

RESEARCH

Open Access



# A residential maintenance-free long-term activity monitoring system for healthcare applications

Xenofon Fafoutis<sup>\*</sup>, Evgeny Tsimballo, Evangelos Mellios, Geoffrey Hilton, Robert Piechocki and Ian Craddock

## Abstract

Demographic changes such as the ageing population and the continuous rise of chronic medical conditions such as obesity, diabetes and depression make our healthcare systems economically unsustainable. Sensing technologies are promising solutions that can provide cost-effective answers to these challenges. In this paper, we focus on long-term in-house activity monitoring that aims at early detection and prevention of such conditions. In this context, we present and experimentally evaluate an ultra low-power (less than 100- $\mu$ W long-term average power consumption) on-body activity sensing prototype system that is based on Bluetooth low energy (BLE). As part of a larger smart home monitoring architecture, the role of the presented system is to collect and reliably deliver acceleration data to the upper layers of the architecture. The system evaluation incorporates a thorough power consumption study that facilitates meaningful battery lifetime estimations, an insightful coverage study in an actual residential environment, and the investigation of energy-efficient packet loss mitigation techniques.

**Keywords:** eHealth, mHealth, Wearable technologies, Healthcare technologies, Internet of things, Bluetooth low energy

## 1 Introduction

Our healthcare systems are challenged by demographic changes. The United Nations (UN) predicts that in 2050, the proportion of the population aged over 60 in the developed world will be 33 %, following an increasing trend [1]. In addition to continuously ageing populations, the rise of chronic illness pushes the limits of our healthcare systems which makes them unsustainable [2]. As medical professionals report that early detection and prevention is more cost-effective than treatment [3], there is a shift towards encouraging people to manage their own well-being at home. Residential healthcare is not only cost-effective, but it is also supporting the dignity and independence of the elderly [4]. Advances in Ambient Assisted Living (AAL) [5], wireless sensor networks (WSN) and wearable technologies [6] provide the necessary infrastructure to support such a shift in healthcare provision.

In this context, the focus of this paper is on long-term activity monitoring in a residential environment. The key challenge of long-term activity monitoring is the energy constraints of wearable hardware. Low power consumption is a key goal of all wireless sensing systems, primarily due to the associated costs of replacing or recharging batteries. To make matters worse, in the context of our application of interest, the users of the system are either elderly or suffering from medical conditions, making the option of user-based system maintenance unrealistic. For these reasons, our aim is an ultra low-power system with a long-term average power consumption of the order of tens of microwatt. Such levels of power consumption correspond to a long battery lifetime, hence, a system that can support maintenance-free residential monitoring for several months. Although only battery-powered solutions are considered in this paper, it should be noted that such low levels of power consumption also facilitate the realisation of completely maintenance-free energy harvesting sensing systems [7], where sensors are powered by the—typically very limited—energy that is available in the surrounding

<sup>\*</sup>Correspondence: xenofon.fafoutis@bristol.ac.uk  
Department of Electrical and Electronic Engineering, University of Bristol,  
Woodland Road, BS8 1UB Bristol, UK

environment. In our proposed platform, the energy efficiency of the wearable device is also enhanced by exploiting the asymmetry of the network resources. Instead of energy-consuming retransmissions, reliable communication is achieved through incorporating energy-efficient packet loss mitigation techniques that are primarily supported by the mains-powered infrastructure.

The presented activity monitoring system is based on Bluetooth low energy (BLE) and is part of a larger platform that we call the SPHERE (a Sensor Platform for HHealthcare in a Residential Environment) architecture [8]. The SPHERE architecture is a residential platform that is armed with environmental sensors, video cameras and on-body sensors for sensing, as well as high-performance computing systems for real-time reasoning and decision-making with machine learning algorithms. A prototype of the SPHERE architecture is being engineered and realised in a furnished house in the city of Bristol, UK, which, for the remainder of the paper, will be referred to as the *SPHERE house*. The purpose of the SPHERE house is to host volunteers, so that the SPHERE technologies can be tested on real subjects. As part of the SPHERE architecture, the scope of the presented activity monitoring system is on-body sensing and the provision of the sensed data to the higher layers responsible for their interpretation.

The contribution of this work lies in the unique combination of energy-efficient techniques and practices for the realisation of a reliable maintenance-free smart home system with wearable sensors for long-term activity monitoring. It also provides insight for researchers and engineers who are developing similar platforms. In particular, we perform a thorough power consumption study that is based on actual measurements and considers each key part of the system separately, in an attempt to be the basis of practical battery lifetime estimations. In addition, we investigate the issue of antenna design for body-centric communications. We provide a thorough wireless coverage study focusing on the challenges of full-house coverage and the effects of the body on the path loss and the packet error rate (PER) in a residential environment. We investigate the trade-off between power consumption and wireless coverage. Finally, we quantify and evaluate the performance of energy-efficient packet loss mitigation techniques, including cyclic redundancy code (CRC) error correction, application-layer coding and antenna polarisation diversity. Overall, this work documents experiences and practical lessons on the use of BLE for streaming data, body-centric communications in residential environments and ultra low-power system design.

The remainder of the paper is structured as follows. Following Section 2 which discusses the related work, Section 3 introduces on-body sensing as part of the SPHERE architecture. Sections 4 and 5 focus on the design

and implementation of the wearable sensor and smart home infrastructure, respectively. Section 6 evaluates the system, primarily focusing on power consumption, wireless coverage and the performance of the energy-efficient packet loss mitigation techniques. Section 7 discusses prototyping and miniaturisation aspects of the wearable platform. Lastly, Section 8 summarises and concludes the paper.

## 2 Related work

In the recent years, fashionable activity monitors, such as Fitbit, Jawbone UP and Nike+ Fuelband SE, have appeared in the consumer electronics market [9]. These wearable gadgets demonstrate the rise of a trend towards self-monitoring, as well as the user acceptability of wearable technologies [6]. However, they are of limited use for medical applications due to the limited accessibility of the raw data, their lack of interoperability with other healthcare systems and their limited expandability to new sensor technologies. Furthermore, their need for regular recharging (typical battery lifetime of less than a week) hinders their suitability for target groups that are uncomfortable with managing modern technologies.

The research community has also proposed and used several wearable devices for activity monitoring systems and AAL healthcare architectures, such as [10] and [11], a few of which are briefly reviewed in this section. We refer the reader to [6] for an exhaustive survey on smart wearable technologies. Verity [12] is an AAL platform in which the wearable device is equipped with an accelerometer and a piezo-resistive sensor for fall detection and heart rate monitoring. In [13], an AAL platform is proposed based on a waist-worn accelerometer that is able to identify basic activities, such as sitting, walking, running and jumping. In the same direction, [14] and [15] describe identification of basic activities using multiple on-body accelerometers and gyroscopes. KNOWME [16] consists of a network of multiple on-body sensors, including accelerometers and an electrocardiogram sensor, that are streaming on a smart phone which estimates the user's energy expenditure. Many of these platforms do not pay the necessary attention to the energy consumption of the hardware used, resulting in wearable devices that require regular recharging, similarly to commercial products. Other works present low-power wearable systems with different types of sensors, such as bio-impedance [17], sound [18], electroencephalogram (EEG) [19] and inertial sensors [20]. Although relatively low power, these platforms target different body sensing applications that require sensing elements that are significantly energy-consuming compared to the sensing elements employed in our platform. The study by Kan and Chen [20] presents a system with similar characteristics to our platform, albeit the paper reports a much inferior power

consumption profile, which is mainly because of the high power consumption of the radio and sensors used, along with the protocols employed for wireless communication.

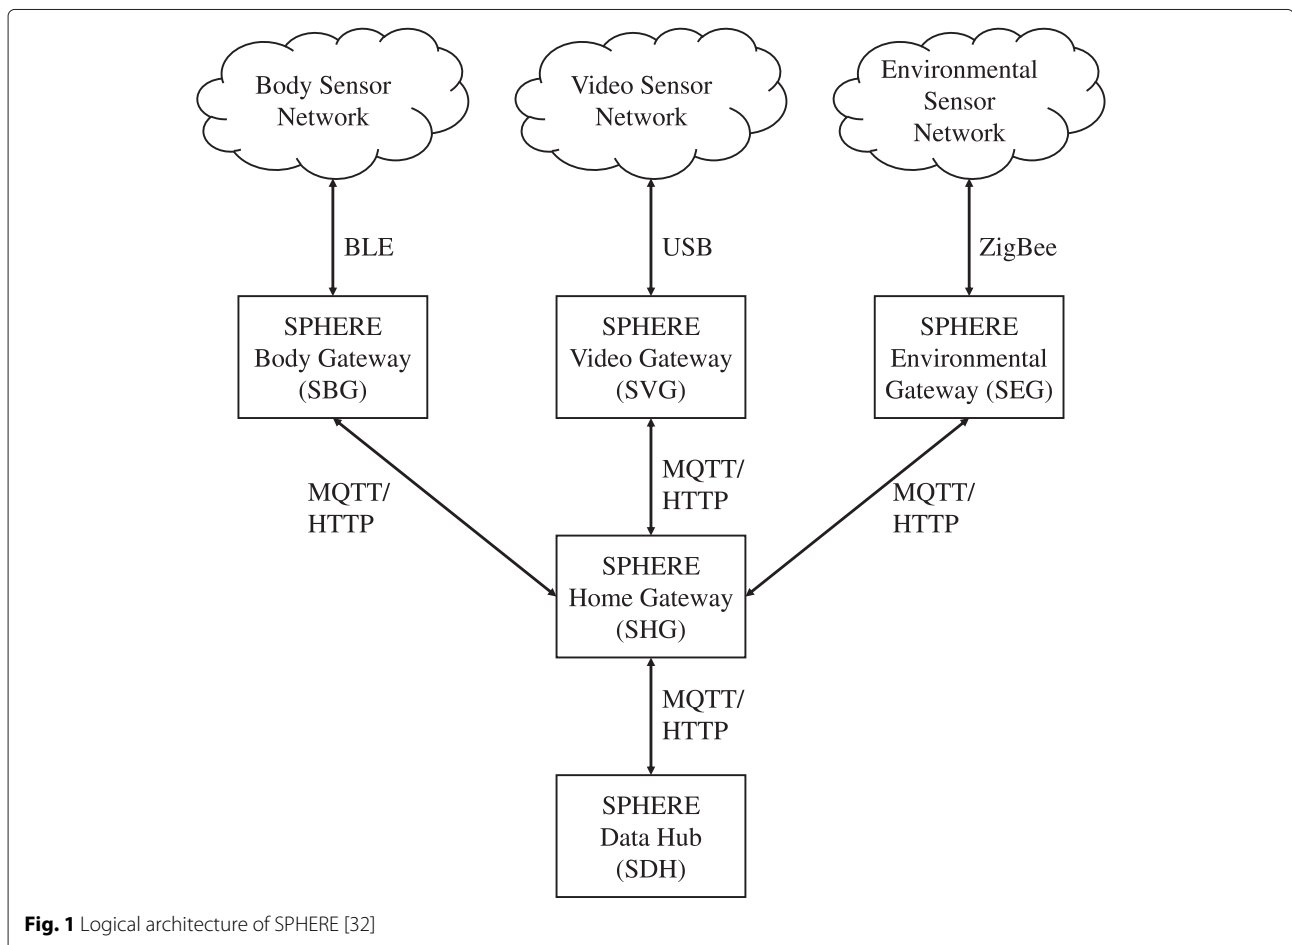
On a different perspective, related work on wireless body area networks (WBANs) typically focuses on wireless communications and the consumption of the radio, often disregarding the power consumption of other parts of the system (for instance, see [21] and [22]). In cases where energy harvesting technologies are employed, related works have proposed adaptive protocols that deal with the challenge of providing the required quality of service under spatio-temporal uncertainty in the energy input [23–26]. For an in-depth survey on WBANs, we refer the reader to [27] and [28]; whereas [29] presents a survey on WBANs with special interest in eHealth. For a thorough investigation on low-power medium access control (MAC) protocols for energy-constrained wireless networks, we refer the reader to [30] and [31].

