



Woznowski, P., Fafoutis, X., Song, T., Hannuna, S., Camplani, M., Tao, L., Paiement, A., Mellios, E., Haghghi, M., Zhu, N., Hilton, G., Damen (Aldamen), D., Burghardt, T., Mirmehdi, M., Piechocki, R., Kaleshi, D., & Craddock, I. (2015). A multi-modal sensor infrastructure for healthcare in a residential environment. In *2015 IEEE International Conference on Communication Workshop, ICCW 2015* (pp. 271-277). [7247190] (ICCW). Institute of Electrical and Electronics Engineers (IEEE). <https://doi.org/10.1109/ICCW.2015.7247190>

Peer reviewed version

Link to published version (if available):
[10.1109/ICCW.2015.7247190](https://doi.org/10.1109/ICCW.2015.7247190)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via IEEE at <http://ieeexplore.ieee.org/document/7247190/>. Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available: <http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

A Multi-modal Sensor Infrastructure for Healthcare in a Residential Environment

Przemyslaw Woznowski, Xenofon Fafoutis, Terence Song, Sion Hannuna, Massimo Camplani, Lili Tao, Adeline Paiement, Evangelos Mellios, Mo Haghighi, Ni Zhu, Geoffrey Hilton, Dima Damen, Tilo Burghardt, Majid Mirmehdi, Robert Piechocki, Dritan Kaleshi and Ian Craddock
Faculty of Engineering, University of Bristol, Bristol, BS8 1UB, UK

Abstract—Ambient Assisted Living (AAL) systems based on sensor technologies are seen as key enablers to an ageing society. However, most approaches in this space do not provide a truly generic ambient space – one that is not only capable of assisting people with diverse medical conditions, but can also recognise the habits of healthy inhabitants, as well as those with developing medical conditions. The recognition of Activities of Daily Living (ADL) is key to the understanding and provisioning of appropriate and efficient care. However, ADL recognition is particularly difficult to achieve in multi-resident spaces; especially with single-mode (albeit carefully crafted) solutions, which only have limited capabilities. To address these limitations we propose a multi-modal system architecture for AAL remote healthcare monitoring in the home, gathering information from multiple, diverse (sensor) data sources. In this paper we report on developments made to-date in various technical areas with respect to critical issues such as cost, power consumption, scalability, interoperability and privacy.

Index Terms—Ambient Intelligence, Ambient Assisted Living, eHealth, Internet of Things

I. INTRODUCTION

The UK, like many other countries, is facing the problem of an ageing society. In 2010, ten million people in the UK were over 65 years old and the projections are for 5.5 million more in 20 years time to reach 19 million by 2050. “65% of Department for Work and Pensions benefit expenditure goes to those over working age, equivalent to £100 billion in 2010/11 or one-seventh of public expenditure” [1]. The average spending of the NHS (National Health System) for a retired household is nearly double of what is spent for non-retired households.

One possible solution to this problem is Ambient Assisted Living (AAL) technologies. Advances in technology have made sensors smaller, cheaper, and more viable for large-scale deployment in residential environments to monitor activities of daily living (ADL). Our vision is to develop a multi-purpose, multi-modal platform of home sensors, which we call SPHERE (Sensor Platform for HEalthcare in Residential Environment). SPHERE aims not only to extend the state of the art in a number of technology domains, but to engage with stakeholders across disciplines and across sectors to illuminate the applications of current technologies to known health needs and of new technologies to emerging health needs.

This generic, multi-modal sensor-based platform, which has been built on a cutting edge platform made up of commercial and prototype components, will be used to test clinical and

health related hypotheses in a real life environment. For this purpose, we have acquired a two storey, two bedroom house from the University of Bristol’s Accommodation Office and have converted it into a fully-instrumented living lab. In this paper, we will refer to this instrumented house as *the SPHERE house*. The remainder of this paper provides an overview of the SPHERE infrastructure; a novel platform that integrates multiple sensing technologies to provide an intelligent ambient space in residential environments.

