of energy-efficient mechanisms

Energy conservation in wireless sensor networks (WSN) in general, and in wearable computing in particular has attracted many research works that have already been surveyed. For WSN, in [36], Anastasi et al. present a valuable taxonomy of energy-conservation schemes focusing on duty cycling and data-reduction approaches. There also exist surveys that concentrate on only one energy-efficient mechanism (like energy-efficient routing protocols or data aggregation techniques [37, 38]) since every category of solution often represents a whole research area in itself. Regarding wearable computing, Lane et al. [39] review many challenges related to mobile phone sensing, and Hoseini-Tabatabaei et al. [19] study smartphone-based user context recognition, but they do not focus on energy-efficiency. Our survey is different since we focus on energy-efficient solutions in the sense that they trade-off energy consumption for latency and accuracy. This is a central characteristic of human context recognition systems since they require low data delivery delay, high accuracy and long-term monitoring capacities. Moreover, we decided to concentrate on recent studies to provide the reader with up-to-date references. Before discussing our classification of energy-efficient mechanisms developed for human context recognition applications, we highlight the main sources of energy consumption in sensor nodes and energy-constrained base station.

We represent the energy consumption profile of different devices, including sensors and mobile phones in figures 4a to 4e. As can be seen, the main draining tasks are sensing, communication and computation operations. Indeed, we can observe in figures 4b and 4c that the sensing task due to power-hungry sensors such as gyroscope can consume a large amount of energy, dominating radio transmissions. This is why some solutions seek to minimize the wakeup time of sensor nodes. The communications also represent a high source of energy expenditure in mobile phones and sensor nodes. For example, in Figure 4a the Wi-Fi communications consume six times as much energy as the sensing task on accelerometer. Similarly, in Figure 4e the continuous transmission of raw data drains twelve times as much energy as sampling the

accelerometer. Therefore, many existing solutions aim to reduce the amount of exchanged data between devices. Finally, we can observe that it is worthy to reduce the computation load, especially on mobile phones since the CPU energy consumption at high frequency can represent a non-negligible fraction of the energy consumption as illustrated in Figure 4d. Regarding sensor nodes, it has been observed in [4] that computing certain feature on sensor node is relatively inexpensive since the additional CPU energy cost can be compensated by the reduction of packet transmissions. However, it does not make sense to compute more computationally-demanding feature. For example, it appears in Figure 4c that the computation cost of a Fourier transform is of the same order of magnitude as the cost of transmitting raw data. In addition to these considerations, battery charging techniques scavenge energy from the environment or human activities, such as lower-limb motion or body heat, which have great potential to become sustainable power sources for wearable sensors.

Based on these observations, our proposed classification of energy-efficient mechanisms for HCR systems is summarized in Figure 5. We have distinguished four main categories of energy-efficient solutions, namely, *power-on time reduction*, *communication reduction*, *computation reduction* and *battery charging*. In the next sections, we survey these different approaches. For each technique, we provide a description of its principle, followed by a review of related solutions.

5. Power-on time reduction

In this section, we survey power-on time reduction schemes whose aim is to adapt the sensor operations in order to save energy by switching the nodes to the sleep mode. The radio, CPU and sensing unit are all turned off to reduce the energy spent for communication, computation and data acquisition.

5.1. Sensor set selection

The objective of sensor set selection techniques is to achieve a good trade-off between the number of sensors activated and the classification accuracy. Indeed, the usage of many sensors may increase the recognition accuracy at the expense of additional energy consumption at the node level. Therefore, sensor set selection activate a subset of nodes depending on their contribution to keep a desired classification accuracy. Sensor set selection can be static [5] or dynamic [49]: the static approach makes the selection prior to deployment; the dynamic solution performs the selection at runtime, depending on some parameters such as the nodes residual energy, the required accuracy or the user activity.

Zappi et al. [49] developed a dynamic sensor selection scheme that selects a new cluster of sensors each time a sensor runs out of energy or fail. The cluster must be able to achieve a minimum desired classification accuracy and is such that even if a node fails, the other active sensors are still able to maintain the given accuracy. Noshadi et al. [5] consider the problem of selecting a subset of sensors, called predictors, so that it is still possible to infer the data of the not-selected sensors through prediction. These two selection schemes assume sensor redundancy (e.g Noshadi et al. consider up to 19 sensors attached to the arms of a user, Zappi et al. applied their method to a plantar pressure system with 99 sensors), but in body sensor network this assumption may not be practicable for sake of wearability. In SeeMon [52], Kang et al. dynamically select a small subset of sensors, called Essential Sensor Set (ESS), sufficient to answer the registered queries. Gao et al. [50] consider multi-modal sensing network for activity recognition and propose a framework that selects in real-time a subset of sensors to reduce energy consumption due to data transmission. They use convex optimization to minimize the number of active nodes while guaranteeing that the probability of misclassification using this subset is under a given threshold. Gordon et al. [51] use prediction to identify the future activities which are likely to occur and adapt the sensor subset accordingly. In a training phase, they first evaluate the dependency of each activity on the sensors, where the degree of dependency is the loss in accuracy compared to the usage of every sensor. At runtime phase, based on the classification history, future activities which are likely to occur are predicted. Then, only a subset of sensors is activated to distinguish the likely activities at little or no recognition cost.

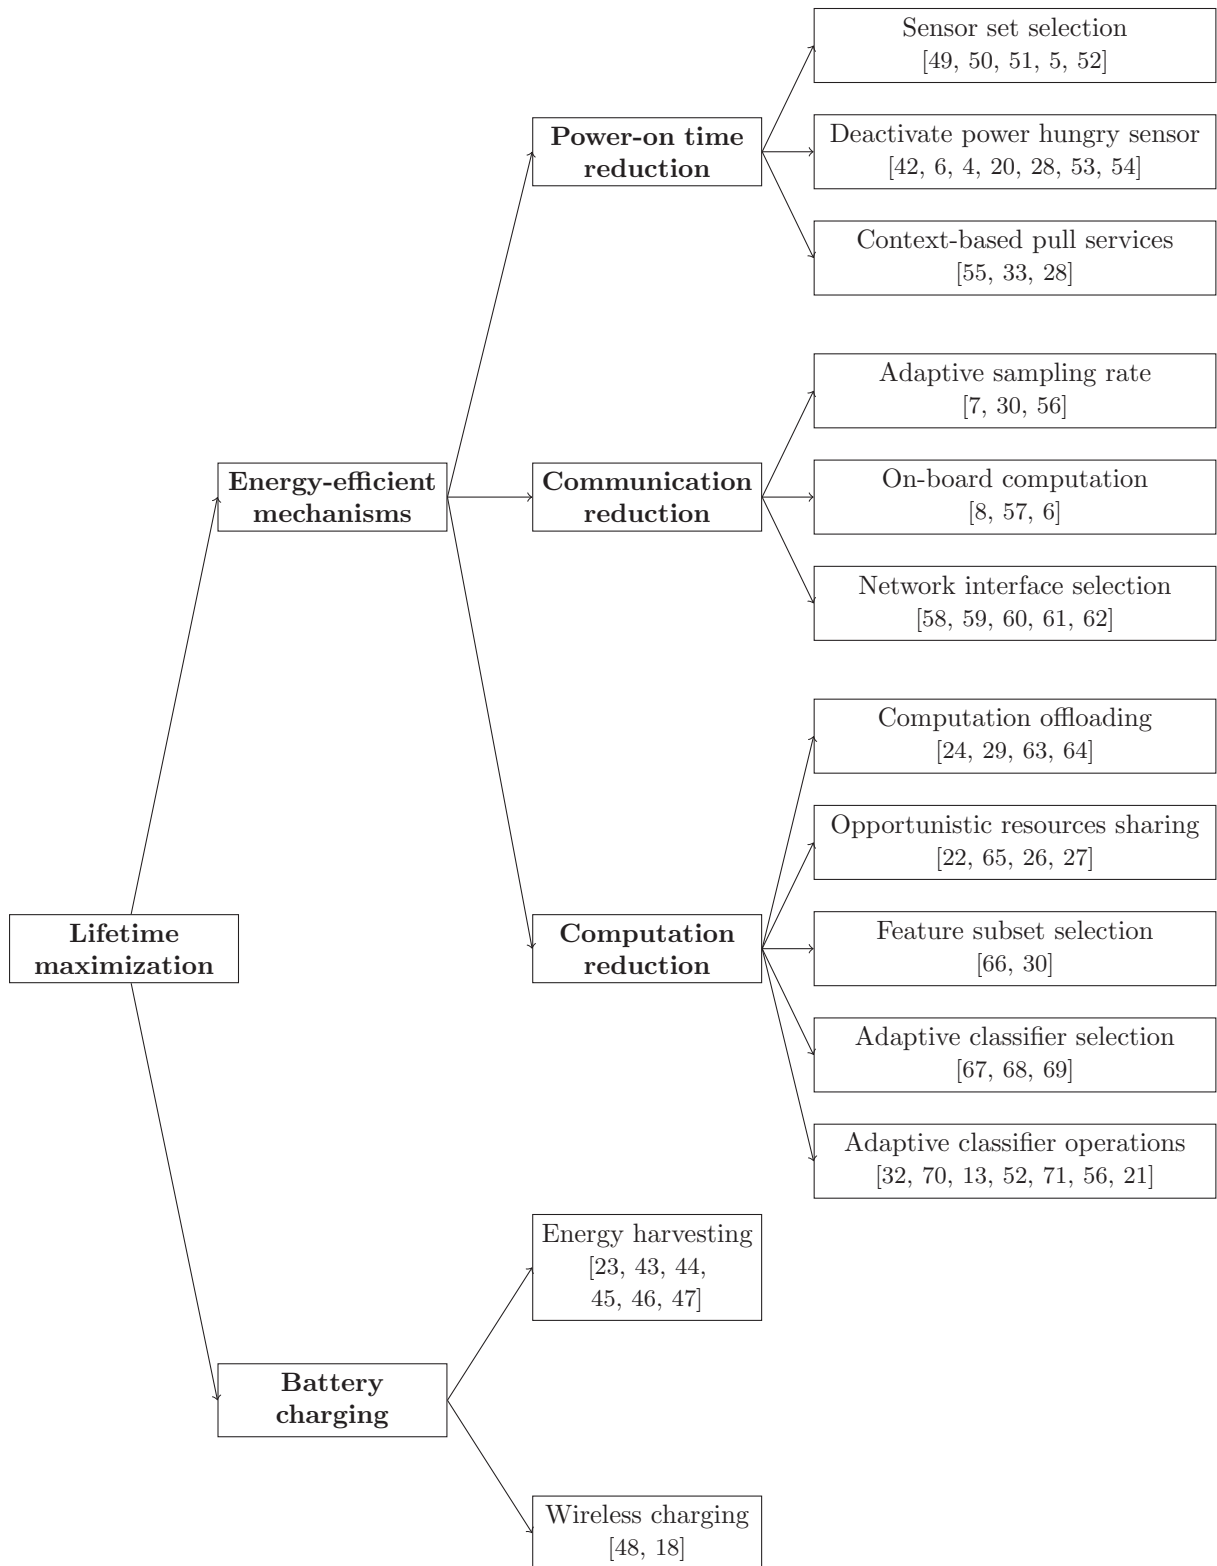


Figure 5: Classification of the energy-efficient mechanisms and battery-charging mechanism for the lifetime maximization of human context recognition applications.