Overall, we observe that activity monitoring systems tend to rely on standardised protocols for communication, such as BLE, ZigBee and legacy Bluetooth. Unlike this paper, none of the surveyed systems pays the necessary

attention to wireless coverage and antenna design. Even in cases of custom hardware, details of the employed solutions are often not provided. Although it can be argued that for certain scenarios, coverage is not a challenge (e.g. streaming to a smart phone), careful antenna design and transmission power optimisation translate to significant energy savings. Unlike wireless coverage, energy consumption attracts more attention in the reviewed literature. Nevertheless, we often observe that it is not documented in depth.

### 3 On-body sensing in SPHERE

The SPHERE system [32] is a multi-modal architecture that is composed of three sub-systems, as shown in Fig. 1. The first sub-system of the SPHERE architecture is associated with on-body sensing. It is composed of a wearable activity sensor and a set of mains-powered receivers that are distributed across the house. Using BLE, the receivers collect the acceleration data from the wearable and forward them to the SPHERE Body Gateway (SBG). The second sub-system is a network of video cameras, which similarly stream their recordings



**Fig. 1** Logical architecture of SPHERE [32]

to the SPHERE Video Gateway (SVG). The third sub-system is a network of mains-powered nodes, equipped with various sensors, including temperature sensors, humidity sensors, noise sensors, proximity sensors and power sensors (for monitoring the consumption of appliances). Using ZigBee for wireless communications, the sensors relay their data to the SPHERE Environment Gateway (SEG).

The three gateways (SBG, SVG and SEG) have a dual role within the SPHERE architecture. Primarily, they are responsible for formatting the data accordingly and forwarding them to the SPHERE Home Gateway (SHG) for permanent data storage in the SPHERE Data Hub (SDH). Message Queue Telemetry Transport (MQTT) and Hypertext Transfer Protocol (HTTP) are used for communication between the gateways. Moreover, the gateways implement online artificial intelligence algorithms that interpret the data and infer conclusions. In practice, when a particular algorithm is fusing data from the three networks, it is implemented in the SHG. Otherwise, it is implemented in the gateway of the respective sub-system.

In the context of the SPHERE architecture, the focus of this paper is on the sub-system of the on-body sensing network. Sections 4 and 5 discuss the design and implementation considerations of the wearable device and the smart home infrastructure, respectively.

#### 4 The wearable device

On-body sensing in the SPHERE architecture is primarily realised through a wrist-worn wearable device that is equipped with an accelerometer, a micro-controller unit (MCU) and a radio for wireless communication.

The system design is guided by the unique requirements of long-term residential monitoring. Aiming at a maintenance-free monitoring period of several months or years, user acceptance and comfort is a major requirement. For this purpose, the wrist is selected as the body position for the wearable device. A wrist-mounted device is widely considered as the most socially acceptable and least invasive choice for the subject's everyday routine, given that people of both sexes commonly wear wrist-worn gadgets, such as watches and bracelets. Alternative positions, such as on the chest or waist, have the risk of being removed by the subjects, compromising the effectiveness of the system. Social studies [33, 34] have shown the importance of wearable devices being comfortable and not intrusive to the daily life activities. Moreover, a wrist-worn device is able to detect everyday activities that do not involve the torso moving, such as eating or brushing teeth. In [35], the authors assess various body positions and present comparison results in which the wrist ranks high in all the considered activities in terms of classification accuracy.

Battery lifetime is another aspect of long-term monitoring that presents a major challenge to our design. Due to the nature of our target group, the system cannot depend on the user for maintaining the system and regularly recharging the batteries. Indeed, our vision is a completely maintenance-free solution for the whole monitoring period or indefinitely through energy harvesting. For this vision to be materialised, the system is designed with energy efficiency as the primary goal. Such a goal has direct effect on the design of every part of the system. The most energy-efficient hardware components need to be employed; the circuit board must be stripped down from any unnecessary components that lead to unwanted leaking currents; the antenna needs to provide the necessary coverage at the lowest transmission power level possible; and the firmware must be stripped down to only the absolutely necessary protocols and algorithms.

On the other hand, an energy-conscious design leads to compromises at the quality of the data that the wearable sensor provides to the reasoning algorithms of the higher layers of the architecture. Although long-term activity monitoring requires relative low sampling rates (10 Hz is sufficient according to [35]), tight energy budgets introduce a trade-off between the quantity (i.e. frequency) and quality (i.e. resolution) of the sensed data, and the battery lifetime. When raw data are not necessary, on-board processing and feature extraction can reduce the energy consumption without compromising the efficacy of the data.

In this section, we present the wearable design for ultra low-power long-term monitoring. Section 4.1 presents the hardware, Section 4.2 presents our antenna design and Section 4.3 presents the firmware.

##### 4.1 Hardware design

The core component of the wearable device is the nRF51822 SoC (system-on-chip). The nRF51822 is an ultra low-power solution for wireless applications. It incorporates a 2.4-GHz, BLE-compatible radio [36], and a 32-bit ARM Cortex M0 MCU. Additionally, it has 31 flexible GPIO (General Purpose Input/Output) ports that support several communication protocols, including Universal Asynchronous Receiver/Transmitter (UART), Synchronous Peripheral Interface (SPI) and I<sup>2</sup>C (integrated circuit). The nRF51822 operates at 1.8 V, whilst it also includes a linear regulator that allows a wider range of supply voltage from 1.8 to 3.6 V. BLE is primarily selected for its high energy efficiency [37–39]. Related work suggests that its energy consumption is less than half of ZigBee for packets of 39 bytes [37]. Moreover, its compatibility with smart phones is an additional merit that can be exploited to increase the coverage of the activity monitor beyond the limits of the house.

The nRF51822 is interfaced over the SPI bus with the ADXL362 accelerometer. The ADXL362 is an ultra low-power 3-axis accelerometer that provides measurements of 12 bits of resolution; 8-bit formatted data are also provided for more efficient single-byte transfers when a lower resolution is sufficient. It supports up to 400 Hz of sampling rate and the measurement range can be configured at either  $\pm 2$ ,  $\pm 4$  or  $\pm 8$  g. Additionally, ADXL362 incorporates a temperature sensor (12 bits of resolution) and a first-in-first-out (FIFO) buffer. Lastly, it has two input/output pins than can be configured to generate interrupts or to synchronise the measurements to an external clock. In addition to the standard SPI wiring, we power the ADXL362 directly from a GPIO of the nRF51822. This allows us to completely disable the accelerometer from software when it is not needed.

**4.2 Antenna design**

Antennas are of great importance for any communication system as they determine the quality and the reliability of the wireless link. In body-centric communications, in particular, the human body imposes a number of extra challenges such as electromagnetic energy absorption, polarisation misalignment and propagation shadowing. Furthermore, the need for small wearable devices makes the design of efficient antennas even more challenging.

Many studies have investigated the challenges and the impact of antennas in body-centric communications [40–42]. Measurements in [43] demonstrated (and quantified) the advantage of using a directional patch antenna for the on-body node in an indoor off-body communication scenario, in which a wearable device is connected to a nearby access point. We therefore designed a low-profile (dimensions:  $17.8 \times 18.5 \times 1.3$  mm) patch antenna on RT/Duroid 6010 substrate (dielectric constant  $\epsilon_r = 10.2$ ) for our wearable device. A picture of the prototype antenna is shown in Fig. 2. The directional nature of the patch ensures that the energy is radiated towards the


room and not to the body of the user, whereas the ground plane minimises the effect of the body on detuning. Two additional issues that needed to be addressed were adequate coverage of the whole 2.4 GHz, which is challenging because of the small form factor of our antenna, and robustness to polarisation misalignment due to random orientation of the wrist-worn wearable device. In order to achieve this, the dimensions of the patch were not square but slightly rectangular and it was fed along its diagonal dimension. In this way, two modes with resonant frequencies very close to each other were excited and, hence, the full band could be covered. The two modes are radiating in orthogonal polarisations, as the sources of radiation are the orthogonal patch edges.

Figure 3 shows the reflection and the transmission coefficients of our patch antenna. The two coefficients are related as follows:

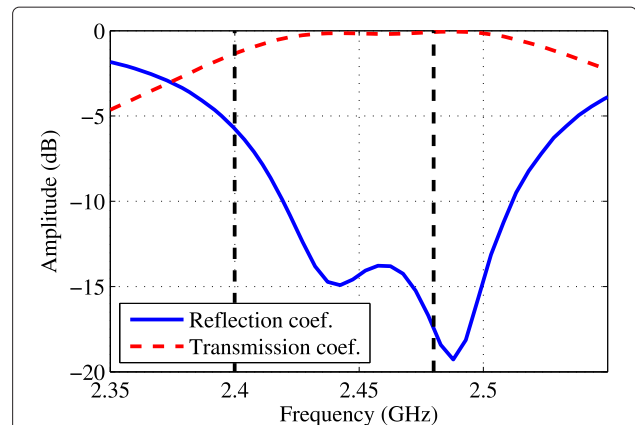
$$T_{11} = 10 \log_{10} (1 - |S_{11}|^2) . \tag{1}$$

$T_{11}$  is the transmission coefficient in decibel and  $S_{11}$  is the reflection coefficient in linear values. The former essentially shows how much energy is transmitted by the antenna (without accounting for dielectric/conduction losses). We can notice that a small amount of fine tuning is required in order to minimise the mismatch losses at the lower frequencies of the 2.4-GHz band (up to approximately 1 dB).