II. BACKGROUND

Existing AAL systems make use of sensor networks, wearable technologies and computer vision technologies. However, they tend to provide solutions that only address specific needs. In contrast, the SPHERE architecture attempts to combine different sensing technologies to provide a generic platform for ADL recognition.

A smart home uses Ambient Intelligence (AmI) technologies to sense, monitor, and control residents’ living environment to enable AAL in an unobtrusive manner [2]. In home environment monitoring, a variety of sensors are used to gather data that enable various activities to be recognised and tracked. Indoor localisation and tracking is a crucial component in AAL applications, and it can be provided by either active or passive sensing technologies [3]. Many other sensing technologies have been employed to detect falls [4], monitor individual daily activities such as sleep measures [5], monitor patients with chronic conditions such as type 2 Diabetes [6] and Alzheimer’s Disease [7], and mental and emotional health [8].

Video based technology has been widely used for smart healthcare systems in home environments, as they have the potential to address several limitations of other sensing modalities [9]. However, having a general vision system is still difficult to achieve hence most systems are designed for specific problems and have rarely been combined with other sensors. A real-time context-aware sleeping-respiration measurement system is proposed in [10] that accurately measures the sleeper’s respiration information. Computer vision techniques for monitoring and clinical evaluation of Parkinson’s disease and stroke patients have been proposed recently in [11]. Since the use of staircases can directly reflect musculoskeletal problems and the progress of recovery, a system has been developed to estimate the quality of movement on stairs in

[12]. Video data is also used to detect falls, especially in the monitoring of the elderly, as in [13].

With regard to on-body sensing, fashionable wearable gadgets, such as Fitbit and Jawbone UP, have appeared in the consumer electronics market in recent years. However, with little capacity for expansion and scant access to raw data, such gadgets are of limited use in medical applications. The research community has also proposed wearable devices for activity monitoring. Verity [14] is an AAL platform that is using a wearable sensor equipped with an accelerometer and a heart rate monitor. Reference [15] proposes an AAL platform based on a waist-worn accelerometer. These platforms use off-the-shelf hardware and do not focus on their power consumption, resulting in wearable devices that need regular charging, similarly to commercial products.

Efforts were also made to implement intelligent spaces through the medium of multiple sensing technologies. The activity recognition system [16] integrated with ICS-FORTH “AmI Sandbox” [17] used a variety of ambient sensors (power consumption monitors, pressure sensors, etc.). Unfortunately, only few home activities have been considered in a very controlled environment and no datasets are available. In the CASAS project [18] a smart apartment was instrumented with motion, ambient temperature, water, cooker, phone usage and contact switch sensors on key objects (medicine container, cooking pot, etc.). Some CASAS datasets are based on 20 people performing a set of activities, other experiments were carried out with a single person occupying the apartment for several months. The lack of rich meta-data (or video data) to annotate these publicly available datasets limits their usability. In [19] wireless PIDs (Presence Infrared Detectors), wireless weight scale, door contact (in cupboards, fridge, etc.), oxymeter and tensiometer sensors were utilised alongside microphones and on-body worn three axis accelerometer. However, such sensor-rich setup is not only labour- but also money-expensive and difficult to maintain in the long run. Other projects such as MIT’s PlaceLab smart home [20], GER’HOME [21] by the INRIA and CSTB, and Kasteren et al. [22] made advancements in this area and provided publicly available datasets. Each of the aforementioned projects adopted a different combination of sensor technologies, however none fully succeeded in solving the problem of reliably recognising all ADLs in a natural, ‘scenario-free’ home environment. Therefore, as pointed out in [23] there is no perfect solution to the problem of activity recognition, as it all depends on the type of recognised activities and application’s requirements.

The collection of data from large-scale sensor-based distributed systems presents a significant challenge. Data streams from diverse sources must be aggregated and linked with relevant contextual data and metadata, and then transformed into a form that is easily accessible and useable. It may be necessary to support multiple data formats, protocols, and semantics when developing applications that convert data streams into information (e.g., intelligent lighting, people mobility support, energy information, etc.). Also, rather than impose stringent

requirements on data to satisfy design and implementation decisions, schema-less data models are increasingly being employed by data aggregation systems, so that they can accept any data record and adapt to data schema changes [24].