5.2. Deactivate power-hungry sensors

The sensors have different energy consumption profiles, so it is worthwhile to deactivate the power-hungry devices when possible. Indeed, in [4], Lorincz et al. compare the energy consumption of different sensors and they measure a factor of 19 between sampling gyroscope (53163 μJ) compared to sampling accelerometer (2805 μJ). Thus, designers may consider disable the gyroscope when the user is performing activities with low intensities and turn off the GPS when the user is sitting or standing still [2].

Several solutions wake up power-hungry sensors on behalf of cheaper accelerometer measurements. Sun et al. [42] consider an ECG-based application where high physical activity can confounds the signal. They propose the PEAR (Power Efficiency through Activity Recognition) architecture, which runs an on-board activity classification algorithm that determines whether the user is active or not by processing accelerometer data. If the user is in active state, the ECG stops sampling, transmitting and switches off its radio. In inactive mode, ECG data are sampled at 100 Hz and features are extracted from these data every 5 s and sent to the base station. Park et al. [6] propose the E-Gesture architecture for hand gesture recognition using hand-worn sensor device and a mobile phone. Their objective is to reduce the power-on time of the gyroscope while achieving high recognition accuracy under dynamic situations. Thus, the accelerometer continuously performs gesture recognition, and when it detects a gesture, it wakes up the gyroscope. Then, if the gyroscope confirms that the user is performing hand gesture, the sensor sends the data to the mobile phone. Lorincz et al. [4] proposed Mercury, a wearable sensor network platform for neuromotor disease studies, which uses accelerometer data to determine whether the node is moving or not. If the node is not moving, power-hungry sensors like gyroscopes are disabled and accelerometer data are dropped. Kim et al. present SensLoc [20], a place detection system which employs a movement detector to control active cycle of power-hungry sensors. Accelerometer data are used to awake WiFi beacons scans for place detection, while the GPS is awoken when the user leaves a place. The Energy Efficient Mobile Sensing System (EEMSS) proposed by Wang et al. [53] activates high-resolution power-hungry sensors only when low-consumption ones sense an interesting event. With a different frame of mind Laerhoven and Gellersen [54] suggest to replace a power-hungry gyroscope with nine cheaper and low-consuming accelerometers. They show that by increasing the number of sensors while decreasing their individual accuracy, it is possible to lower power consumption. Similarly, when the locality of a user is required, one can take advantage of the energy-accuracy trade-off between different techniques, where energy consumption increases from WiFi-based localization to GPS schemes, while the accuracy decreases from GPS to WiFi methodologies.

The solutions that consist in deactivating power-hungry sensors adapt the sensors operations to the user's context. Thus, the energy gain is dependent on the user activity profile. For instance, PEAR is designed in such a way that more energy is saved when the user is active because the solution switches off the ECG during high physical activity as it can confounds the signal. In contrary, Mercury saves more energy when the user is inactive as power-hungry sensors are disabled. Similarly, SensLoc saves more energy when the user mobility is low because both GPS and WiFi scans are deactivated.

5.3. Context-based pull services

An alternative approach towards reducing energy consumption exploits the correlation between contexts. Instead of continuously collecting data from sensors, these solutions decide whether or not acquire more data depending on statistical correlation across different forms of contexts.

The Acquisitional Context Engine (ACE) [28] uses intelligent sensor caching and activity correlation to reduce power consumption by using already inferred attributes or choosing less power hungry sensors to infer correlated activities. For example, when the system already knows that the user is driving and has learned the rule: 'if the user is driving, the user is not at home', ACE can deduce that the user is not at home without acquiring any sensor data. Similarly, ACE can first probe less energy consuming sensor – accelerometer – to determine whether the user is driving or not, and if the result imply that the user is not driving, it turns the GPS on and infers the actual user's location. Nath argues that ACE can be complemented with SeeMon [52] to integrate the temporal continuity of context in order to extract sophisticated rules such as 'if

a user is in his/her office now, he cannot be at home in next 10 minutes’. MediAlly [55] adopts an Activity Triggered Deep Monitoring (ATDM) paradigm as an energy saving mechanism, where medical data streams are collected and relayed to a central server only when monitored context information is evaluated and satisfies given predicates. As an example, MediAlly can collect ECG data whenever the subject’s context indicates physical or emotional stress. The ACQUA (Acquisition Cost-Aware Query Adaptation) [33] system avoids retrieving unnecessary sensor data by sequentially acquiring the data which has the highest probability to determine the result of the query. This choice is based on a-priori knowledge of energy consumption costs and probabilities, which can be inferred based on historical traces obtained for previous query executions. These approaches can save energy in the case of repetitive context, but are not applicable to detect more spontaneous contexts.

6. Communication reduction

In this section, we review techniques that aim to reduce the data communication costs, either by reducing the amount of data to be transmitted in the network or by carefully selecting the communication technology.

6.1. Adaptive sampling rate

In [72], Maurer et al. studied the impact of the accelerometer sampling rate on the recognition accuracy and they found that no significant gain in accuracy is achieved above 20 Hz for ambulation activities. Therefore, depending on the user activity and the type of sensor, it may not be necessary to sample at the maximum rate in order to achieve a desired accuracy. The key idea behind adaptive sampling rate is to adapt the frequency of data acquisition depending on the recognized context. This technique leads to fewer unneeded samples, which will result in less data to transmit.

In A3R (Adaptive Accelerometer-based Activity Recognition), Yan et al. [7] propose to adapt both the sampling rate (SR) and the classification feature (CF) in real-time. For each activity they compute offline the accuracy that can be achieved for different tuple $\langle SR, CF \rangle$ of sampling rate and classifier feature. Then, depending on the strategy, they determine per-activity the optimal tuple that can achieve a user-specified accuracy, or the tuple that have the highest ratio of accuracy over energy consumption. At runtime, when an activity is recognized, the application selects the corresponding tuple $\langle SR, CF \rangle$ which enables to tradeoff energy and accuracy. Afterwards, in [30], Qi et al. introduced the idea that it requires a lower sampling rate to determine whether a specific activity happens or not than classifying multiple activities. Thus, their AdaSense architecture periodically decides whether the current activity changes or not through single activity event detection with a lower sampling rate. If a change is detected, AdaSense switches to the multi-activity classification at a higher sampling rate. Then the solution performs event detection of the new activity at a lower sampling rate. The major drawback of these solutions is the lack of flexibility as configurations are selected once and for all for a given activity, which requires a new training phase when considering new activities. In order to reduce the energy consumption and keep a low localization error, Liu et al. [56] consider that activities associated with faster speed need a higher GPS sampling rate and vice versa. The Jigsaw solution dynamically learns the optimal GPS duty cycle scheme given the remaining battery budget level, the expected duration of the tracking application and the user’s runtime mobility level provided by the accelerometer. Jigsaw comprises a set of sensing pipelines for the accelerometer, microphone and GPS sensors, which are built in a plug and play manner to support: i) resilient accelerometer data processing, which allows inferences to be robust to different phone hardware, orientation and body positions; ii) smart admission control and on-demand processing for the microphone and accelerometer data, which adaptively throttles the depth and sophistication of sensing pipelines when the input data is low quality or uninformative; and iii) adaptive pipeline processing, which judiciously triggers power hungry pipeline stages (e.g., sampling the GPS) taking into account the mobility and behavioral patterns of the user to drive down energy costs.

6.2. On-board computation

On-board computation is used to reduce the amount of data to be transmitted towards the base station. The rationale behind this method is that computation requires less energy than communication. For example, a node can perform signal processing tasks such as feature extraction in order to transmit the result instead of raw data. Similarly, compression encodes information in such a way that the number of bits needed to represent the initial message is reduced. This is energy-efficient because it reduces transmission times as the packet size is smaller. Another possibility that can be exploited jointly is to send burst of data instead of one data per packet.

Wang et al. [8] present a hierarchical architecture for real-time activity recognition, composed of accelerometer body sensors, two RFID readers positioned in each hand to detect handling of tagged objects and an UHF RFID reader in each room to locate the user who wears a UHF tag. The system operates in two stages. First, the accelerometers process their data to perform gesture recognition (e.g. right hand moving up, sitting down) at the node level. Then, recognized gestures, tagged objects and user locations information are transmitted to a PDA device, which runs the activity recognition (e.g. eating, making coffee, shaving) algorithm. Energy is saved at the sensor level thanks to reduced communication overhead, as nodes only transmit 1 byte data containing the gesture label every second. They demonstrate that the hierarchical approach enables to reduce the communication cost by 60% compared to a solution where all the sensor readings are transferred to the base station. The E-Gesture architecture [6] reduces packet transmissions by sending data only when a gesture is detected by both the accelerometer and the gyroscope sensors. Wang et al. [57] compare the energy/performance trade off between an on-node scheme where the classification is performed on the accelerometer, and an off-node scheme where the accelerometer sends its raw data to the base station. They show that with the on-node scheme 40% energy can be saved with 13% reduction in classification accuracy. So, their results suggest that it is more energy-efficient to perform all the tasks on the sensors. However, computation offloading is sometimes more energy-efficient than local computation without loss in accuracy (see subsection 7.1).

6.3. Network interface selection

Sensors and mobile devices are often equipped with multiple network interfaces, such as ZigBee, Bluetooth, WiFi, and 3G. Each communication technology has different characteristics in terms of bandwidth availability, service area, economic costs and energy-consumption pattern. Thus, it is necessary to efficiently select the network interface in order to extend the lifetime of the mobile devices. This selection depends on the environment opportunities, that is, on the connection availability and on the application requirements (bandwidth, link quality, energy). Energy-efficient radio selection protocols have been proposed for both mobile phones and sensors.