Figure 4 shows the 3D far-field radiation patterns of the patch antenna at the lower end (2400 MHz), the middle (2440 MHz) and the higher end (2480 MHz) of the 2.4-GHz frequency band. The patterns refer to electric field levels and were measured in an anechoic chamber. In the far-field region, where the propagating wave has a plane wavefront, the electric field can be resolved in two components, vertical and horizontal (with respect to

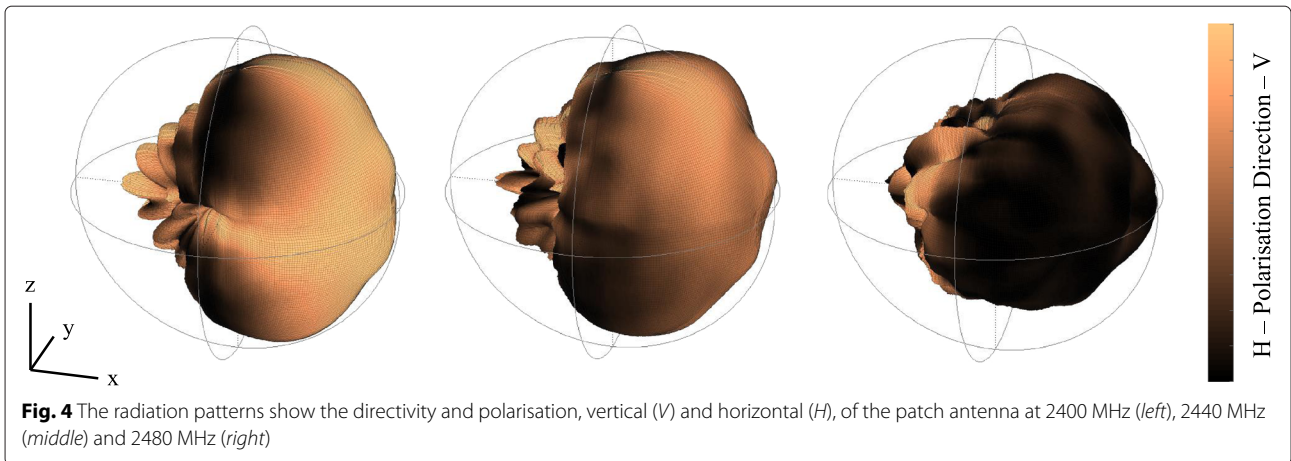


**Fig. 2** The wearable sensor prototype



**Fig. 3** Reflection and transmission coefficients of the patch antenna. The vertical lines indicate the limits of the 2.4-GHz frequency band





the coordinate system shown in Fig. 4). These components are orthogonal to each other and to the direction of propagation. The third component, which lies along the direction of propagation, has a level that falls off at a higher rate with distance and hence in the far-field region is negligible. The distribution of the radiated power between the vertical and horizontal components is shown in Fig. 4 using the colour code and is also summarised in Table 1. We can initially notice that the polarisation split changes with frequency (the pattern is essentially rotated around the  $x$ -axis). Seventy-seven percent of the power is concentrated in the vertical polarisation component at 2400 MHz, whereas this percentage drops to only 9 % at 2480 MHz. The split is more balanced at 2440 MHz, with 47 % of the power in the vertical polarisation component. We can also see that the antenna radiates directionally towards the  $+x$ -axis with a maximum directivity of 6.2, 6.8 and 8 dBi at 2400, 2440 and 2480 MHz, respectively. The higher directivity at the high frequencies is due to ripples in the pattern, caused by imperfections during the manual fine tuning of the antenna (scratched edges).

We should mention that the total efficiency (including mismatch and dielectric/conduction losses) of the antenna was measured (at 2440 MHz) to be 55 %, following the method described in [44]. The measurement was relative to a highly efficient reference antenna radiating at the same frequency as the antenna-under-test, and multiple repetitions showed the accuracy of  $\pm 5$  %. We should note that all antenna measurements reported in this section were performed with the antenna

element mounted on a large ground plane. Ongoing work focuses on developing a phantom that will allow the measurement of the impact of the human body on the antenna characteristics.

Lastly, a system-level antenna comparison study that includes our patch antenna is documented in [45] and demonstrates a performance comparable to an external bulky monopole, whilst yielding far superior performance than a printed alternative.

### 4.3 Firmware design

According to the BLE specifications [36], a device can operate in either of four distinct roles, namely *broadcaster*, *observer*, *peripheral* and *central*. The two former roles are based on a connectionless unidirectional communication model where the *broadcaster* is the transmitter and the *observer* is the receiver. The two latter roles are based on a connection-oriented model, where the *Central* is the master and the *peripheral* is the slave. Due to the nature of long-term residential monitoring (i.e. ultra low power consumption and full-house mobility), we opt for the connectionless mode, in which the wearable device acts as a *broadcaster*. The connectionless mode of BLE operates as follows. At a minimum period of 100 ms, the *broadcaster*, i.e. the wearable device, transmits a series of packets that are called BLE advertisements. In BLE, the 2.4-GHz band is divided in 40 frequency channels of 2 MHz, three of which are dedicated for advertisements. The advertisement channels are the channels 37 (at 2402 MHz), 38 (at 2426 MHz) and 39 (at 2480 MHz). On each transmission event, the BLE specifications allow the *broadcaster* to transmit on any combination of the three advertisement channels. Each advertisement is a packet of up to 39 bytes. Part of these 39 bytes are used for addressing and other headers. Up to 26 bytes per advertisement are available for the application layer. On the infrastructure side, the *observer* is scanning, i.e. listening to, the advertisement channels and collects the advertisements.

**Table 1** Patch antenna measured pattern statistics

Frequency (MHz)	Maximum directivity (dBi)	Power in vertical (%)	Power in horizontal (%)
2400	6.3	77	23
2440	6.8	47	53
2480	8.0	9	91

The simplicity of the connectionless mode of BLE makes it an ideal choice for our application of interest. Its advantages can be summarised as follows. First, the three advertisement channels are carefully selected to fall in between the popular IEEE 802.11 channels 1, 6 and 11. Since these WiFi channels are most commonly used [46, 47], the interference between WiFi and BLE is mitigated. Most importantly, when the user is moving, handovers to other *observers* are seamless, quick and without overhead. In fact, as shown in the following sections, the smart home infrastructure is leveraging parallel receptions from multiple receivers to mitigate the packet loss. Lastly, the minimalistic design of connectionless BLE implies low overhead, with direct benefits on energy consumption. In short, there are no transmissions of acknowledgements, no idle listening overheads for receiving data and no connection establishment overheads. The main weakness of connectionless BLE is packet loss due to the lack of retransmissions. In Section 5.1, we introduce a set of alternative techniques that the infrastructure employs to mitigate packet loss.

Whilst the system can fundamentally support a vast amount of firmware variations, in this paper, we consider raw data streaming, as well as on-board feature extraction with histograms [48, 49]. The latter is a means to keep both the power consumption low and the measurement frequency high. The idea is to monitor a time window and generate histograms, i.e. an empirical approximation of the probability density function, of the acceleration profile in the three axes. Histogram generation is a function of several parameters, such as the window size, the bin size and the number of bins. As the scope of this paper is on the infrastructure of the system, we do not fine-tune these parameters. Instead, we provide an indicative example.

We generate histograms of 8 bins of 1 byte for each of the 3 axes; a total of 24 bytes. Each bin is, practically, counting the number of samples that fall within its range. Each particular sample is shifted appropriately, so that the 3 most significant bits are used as a bin index. At 50-Hz sampling rate and with 8 bits of resolution, the FIFO buffer fits 128 samples (2.56 s). The time window is set to 5.12 s. During the time window, we calculate histograms of 256 samples (2 FIFO buffers). At the end of the window, the 24 bytes of the histograms, along with a sequence number, are packed in an advertisement and transmitted over BLE. We note that this solution compresses 256 3D samples into a single advertisement. In other words, it keeps radio duty cycle low, despite the high sampling frequency.

## 5 The smart home infrastructure

The key idea of the infrastructure design is to maintain the wearable sensor simple and energy-efficient, by applying mechanisms that mitigate packet loss in the infrastructure. Indeed, a network with wearable sensors

is highly asymmetrical. In contrast to the wearable device discussed in the previous sections, the receiver does not have power constraints, as it is mains-powered. Therefore, the design principle is to incorporate into the receiver features that unburden the wearable device from power-consuming tasks.

In this section, we present the smart home infrastructure composed of a distributed network of receiver units. In particular, Section 5.1 presents four packet loss mitigation techniques, whereas in Section 5.2, we combine these techniques in a distributed network of receiver units that deliver the data to the SHG [50].

### 5.1 Packet loss mitigating techniques

In a standard BLE-based communication system, data bits from the upper layers are first encapsulated into physical (PHY) layer packets. Additional 24 redundant bits are then appended to each packet by the CRC encoder and are used to detect errors at the receiver (hence the name of the code, CRC-24). The encoded data is then modulated and up-converted onto a carrier frequency. At the receiver, the arrived signal is down-converted and demodulated. The CRC check is then performed for each packet to detect any bit errors. If there is at least one bit error, the packet is discarded. Otherwise, the packet is forwarded to the upper layers of the receiver. In the connectionless mode, dropped packets are experienced as application-layer data loss. It is clear that an off-the-shelf connectionless BLE system performs well only if the probability of bit error is very low.

#### 5.1.1 Orthogonally polarised antennas

One of the simplest ways to improve the reliability of a practical BLE system is to introduce two orthogonally polarised antennas at the receiver [51]. In practice, each receiver unit is composed of a Raspberry Pi B+ micro-computer interfaced, over UART, with two nRF51822 BLE receivers having antennas with orthogonal polarisations (i.e. a vertical and a horizontal dipole). For details on the receiver dipole antennas, we refer the reader to [43]. Considering that one end of the link, i.e. the wearable sensor, is unpredictably polarised, as it is following the movements of the user, such antenna arrangement would increase the probability of successful reception. The input streams from the two radios are required to be combined into one stream after the CRC check. The cost of the solution is a minimal increase in computation complexity at the receiver. However, there are no costs for the energy-constrained transmitter.

#### 5.1.2 CRC error correction

Whilst the CRC code employed by BLE is used for error detection only, it has an inherent error correction potential due to redundancy it introduces to transmitted data.

CRC error correction would directly improve the packet reception rate at the receiver, thus, for example, allowing it to be placed further from the transmitter.

In our previous works [52, 53], we introduced a novel approach to the error correction of CRC codes based on iterative decoding techniques. In particular, it was shown how the BLE CRC code can be converted to an equivalent low-density parity check (LDPC) code by making its parity check matrix sparse using the algorithm proposed in [54]. The most widely used decoding approach for LDPC codes is based on belief propagation (BP) [55]. As an alternative to BP, a linear programming approach to the decoding of LDPC codes was proposed in [56]. The resulting algorithm known as the alternating direction method of multipliers (ADMM) was developed in [57] and [58]. In [52] and [53], we studied both BP- and ADMM-based decoding algorithms in the context of the BLE CRC code using simulations and real data, respectively. The ADMM-based approach showed a superior performance to that of BP, so it is employed in the SPHERE infrastructure.

The iterative decoding methods described above use soft information about the bits to be decoded in the form of log-likelihood ratios (LLRs). The LLR of a bit indicates how close the bit value is to 1 or 0 based on channel measurements. In a practical BLE system, the LLRs are not directly available at the receiver, since the demodulator produces hard decisions in the form of 1 and 0 s. In [52], it is shown that the transmitter, channel and receiver can be modelled as an equivalent binary symmetric channel (BSC) that flips some of the bits encoded by the CRC code with a probability  $p$ . The LLR for the  $i$ th received bit  $r_i$  can be calculated as

$$\gamma_i = (2r_i - 1) \ln \left( \frac{1-p}{p} \right). \quad (2)$$

In a practical BLE system, some indication of the quality of the received signal is usually available. For example, in our platform, the received signal strength indication (RSSI) is available for every received packet. Therefore, a statistics on the bit flipping probability (or the bit error rate, BER)  $p$  as a function of the RSSI level can be measured in advance by sending some known packets. This statistics can then be used as a look-up table to calculate LLRs in a real scenario.

### 5.1.3 Application-layer coding

The packet loss mitigation techniques considered so far did not change the way the information is transmitted. At the same time, it is clear that if some error correction coding was introduced amongst the data bits prior to transmission, it would allow more packets to be corrected than by means of CRC correction only. We call this approach application-layer coding (ALC), since it applies to the actual application data bits, before sending them

to the PHY layer. For a general communication system, error correction coding offers a compromise between the amount of data that can be corrected and the level of redundancy. In one extreme case, no redundancy offers the maximum bandwidth utilisation but at the same time the highest error rate. In another extreme case, data can be transmitted repetitively (using the so-called repetition coding), which provides excellent error correction capabilities, but the bandwidth utilisation reciprocal to the repetition factor.