III. SPHERE SYSTEM ARCHITECTURE

The SPHERE platform is based on three sensing technologies: a Body Sensor Network made up of ultra low-power wearable sensors; a Video Sensor Network focusing on recognition of activities through video analysis of home inhabitants; and an Environment Sensor Network made up of hardware sensing the home ambience. Fig. 1 provides a high-level view of the SPHERE hub and data sharing system, which aims to: a) collect, organise, and store data captured by environment, video, and on-body sensors installed in the SPHERE house; b) enable data owners to view and manage the access to their data, which is necessary to enhance trust in the SPHERE monitoring system; c) provide end-users (e.g. clinicians) with catalogue-based services for searching and retrieving sensor information and sensor-generated data at different levels of granularity (ranging from raw data to detected activities and events) in a form that is easily accessible and usable by domain experts and end-users; d) enable dynamic management of intermediate processing of data, to create value-added streams of data.

The SPHERE platform is made up of three basic architectural layers: sensing networks, in-home data aggregation, and central data aggregation with data analytics. Data captured from the sensing networks are collected by their respective sensor gateways, either in raw form or as detected activities and events, and provided to the SPHERE Home Gateway (SHG) using secure communications. Sensor gateways may also accept commands from the SHG, allowing the system to query the sensing environment for additional data, or to configure settings on sensor nodes.

The SPHERE Home Gateway (SHG) brings together sensor data outputs from home environment sensors, video sensors, and on-body sensors (via sensor gateways) and manages data access using suitable secure communications to the SPHERE Data Hub (SDH). Sensor data collection is done using MQTT and HTTP protocols, either directly with the sensors or with the corresponding cluster gateway. The SHG can also perform additional processing on data (on top of that already managed by sensor gateways) and provide sensor networks with integrated data streams (across different sensor networks within the same home, e.g. to cross-correlate data, aid activity detection, conflict resolution, handle uncertainty, improve accuracy, etc.). Multiple home residents can be differentiated between by correlating data from multiple sensor networks e.g. using location information from wearable sensors and/or face recognition in video sensors. The SHG also provides a management console for home owners to view, control, and manage data collection, as well as access to collected data. The favoured data management device identified by target users (in multiple consultation sessions) is a dedicated tablet, which is also preferred for issuing alerts and reminders.

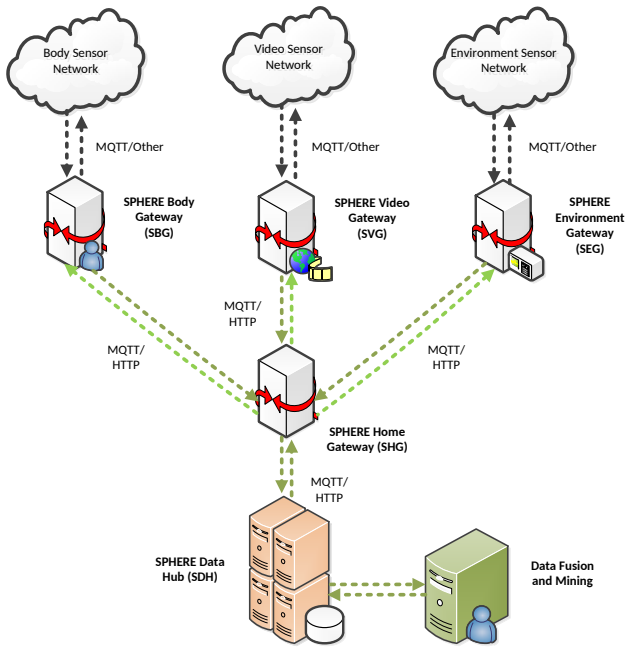


Fig. 1. An overview of the SPHERE system architecture.