In [59], Kim et al. privilege WiFi communications because they are more energy-efficient than 3G. Since 3G networks have wider service areas, their objective is to detect rapidly available WiFi access point while limiting the number of power-consuming WiFi scans. Similarly, Rhamati and Zhong [60] are interested in estimating the WiFi availability without turning on the interface because the energy cost for establishing a connection is very high, while it consumes less energy than cellular networks. They propose different selection policies based on the time, history, cellular network conditions (available access point, signal strength), and device motion. In [58], Paschou et al. select the most suitable data transfer technology to be used (e.g. WiFi, 3G, SMS), depending on the volume and type of health data to be sent. The authors minimize the economic cost for the customer, but their work can be extended to consider other metrics.

Some studies consider sensors equipped with two radios: a low-power, low throughput radio, and a high-power, high throughput radio. Indeed, higher throughput radios have a lower energy-per-bit cost, but they also have higher start up time and cost. Therefore, high bandwidth radios become more energy efficient only when a large number of bytes have to be transmitted, which compensate the wakeup energy overhead. The exact number of bytes above which a radio becomes more energy-efficient is called the *break-even point*. Nonetheless, not every healthcare application requires such a high data rate. Moreover, the energy

consumption of high-power radio when idle is prohibitive, so it is imperative to turn it off when not in use. The transition overhead due to start-up cost can be amortized by buffering a minimum amount of data. In this context, Sengul et al. [62] propose the Bulk Communication Protocol (BCP) where the source node accumulates data in its buffer, and once it reaches a threshold (superior to the break-even point), it sends a message to the receiver over the low-power radio to ask him to turn on its high-power radio. Loseu et al. [61] accumulate data at local nodes and transmit them in bulk while not violating the timing constraint. The drawback of these solutions is the delay induced by the buffering process. Therefore, for strict deadline (e.g 20ms – 1s) healthcare applications it is preferable to use a low-power low-throughput radio since there is no time to accumulate enough packets in buffer.

Multi-radio devices offer new opportunities to save energy at both the sensors and the base station. However, even though the break-even point is indicative of the relative energy efficiency between two radios, it does not capture every parameters. For example, the radios may have different transmission range which will require more or less relay nodes depending on the used interface. Similarly, all the radios may not be available at the same time, or they may experiment different packet loss and interference profile, which directly impacts the energy consumption and robustness. Thus, there is a lot of trade-off to be explored between energy-consumption, latency and robustness when designing network interface selection protocols in multi-radio systems [73].

7. Computation reduction

In this section, we focus on mechanisms that target the reduction of computation costs, as human context recognition applications necessitate power-hungry and complex algorithms to process the signals and perform the classification task.

7.1. Computation offloading

For a node, computation offloading strategy consists in offloading a part of its computation tasks to another device. For example, a sensor can execute (a part of) its tasks on a more powerful base station (e.g. laptop, PDA), while a base station can offload its tasks on a back-end server. Thus, intensive calculation are performed on nodes with higher resources. This approach has proven to be energy-efficient under some conditions [74, 24, 27], and the usage of devices with different computing, storage and energy capabilities allows incorporating more complex algorithms and models. Thus, computation offloading leads to a trade-off between energy, latency and recognition accuracy.

In [24], Rachuri et al. propose SociableSense, an architecture dedicated to activity recognition for sociable interactions, where a phone with integrated sensors can offload its tasks to the cloud. SociableSense dynamically decides whether to perform the computation locally on the phone or remotely on back-end servers by maximizing a weight additive utility function that integrates the energy consumption, the latency and the amount of data sent over the network. The results show that it is more energy-efficient to offload the speaker identification task to the cloud when using the phone’s microphone. This is also the configuration that achieves the lowest latency. CenceMe [29] performs audio and activity classification on the phone, while higher level forms of context such as social interaction and location are classified on a remote server. Liu et al. [63] developed a localization technique based on GPS sensing that offloads some calculations to the cloud. With few milliseconds of raw satellite’s data, the server can estimate the device’s past locations by exploiting information from public, on-line databases. Therefore, uploading data on a backend server can ease the execution of burdensome tasks when an appropriate strategy controls the impact of the communication load. However, custom-made application execution partitioning, such as the one used in CenceMe and SociableSense, requires significant effort from the developer’s side. More general solutions allow an application developer to delegate the partitioning task to a dedicated middleware. MAUI [64], for example, supports fine-grained code offloading to a cloud in order to maximize energy savings on a mobile device.

7.2. Opportunistic resources sharing

High density of sensing devices including mobile phone, wearable and ambient sensor networks has accelerated the rise of innovative collaborative sensing applications. Example of such application is Ear-Phone [71], which constructs a noise map from samples obtained by crowdsourcing data collection. On the one hand, the presence of other devices in close proximity to the user can be employed to increase the sensing capabilities of a unique device. On the other hand, many opportunities for establishing opportunistic connections between devices can be used for energy-efficient resources sharing.

The Remora architecture introduced by Keally et al. [26] exploits the physical proximity of several users to share resources among neighbouring Body Sensor Networks (BSN). Indeed, colleagues, family members or sport team may spend time together performing the same activity, in such a way that their respective BSN can overheard data transmission and collaborate. In this situation, if the proximity duration is long enough, resources sharing can provide an energy benefit. Remora implements a collaborative sensor node selection to improve accuracy while disabling as much sensors as possible, and duty-cycles classifiers so that one phone make the classification while other phones go into sleep mode. The CoMon platform [22] also addresses the energy problem through opportunistic cooperation among nearby users. CoMon first discovers nearby devices by periodic Bluetooth scans and selects candidates that will potentially stay in the vicinity for long period. This is to ensure that the cooperators will stay longer enough to breakeven the overhead due to the initiation of the cooperation. Then the devices negotiate a cooperation plan than will provide a mutual energy benefit. The authors highlight the difficulty of providing a fair negotiation process taking into consideration available sensing devices and their residual energy, user policies, and different sampling rate or accuracy requirements. ErdOS [65] is a mobile operating system that extends the battery life of mobile handsets by exploiting opportunistic access to resources in nearby devices using social connections between users. Notably, they distribute the sensing burden of shared ambient context by sharing GPS reads with nearby devices. Rachuri et al. [27] propose to offload some sensing tasks from user phone to remote sensors embedded in the environment. Their algorithm, called METIS, dynamically distributes the sensor tasks between the user phone and the infrastructure sensors based on the sensing power, sampling rate and user mobility. METIS first performs a discovery of the sensing devices that are available in the environment and estimates the minimum amount of time the phone should suspend local sensing and use remote sensing to achieve energy gain.

Opportunistic resource sharing approaches are particularly suitable for *environmental* context monitoring, that is for contexts that are common to multiple neighboring nodes (e.g. noise level, speed movements). In this case, incentive mechanisms are likely to be necessary for crowdsourcing applications in order to enable collaboration between unknown users, and to obtain the requisite data. With another point of view, Rodriguez et al. [65] argue that if users notice a personal benefit in terms of energy savings and usability by sharing their resources in the long-term with subjects they personally know and trust, there is no need to implement potentially complex or costly incentive schemes to enforce cooperation.

7.3. Feature subset selection

Feature selection is generally used to select the features that are most discriminative and contribute more to the performances of the classifier. There exist a trade-off between delay and accuracy since using a lot of features certainly enhances the accuracy but it is energy and computation consuming. Inversely, reducing the number of features allow saving energy and reducing computation at the cost of reduced accuracy.

Cilla et al. [66] proposed a Pareto-based Multiobjective Optimization that selects feature subsets whose size is minimize while maximizing the accuracy of the classifier. Their approach provides the designer with diversified solutions closed to the Pareto-front. The advantage is that the expert does not have to set a priori tradeoff between accuracy and number of feature, but he will choose a posteriori the subset that best fits the application requirement knowing that these solutions are closed to the optimal. In AdaSense [30], Qi et al. design the offline SRO (Sampling Rate Optimization) algorithm that selects the optimal sampling rate and feature set of both single activity event detection and multiple activity classification under any

accuracy requirements. It is a genetic programming based algorithm that explores the feature set with the objective to minimize the sampling rate while achieving a given accuracy.

7.4. Adaptive classifier selection

There exist different types of classifiers based on Decision trees, Bayesian approaches, Neural Networks, Markov models and so on. They differ from each other in terms of delay, memory size and computation requirements. Indeed, lightweight decision tree methods are usually preferred for embedded systems while distance-based classifiers that measure a similarity between instances stored on a database are not convenient to be implemented on a mobile device. Similarly, ensembles of classifiers that combine the output of several classifiers to improve the classification accuracy are computationally expensive. Thus, every classifier has different energy, delay and accuracy properties. This is why some works propose to adapt the classifier to the available resources and application requirements, either in a static or dynamic way, in order to improve the system performances.

Martin et al. [67] developed a fuzzy-based on-line classifier selection that selects the best-suited classifier given *i*) offline classifier performance evaluation (trained accuracy, delay, memory need, complexity) *ii*) online accuracy and delay requirements (low, medium, high) and *iii*) the current device state (battery level, available memory, CPU load). In practice, given the current application requirements and device state, the fuzzy selector outputs a score indicating the desired quality level of the classification (i.e the desired accuracy, response time, complexity and memory use of an ideal classifier). Then, a distance-based algorithm selects the best classifier in the database that fulfils these scores. The results show that this approach enables to periodically choose the classifier that best suits the application requirements and device context. However, the difficulty of this study is to define the fuzzy rules used to score classifiers quality. Chu et al. [68] have developed the Kobe toolkit that performs profiling and optimization of mobile embedded classifiers to achieve an optimal energy-latency-accuracy trade-off. The developers have to specify the accuracy, latency and energy constraints. Then, Kobe tunes the classifier parameters (e.g. sample windows size, number of FT sample points). Kobe also adopts computation offloading at runtime: it decides whether to partially or entirely offload the computation to the cloud in response to system changes (phone battery levels, network bandwidth and latency, processor load). Their results show that for comparable levels of accuracy, traditional classifiers suffer between 66% and 176% longer latencies and use between 31% and 330% more energy. Orchestrator adaptively changes sets of features and classifiers for context recognition depending on current set of available sensors and pre-defined system policy [69].

7.5. Adaptive classifier operations

In traditional architectures, classifiers are usually triggered periodically. However, it is possible to adjust the classifier behaviour depending on the recognized context in order to reduce the computation time due to classification tasks. For example, some solutions propose not to run the classifier if the sensed signal is not accurate or if the user does not require a close monitoring. It is also possible to reduce the frequency at which the classification is performed depending on the recognized activity.