Compared with a traditional, non-constrained communication system, low-energy systems have an additional factor to take into account—the energy consumed to transmit data. It is therefore natural to define the energy efficiency of a system as a ratio between the energy available at the transmitter and the amount of information reached the receiver. If error correction coding is applied, the encoding process will consume some additional energy at the wearable device. Let  $n$  be the total number of bits after encoding. The number of original message bits can be expressed as  $m = nr$ , where  $r \leq 1$  is the code rate. The energy efficiency can now be defined as

$$\eta = \frac{E + \delta}{m(1 - \text{MER})} = \frac{E + \delta}{nr(1 - \text{MER})}, \quad (3)$$

where  $E$  is the energy required to transmit an uncoded message,  $\delta$  is the additional energy consumed by the encoding process and MER is the message error rate after error correction. By fixing  $E$  and  $n$ , (3) can be used to benchmark and compare the performance of different coding techniques: the lower the  $\eta$ , the more efficient the technique is.

In our context, we exploit the fact that the transmitter sends three advertisement packets on different carrier frequencies, which can be viewed as three independent channels. The natural options are therefore to encapsulate *unique* message bits into each packet (no ALC,  $r = 1$ ), to encapsulate the *same* message bits into each packet (repetition coding,  $r = 1/3$ ) or to employ some encoding between message bits before sending them to the PHY layer. For the latter case, we employ a simple encoding scheme based on the XOR operation: a message is split into two parts,  $A$  and  $B$ , with  $A$  being sent on the first channel,  $B$  on the second and  $A \oplus B$  on the third, where  $\oplus$  denotes a bit-wise XOR operation. The resulting code rate is  $r = 2/3$ . At the receiver, it is clear that to receive the message correctly, it is sufficient to successfully receive any two out of the three packets. Let  $q$  denote the PHY PER. Assuming that each packet experiences an independent loss, the MER for the XOR-based ALC can be calculated as follows:

$$\text{MER}_{\text{XOR}} = q^3 + 3(1 - q)q^2 = 3q^2 - 2q^3. \quad (4)$$



The MER for repetition coding and the uncoded systems can be similarly obtained as follows:

$$\text{MER}_{\text{rep}} = q^3; \quad (5)$$

$$\text{MER}_{\text{unc}} = 1 - (1 - q)^3 = 3q - 3q^2 + q^3. \quad (6)$$

By substituting (4), (5) and (6) to (3) and using the corresponding code rates, the three techniques can be compared in terms of the energy efficiency.

Application-layer coding within the three advertisement channels provides resilience not only against interference but also against polarisation mismatch by utilising the fact that the wearable patch antenna has different polarisation components on different channels (Fig. 4).

#### 5.1.4 Multiple receivers

Since the transmission power of BLE radios is typically low, multiple receivers are necessary to provide signal coverage over a residential house. Due to the broadcasting nature of connectionless BLE, traffic combining can be performed at the SBG. Indeed, not only the coverage is increased but also different pieces of a broken message (if ALC is enabled) may reach different overlapping receivers. With no additional energy overheads for the energy-constrained wearable device, this would increase the probability of successful reception and allow the wearable to operate at smaller transmission power levels compared to a single receiver scenario.

#### 5.2 Distributed network of receiver units

In this section, we combine all error correction techniques considered so far and implement them in a distributed system consisting of multiple receivers and the SBG that is responsible for delivering the data to the SHG. To this end, we establish two performance goals. Primarily, we aim at minimising the PER, thus maximising the reliability of the system. Secondly, we spread the computational burden across the network and reduce the traffic between the devices.

Combining packets from two antennas with orthogonal polarisations, due to its simplicity, is implemented on each receiver. The combining process can be described as follows. If the same packet is received in both polarisation and if at least one copy is correct (i.e. it passed the CRC check), the other copy is discarded. In this way, the amount of forwarded traffic is reduced. If both copies are incorrect, they are both forwarded further, so that the CRC error correction algorithm can combine the reliability information. Combining packets from multiple receivers in the SBG is implemented in the same manner.

Of all four techniques, CRC error correction is the most computationally intensive. Since the same packet can potentially be received by all receivers and in both polarisations, attempting to correct errors in each copy of the packet would lead to unnecessary processing. On

the other hand, combining all copies of the packet and removing unnecessary duplicates in the central device can significantly reduce the total number of corrupted packets. Therefore, it is reasonable to perform CRC error correction block in the SBG. By contrast with the single receiver scenario, soft information about bits should now be computed based on multiple copies of the same packet. Assuming that all received copies are statistically independent, (2) can be rewritten as

$$\gamma_i = \sum_j (2r_{i,j} - 1) \ln \left( \frac{1 - p_j}{p_j} \right), \quad (7)$$

where  $r_{i,j}$  denotes the  $i$ th bit of the  $j$ th copy of the packet and  $p_j$  is the RSSI-based bit-flip probability. By summing up the LLRs for statistically independent observations, the reliability estimation is averaged out.

Application-layer (AL) decoding, by definition, is normally applied at the end of a signal processing chain. Following the strategy described in Section 5.1.3, let the encoded sequence of packets be  $A$ ,  $B$  and  $A \oplus B$ . It can be observed that such decoding procedure, due to its simplicity, could be applied not only after CRC error correction but also after polarisation combining, with the aim of reducing the amount of traffic sent to the SBG: if any two out of three packets are correct, the third packet becomes redundant. Similarly, after packets from all receivers are combined in the SBG, it may be possible to reconstruct both  $A$  and  $B$  without resorting to CRC correction. In total, AL decoding can be applied after each of the other three techniques described above.

Algorithm 1 summarises the order and brief description of the error correction techniques, which is also visualised in Fig. 5.

---

#### Algorithm 1 Distributed error correction

---

**Input:** Encoded message comprising of packets  $A$ ,  $B$  and  $A \oplus B$  fully or partly received by multiple receivers.

**Output:** Decoded packets  $A$  and  $B$ .

*Local receiver processing:* on each receiver do:

Step 1: **Polarisation combining:** for each unique received packet:

**if** only one copy is received, forward it;  
**elseif** both copies are not correct, forward both;  
**else** forward the correct copy and discard the other.

Step 2: **AL decoding:** attempt to decode  $A$  and  $B$ . If successful, forward both and discard redundant copies. If not, forward all copies.

*SBG:*

Step 3: **Combining traffic from multiple receivers:** for each unique received packet:

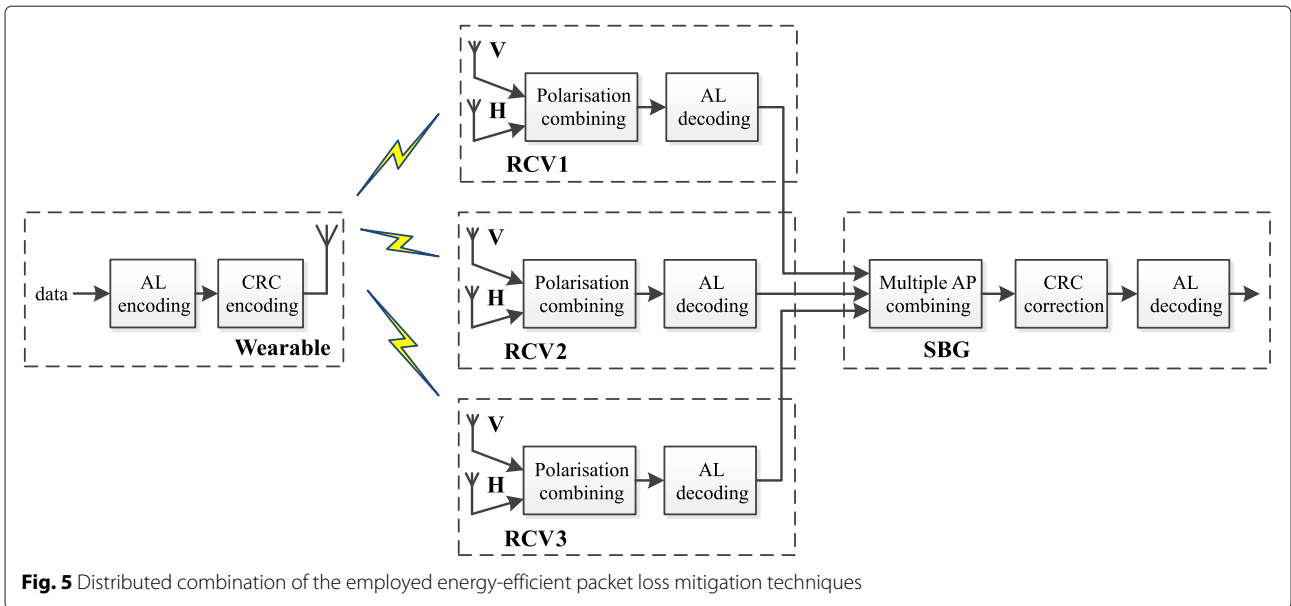
**if** only one copy is available, forward it;  
**elseif** all copies are not correct, forward all;  
**else** forward just one correct copy and discard others.

Step 4: **AL decoding:** repeat Step 2.

Step 5: **CRC correction:** apply Algorithm 1 from [52].

Step 6: **AL decoding:** repeat Step 2.

---




### 6 System evaluation

The evaluation of the system is focused on the energy consumption of the wearable, the provision of full-house wireless coverage, and the reliability of the system, i.e. the effectiveness of the packet loss mitigation techniques that are employed by the infrastructure. Additionally, we evaluate the accuracy of the RSSI and the transmission power reported by the nRF51822. For the evaluation experiments, we use a prototyped wearable device (Fig. 2) and three prototyped receiver units (Fig. 6) that are based on the nRF51822 development kits.

#### 6.1 Energy consumption of the wearable device

In Section 4.1, we selected state-of-the-art components to design our wearable device. In this section, we present the power consumption evaluation of the wearable device in contrast to the alternative candidates. In particular,


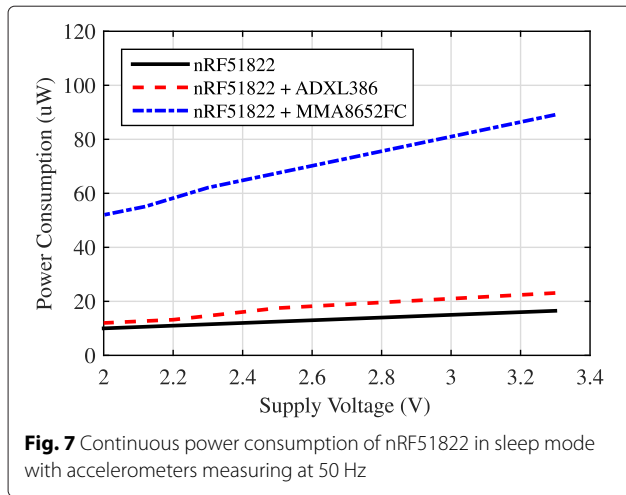


**Fig. 6** The receiver prototype that employs two orthogonally polarised antennas

we will compare the ADXL362 accelerometer to the MMA8652FC which is an I<sup>2</sup>C-based 3D accelerometer. This also implies comparison between SPI and I<sup>2</sup>C for data transfer between the accelerometer and the MCU. The nRF51822 is compared to another high-end off-the-shelf 1.8-V BLE SoC, which, for the remainder of the paper, will be referred to as the BLE SoC.