The SPHERE Data Hub (SDH) receives sensor data from participating homes (via SHGs), stores them in a data warehouse, and enables easy and secure access to the data through appropriate graphical and programmatic interfaces for various stakeholders e.g. clinicians, carers, family. The types of services that the SDH offer users depend on privacy constraints and levels of access granted by the data owner, which determine what data is available and where data can be moved. The sensing task in the presented architecture is distributed across the three sensing networks with a single point for data storage and processing. The processing involves automatic recognition of home activities and data mining techniques to learn habits and daily routines of home occupants. The infrastructure does not only provide a complete AAL environment in which various longitudinal studies can be undertaken but more importantly, it can also connect people with various medical conditions to care providers.

A. Home Environment Sensing

Home environment monitoring in SPHERE, enabled by passive ambient sensing, utilises an assortment of data dedicated to describing the characteristics of surrounding ambience. It establishes linked data sets and builds profiles of ADL that enable further data applications (i.e., data mining / fusion) as well as clinical or non-clinical studies, investigating links between environmental variables to health needs. An environment sensing platform has been designed and developed to provide real-time ambient sensing data acquisition, processing, transmission and distribution over wireless sensor networks (WSN). As part of the SPHERE architecture, this platform collects structured data from multi-cluster sensors, monitoring temperature, humidity, luminosity, noise level, air




Fig. 2. Ambient sensor nodes (left), and a wearable device with patch antenna (right).

quality, room occupancy, door contact, cold and hot water consumption, as well as electricity metering.

Our environment sensing platform employs WSN-enabled off-the-shelf products and development kits. In particular, we use the Internet of Things (IoT) development kits from Libelium™ to build a ZigBee-based WSN and collect multi-cluster environment sensor data – prototype in Fig. 2 (left). The Libelium sensor nodes are the basic units for collecting information from the physical world. Integrated in a credit card sized platform, they consist of three main modules: a main processing board, a sensor board and a wireless module. We have developed three different variations of the sensor nodes, each equipped with different sensors and their respective analogue signal conditioning circuits. The main processing board receives these pre-processed signals and digitises them as sensor data. These data are structured into a pre-defined data format for wireless transmission over the WSN. The wireless module uses ZigBee for communications. These sensor nodes are of several fundamental capabilities, providing multi-channel sensor data acquisition, signal processing as well as wireless transmission in an energy-constrained mechanism. Additionally, a set of smart metering products by CurrentCost™ are used to gather electricity consumption information from different household electrical appliances.

Sensed data are processed and aggregated into a single message and then wirelessly transmitted to the SEG. Messages are either event-driven or periodic notifications. Event-driven messages are generated from the door contact sensors, the occupancy sensors, and the water metering sensors; and are sent immediately when either the sensors detect certain events, or the status of the sensor reading reaches thresholds. Periodic data are generated by the other sensors (i.e. temperature, humidity, luminosity, air quality, noise level and power consumption sensors) in a selected reading frequency.

The environment sensor data are described in a predefined name schema to format the data and identify the context of the data, where the name schema describes which, when, where the sensor readings are taken, as well as their unit of measurement. Moreover, a time-stamp is generated at the time the sensor readings are collected and attached to the messages. When the SEG receives the data it parses it into the JSON (Javascript Object Notation) format for further distribution using a publish-subscribe model. Time synchronisation is required for providing temporal relationship between the sensor data. When a new sensor node joins the network, the

mechanism of time synchronisation is applied during setup. Afterwards, periodic messages containing the time reference are sent to all the nodes in the WSN to synchronise regularly. Thirty sensor nodes are deployed in the SPHERE house, providing 90 environmental and ambient data items.