Au et al. [32] introduce the episodic sampling technique which aims at adapting the classification rate depending on the activity variation. The proposed solution additively increases the interval between activity recognition when the classification result remains the same, and multiplicatively decreases the interval when a change in the activity is detected. It trades-off accuracy with energy as nodes can go to sleep between two classification episodes. Nevertheless, the algorithm may exhibit lower performances when context changes experience high dynamics which necessitates a careful parameters setting. Srinivasan and Phan [70] focus on reducing the amount of time the activity recognition system stays awake while guaranteeing a high confidence prediction of the current activity being performed. The idea is that differentiating between idle and non-idle activities requires a smaller window of accelerometer samples compared to identifying specific user's physical activities. Thus, their two-tier classifier first uses a small window size to determine if the user is idle or not; if the user is not idle, the two-tier classifier acquires more sensor samples to identify the user's

fine-grained activity. In [13], Lin et al. propose the BeWell+ architecture to continuously monitor user behaviour along three health dimensions: sleep, physical activity and social interaction. In order to save energy, the solution dynamically tunes the rate at which sampling, feature extraction and activity inference are performed, each time there is a change in the wellbeing scores. Thus, BeWell+ allocates energy to provide more accurate monitoring and responsive feedback for wellbeing dimension of highest concern, by giving priority on specific health dimension (e.g. sleep) where the user needs most help. In contrast, user with healthy behaviours are monitored less closely, with a feedback provided on a lower scale - therefore using less resources on the smartphone. As a result, the gain in energy-efficiency will depend on the user behaviour. That is, a user with poor wellbeing scores across all dimensions will consume more energy than a consumer who exhibits comparatively high wellbeing scores, allowing the adaptive scheme to lower energy used for these dimensions.

In order to spare computation resources, some researches aim at filtering out samples that are not likely to contain information at an early stage of the context monitoring process. SeeMon [52] makes early decision in the processing pipeline in order to significantly save computational overhead when no changes in the context are detected. SoundSense [21], EAR-phone [71] and Jigsaw [56] duty-cycle classifiers depending on the quality of sensed data: noise measurements that are not accurate enough for the target application are discarded. For example, Ear-Phone is a noise-mapping application that first determines the phone sensing context (in hand or in pocket / bag), and measurements that are detected as measured in the pocket or bag are not taken into account in the calculation of the user’s noise exposure. The Jigsaw’s admission control block throw away packets indicative of silence.

8. Battery charging

8.1. Energy harvesting

New technologies have been developed to enable sensors to harvest energy from their environment like sun, heat or movement [75, 76, 77]. Compared to traditional sensors, rechargeable motes can operate continuously and, theoretically, for an unlimited length of time. They convert ambient energy to electrical energy and then either consume it directly or store it for later use. For healthcare applications, body heat and body movement are two main sources of energy harvesting.

Body heat

The human body heat represents an interesting source of energy since the human metabolism generates and dissipates heat. Indeed, a human being generates more than 100W of heat on average, and about 1–2% of this heat can be used to obtain an electrical power of the order of milliwatts. At typical indoor conditions, the heat flow in a person mainly stays within the 1–10 mW/cm². In this conditions, powering of a low-power health-monitoring sensor by using energy harvesters become feasible. In [23], Leonov study the efficiency of a thermoelectric energy harvester integrated into an office-style shirt. The system can generate power in 5-0.5 mW range at ambient temperatures of 15-27 °C respectively. The author also highlights that the power generated by wearable thermal devices does not depends on the user metabolic rate, but on the overall body heat, air speed and sweating rate. So, during exercise the harvesting power will not necessarily increase. In [43], Leonov et al. are able to power a watch-style sensor module that measures and transmits the battery voltage, the temperature and/or light intensity every 2s. The thermal energy generator successfully delivers 100 μ W at 22 °C, which is enough to power the system as it requires 50-75 μ W to operate. As can be seen, given the small efficiency of thermal harvesting, only low-power applications (below 1-2 mW) can run efficiently on human body heat.

Body movement

Human activities, such as walking or running are potential sources of energy harvesting since mechanical vibrations can be converted into electrical energy by different mechanisms: electromagnetic, electrostatic, and piezoelectric. In [44], Gorlatova et al. quantify the inertial energy that can be harvested from the activities of 40 individuals over periods up to 9 days. The data were collected from three sensing units

located on the shirt pocket, waist belt and trouser pocket. It appears that 20-second walking and running motion can generate up to $202 \mu W$ and $808 \mu W$ respectively. Cycling activity generates only $52 \mu W$, but in this case harvester placements on the lower legs should be considered. The authors highlighted other main findings: Walking downstairs generates more power than going upstairs because of the higher accelerations involved. Taller people generate around 20% more power than shorter people. People are passive most of the time since 95% of the total harvestable energy is generated during less than 7% of the day. In their survey, Paradiso and Starner argue that the most promising way to harvest energy from people is by taping their gait [77]. Therefore, the shoe designed by Shenck and Paradiso [47], which integrates piezoelectric elements beneath a standard running insole, scavenges energy from heel strike. Each shoe produces sufficient energy to transmit a 12-bit code via an on-board radio as the bearer walked. In practice, it is possible to generate more power, but the devices become bulky and heavy. For example, in [45], Donelan et al. present an energy harvester that generates electricity during human walking with minimal user effort. The energy harvester is mounted at the user knee and is only engaged during the decelerating portion of the stride cycle. In this way, the system assists the leg motion at the end of the swing phase. The authors demonstrate that the metabolic cost, that is the additional metabolic power required to generate 1 Watt, is reduced by the generative braking action of the device. An output power of 5 W is obtained when the subjects wear one device on each leg during normal walking motion. In [46], Rome et al. exploit the mechanical energy produced by the vertical movement of a backpack during walking and convert it to electricity. They observe that generally the electrical power increases with walking speed and the weight of the load in the backpack. Their system is able to generate up to 7.4 Watts when the user is wearing a 38 kg backpack during fast walking. For the time being, the weight of the backpack (between 20 kg and 38 kg) is a limitation as only serious hikers or military personnel will wear such a load.

Mitcheson [78] compare the performance of both kinetic and thermal energy harvesters. Theoretically, kinetic harvesters achieve higher power density of around $300 mW/cm^3$ compared to around $20 mW/cm^3$ with a thermal device. However, given the current state development of kinetic and thermal devices, Mitcheson shows that for equivalent volume, current thermoelectric devices achieve greater densities. Moreover, it appears that kinematic generators are more dependent on the user mobility. As a consequence, if the user is immobile or if the sensor is not worn on a limb, the generated power would be lower.

8.2. Wireless charging

Recent breakthroughs in wireless power transfer are expected to increase the sustainability of sensor networks and make them perpetually operational, since these techniques can be used to transmit power between devices without the need of any contact between the transmitter and the receiver. Wireless charging in wireless sensor networks can be achieved in two ways: electromagnetic (EM) radiation and magnetic resonant coupling. Xie et al. [79] show that omni-directional EM radiation technology is applicable to a WSN with ultra-low power requirement and low sensing activities (like temperature, light, moisture). This is because EM waves suffer from rapid drop of power efficiency over distance, and active radiation technology may pose safety concerns to humans. In contrast, magnetic resonant coupling appears to be the most promising technique to address energy needs of WSNs thanks to an higher efficiency within several-meter range.

Wireless energy transfer has already been applied to power medical sensors and implantable devices. In [48], Damm et al. use inductive power supply in order to power *in vivo* hip prosthesis. The prosthesis is instrumented with 6 strain gauge sensors in order to measure the forces and moments acting in the hip joint prosthesis. An external coil is fixed around the patient implant during the experimentation in order to deliver the power of 5mW, required by system to acquire data and send them to the base station. Still, some challenges remain to be addressed: the energy transmission efficiency decays rapidly with the distance. Moreover, the system requires the usage of a large induction coil wired to an external device [80]. In a more pervasive manner, Gummesson et al. [18] designed a wearable ring for gesture recognition that opportunistically harvests energy from a NFC-enabled phone. The achieved scavenging performances depend on the relative positioning of coils, that is on the position of the user's finger when holding the phone.

Table 3: Energy harvesting opportunities for wearable sensors.

Power source	Reference	Power
Thermal	watch [43]	100 μW at 22 $^{\circ}C$
	shirt [23]	5-0.5 mW at 15-27 C
Kinetic	cycling [44]	52 μW
	walking [44]	202 μW
	running [44]	813 μW
	insole	1-2 mW per step
	knee [45]	5 W
	backpack [46]	7.4 W
Wireless charging	inductive charging [48]	5 mW
	wearable ring [18]	18 mW
Solar for comparison	indoor light [44]	50-100 μW per cm^2

Table 4: Power needs of different platforms.

Shimmer [82]	Hip prosthesis [48]	Watch [43]
70 mW	5 mW	50-75 μW

Magnetic coupling, kinetic and thermal energy sources have been studied for both external and in-vivo sensors. While wireless charging is more controllable, it requires an external device or installation to power the sensor. This can be inconvenient and trouble the activities of the patients, but efforts are being made to develop small-size wearable device that can be placed over the implantation area to power implanted sensors and communicate with them [81]. We summarize in Table 3 the power that can be harvested from different sources of energy. For comparison, we give in Table 4 the power consumption of the different platforms. As can be seen, the power that can be harvested is still low, so that for the moment these solutions can only power specifically designed sensors with low duty-cycle.

9. Summary and discussion

In the previous sections, we have detailed four main strategies for energy optimization, namely, *power-on time reduction*, *communication reduction*, *computation reduction* and *battery charging*. Ideally, a designer will need to apply strategies from all four categories to maximize the lifetime of his/her human context recognition system. Examples of studies that implement several mechanisms are Mercury [4], which combines power-on time reduction through deactivation of power-hungry sensors, with communication reduction through on-board processing; A3R [7] and AdaSense [30] architectures implement a per-activity feature selection to reduce computation, in combination with adaptive sampling mechanisms in order to reduce communications. However, depending on the application or on the selected architecture, it may not be possible to implement or combine some solutions. A cardiac monitoring application using a unique ECG sensor that transmits its data towards the user’s mobile phone will not use sensor set selection nor adaptive sampling. It will more likely employ both feature subset selection and on-board computation in order to extract relevant characteristics of the signal and to reduce the number of transmissions. Moreover, energy harvesting is not sufficient to solve the energy problem. Indeed, between two harvesting operations, nodes may deplete their energy, so it is important to implement energy saving mechanisms in addition to energy harvesting solutions. For example, the wearable ring [18] scavenges energy opportunistically when the user interacts with his/her mobile phone. As a consequence, the ring platform also integrates energy-efficient mechanisms such as deactivation of power-hungry audio sensor, and classifier selection. Such an energy-aware design helps to reduce harvesting needs, which in turn enables embedded devices to conduct more useful tasks with the limited power they can scavenge.