We begin the evaluation with the continuous power consumption of the wearable device. This includes all the sources of power consumption that are not duty-cycled, such as the accelerometer sampling at 50 Hz and the BLE SoC in the sleep mode. We measured the current drain of the system by measuring the voltage drop across a shunt resistor  $R_S = 100 \Omega$  placed on the positive side of the power supply. As the system is linearly regulated, we measured a constant 5  $\mu\text{A}$  current drain over the whole supply voltage range for the nRF51822 alone. The ADXL362 and the MMA8652FC increase the continuous current by an additional 2 and 22  $\mu\text{A}$ , respectively (notice one order of magnitude of difference). Figure 7 shows the continuous power consumption of the nRF51822 with and without an accelerometer for different supply voltages ( $V_S$ ). Figure 8 compares the continuous power consumption of the nRF51822 to that of the BLE SoC without any accelerometers. Notice that the BLE SoC is not linearly regulated, but even if we compare them at  $V_S = 2 \text{ V}$ , the BLE SoC idle current is one order of magnitude higher, 52  $\mu\text{A}$ .

Next, we measure the power consumption of using the MCU, i.e. transferring and processing acceleration data. In these measurements, we use a  $R_S = 10 \Omega$  shunt resistor. As shown in Fig. 9, the consumption of both chips is of the same order of magnitude (6 mA). Additionally, we





measured how much time the MCU needs to be active in order to collect a full FIFO buffer for the accelerometer. The ADXL362 (SPI clocked at 4 Mbps) needs  $\approx 13$  ms to transfer a 128-sample FIFO (total charge of 46.1  $\mu$ C). The MMA8652FC, on the other hand, needs  $\approx 2.5$  ms to transfer a 32-sample FIFO (total charge of 12.1 or 48.4  $\mu$ C to repeat it for 128 samples). The measurements show no significant difference between the two accelerometers in terms of data transfer.


Lastly, we measure the consumption during a triple advertisement transmission event, again using a  $R_S = 10 \Omega$  shunt resistor. Figure 10 shows the current profile during the transmission of a triple advertisement. We can clearly see in the profile an initialisation stage, followed by three peaks of high current, one for each advertisement. For both chips, the current profiles did not change with the supply voltage, which suggests linear regulation for both radios. Figure 11 experiments with different transmission power levels, ranging from  $-20$  to 4 dBm. In particular, we measure the total energy consumption of the transmission of a triple advertisement

( $V_S = 3$  V). The measurements suggest that the nRF51822 is approximately twice better than the alternative solution. Additionally, the measurements demonstrate how the transmission power affects the consumption, suggesting non-negligible energy savings at  $-4$  dBm. Lower transmission power settings appear to have limited value, as they compromise the wireless coverage without offering significant energy savings.

The consumption levels of sleeping, processing and transmitting can be combined to estimate the long-term average power consumption of different scenarios. For reference, Table 2 shows battery lifetime estimations assuming the transmission of raw accelerometer data for various sampling frequencies ( $f_s$  in Hz) and transmission power levels ( $P_{tx}$ ). The battery lifetime estimations are based on the following equation:

$$T = \frac{E_{BAT}}{P_I + (E_{SPI} + E_{BLE}) \times f_s}, \tag{8}$$





where  $E_{BAT}$  is the total energy of the battery;  $P_I$  is the continuous power consumption given in Fig. 7 (nRF51822 + ADXL386);  $E_{SPI}$  is the energy consumed for transferring a single sample over SPI from the accelerometer to the MCU;  $E_{BLE}$  is the energy consumed for the transmission of a single sample over BLE given in Fig. 11 divided by the number of samples inside a packet;  $f_s$  is the sampling frequency of the accelerometer. In raw data scenarios, the application layer populates the payload of a BLE advertisement with either 4 triaxial samples of 12-bit resolution or 8 triaxial samples of 8-bit resolution, respectively. The estimations assume the use of a 210-mAh coin cell battery (3 V). In [35], the authors use accelerometers with 8-bit resolution to perform activity classification. Experimenting with different sampling frequencies, the authors show that the performance of the classifier reaches a high level at approximately  $f_s = 10\text{ Hz}$  with only marginal improvement at higher frequencies. In this configuration, the expected battery lifetime of our system is approximated at 172 days at 4 dBm and 213 days at -4 dBm.

On-board processing can yield even longer battery lifetimes. Next, we evaluate the firmware implementation described in Section 4.3, in which the accelerometer sampling frequency is 50 Hz and the transmission period of the advertisements containing the histogram is 2.54 s. The transmission power is set at 4 dBm.

To increase our confidence in the measurements, we measure the overall consumption of the wearable system using an alternative method. We charge a 18-mF capacitor using a 3-V battery and then we use the capacitor to power the wearable device. Then, we can approximate the long-term average power consumption by measuring the voltage drop ( $V_2 - V_1$ ) of the capacitor over a known time interval (1 min). The energy stored in the capacitor was enough to power the system for over 7 min. Table 3 summarises the results. Note that the table overestimates the consumption as it includes the leakage of the capacitor. The measurements suggest that, with an efficient voltage converter, the wearable device yields approximately 60- $\mu\text{W}$  long-term average power consumption. Assuming a 210-mAh coin cell battery, this consumption level translates to an approximate lifetime of 437 days without battery replacement.

### 6.2 Evaluation of RSSI and transmission power levels

Our next goal in the evaluation of the wearable device is to measure the inaccuracy of transmission and reception power levels reported by the nRF51822. The nRF51822 supports 7 transmission power levels (from -20 to 4 dBm with a step of 4 dBm). In the next experiment, we connect the radio output of the wearable device directly to the input of a spectrum analyser, and we measure the power of the transmitted signal. We note that this experiment incorporates the losses in the matching circuit of the wearable device. Table 4 summarises the results for the 3 advertisement channels. We notice an approximately 1-dB difference between the two ends of the 2.4-GHz band and an error between the nominal and the measured power that becomes significant at the higher transmission power levels. Cable losses were measured and accounted for in the calculations.

Armed with the actual transmission power levels, we next measured the accuracy the reported RSSI by connecting the wearable device directly to the receiver (using low-loss wires and 30-dB attenuators). Figure 12 is a conversion figure that can be used to translate a reported RSSI to the actual received power level. We notice that we never receive a packet with the RSSI less than -102 dBm and that the maximum reported level is -44 dBm. Between these two thresholds, the error remains approximately

**Table 2** Battery lifetime approximations in days

$P_{tx}$	$f_s$	1	5	10	20	30	40	50
4 dBm	12-bit	591	195	106	55	37	28	22
	8-bit	749	301	172	92	63	48	38
-4 dBm	12-bit	680	250	139	74	50	38	30
	8-bit	817	361	213	117	80	61	49

**Table 3** System powered by a charged capacitor for 1 min

$V_1$ (V)	$V_2$ (V)	$\Delta Q$ (mC)	$I_{avg}$ ( $\mu A$ )	$P_{avg}$ ( $\mu W$ )
2.85	2.68	3.02	50.4	139
2.68	2.54	2.48	41.4	108
2.54	2.42	2.16	36	89
2.42	2.31	2.3	33.9	80
2.31	2.21	1.77	29.4	66
2.21	2.12	1.71	28.5	62
2.12	2.02	1.69	28.2	58

constant. The losses of the attenuators and the wires were measured and considered in the calculations.


**6.3 Wireless coverage and polarisation diversity**

In the next series of experiments, we evaluate the wireless coverage of the wearable device in the SPHERE house, providing insight on the benefits of employing multiple orthogonally polarised antennas. The SPHERE house is a two-storey building with concrete brick walls. For the following experiments, we consider three receivers in different locations, two in the lounge and one in the study room, and seven different wearable locations, one in each of the rooms, as shown in the 3D model of the SPHERE house in Fig. 13. In each of the two following experiments, a subset of these links is evaluated.

The first experiment focuses on measuring the off-body wireless channel gain in the 2.4-GHz band using a vector network analyser (VNA). We should note that the main difference between the channel gain and path loss is that the former includes the transmit and receive antenna gains. All channel measurements were performed with the patch antenna mounted on the wrist of a human subject (male, 176-cm height, 78-kg weight) in locations W1, W2 and W3. All three receivers were considered, and the links are classified into the following four categories: same room (RCV1-W1, RCV2-W1, RCV3-W2); one wall (RCV1-W2, RCV2-W2, RCV3-W1); two walls (RCV3-W3); and three walls (RCV1-W3, RCV2-W3).

**Table 4** Nominal vs. measured transmission power

Nominal	Measurements (channel 37)	Measurements (channel 38)	Measurements (channel 39)
4	1.3	1.5	1.8
0	-1.6	-1.1	0.6
-4	-6.1	-5.5	-5
-8	-9.5	-8.9	-8.4
-12	-13.3	-12.8	-12
-16	-17.3	-16.7	-15.9
-20	-20.7	-20.6	-19.4

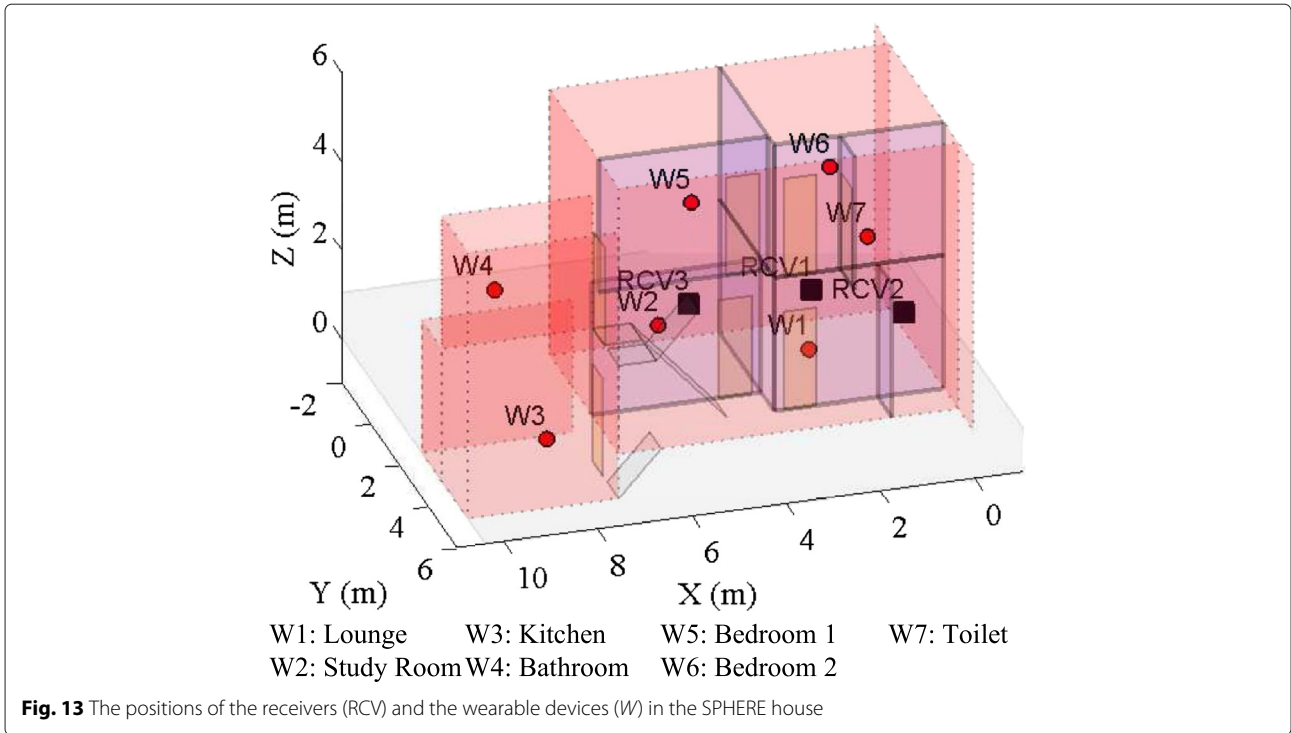


**Fig. 12** Mapping the reported RSSI to measured received signal power

Each transmitter-receiver link was measured separately. In order to investigate the effect of the body, the subject was sitting on a stool mounted on a turntable and being rotated in azimuth through 360° with the help of motors. Two different arm positions were considered, as shown in Fig. 14. The location and the orientation of the receiver remained fixed each time on a wooden table, and 801 samples of the channel gain were recorded every 30° using the VNA over the whole 2.4-GHz band (2400 to 2480 MHz). All measurements were performed twice (clockwise and anticlockwise rotation of the subject) and results have been averaged. We should also note that in this section, we report the channel gain value averaged over the whole frequency band (after having compensated for any extra gains and losses introduced for example by the amplifier and the cables used). Figure 15 shows an example of a channel gain measurement as a function of frequency and body azimuth rotation. The impact of the body shadowing and fast fading is clear, as large channel gain variations can be noticed (up to about 40-dB difference between the maximum and the minimum recorded value in this example).