B. Video Monitoring

The video monitoring component of the SPHERE architecture is tasked with developing a real-time multi-camera system for activity and health monitoring within the home environment. Multi-camera architectures have been used for activity monitoring in indoor and outdoor environments [25][9]. Two main different architecture schemes have been presented so far: centralized and distributed networks. In the SPHERE platform, integration with other sensing modalities, user acceptance and deployment budget are key factors. In fact, in order to make deployment into the local community financially feasible, it has been necessary to limit hardware selection to low cost consumer RGB-D cameras and web cams such as the Asus Xtion, Microsoft Kinect (v2), and Genius WideCam 1050. We believe it is vital that any algorithms developed are based on the actual hardware to be deployed in people's homes. All of these cameras require USB 2.0 or 3.0 connections (on a dedicated bus) and, in case of the Kinect v2, a high-performance computer running Windows 8 or higher. The use of these USB based devices, and the relatively high computational burden of the computer vision techniques used to track and assess movement within the home lead us to develop a centralized network where a central node (see SVG module in Fig. 1) collects and locally processes the data provided by the cameras. The number of video devices is limited (between five and ten in a typical house scenarios), hence the lack of scalability that generally affect centralized architectures is not an issue for the SPHERE sensing platform. In particular, video devices will be positioned in relevant areas where important activity and actions take place: kitchen, living room, corridor, stairs, etc., while blind areas will be covered by integrating other sensors, such as on-body or environmental sensors. Currently, four Asus Xtions and one Kinect v2 are deployed in the SPHERE house.

Tracking individuals moving around home environments feeds into a number of research work flows. For example, when combined with prior knowledge of the environment, such as the location of white goods activities, activities like washing up, cooking and watching television can be readily identified. Furthermore, general activity levels including the amount of time spent sedentary can be extracted. Tracking is based on the state-of-the-art people trackers such as [26]. From the silhouette provided, a bounding box can be fit to a tracked person, to obtain a 3D trajectory over time that acts as a coarse shape descriptor for subsequent analysis. For example, a near-square bounding box indicates a sitting individual, whereas a tall thin bounding box suggests a person standing.

In our different application scenarios, tight camera synchronization is not required, and we use a time-stamping system based on the recording device clock and on the central

node acquisition time. This software based synchronization is affected only by the USB latency. Also, since the cameras will be positioned with non-overlapping views (to maximize the house coverage), a weak synchronization is sufficient to manage tracking hand-over or higher level tasks, such as activity labeling or high level feature fusions.

The SPHERE platform allows us to obtain an accurate 3D description of both human motion and environment. The main shortcomings of the depth sensors, such as limited range and depth of view and interference with natural light, are limited in the SPHERE indoor scenario. One of the main advantages of depth sensors is that they allow us to employ detailed, but anonymous, data thus guaranteeing that these systems are accepted by the users and regulatory committees, and hence enabling the system's deployment in large scales. To aid privacy, only results of video analysis and no video footages are allowed to leave the home boundaries.

Much of our recent work has focused on developing algorithms, based on depth data analysis, for assessing the *quality* of particular activities as opposed to identifying *which* activities are taking place. We argue that identifying the actions comprising functional mobility, such as walking, sitting to standing and ascending the stairs are best determined via context as opposed to low level analysis incorporating computer vision algorithms. Currently, we are particularly interested in determining the quality of stair climbing and sitting to standing within the SPHERE house.

Movement quality assessment is based on the skeletons provided by the PrimeSense middleware and the Kinect SDK. We normalise these skeletons for global positioning and orientation of the camera and height variation. The normalised skeletons are of high dimensionality (60D) and often contain outliers. Thus, their dimensionality is reduced using a modified version of Diffusion Maps [27], where Gerber's [28] method for addressing outliers in Laplacian Eigenmaps is exploited. The resulting high level feature vector, obtained from the normalised skeleton at one frame, represents individual poses and is used to build a statistical model of normal movement. The camera system does not require a calibration (in terms of relative camera positions etc.) as the camera views are not overlapping. For this reason bounding box positions for tracked individuals are provided relative to each individual camera position. However, the cameras have been positioned to maximise the potential for obtaining useful skeleton data from the OpenNI/Primesense middleware. Skeleton data is then normalised using Procrustes analysis.