	Environmental context		Personal context	
	Patient location detection (coffee shop, home, office, ...)	Social interaction detection	Physical activities recognition	Physiological parameters supervision
Power-on time reduction	Deactivate power-hungry sensors <i>e.g. promote WiFi-based localisation over GPS</i> impact accuracy	Sensor set selection <i>e.g. select between Bluetooth traces, WiFi traces or microphone to detect users co-location</i> impact accuracy	Deactivate power-hungry sensors <i>e.g. deactivate gyroscope or ECG when the patient is inactive</i> impact accuracy	
Communication reduction	Network interface selection <i>e.g. as the user is moving, select the communication technology that minimizes the energy consumption</i> impact delay		Adaptive sampling rate <i>e.g. increase or decrease the sampling rate depending on the current activity</i> impact accuracy	On-board computation <i>e.g. perform some "light" computing and send only the result</i> impact accuracy
Computation reduction	Opportunistic resources sharing <i>e.g. in a limited geographical area, share its localisation information</i> impact delay	Computation offloading <i>e.g. in a social environment, offload computation in order to merge data collected from different sources</i> impact delay and accuracy Adaptive classifier operation <i>e.g. discard noisy measurements</i> impact accuracy	Adaptive classifier <i>e.g. adapt the classification rate to the user activity variation</i> impact accuracy	Feature subset selection <i>e.g. minimize the number of extracted features given an accuracy lower bound</i> impact accuracy
Battery charging	Wireless charging <i>e.g. recharge batteries thanks to devices available in the user environment</i>		Body movement energy scavenging <i>e.g. as user is moving scavenge energy thanks to harvesters located on lower limbs or shoes</i>	Body heat energy scavenging <i>e.g. scavenge energy thanks to devices located on user's skin or clothes</i>

Table 5: Applications and energy-efficient mechanisms.

In Table 5, for each type of identified applications (user location detection, social interaction detection, physical activity recognition and physiological parameters supervision), we give suitable approaches for these applications to save energy by specifying illustrative examples, and we also highlight the influence of the technique on the delay and accuracy requirements. Therefore, we guide the designer so that he/she can take the final decision in choosing the efficient solutions for global energy saving.

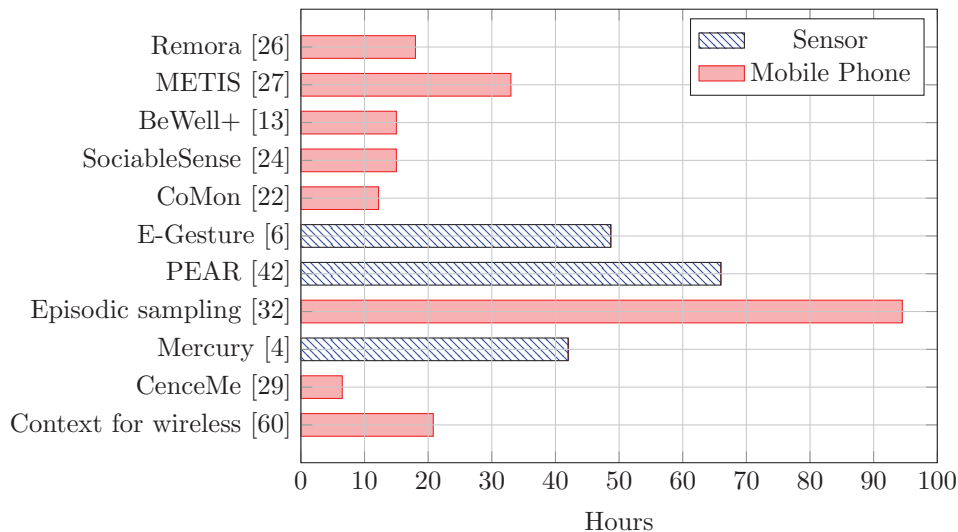
It is hard to estimate the lifetime we can expect from a system that combines the best techniques. We represent in Figure 6 the lifetime achieved by different architectures in hours of operation. We are aware that those results (directly taken from the reviewed papers) are not comparable since they were obtained for different targeted applications, number of type of sensors, sampling frequency, used platforms, etc. As a consequence, we do not draw conclusion on which solution achieves better performances. It is just interesting to notice that globally the lifetime of mobile phones are shorter than the lifetime of sensors. This means that although mobile phones have higher battery capacity than sensor nodes, they deplete their energy faster. This can be explained by the fact that mobile phone have to ensure services such as text messaging, phone call or web browser in addition to the healthcare or well-being application. As a consequence, when designing energy-efficient heterogeneous architectures like the one introduced in Figure 2, the mobile phones used as base station may become serious bottlenecks. Therefore, there is a need to take into consideration the energy-limitation of every device involved in the activity recognition chain.

We plot in Figure 7 the recognition accuracy achieved by different architectures. As we can see, most of the solutions experiment an accuracy greater than 90%. Two of the three architectures which have a recognition accuracy lower than 90% carry their classification task on sensor node. This may be explained by the fact that sensor nodes have fewer resources, so that the classification is limited to lightweight classifiers which may have lower classification performance. It is interesting to note that some solutions, such as Remora [26] or CoMon [22], save energy while improving the accuracy. These solutions carefully determine *when* and *what* resources to share among users, depending on the costs and benefits of sharing a resource.

Although latency is a requirement of importance, few studies report their achieved latency. Therefore, it is not clear how long complex architectures will take to recognize current context. For example, solutions based on opportunistic resource sharing need to detect available sensors, and estimate if the co-location time will be long enough to provide an energy benefit. The overhead induced by the initiation of the collaboration may be affordable in some applications, particularly for environment context recognition.

Finally, it is possible to notice that in the proposed solutions the energy consumption is generally related with the user’s context. For example, in some solutions the energy consumption of the system is directly related to the time spent by the user in active state. In other solutions, the energy consumption depends on the user location. This ability of a system to reduce its energy consumption in function of external parameters can be characterized as follows:

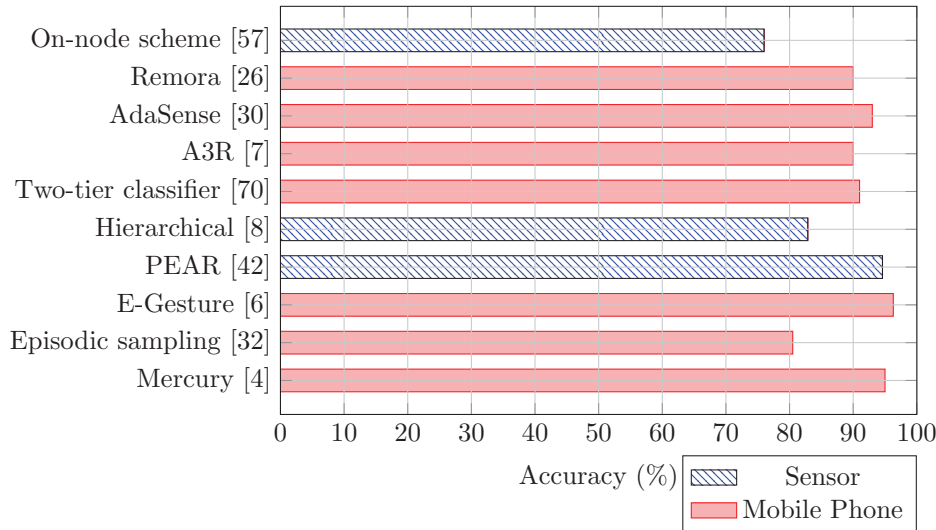
- **Agnostic:** The solution does not implement any adaptive mechanism and behaves in the same way whatever the user activity or environment are. The energy consumption is constant over the time. This is the case of the Hierarchical architecture [8], in which the sensors transmit 1 byte of data every seconds.
- **User level of activity:** The energy consumption of the architectures depends on the user level of activity. In Mercury [4], PEAR [42] and Two-tier classifier [70], the system operates differently depending on the user state (idle or not).
- **User type of activity:** The energy consumption of the system depends on the type of activities being performed by the patient. Typically, A3R [7], AdaSense [30] and E-Gesture [6] architectures adapt their policy in function of the current user activity.
- **User frequency change of context:** The solution consumes more energy when the user change from one activity to another frequently. For example, the On-node scheme [57] sends a report only when the activity changes while the Episodic sampling technique [32] spaces signal processing when no changes are detected.



Architecture	Year Published	Mobile/Mote Model	Sensors	Battery
Context for wireless [60]	2007	HTC Wizard PDA phones (T-Mobile MDA and Cingular 8125)	3-channel ECG, 3axis ACC	1250mAh
CenceMe [29]	2008	Nokia N95	-MIC, ACC, BT scans, GPS, camera	1400mAh
Mercury [4]	2009	SHIMMER	3-axis ACC, GYR	250 mAh Li-polymer rechargeable battery
Episodic sampling [32]	2009	-ADInstrument's MLT1132 -Wearable device. -Nokia N770 PDA	- HR, RR -ACC ADXL330, GYR IDG-300	-NS -NS - Li-po 1500 mAh
PEAR [42]	2011	SHIMMER	-ACC, ECG	280 mAh
E-Gesture [6]	2011	- Wrist-watch sensor compliant with TinyOS. - Nokia N96, Google Nexus One	-3-axis ACC (ADXL335), 3-axis GYR (3 XV-3500CB)	-A coin-size Li-Ion 250 mAh -1400 mAh Li-ion
CoMon [22]	2012	- Sensor motes running TinyOS - Google Nexus One	-Dust, CO2 sensors -GPS, MIC	-NS - 1400mAh
SociableSense [24]	2011	Nokia6210 Navigator phone, Samsung Galaxy S	ACC, MIC, BT traces	750mAh
BeWell+ [13]	2012	Nexus S	ACC, MIC	3200mAh vs. 1400mAh
METIS [27]	2013	-Nokia 6210 Navigator phones -imote2 - Samsung Galaxy S, HTC Desire	-MIC, BT -3-axis ACC -BT detection, MIC, ACC	-NS -3xAAA batteries -NS
Remora [26]	2013	-IRIS motes -Android HTC Hero	- 2-axis ACC, MIC, light and temperature - ACC, velocity from GPS and WiFi	-2xAA Batteries, 2200 mAh