Figure 16 shows the cumulative distribution function (CDF) of the difference (in dB) of the channel gain between the two antennas of each receiver for the four different link categories. We can see that when the transmitter and the receiver are in the same room, the difference between the two orthogonal receiving antennas is up to approximately 15 dB, with an average of about 6 dB. As the number of walls between the wearable device and the receiver increases, the difference becomes smaller but is still significant, with an average value of about 2.7 dB for the three investigated link categories. This highlights the advantage of exploring polarisation orthogonality at the receiver, especially for an off-body scenario in which the orientation (and hence the polarisation) of the wrist-worn antenna is continuously changing following the






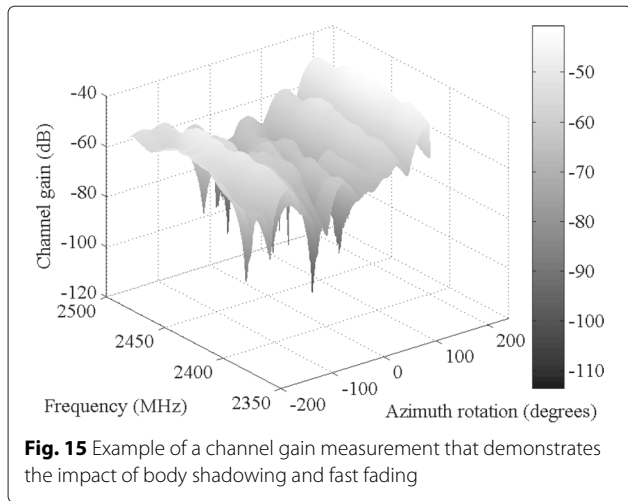
movements of the user. With both antennas at the receiver, we are able to use the one that gives the best reception at a per-packet basis.

Figure 17 shows the CDFs of the measured channel gain for the four link types. It should be noted that each CDF graph is over all the body azimuth rotations and the two arm positions, and it refers to the best of the two orthogonal antennas at the receiver each time. We can observe that the first (brick) wall introduces an extra loss of approximately 15 dB on average, with about 10-dB additional loss for each extra wall. Note that these numbers

include the additional path loss due to the longer distance between the transmitter and the receiver and do not refer to losses due to the wall only. Furthermore, the dynamic range of each CDF graph demonstrates the effect of body shadowing. We can see that when the wearable and the receiver are in the same or in adjacent rooms, the dynamic range of the channel gain is about 20 to 22 dB. Additionally, we can observe that the influence of the body decreases as we move towards a more difficult link and it drops to approximately 12 to 13 dB when the transmitter and the receiver are separated by two or three walls. This








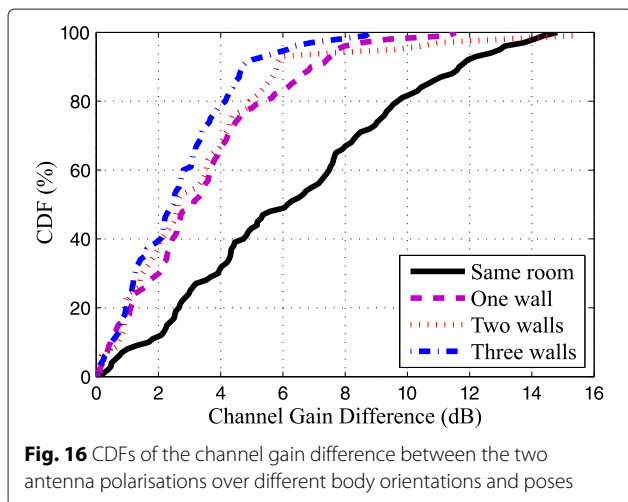
is because in scenarios with a strong dominant signal component, the reception is more vulnerable to changes in the body position. On the other hand, through-wall links depend less on a single dominant component and more on multipath ones, and hence are less sensitive to the body orientation.


The following experiments investigate the PER at the system level, aiming to identify the limits of wireless radio coverage of our BLE system in the house. The receiver is placed in the study room (RCV3), and we investigate all 7 links for the wearable considering all possible 7 transmission power levels. For each combination, we calculate the average RSSI and PER.

The system-level measurements verify the previous results, yielding a significant difference in the PER for the two orthogonal receiving antennas of the same link. In an attempt to highlight these differences, Fig. 18 demonstrates the PER of 5 characteristic links that result in large differences between the two antennas. For

example, in the most extreme case of the RCV3-W5 link, the horizontal dipole results in a PER of approximately 90 %, whereas the PER for the vertical one is only about 8 %. These differences are due to the orientation of the wearable antenna and the properties of the propagation environment (multipath propagation). Indeed, when a mobile link operates close to the sensitivity threshold, even small differences in the RSSI, caused by polarisation mismatch, can lead to a very high PER.

Figure 19 shows the cumulative distribution function (CDF) of the RSSI of the received packets that correspond to the horizontally and vertically polarised antenna, respectively, in measurements that involve a user performing random walks in all 7 areas of the house. The figure also plots the CDF curve after the two orthogonally polarised antennas are combined. Similarly





to the previous experiment, the unpredictably polarised wearable antenna occasionally matches the polarisation of the vertical receiver antenna, whilst other times, it matches the horizontal receiver antenna. By combining two orthogonally polarised antennas in a per-packet basis, the system increases the probability that one of them will match the unpredictable polarisation of the wearable antenna; thus improving the overall performance. Without any change in the energy-constrained wearable sensor, this technique accounts for an approximately 3-dB median improvement in the RSSI.

Table 5 summarises the results of the 7 links over the 7 different transmission power levels. For each combination, we show the PER (as a percentage) and the average RSSI (in dBm) for the best of the two receiving antennas. In the same room, we notice excellent reception at all transmission power levels. In the lounge and the first bedroom, which are one wall away from the receiver, we observe excellent reception at high transmission power levels that gradually drops as we decrease the transmission power level. Two walls away from the receiver, reception is a challenge: the results demonstrate a high PER and several links with no reception. We encourage the reader to compare the results of Table 5 and Fig. 11, which together demonstrate the trade-off between wireless coverage and

energy consumption. The combined measurements provide valuable information on how to tune the transmission power to different specifications, such as the size of the house, the battery lifetime requirements and the PER requirements. It can be seen, for example, that for a relatively small flat (studio or 1-bedroom), a transmission power of  $-4$  dBm would be able to provide sufficient in-house coverage, whilst requiring 33 % less energy for BLE packet transmissions. According to Table 2, this reduction corresponds to a 23.8 % increase in the battery lifetime (8-bit resolution,  $f = 10$  Hz).

The results presented in this section can be combined for additional insight into the behaviour of the system. As an example, we conclude this section by demonstrating how the PER in an unknown topology can be estimated. Table 5 suggests that errors begin at an RSSI of approximately  $-80$  dBm. The PER is very high at an RSSI of approximately  $-95$  dBm until an RSSI of  $-102$  dBm where the reception is completely lost. Using Fig. 12, we can map these RSSI thresholds as reported by the receiver to the actual received signal strength of  $-76$ ,  $-88$  and  $-94$  dBm, respectively. Using Fig. 17 and the link budget equation, i.e. adding the channel gain to the transmission power (Table 4), we can anticipate the PER in an unknown scenario. For instance, at a transmission power of 4 dBm and for a link between adjacent rooms, we expect a received signal between  $-79$  and  $-55$  dBm, which translates to the zero PER generally and a small PER ( $<5$  %) in the worst case body orientation scenarios.

#### 6.4 CRC error correction and application-layer coding

The evaluation of the CRC error correction is based on a large set of corrupted packets that is collected in an office environment. In this case, the receiver was fixed on a desk, whilst the wearable transmitter was positioned in different locations within the office and the surrounding areas. To minimise interference, the data collection sessions were performed at off-peak hours when the office was empty. In a period of over 10 h, we collected a total of approximately 400,000 packets, 6000 of which were corrupted.


**Table 5** PER (%) and average RSSI (in dBm) for different links and transmission power levels (NR corresponds to no reception)

ID	Room	Walls	+4 dBm	0 dBm	-4 dBm	-8dBm	-12 dBm	-16 dBm	-20 dBm
W2	Study	0	0.0 ( $-51$ )	0.0 ( $-52$ )	0.0 ( $-58$ )	0.2 ( $-61$ )	0.0 ( $-65$ )	0.2 ( $-70$ )	0.4 ( $-71$ )
W5	Bedroom 1	1	0.3 ( $-70$ )	0.7 ( $-73$ )	0.2 ( $-75$ )	3.9 ( $-83$ )	5.8 ( $-84$ )	4.6 ( $-90$ )	7.4 ( $-95$ )
W1	Lounge	1	0.0 ( $-74$ )	0.5 ( $-77$ )	0.4 ( $-83$ )	1.1 ( $-85$ )	6.1 ( $-87$ )	15.7 ( $-94$ )	NR
W4	Bathroom	2	0.0 ( $-80$ )	1.2 ( $-85$ )	21.4 ( $-93$ )	NR	NR	NR	NR
W3	Kitchen	2	0.8 ( $-82$ )	2.3 ( $-84$ )	15.8 ( $-93$ )	NR	NR	NR	NR
W6	Bedroom 2	2	4.5 ( $-88$ )	4.5 ( $-88$ )	69.6 ( $-95$ )	58.0 ( $-92$ )	NR	NR	NR
W7	Toilet	3	31.4 ( $-96$ )	64.0 ( $-99$ )	NR	NR	NR	NR	NR

Figure 20 shows the PER for each RSSI level before any correction algorithm is applied (solid line). We can observe that without error correction, the 1 % PER threshold is at  $-91$  dBm. On the other hand, when ADMM is enabled, a gain of up to 3 dB can be achieved, i.e. the 1 % PER threshold is at  $-94$  dBm. In terms of the PER, error correction significantly improves the reliability, correcting up to 60 % of corrupted packets at  $-97$  dBm. Such improvements come without any change at the transmitter. Hence, CRC error correction is a means for improving the reliability of the system without compromising its low power consumption.


In contrast to CRC error correction and the other packet loss mitigation techniques, application-layer coding does affect the energy consumption of the transmitter. Considering the maximum permitted packet size, i.e. 39 bytes, each of the three BLE advertisement has a payload of 24 bytes. Therefore, a triple advertisement carries 72 payload bytes ( $n = 576$ ). The energy consumption of a triple advertisement event ( $E$ ) is given in Fig. 11. In the case of XOR-based ALC, an additional  $0.3 \mu\text{J}$  was recorded to be consumed by the encoding process. Assuming an error-free channel ( $q = 0$ ) and based on (3), Fig. 21 shows the energy consumption per bit for the cases of no redundancy, packet repetition and XOR-based ALC, for various transmission power settings. It can be seen that the use of XOR-based ALC increases the energy per bit required for transmission, primarily because of the 33 % reduction in channel utilisation that it introduces. Whilst ALC compromises the energy efficiency in error-free channels, in a real scenario, it can improve the reliability of wireless communication.