Our statistical model [12] is made up of two components that describe normal poses and the normal dynamics of the movement, respectively. The first pose model is in the form of the probability density function (PDF) of the poses, and it is learnt from normal movement frames. The quality of a new pose at each frame is then assessed as the log-likelihood of being described by the pose model. The dynamics model is represented as the PDF that describes the likelihood of a pose at a new frame given the poses at the previous frames. The dynamics quality is then assessed as the log-likelihood

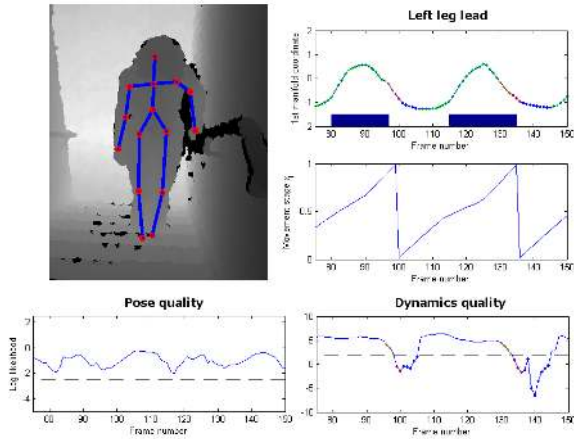


Fig. 3. Analysis of gait on stairs. From top left: raw skeleton data, high level description of the gait, pose and dynamics quality measures.

of the model describing a sequence of poses within a fixed size window. Each frame is classified as normal or abnormal using an empirically determined threshold on the likelihood output from the dynamics model. An example of the quality assessment pipeline is shown in Fig. 3. In the example, a person is simulating an abnormal gait pattern where she only climbs using her left leg leading as opposed to a normal reciprocal motion pattern. The top-left image shows the raw skeleton output and the top-right shows the periodic nature of the high level descriptor derived from the dimensionality reduction method described above. Note how in the dynamic quality measure graph (bottom-right), the quality of movement drops below the threshold. This is where the subject fails to lead with their right leg. For further details, please see [12].

C. Body-Area Sensing

On-body sensing in the SPHERE architecture is realised through a wrist-worn wearable device. The wrist is selected as the most user-acceptable and least invasive body position. The core component of the wearable device is the nRF51822 radio by Nordic Semiconductor. nRF51822 is an ultra low-power Bluetooth Low Energy (BLE) [29] solution for wireless applications. We selected BLE primarily for its high energy efficiency; that is less than half of ZigBee [30]. nRF51822 is interfaced with two ADXL362 3-axis accelerometers by Analog Devices. A low-profile directional patch antenna on RT/Duroid 6010 substrate (dimensions: 18.5×19 mm; measured radiation efficiency: 55%) was designed to enable the wireless communication of the wearable device with a nearby receiver unit. Measurements reported in [31] have shown that a directional antenna is the best choice for the on-body node of an off-body communication system. The wearable system was prototyped as shown in Fig. 2 (right). Using an 100Ω shunt resistor on the power supply, we measured the idle current drain of nRF51822 at $5\mu\text{A}$. Each ADXL362 accelerometer contributes an additional $2\mu\text{A}$ when sampling at 50Hz.

For communication, we use the connectionless mode of BLE, in which the wearable device periodically broadcasts

advertisements of 26 bytes of payload. Fig. 4 (top) shows the peak current drain of an advertisement transmission at different transmission power levels, ranging from -20 to $+4\text{dBm}$. We can see that for transmit powers higher than -4dBm the energy consumption increases considerably. This highlights a challenging trade-off between energy efficiency and wireless coverage. Using the connectionless BLE for communication, the firmware operates as follows. Before going into sleep mode, the micro-controller unit (MCU) sets the accelerometer to sense at 50Hz. Once the accelerometer buffer is full, the MCU gets into active mode to transfer the raw data into its memory and process them into histograms, i.e. empirical probability density functions [32]. The cycle continues until enough data is collected to fill up an advertisement. The transmission follows and the cycle restarts. We measure the overall consumption using a charged 18mF capacitor to power the wearable device. By measuring the capacitor's charge difference over a period of 1 minute, we derive to a long-term average power consumption of $60\mu\text{W}$. Assuming a 210mAh coin cell battery, this consumption level translates to an approximate lifetime of 10 months without battery replacement.