Figure 6: The estimated battery lifetime in hours of different architectures. The table provides hardware configuration of the reviewed solutions (ACC: Accelerometer, GYR: Gyroscope, HR: Heart rate, RR: Respiratory rate, ECG: Electrocardiogram, BT: Bluetooth, MIC: Microphone)

- **User location:** The energy saving opportunities depend on the user environment. As a direct illustration, SociableSense [24], METIS [27] and Context-for-wireless [60] monitor the user mobility and location in order to take advantage of the existence of external devices in the environment.
- **User co-location:** The monitoring system can take advantage of the existence of another system in the environment, and they can initiate a collaboration. For instance, the Remora [26] architecture shares sensing and computational resources among neighboring personal sensor networks in order to save energy.
- **User needs:** The system adapts to the user needs in terms of monitoring since some patients does not require to be monitored as closely as others. For example, BeWell+ [13] provides a lower feedback to patient that need less help in order to save resources. MediAlly [55] can trigger medical data



Architecture	Year Published	Mobile/Mote Model	Sensors	Battery
Mercury [4]	2009	SHIMMER	3-axis ACC, GYR	250 mAh Li-polymer rechargeable battery
Episodic sampling [32]	2009	-ADInstrument's MLT1132 -Wearable device. -Nokia N770 PDA	- HR, RR -ACC ADXL330, GYR IDG-300	-NS -NS - Li-po 1500 mAh
E-Gesture [6]	2011	- Wrist-watch sensor compliant with TinyOS. - Nokia N96, Google Nexus One	-3-axis ACC (ADXL335), 3-axis GYR (3 XV-3500CB)	-A coin-size Li-Ion 250 mAh -1400 mAh Li-ion
PEAR [42]	2011	SHIMMER	-ACC, ECG	280 mAh
Hierarchical [8]	2012	-IMOTE2 motes -MICA2Dot mote - HTC G7	-3-axis ACC - short-range RFID reader	NS
Two-tier classifier [70]	2012	Galaxy S2 vs. Nexus S	3-axis ACC	NS
A3R [7]	2012	Nokia N95 vs. Samsung Galaxy S2	ACC	NS
AdaSense [30]	2013	HTC Google Nexus One	ACC	Removable Li-Ion 1400 mAh
Remora [26]	2013	-IRIS motes -Android HTC Hero	- 2-axis ACC, MIC, light and temperature - ACC, velocity from GPS and WiFi	-2*AA Batteries, 2200 mAh
On-node scheme [57]	2013	Texas Instruments eZ430-RF2500	3-axis ACC (ADXL335)	NS

Figure 7: The accuracy achieved by different architectures in percent. The table provides hardware configuration of the reviewed solutions (ACC: Accelerometer, GYR: Gyroscope, HR: Heart rate, RR: Respiratory rate, ECG: Electrocardiogram, BT: Bluetooth, MIC: Microphone)

acquisition whenever the subject's context indicates physical or emotional stress.

10. Conclusion and open research issues

In conclusion, we carried out a novel comparative study, in which we considered how each category of energy efficient mechanisms (power-on time reduction, communication reduction, computation reduction, and battery charging) can be applied to different techniques used in health wellbeing applications (user location detection, social interaction detection, physical activity recognition and physiological parameters supervision). We give suitable approaches for these applications to save energy. We also highlight the influence of the technique on the delay and accuracy requirements. Therefore, we guide the designer so that he/she can take the final decision in choosing the efficient solutions for global energy saving. In this section,

we aim to shed some light on open research issues relating to energy-efficient human context recognition based on wearable sensing through the following points:

10.1. Unbalanced residual energy

Devices involved in context recognition applications may have unbalanced residual energy due to different energy consumption rates, and energy harvesting rates. Indeed, depending on their role, sampling frequency or communication technology, nodes will deplete their battery more or less rapidly. Furthermore, nodes may have an uneven residual energy distribution due to the difference in the quantity of the collected energy. For instance, the energy that can be harvest from human activities will not be the same for two identical sensors if one is located on the chest and another is located on the leg. As a consequence, nodes with low residual energy may be assigned lower sensing periods, while those with high residual energy may be preferred for running context recognition tasks. In addition, the same battery level at two different devices could have different implications: for example, 50% battery left at a sensor dedicated to a unique application represents a longer operating time compared to 50% battery left at a mobile phone running multiple applications. Thus, the fraction of the phone’s energy dedicated to the context recognition application must not be as high as the one of the sensors: a part of phone’s energy must be preserved to be dedicated to other tasks, such as calling or text messaging. Therefore, it is possible to consider new decision metrics that include the desired involvement of a device in the application in terms of energy usage. It is also important to take into consideration that the same battery level at a different time could mean a totally different thing. Indeed, 50% battery left in the early morning may necessitate that the device be recharged before the end of the day, which may not be the case if 50% battery is left in the early evening [56]. These discrepancies between devices regarding unbalanced residual energy have to be taken into account when designing context recognition systems based on wearable sensors.

10.2. Mutual and cross technology interference mitigation

As the popularity of sensing devices grows, we can expect that a single user will carry a number of context sensing devices. In this context, multiple devices will work together towards improved context sensing. We showed in section 7.2 that the cooperation of multiple body area networks can save energy through opportunistic resources sharing. However, the collocation of several wearable sensor networks sharing the same unlicensed band dramatically increases the level of interference, which in turn negatively affects the network performance [83]. Indeed, mutual (e.g. WiFi-WiFi, ZigBee-ZigBee) and cross (e.g WiFi-ZigBee) technology interference arise when several body sensor networks operate in the same vicinity, due to the broadcast nature of the wireless channel. For instance, it is shown that the throughput of WiFi and Zigbee degrade heavily when they are co-located [84]. As a consequence, the increase of re-transmissions due to collisions, and multiple carrier sensing, result in increased energy consumption. Therefore, there is a need to develop interference-aware solutions in order to ensure an energy-efficient radio spectrum utilization.

10.3. Devices heterogeneity

Mobile sensing is challenged by a growing number of devices used for context-aware applications. This combination of sensors leads to heterogeneous architectures composed of devices with different capabilities in terms of sensing, accuracy, computation, communications and battery power. Heterogeneous architectures offer a certain flexibility since the devices have different resources whose usage can be optimized, and the use of multiple sensors allow improving the context recognition. However, heterogeneity raises programmability issues since devices are provided by different manufacturers, which make it very hard to build a common sensing platform that can work across different devices and OS. Envisioned solutions to the problem of heterogeneity are cloud sensing, which combine virtualization of nodes and semantic-based query processing. The objective of this approach is to reduce the configuration time by providing a common framework to represent and access information, while allowing optimization across multiple devices. However, this approach induces a communication overhead to the data transmission between virtualized devices and the cloud.

10.4. Collocation of concurrent application

Until recently, most of the recognition systems were developed for a unique application, and researchers adopted an application-driven approach. In this approach, applications determine the types and amount of resources required to execute the program, and request the needed resources through the provided APIs. However, future systems are likely to run multiple context recognition applications at a time. In this context, if the device is not equipped with an appropriate OS for concurrent resource management, energy performance will degrade. For example, a user may be willing to run both a wellbeing application that encourages healthy behavior such as BeWell+ [13], and a noise exposure mapping application such as Ear-Phone [71], which both use the microphone sensor on phone. In order to save energy, BeWell+ dynamically tunes the rate at which sampling, feature extraction and context inference are performed, while Ear-Phone does not take into consideration measurements taken when the phone is carried in the bag or pocket. If these two applications are used simultaneously, it is not clear how energy gain of both solutions will be offset by the behavior of the other application. In Orchestrator [69], Kang et al. defend the need for a system-driven approach to tackle the complexity of resource management in dynamic and concurrent environment. The LiveLabs [85] experience confirms the difficulty of satisfying multiple applications accessing a sensing platform when considering different requirements in terms of type of information, sampling frequency, level of accuracy, or latency. Although some studies propose to manage resources of a mobile device or sensor node [86, 87], a general framework for multi-context recognition across heterogeneous devices is missing.

10.5. Anticipatory energy saving

Beyond context recognition, *context anticipation* represents a powerful tool for tackling a number of health and well-being issues, from obesity to stress and addiction [88]. Based on the predictions of the future state of the context, these applications will anticipate changes in user's health and behaviour and take proactive decisions in order to impact the future state. For instance, an application can proactively tackle depression by detecting decreased movement, the lack of socializing, irregular sleep patterns of the patient, and then encourage the patient to go out and socialize, for example by sending a link to two discounted theater tickets. These systems will offer great benefit for the patient follow up by allowing the development of innovative ways of delivering interventions. In addition to the social benefit, such anticipatory applications will help to proactively manage energy and wireless resources. Indeed, the ability to predict users' location, social encounters or health hazards will enable devices to predict future harvesting opportunities, resource sharing opportunities or application demands. For example, given good weather forecasts, a system can anticipate that the user will jog for one hour this weekend, thus expecting to recharge a part of its battery through body movement.

To conclude, we anticipate that energy management will remain a huge challenge because it is a key enabler of real-world applications development and market penetration. In this context, we hope that this survey will be helpful to the pervasive and mobile research community willing to address the energy issues of health-related wearable sensing systems. With this objective, this paper surveys existing energy-efficient approaches designed for human context recognition based on wearable sensors for healthcare and wellbeing applications. More particularly, we have proposed a new classification of these mechanisms, and we have reviewed up-to-date references in details.

Acknowledgments

This work was carried out and funded in the framework of the Labex MS2T. It was supported by the French Government, through the program "Investments for the future" managed by the National Agency for Research (Reference ANR-11-IDEX-0004-02).

References

- [1] <http://www.un.org/esa/population/publications/worldageing19502050/>, World Population Ageing: 1950-2050, online; Chapter II, accessed december 2014.