Assuming the maximum available transmission power setting (4 dBm), we next consider channel errors. In particular, we use the PER measured and shown in Fig. 20 (without CRC error correction) to calculate the energy

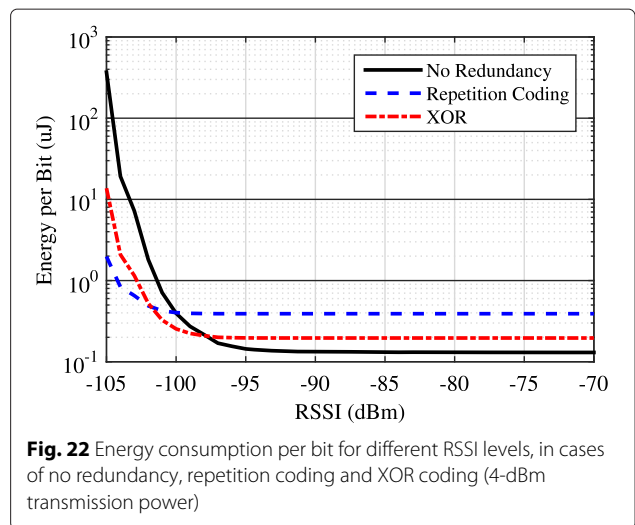


**Fig. 21** Energy consumption per bit for different transmission power levels, in cases of no redundancy, repetition coding and XOR coding ( $p = 0$ )

efficiency per bit (3) based on the MER of all three schemes given in (4), (5) and (6). Figure 22 plots the results as a function of the RSSI. It can be observed that each method performs best in a separate RSSI range. As expected, at high signal levels, any coding is redundant, resulting in some increase in the energy consumption. Repetition coding has the highest energy efficiency only when the signal level is extremely low, due to its superior error correction capability. The XOR-based method provides the lowest energy per bit between  $-102$  and  $-97$  dBm, whilst being more energy efficient than the repetition code at high signal levels. We note that the long-term optimum choice in a real scenario depends on the distribution of the received signal level. However, without a priori knowledge of such statistics, coding methods similar to XOR can be a middle-ground solution.



**Fig. 20** PER for various RSSI levels before and after CRC error correction



**Fig. 22** Energy consumption per bit for different RSSI levels, in cases of no redundancy, repetition coding and XOR coding (4-dBm transmission power)

### 6.5 Overall performance of packet loss mitigation

The performance of the distributed system, introduced in Section 5.2, is evaluated next. The evaluation is based on measurements in the SPHERE house in which the user is performing random walks in all 7 areas of the house. The transmission power of the wearable sensor is set at  $-4$  dBm. Three receiver units are deployed in the three considered locations (RCV1, RCV2 and RCV3).

Figure 23 illustrates the performance of the system in terms of the PER after each stage of error correction. The performance of the system without ALC is also shown as a reference. The initial PER prior to error correction was recorded at 7 % and is shown on the graph as a horizontal line. It can be observed that polarisation combining reduces the PER by 10 %, which suggests some degree of correlation between the two polarisations. At the same time, it can be seen that combining packets from multiple APs has a dramatic effect on the PER, reducing it further by 54 % without ALC and by 42 % with ALC. CRC error correction provides an additional reduction of 20 % in the PER.

When ALC is enabled and decoding is performed after each of the three other stages, it has a profound effect on the performance and the amount of traffic sent from the APs to the central device. Indeed, the final PER with ALC is only 1.4 %, 40 % lower than without ALC. In terms of the traffic flow, it was recorded that the number of packets sent to the central device was reduced by 30 % when AL decoding was enabled locally at each AP. Finally, the number of packets processed by the CRC error correction block was reduced by 45 %, which significantly decreases the computational load in the central hub. Overall, all techniques combined achieve 80 % reduction in the PER.

### 7 Discussion on miniaturisation


Armed with the results obtained so far, our next goal is to miniaturise the prototype and manufacture a compact circuit board for the wearable device. A miniaturised

version of the wearable prototype would employ the very same key components of the prototype, which are, by design, suitable for a tiny wearable board. Providing that the proposed (and evaluated) elements are employed, the miniaturised circuit board is expected to yield similar performance characteristics to the prototype, as documented in Section 6. These include (i) the patch antenna ( $17.8 \times 18.5 \times 1.3$  mm), (ii) a CC2032 coin cell battery (diameter 20 mm, height 3.2 mm) and (iii) ultra low-power electronic components, including the nRF51822 ( $3.5 \times 3.83$  mm) BLE radio and the ADXL362 ( $3 \times 3.25$  mm) accelerometer. Considering the height of an off-the-shelf battery holder and the height of the components on the main circuit board, the end system fits in an enclosure with dimensions of  $20 \times 20 \times 10$  mm. Therefore, the end platform can be as small as a typical wristwatch. In addition, both the wearable device and the receivers use low-cost components that are suitable for low-cost consumer electronics.

### 8 Conclusions

Smart homes and wearable technologies constitute a promising answer to the challenges introduced by changes in demographics and the rise of chronic medical conditions on our healthcare systems. In this paper, we present and evaluate the infrastructure of a long-term activity monitoring system that aims the early detection and cost-effective intervention of such emergencies. The prototype system is realised by a wearable BLE-enabled activity sensor. We evaluate the power consumption and wireless coverage of the system. Specifically, we provide detailed consumption measurements for all the key components of the system (i.e. accelerometer, radio, processor), comparing them to alternative components. Moreover, we evaluate the radio coverage of the system using both channel-level and system-level measurements, which quantify the coverage challenges and the effect of the body in the quality of the wireless link.

The coverage study suggests that a single receiver does not suffice for full-house coverage even at the highest transmission power level. For this reason, the SPHERE architecture incorporates several receivers. In fact, we employ multiple overlapping receivers which, together with other packet loss mitigation techniques (orthogonally polarised antennas, CRC error correction and ALC), provide an energy-efficient alternative to retransmissions. Overall, the recorded PER is measured at 1.4 % in an experiment where a user is performing random walks in all areas of the house. In addition to full-house coverage, this approach allows us to turn the transmission power down to  $-4$  dBm and save approximately 33 % energy in each transmission, respectively. Moreover, the broadcasting nature of connectionless BLE allows seamless, overhead-free handovers between receivers.



Our energy-conscious wearable design is exploiting the connectionless advertising mode of BLE to achieve battery lifetimes of several months, enabling long-term maintenance-free data collection campaigns. For example, the lifetime estimation when streaming raw data (accelerometer sampling rate of 10 Hz) is 213 days; and the lifetime estimation when streaming histograms (accelerometer sampling rate of 50 Hz) is 437 days. Overall, by exploiting the broadcasting nature of connectionless BLE, incorporating efficient wearable antennas, employing energy-efficient packet loss mitigation techniques and deploying multiple receivers, we are able to provide full-house coverage and reliable links, whilst moving power-intensive tasks from the wearable to the infrastructure.

#### Competing interests

The authors declare that they have no competing interests.

#### Acknowledgements

This work was performed under the SPHERE IRC funded by the UK Engineering and Physical Sciences Research Council (EPSRC), Grant EP/K031910/1.