To ensure full-house reliable coverage, multiple Access Point (AP) units need to be deployed around the house. An AP is a Raspberry Pi B+ micro-computer interfaced with two nRF51822 BLE receivers, which employ antennas in orthogonal polarisations (a horizontal and a vertical dipole). In contrast to the wearable device, the AP does not have power constraints as it is mains-powered. Using two orthogonal polarisations at the AP, we minimise the impact of the polarisation change of the on-body antenna due to random orientation of the arm. Fig. 4 (bottom) shows as an example the measured received signal strength indication (RSSI) in the SPHERE house. In this scenario, the wearable device was mounted on the wrist of a user and it was separated from the receiver by two concrete walls. We also considered a number of body rotations through 360° and two different arm positions. The transmission power was set to 0dBm . The RSSI is shown as cumulative distribution function (CDF) for the *best* and the *worst* of the two receiver antennas (the horizontal dipole was the best antenna in 63% of the cases in this scenario). We can notice an approximately 15dB dynamic range on the received signal strength that solely depends on the body orientation and arm position with respect to the receiver. We can also see a difference of up to approximately 7dB (with an average of 3dB) between the two receiver antennas. To translate the RSSI levels to packet error rate (PER) and provide more intuition on the practical benefits of the two orthogonal antenna polarisations at the receiver and on the impact of body orientation, Fig. 4 (bottom) also shows the respective measured PER at different RSSI levels. We notice that errors begin to occur at -80dBm , while at -102dBm the reception is completely lost.

Early prototyping suggests that our energy-conscious wearable sensor design offers a battery lifetime of several months facilitating the use of energy harvesting, while with the use of multiple receivers and with careful antenna design we provide reliable and robust full-house coverage.




Fig. 4. Peak current drain (top) for various transmission events. CDF of the RSSI (bottom) for various body orientations and arm positions for the best and worst receiver antennas in a link through two concrete walls.

D. Integration and Monitoring

The environment and wearable sensors are integrated as follows. The environment sensor network, as described in Sec.III-A, is functional with a Raspberry Pi model B constituting the SEG (running an MQTT client). Since the on-body sensing system requires multiple APs for maximum coverage, multiple Raspberry Pi APs have been deployed in the SPHERE house with a distributed middleware for capturing data from the on-body sensors. APs cooperatively negotiate the collected data packets in order to reduce redundancy and optimise data gathering. The middleware on the APs is based on SENSOMAX [33], which is an agent-based distributed RTOS (Real Time Operating System) supporting multiple concurrent applications. The concurrency feature on the APs allows them to integrate and apply distributed computational algorithms for various purposes such as error correction of the received packets. Over-The-Air (OTA) updating and reprogramming of the on-body sensors is another important features of APs, by which the firmware and systematic configuration of devices can be modified dynamically. The APs act as the SBG, being responsible for conveying data and carrying out required functionalities in a reliable manner. The agent-based nature of SENSOMAX can be utilised for fault-detection, by sending smart agents around the network, looking for anomalies in the access points. Such a fault-detection feature promotes reflectivity in the network, whereby the performance of the APs can be optimised proactively with regards to their occupied memory and processing. The same physical Raspberry Pies implement the distinct logical roles of both SEG and SBG.

The House's ambient data and on-body data are streamed over an MQTT stream to the Next Unit of Computing (NUC) Intel computer serving as the SHG. The SPHERE Home Gateway is a Linux server hosting a number of software components, including an ActiveMQ open source messaging and integration patterns server (supporting MQTT v3.1 broker) and an Apache web server.