- [2] O. D. Lara, M. A. Labrador, A Survey on Human Activity Recognition based on Wearable Sensors, *IEEE Communications Surveys and Tutorials* 15 (3) (2013) 1192–1209.
- [3] G. Bleser, D. Steffen, M. Weber, G. Hendeby, D. Stricker, L. Fradet, F. Marin, N. Ville, F. Carré, A personalized exercise trainer for the elderly, *Journal of ambient intelligence and smart environments* 5 (2013) 547–562.
- [4] K. Lorincz, B. rong Chen, G. W. Challen, A. R. Chowdhury, S. Patel, P. Bonato, M. Welsh, Mercury: a wearable sensor network platform for high-fidelity motion analysis, in: *7th ACM Conference on Embedded Networked Sensor Systems*, Berkeley, California, 183–196, 2009.
- [5] H. Noshadi, F. Dabiri, S. Meguerdichian, M. Potkonjak, M. Sarrafzadeh, Energy optimization in wireless medical systems using physiological behavior, in: *Wireless Health Conference*, San Diego, CA, 128–136, 2010.
- [6] T. Park, J. Lee, I. Hwang, C. Yoo, L. Nachman, J. Song, E-Gesture: a collaborative architecture for energy-efficient gesture recognition with hand-worn sensor and mobile devices, in: *9th ACM Conference on Embedded Networked Sensor Systems*, Seattle, WA, 260–273, 2011.
- [7] Z. Yan, V. Subbaraju, D. Chakraborty, A. Misra, K. Aberer, Energy-Efficient Continuous Activity Recognition on Mobile Phones: An Activity-Adaptive Approach, in: *16th International Symposium on Wearable Computers*, Newcastle, UK, 17–24, 2012.
- [8] L. Wang, T. Gu, X. Tao, J. Lu, A hierarchical approach to real-time activity recognition in body sensor networks, *Pervasive and Mobile Computing* 8 (1) (2012) 115–130.
- [9] R. Fensli, E. Gunnarson, T. Gundersen, A wearable ECG-recording system for continuous arrhythmia monitoring in a wireless tele-home-care situation, in: *18th IEEE Symp. on Computer-Based Medical Systems*, 407–412, 2005.
- [10] J. Chen, K. Kwong, D. Chang, J. Luk, R. Bajcsy, Wearable sensors for reliable fall detection, in: *27th IEEE Conference of the Engineering in Medicine and Biology Society*, 3551–3554, 2005.
- [11] K. Zhan, S. Faux, F. Ramos, Multi-scale Conditional Random Fields for first-person activity recognition on elders and disabled patients, *Pervasive and Mobile Computing* (In Press) 1–27.
- [12] Q. Lin, D. Zhang, K. Connelly, H. Ni, Z. Yu, X. Zhou, Disorientation detection by mining GPS trajectories for cognitively-impaired elders, *Pervasive and Mobile Computing* (In Press) 1–15.
- [13] M. Lin, N. D. Lane, M. Mohammod, X. Yang, H. Lu, G. Cardone, S. Ali, A. Doryab, E. Berke, BeWell+: multi-dimensional wellbeing monitoring with community-guided user feedback and energy optimization, in: *Wireless Health Conference*, San Diego, USA, 1–5, 2012.
- [14] Y. Lee, C. Min, C. Hwang, J. Lee, I. Hwang, Y. Ju, C. Yoo, M. Moon, U. Lee, J. Song, SocioPhone: Everyday Face-to-face Interaction Monitoring Platform Using Multi-phone Sensor Fusion, in: *Proceeding of the 11th Annual International Conference on Mobile Systems, Applications, and Services, MobiSys '13*, ACM, New York, NY, USA, ISBN 978-1-4503-1672-9, 375–388, doi:10.1145/2462456.2465426, URL <http://doi.acm.org/10.1145/2462456.2465426>, 2013.
- [15] J. K. Aggarwal, Q. Cai, Human Motion Analysis: A Review, *Computer Vision and Image Understanding* 73 (1999) 428–440.
- [16] R. E. Mayagoitia, A. V. Nene, P. H. Veltink, Accelerometer and rate gyroscope measurement of kinematics: an inexpensive alternative to optical motion analysis systems, *Journal of Biomechanics* 35 (2002) 537–542.
- [17] L. Bao, S. S. Intille, Activity recognition from user-annotated acceleration data, *Springer Berlin Heidelberg* (2004) 1–17.
- [18] J. Gummeson, B. Priyantha, J. Liu, An energy harvesting wearable ring platform for gesture input on surfaces, in: *12th annual Int. Conf. on Mobile systems, applications, and services*, Bretton Woods, NH, USA, 162–175, 2014.
- [19] S. A. Hoseini-Tabatabaei, A. Gluhak, R. Tafazolli, A survey on smartphone-based systems for opportunistic user context recognition, *ACM Computing Surveys* 45 (3) (2013) 1–51.
- [20] D. H. Kim, Y. Kim, D. Estrin, M. B. Srivastava, SensLoc: sensing everyday places and paths using less energy, in: *8th ACM Conference on Embedded Networked Sensor Systems*, Zurich, Switzerland, 43–56, 2010.
- [21] H. Lu, W. P. N. D. Lane, T. Choudhury, A. T. Campbell, SoundSense: scalable sound sensing for people-centric applications on mobile phones, in: *7th international conference on Mobile systems, applications, and services*, Krakow, Poland, 165–178, 2009.
- [22] Y. Lee, Y. Ju, C. Min, S. Kang, I. Hwang, J. Song, CoMon: cooperative ambience monitoring platform with continuity and benefit awareness, in: *10th Int. Conf. on Mobile systems, applications, and services*, Low Wood Bay, UK, 43–56, 2012.
- [23] V. Leonov, Thermoelectric Energy Harvesting of Human Body Heat for Wearable Sensors, *IEEE Sensors Journal* 13 (6) (2013) 1–8.
- [24] K. K. Rachuri, C. Mascolo, M. Musolesi, P. J. Rentfrow, SociableSense: exploring the trade-offs of adaptive sampling and computation offloading for social sensing, in: *ACM International Conference on Mobile computing and networking - MobiCom*, Las Vegas, Nevada, USA, 73–84, 2011.
- [25] D. Roggen, G. Troster, P. Lukowicz, A. Ferscha, J. del R. Millan, R. Chavarriaga, Opportunistic human activity and context recognition, *Computer* 46 (2) (2012) 36–45.
- [26] M. Keally, G. Zhou, G. Xing, J. Wu, Remora: Sensing resource sharing among smartphone-based body sensor networks, in: *21st International Symposium on Quality of Service*, Montreal, QC, 1–10, 2013.
- [27] K. K. Rachuri, C. Efstratiou, I. Leontiadis, C. Mascolo, P. Rentfrow, METIS: Exploring mobile phone sensing offloading for efficiently supporting social sensing applications, in: *IEEE International Conference on Pervasive Computing and Communications*, San Diego, CA, 85–93, 2013.
- [28] S. Nath, ACE: exploiting correlation for energy-efficient and continuous context sensing, in: *10th Int. Conf on Mobile Systems, Applications and Services*, Low Wood Bay, UK, 29–42, 2012.
- [29] E. Miluzzo, N. D. Lane, K. Fodor, R. Peterson, H. Lu, M. Musolesi, S. B. Eisenman, X. Zheng, A. T. Campbell, Sensing meets mobile social networks: the design, implementation and evaluation of the CenceMe application, in: *6th ACM conference on Embedded network sensor systems*, Raleigh, NC, USA, 337–350, 2008.

- [30] X. Qi, M. Keally, G. Zhou, Y. Li, Z. Ren, AdaSense: Adapting Sampling Rates for Activity Recognition in Body Sensor Networks, in: 19th IEEE Real-Time and Embedded Technology and Applications Symposium, Philadelphia, PA, 163–172, 2013.
- [31] A. Bulling, U. Blanke, B. Schiele, A tutorial on human activity recognition using body-worn inertial sensors, *ACM Computing Surveys* 46 (3) (2014) 1–33.
- [32] L. Au, M. Batalin, T. Stathopoulos, A. Bui, W. Kaiser, Episodic sampling: Towards energy-efficient patient monitoring with wearable sensors, in: IEEE Engineering in Medicine and Biology Society, Minneapolis, Minnesota, USA, 6901–6905, 2009.
- [33] L. Lim, A. Misra, T. Mo, Adaptive data acquisition strategies for energy-efficient, smartphone-based, continuous processing of sensor streams, *Distributed and Parallel Databases* 32 (2) (2013) 321–351.
- [34] N. Balasubramanian, A. Balasubramanian, A. Venkataramani, Energy consumption in mobile phones: a measurement study and implications for network applications, in: 9th ACM SIGCOMM conference on Internet measurement conference, Chicago, Illinois, USA, 280–293, 2009.
- [35] J. A. Ward, P. Lukowicz, H. W. Gellersen, Performance metrics for activity recognition, *ACM Transactions on Intelligent Systems and Technology* 2 (2011) 1–23.
- [36] G. Anastasi, M. Conti, M. Di Francesco, A. Passarella, Energy conservation in wireless sensor networks: A survey, *Ad Hoc Networks* 7 (3) (2009) 537–568.
- [37] K. Akkaya, M. Younis, A survey on routing protocols for wireless sensor networks, *Ad Hoc Networks* 3 (3) (2005) 325–349.
- [38] R. Rajagopalan, P. K. Varshney, Data-aggregation techniques in sensor networks: a survey, *IEEE Communications Surveys & Tutorials* 8 (4) (2006) 48–63.
- [39] N. Lane, E. Miluzzo, H. Lu, D. Peebles, T. Choudhury, A. Campbell, A survey of mobile phone sensing, *IEEE Communications Magazine* 48 (9) (2010) 140–150.
- [40] F. B. Abdesslem, A. Phillips, T. Henderson, Less is more: energy-efficient mobile sensing with senseless, in: 1st ACM workshop on Networking, systems, and applications for mobile handhelds, Barcelona, Spain, 61–62, 2009.
- [41] J. Furthmuller, O. Waldhorst, Energy-aware Resource Sharing with Mobile Devices, *Performance Metrics for Intelligent Systems* 56 (2012) 1920–1934.
- [42] F.-T. Sun, C. Kuo, M. Griss, PEAR: Power efficiency through activity recognition (for ECG-based sensing), in: 5th International Conference on Pervasive Computing Technologies for Healthcare, Dublin, Ireland, 115–122, 2011.
- [43] V. Leonov, T. Torfs, P. Fiorini, C. Van Hoof, Thermoelectric Converters of Human Warmth for Self-Powered Wireless Sensor Nodes, *IEEE Sensors Journal* 7 (5) (2007) 650–657.
- [44] M. Gorlatova, J. Sarik, M. Cong, I. Kymissis, G. Zussman, Movers and Shakers: Kinetic Energy Harvesting for the Internet of Things, *Technical Report* (2013) 1–11.
- [45] J. M. Donelan, Q. Li, V. Naing, J. A. Hoffer, D. J. Weber, A. D. Kuo, Biomechanical Energy Harvesting: Generating Electricity During Walking with Minimal User Effort, *Science* 319 (5864) (2008) 807–810.
- [46] L. Rome, L. Flynn, E. Goldman, T. Yoo, Generating Electricity While Walking with Loads, *Science* 309 (5741) (2005) 1725–1728.
- [47] N. Shenck, J. Paradiso, Energy scavenging with shoe-mounted piezoelectrics, *IEEE Micro* 21 (3) (2001) 30–42.
- [48] P. Damm, F. Graichen, A. Rohlmann, A. Bender, G. Bergmann, Total hip joint prosthesis for in vivo measurement of forces and moments, *Medical Engineering & Physics* 32 (1) (2010) 95–100.
- [49] P. Zappi, T. Stiefmeier, E. Farella, D. Roggen, L. Benini, G. Troster, Activity recognition from on-body sensors by classifier fusion: sensor scalability and robustness, in: 3rd International Conference on Intelligent Sensors, Sensor Networks and Information, Melbourne, Qld., 281–286, 2007.
- [50] L. Gao, A. Bourke, J. Nelson, Activity recognition using dynamic multiple sensor fusion in body sensor networks, in: IEEE Engineering in Medicine and Biology Society, San Diego, CA, 1077–1080, 2012.
- [51] D. Gordon, J. Czerny, T. Miyaki, M. Beigl, Energy-Efficient Activity Recognition Using Prediction, in: 16th International Symposium on Wearable Computers, Newcastle, 29–36, 2012.
- [52] S. Kang, J. Lee, H. Jang, Y. Lee, S. Park, J. Song, A Scalable and Energy-Efficient Context Monitoring Framework for Mobile Personal Sensor Networks, *IEEE Transactions on Mobile Computing* 9 (5) (2006) 686–702.
- [53] Y. Wang, J. Lin, M. Annavaram, Q. A. Jacobson, J. Hong, B. Krishnamachari, N. Sadeh, A framework of energy efficient mobile sensing for automatic user state recognition, in: 7th International Conference on Mobile systems, applications, and services, Krakow, Poland, 179–192, 2009.
- [54] K. V. Laerhoven, H.-W. Gellersen, Spine versus Porcupine: A Study in Distributed Wearable Activity Recognition, in: Int. Symp. on Wearable Computers, Washington, DC, USA, 142–149, 2004.
- [55] A. Chowdhury, B. Falchuk, A. Misra, MediAlly: A provenance-aware remote health monitoring middleware, in: IEEE Int. Conf on Pervasive Computing and Communications, Mannheim, Germany, 125–134, 2010.
- [56] H. Lu, J. Y. Z. Liu, N. D. Lane, T. Choudhury, A. T. Campbell, The Jigsaw continuous sensing engine for mobile phone applications, in: 8th ACM Conference on Embedded Networked Sensor Systems, Zurich, Switzerland, 71–84, 2010.
- [57] N. Wang, G. V. Merrett, R. G. Maunder, A. Rogers, Energy and Accuracy Trade-Offs in Accelerometry-Based Activity Recognition, in: 22nd International Conference on Computer Communications and Networks, Nassau, Bahamas, 1–6, 2013.
- [58] M. Paschou, E. Sakkopoulos, E. Sourla, A. Tsakalidis, Health Internet of Things: Metrics and methods for efficient data transfer, *Simulation Modelling Practice and Theory* (34) (2013) 186–199.
- [59] B. Kim, Y. Cho, J. Hong, AWNIS: Energy-Efficient Adaptive Wireless Network Interface Selection for Industrial Mobile Devices, *IEEE Transactions on Industrial Informatics* (10) (2014) 714–729.
- [60] A. Rahmati, L. Zhong, Context-for-wireless: context-sensitive energy-efficient wireless data transfer, in: 5th International