Received: 30 October 2015 Accepted: 18 January 2016

Published online: 28 January 2016

#### References

- Department of Economic and Social Affairs—Population Division—United Nations, World Population Ageing: 1950–2050. Technical report (2011). <http://www.un.org/esa/population/publications/worldageing19502050/>, ISBN 92-1-051092-5
- AJ Cruz-Jentoft, A Franco, P Sommer, JP Baeyens, E Jankowska, A Maggi, P Ponikowski, A Rys, K Szczerbinska, JP Michel, A Milewicz, Silver paper: the future of health promotion and preventive actions, basic research, and clinical aspects of age-related disease—a report of the European Summit on Age-Related Disease. *Aging Clin. Exp. Res.* **21**(6), 376–385 (2009)
- Royal College of Nursing: children and young people's mental health—every nurse's business: RCN guidance for nursing staff (2014). [https://www2.rcn.org.uk/\\_data/assets/pdf\\_file/0005/587615/004\\_587\\_WEB.pdf](https://www2.rcn.org.uk/_data/assets/pdf_file/0005/587615/004_587_WEB.pdf)
- M Haak, A Fänge, S Iwarsson, SD Ivanoff, Home as a signification of independence and autonomy: experiences among very old Swedish people. *Scand. J. Occup. Ther.* **14**(1), 16–24 (2007)
- P Rashidi, A Mihailidis, A survey on ambient-assisted living tools for older adults. *IEEE J. Biomed. Health Inform.* **17**(3), 579–590 (2013)
- Smart wearable systems: current status and future challenges. *Artificial Intell. in Medicine*. **56**(3), 137–156 (2012)
- WKG Seah, ZA Eu, H Tan, in *Proc. 1st Int. Conf. Wireless Commun., Veh. Technology, Inform. Theory and Aerospace Electron. Syst. Technology (Wireless VITAE 2009)*. Wireless sensor networks powered by ambient energy harvesting (WSN-HEAP) - Survey and challenges (IEEE, Aalborg, Denmark, 2009), pp. 1–5
- N Zhu, T Diethel, M Camplani, L Tao, A Burrows, N Twomey, D Kaleshi, M Mirmehdi, P Flach, I Craddock, Bridging e-Health and the Internet of things: the SPHERE project. *IEEE Int. Syst.* **30**(4), 39–46 (2015). doi:10.1109/MIS.2015.57
- TJM Kooiman, ML Dontje, SR Sprenger, WP Krijnen, CP van der Schans, M de Groot, Reliability and validity of ten consumer activity trackers. *BMC Sports Sci. Med. Rehabil.* **7**(1), 24 (2015). doi:10.1186/s13102-015-0018-5
- S Tennina, MD Renzo, E Kartsakli, F Graziosi, AS Lalos, A Antonopoulos, P-V Mekikis, L Alonso, WSN4QoL: a WSN-oriented healthcare system architecture. *Int. J. Distrib. Sens. Netw.* **2014** (2014). doi:10.1155/2014/503417
- E Kartsakli, A Antonopoulos, A Lalos, S Tennina, M Renzo, L Alonso, C Verikoukis, Reliable MAC design for ambient assisted living: moving the coordination to the cloud. *Commun. Mag. IEEE.* **53**(1), 78–86 (2015). doi:10.1109/MCOM.2015.7010519
- J Winkley, P Jiang, W Jiang, Verity: an ambient assisted living platform. *IEEE Trans. Consum. Electron.* **58**(2), 364–373 (2012)
- P Gupta, T Dallas, Feature selection and activity recognition system using a single triaxial accelerometer. *IEEE Trans. Biomed. Eng.* **61**(6), 1780–1786 (2014)
- Y Zhang, S Markovic, I Sapir, RC Wagenaar, TDC Little, in *Proc. 5th Int. Conf. Pervasive Computing Technologies for Healthcare (PervasiveHealth)*. Continuous functional activity monitoring based on wearable tri-axial accelerometer and gyroscope (IEEE, Dublin, Ireland, 2011), pp. 370–373
- J Cheng, X Chen, M Shen, A framework for daily activity monitoring and fall detection based on surface electromyography and accelerometer signals. *IEEE J. Biomed. Health Informatics.* **17**(1), 38–45 (2013)
- U Mitra, BA Emken, S Lee, M Li, V Rozgic, G Thatte, H Vathsangam, D Zois, M Annavaram, S Narayanan, M Levorato, D Spruijt-Metz, G Sukhatme, Knowme: a case study in wireless body area sensor network design. *IEEE Commun. Mag.* **50**(5), 116–125 (2012)
- S Lee, S Polito, C Agell, S Mitra, R Firat Yazicioglu, J Riistama, J Habetha, J Penders, A low-power and compact-sized wearable bio-impedance monitor with wireless connectivity. *J. Phys. Conf. Ser.* **434**(1), 012013 (2013). doi:10.1088/1742-6596/434/1/012013
- D Oletic, B Arsenali, V Bilas, Low-power wearable respiratory sound sensing. *Sensors (Basel)*. **14**(4), 6535–66 (2014). doi:10.3390/s140406535
- KC Tseng, B-S Lin, L-D Liao, Y-T Wang, Y-L Wang, Development of a wearable mobile electrocardiogram monitoring system by using novel dry foam electrodes. *IEEE Syst. J.* **8**(3), 900–906 (2014). doi:10.1109/JSYST.2013.2260620
- Y-C Kan, C-K Chen, A wearable inertial sensor node for body motion analysis. *IEEE Sens. J.* **12**(3), 651–657 (2012). doi:10.1109/JSEN.2011.2148708
- C Buratti, R D'Errico, M Maman, F Martelli, R Rosini, R Verdona, in *Proc. 4th Int. Symp. on Applied Sci. in Biomedical and Commun. Technologies*. Design of a body area network for medical applications: the WiserBAN project (ACM, New York, USA, 2011), pp. 1641–1645
- S Ullah, K Kwak, An ultra low-power and traffic-adaptive medium access control protocol for wireless body area network. *J. Med. Syst.* **36**(3), 1021–1030 (2012)
- E Ibarra, A Antonopoulos, E Kartsakli, C Verikoukis, HEH-BMAC: hybrid polling MAC protocol for WBANs operated by human energy harvesting. *Telecommun. Syst.* **58**(2), 111–124 (2015). doi:10.1007/s11235-014-9898-z
- V Esteves, A Antonopoulos, E Kartsakli, M Puig-Vidal, P Miribel-Català, C Verikoukis, Cooperative energy harvesting-adaptive MAC protocol for WBANs. *Sensors.* **15**(6), 12635 (2015). doi:10.3390/s150612635
- E Ibarra, A Antonopoulos, E Kartsakli, J Rodrigues, C Verikoukis, QoS-aware energy management in body sensor nodes powered by human energy harvesting. *Sens. J. IEEE.* **PP**(99), 1–1 (2015). doi:10.1109/JSEN.2015.2483064
- X Fafoutis, A Di Mauro, C Orfanidis, N Dragoni, Energy-efficient medium access control for energy harvesting communications. *IEEE Trans. on Consum. Electron.* **61**(4) (2015)
- S Ullah, H Higgins, B Braem, B Latre, C Blondia, I Moerman, S Saleem, S Rahman, K Kwak, A comprehensive survey of wireless body area networks. *J. Med. Syst.* **36**(3), 1065–1094 (2012)
- R Cavallari, F Martelli, R Rosini, C Buratti, R Verdona, A survey on wireless body area networks: technologies and design challenges. *IEEE Commun. Surv. Tutorials.* **16**(3), 1635–1657 (2014). doi:10.1109/SURV.2014.012214.00007
- E Kartsakli, AS Lalos, A Antonopoulos, S Tennina, MD Renzo, L Alonso, C Verikoukis, A survey on m2m systems for mHealth: a wireless communications perspective. *Sensors.* **14**(10), 18009 (2014). doi:10.3390/s141018009
- C Cano, B Bellalta, A Sfairpoulou, M Oliver, Low energy operation in WSNs: a survey of preamble sampling MAC protocols. *Comput. Networks.* **55**(15), 3351–3363 (2011)
- X Fafoutis, A Di Mauro, MD Vithanage, N Dragoni, Receiver-initiated medium access control protocols for wireless sensor networks. *Comput. Netw.* **76**, 55–74 (2015). doi:10.1016/j.comnet.2014.11.002
- P Woznowski, X Fafoutis, T Song, S Hannuna, M Camplani, E Mellios, M Haghighi, N Zhu, M Mirmehdi, R Piechocki, D Kaleshi, I Craddock, in *Proc. IEEE Int. Conf. Communications Workshops (ICCW)*. A multi-modal sensor infrastructure for healthcare in a residential environment (IEEE, London, UK, 2015), pp. 271–277



33. M Ziefle, C Rocker, in *Proc. 4th Int. Conf. Pervasive Comput. Technologies for Healthcare (PervasiveHealth)*. Acceptance of pervasive healthcare systems: a comparison of different implementation concepts (IEEE, Munich, Germany, 2010), pp. 1–6
34. JHM Bergmann, AH McGregor, Body-worn sensor design: what do patients and clinicians want? *Ann. Biomedical Eng.* **39**(9), 2299–2312 (2011)
35. U Maurer, A Smailagic, DP Siewiorek, M Deisher, in *Int. Workshop on Wearable and Implantable Body Sensor Networks (BSN)*. Activity recognition and monitoring using multiple sensors on different body positions (IEEE, Boston, USA, 2006)
36. Bluetooth SIG, Specification of the Bluetooth System—Covered Core Package version: 4.0. [Online] Available at (2010). <https://www.bluetooth.org/en-us/specification/adopted-specifications>
37. M Siekkinen, M Hienkari, JK Nurminen, J Nieminen, in *Proc. IEEE Wireless Communications and Networking Conference Workshops (WCNCW)*. How low energy is Bluetooth low energy? Comparative measurements with ZigBee/802.15.4 (IEEE, Paris, France, 2012), pp. 232–237
38. A Dementyev, S Hodges, S Taylor, J Smith, in *IEEE Int. Wirel. Symp. (IWS)*. Power consumption analysis of Bluetooth low energy, ZigBee and ANT sensor nodes in a cyclic sleep scenario (IEEE, Beijing, China, 2013)
39. K Mikhaylov, N Plevritakis, J Tervonen, Performance analysis and comparison of Bluetooth low energy with IEEE 802.15.4 and SimpliciTI. *J. Sens. Actuator Netw.*, vol. 2. (MDPI, Basel, Switzerland, 2013), pp. 589–613
40. PS Hall, Y Hao, in *Proc. 1st Eur. Conf. Antennas and Propagation (EuCAP)*. Antennas and propagation for body centric communications (IEEE, Nice, France, 2006), pp. 1–7
41. A Alomainy, Y Hao, F Pasveer, Numerical and experimental evaluation of a compact sensor antenna for healthcare devices. *IEEE Trans. Biomed. Circuits Syst.* **1**(4), 242–249 (2007)
42. PS Hall, Y Hao, YI Nechayev, A Alomainy, CC Constantinou, C Parini, MR Kamarudin, TZ Salim, DTM Hee, R Dubrovka, AS Owadally, W Song, A Serra, P Nepa, M Gallo, M Bozzetti, Antennas and propagation for on-body communication systems. *IEEE Antennas Propag. Mag.* **49**(3), 41–58 (2007)
43. E Mellios, A Goulianos, S Dumanli, G Hilton, R Piechocki, I Craddock, in *Proc. 9th Int. Conf. on Body Area Networks (BODYNETS)*. Off-body channel measurements at 2.4 GHz and 868 MHz in an indoor environment (ICST, London, UK, 2014)
44. GS Hilton, HWW Hunt-Grubbe, in *Proc. Eur. Conformal Antennas Workshop (EWAC)*. Simulation and practical analysis of a cavity-backed linear slot antenna for operation in the IEEE802.11a band (IEEE, Bristol, UK, 2007)
45. MW Abdullah, X Fafoutis, E Mellios, M Klemm, G Hilton, in *Proc. Loughborough Antennas and Propagation Conf. (LAPC)*. Investigation into off-body links for wrist mounted antennas in Bluetooth systems (IEEE, Loughborough, UK, 2015)
46. G Valadon, F Le Goff, C Berger, in *Proc. 5th Adv. Int. Conf. Telecomm. (AICT)*. A practical characterization of 802.11 access points in Paris (IEEE, Venice, Italy, 2009), pp. 220–225
47. A Akella, G Judd, S Seshan, P Steenkiste, in *Proc. 11th ACM An. Int. Conf. on Mobile Comp. and Networking (MobiCom)*. Self-management in chaotic wireless deployments (ACM, New York, USA, 2005), pp. 185–199
48. T Plötz, NY Hammerla, P Olivier, in *Proc. 22nd Int. Joint Conf. on Artificial Intelligence (IJCAI) - Volume Volume Two*. Feature learning for activity recognition in ubiquitous computing (AAAI Press, Barcelona, Spain, 2011), pp. 1729–1734
49. NY Hammerla, R Kirkham, P Andras, T Ploetz, in *Proc. 2013 Int. Symp. on Wearable Comput. (ISWC)*. On preserving statistical characteristics of accelerometry data using their empirical cumulative distribution (ACM, New York, USA, 2013), pp. 65–68
50. E Tsimbalò, X Fafoutis, E Mellios, M Haghghi, B Tan, G Hilton, R Piechocki, I Craddock, in *Proc. 2nd IEEE World Forum on Internet of Things (WF-IoT)*. Mitigating packet loss in connectionless Bluetooth low energy (IEEE, Milan, Italy, 2015)
51. RG Vaughan, Polarization diversity in mobile communications. *IEEE Trans. Veh. Technol.* **39**(3), 177–186 (1990). doi:10.1109/25.130998
52. E Tsimbalò, X Fafoutis, R Piechocki, in *Proc. 26th IEEE Int. Symp. Personal, Indoor and Mobile Radio Commun. (PIMRC)*. CRC Error correction for energy-constrained transmission (IEEE, Hong Kong, China, 2015)
53. E Tsimbalò, X Fafoutis, R Piechocki, in *Proc. 2nd IEEE World Forum on Internet of Things (WF-IoT)*. Fix It, Don't Bin It!—CRC error correction in Bluetooth low energy (IEEE, Milan, Italy, 2015)
54. V Kumar, O Milenkovic, On graphical representations of algebraic codes suitable for iterative decoding. *IEEE Commun. Lett.* **9**(8), 729–731 (2005). doi:10.1109/LCOMM.2005.1496597
55. TK Moon, *Error correction coding: mathematical methods and algorithms*. (Wiley-Interscience, Hoboken, New Jersey, 2005)
56. J Feldman, MJ Wainwright, DR Karger, Using linear programming to decode binary linear codes. *IEEE Trans. Inf. Theory.* **51**(3), 954–972 (2005). doi:10.1109/TIT.2004.842696
57. S Boyd, Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.* **3**(1), 1–122 (2010). doi:10.1561/2200000016
58. S Barman, X Liu, S Draper, B Recht, Decomposition methods for large scale LP decoding. 2011 49th Annu. Allert. Conf. Commun. Control. Comput. **59**(12), 253–260 (2011). doi:10.1109/Allerton.2011.6120176

**Submit your manuscript to a SpringerOpen<sup>®</sup> journal and benefit from:**

- Convenient online submission
- Rigorous peer review
- Immediate publication on acceptance
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

---

Submit your next manuscript at ► [springeropen.com](http://springeropen.com)