- Conference on Mobile systems, applications and services, San Juan, Puerto Rico, USA, 165–178, 2007.
- [61] V. Loseu, H. Ghasemzadeh, R. Jafari, A wireless communication selection approach to minimize energy-per-bit for wearable computing applications, in: *Int. Conf. on Distributed Computing in Sensor Systems and Workshops - DCOSS*, Barcelona, Spain, 1–8, 2011.
- [62] C. Sengul, M. Bakht, A. Harris, T. Abdelzaher, R. Kravets, Improving Energy Conservation Using Bulk Transmission over High-Power Radios in Sensor Networks, in: *Int. Conf. on Distributed Computing Systems -ICDCS*, Beijing, China, 801–808, 2008.
- [63] J. Liu, B. Priyantha, T. Hart, H. S. Ramos, A. A. F. Loureiro, Q. Wang, Energy efficient GPS sensing with cloud offloading, in: *10th ACM Conference on Embedded Network Sensor Systems*, Toronto, Canada, 85–98, 2012.
- [64] E. Cuervo, A. Balasubramanian, D. ki Cho, A. Wolman, S. Saroiu, R. Chandra, P. Bahl, MAUI: making smartphones last longer with code offload, in: *8th Int. Conf on Mobile Systems, Applications and Services*, San Francisco, USA, 49–62, 2010.
- [65] N. Vallina-Rodriguez, J. Crowcroft, ErdOS: achieving energy savings in mobile OS, in: *6th Int. workshop on MobiArch*, Bethesda, USA, 37–42, 2011.
- [66] R. Cilla, M. Patricio, A. Berlanga, J. Molina, Creating Human Activity Recognition Systems Using Pareto-based Multiobjective Optimization, in: *6th IEEE International Conference on Advanced Video and Signal Based Surveillance*, Genova, 37–42, 2009.
- [67] H. Martin, J. Iglesias, J. Cano, A. M. Bernardos, J. R. Casar, Towards a fuzzy-based multi-classifier selection module for activity recognition applications, in: *Workshop proceedings of the 10th IEEE International Conference on Pervasive Computing and Communications*, Lugano, Switzerland, 871–876, 2012.
- [68] D. Chu, N. D. Lane, T. T.-T. Lai, C. Pang, X. Meng, Q. Guo, F. Li, F. Zhao, Balancing energy, latency and accuracy for mobile sensor data classification, in: *9th ACM Conference on Embedded Networked Sensor Systems*, Seattle, WA, USA, 54–67, 2011.
- [69] S. Kang, Y. Lee, C. Min, Y. Ju, T. Park, J. Lee, Y. Rhee, J. Song, Orchestrator: An active resource orchestration framework for mobile context monitoring in sensor-rich mobile environments, in: *IEEE Int. Conf on Pervasive Computing and Communications*, Mannheim, Germany, 135–144, 2010.
- [70] V. Srinivasan, T. Phan, An accurate two-tier classifier for efficient duty-cycling of smartphone activity recognition systems, in: *3rd International Workshop on Sensing Applications on Mobile Phones*, Toronto, Canada, 1–5, 2012.
- [71] R. Rana, C. T. Chou, N. Bulusu, S. Kanhere, W. Hu, Ear-Phone: A context-aware noise mapping using smart phones, *Pervasive and Mobile Computing* (In Press) 1–22.
- [72] U. Maurer, A. Smailagic, D. Siewiorek, M. Deisher, Activity recognition and monitoring using multiple sensors on different body positions, in: *Int. Workshop on Wearable and Implantable Body Sensor Networks*, Cambridge, MA, 1–5, 2006.
- [73] D. Lymberopoulos, N. Priyantha, M. Goraczko, F. Zhao, Towards Energy Efficient Design of Multi-radio Platforms for Wireless Sensor Networks, in: *Int. Conf. on Information Processing in Sensor Networks (IPSN)*, St. Louis, MO, 257–268, 2008.
- [74] K. Kumar, Y.-H. Lu, Cloud Computing for Mobile Users: Can Offloading Computation Save Energy, *Computer* 43 (4) (2010) 51–56.
- [75] S. Sudevalayam, P. Kulkarni, Energy Harvesting Sensor Nodes: Survey and Implications, *IEEE Communications Surveys & Tutorials* 13 (3) (2011) 443–461.
- [76] P. Mitcheson, E. Yeatman, G. Rao, A. Holmes, T. Green, Energy Harvesting From Human and Machine Motion for Wireless Electronic Devices, *Proceedings of the IEEE* 96 (9) (2008) 1457–1486.
- [77] J. Paradiso, T. Starner, Energy scavenging for mobile and wireless electronics, *IEEE Pervasive Computing* 4 (1) (2005) 1–10.
- [78] P. Mitcheson, Energy harvesting for human wearable and implantable bio-sensors, in: *32th Annual Conference of the IEEE EMBS*, Buenos Aires, Argentina, 3432–3436, 2010.
- [79] L. Xie, Y. Shi, Y. Hou, A. Lou, Wireless power transfer and applications to sensor networks, *IEEE Wireless Communications* 20 (4) (2013) 140–145.
- [80] F. Graichen, R. Arnold, A. Rohlmann, G. Bergmann, Implantable 9-channel telemetry system for in vivo load measurements with orthopedic implants, *IEEE Transactions on Bio-medical Engineering* 54 (2) (2007) 253–261.
- [81] J. Olivo, S. Carrara, G. De Micheli, IronIC patch: A wearable device for the remote powering and connectivity of implantable systems, in: *IEEE International Instrumentation and Measurement Technology Conference - I2MTC*, Graz, 286–289, 2012.
- [82] www.shimmersensing.com, 2014.
- [83] J. Elias, S. Paris, M. Krunz, Cross Technology Interference Mitigation in Body Area Networks: an Optimization Approach, *IEEE Transactions on Vehicular Technology* (99) (2014) 1–14.
- [84] C.-J. M. Liang, N. B. Priyantha, J. Liu, A. Terzis, Surviving wi-fi interference in low power ZigBee networks, in: *8th ACM Conference on Embedded Networked Sensor Systems*, Zurich, Switzerland, 309–322, 2010.
- [85] R. Balan, Y. Lee, T. K. Wee, A. Misra, The challenge of continuous mobile context sensing, in: *6th Int. Conf. on Communication Systems and Networks*, Bangalore, Inde, 1–8, 2014.
- [86] A. Lachenmann, P. J. Marron, D. Minder, K. Rothermel, Meeting lifetime goals with energy levels, in: *5th Int. Conf. on Embedded networked sensor systems*, Sydney, Australia, 131–144, 2007.
- [87] K. Lorincz, B. rong Chen, J. Waterman, G. Werner-Allen, M. Welsh, Resource aware programming in the Pixie OS, in: *6th ACM conference on Embedded network sensor systems*, Raleigh, NC, USA, 211–224, 2008.
- [88] V. Pejovic, M. Musolesi, Anticipatory mobile computing for behaviour change interventions, in: *ACM International Joint Conference on Pervasive and Ubiquitous Computing*, Seattle, Washington, USA, 1025–1034, 2014.