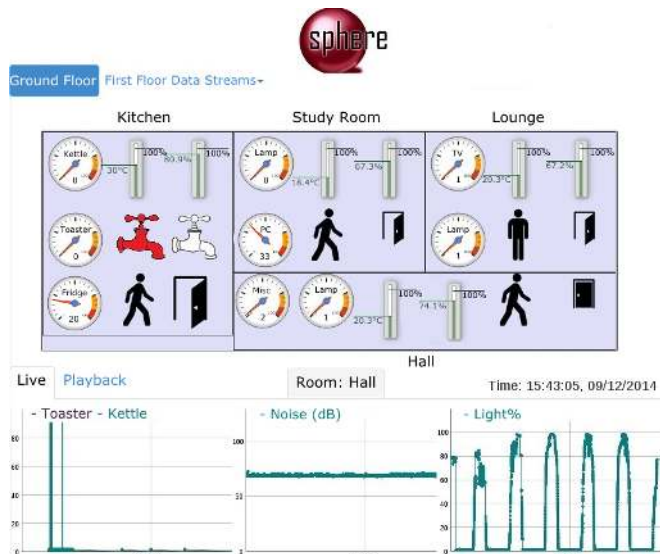


Fig. 5. Dashboard prototype showing the status and readings from sensors in different areas of the SPHERE living lab.

We have developed a prototype dashboard (shown in Fig.5) to display live sensor data using a range of different visualisation techniques. These include simple images (for room occupancy, door contact, water flow), gauges (electricity consumption), thermometer-like meters (temperature and humidity) and graphs (electricity consumption, noise and light levels) representing the current state of sensors deployed in the SPHERE house. In addition to displaying the live state of sensors in the SPHERE house, the dashboard also allows for raw sensor data to be viewed and exported in JSON format.

The dashboard has been implemented using the Bootstrap framework making it well suited not only for computers and laptops, but also for mobile devices. Graphical components have been developed based on open source public Javascript (JS) libraries, such as Dygraphs. MQTT data is supplied to the JS MQTT client over Websockets – natively supported by the ActiveMQ. The dashboard's evolution involves integration of wearable devices' (Sec.III-C) and video sensors' (Sec.III-B) data. Another suit of visual interface will be designed for home owners to view, control and manage access to the collected data. Subsequently, visual tools for doctors, carers and family members are to be provided to enable remote care and in-depth analysis of home inhabitant's activities.

IV. CONCLUSIONS AND FUTURE WORK

Many systems in the area of AAL focus on specific medical conditions, and can only recognise and act on specific activities. Medication intake, fall detection, and activity level monitoring, for example, are just a few of the applications often tackled in separation. To meet the healthcare challenges posed by an ageing society, SPHERE aims to create a comprehensive, multi-modal, system that uses complimentary sensing technologies to provide a more complete view of users' ADL. The goal is to connect doctors, carers, and family members

with people suffering from various medical conditions, allowing them to live their lives independently in the comfort of their own homes, and at the same time significantly reducing the costs associated with healthcare provision.

In this paper, we presented an overview of the SPHERE architecture and described developments made in various technical areas using off-the-shelf and prototype sensing technologies. Through the mixture of these technologies the entire system is envisaged to be affordable for a typical household, costing no more than few thousands GBP. The work in SPHERE is addressing critical issues such as cost, power consumption, scalability, interoperability, and privacy to enable large-scale deployment of the system for healthcare studies and real-life AAL applications. We believe that sharing research data can help accelerate the progress of research and its application for the public good. As such, we aim to make data collected from studies carried out in the SPHERE living lab publicly available (subject to appropriate safeguards), together with rich metadata and contextual information to enable data mining and activity recognition algorithms to be developed.

V. ACKNOWLEDGEMENTS

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

REFERENCES

- [1] R. Cracknell *et al.*, "The ageing population," *Key issues for the new parliament*, p. 44, 2010.
- [2] M. R. Alam, M. B. I. Reaz, and M. A. M. Ali, "A review of smart homes: Past, present, and future," *IEEE Trans. Syst. Man Cybern. C, Appl. Rev.*, vol. 42, no. 6, pp. 1190–1203, 2012.
- [3] G. Deak, K. Curran, and J. Condell, "A survey of active and passive indoor localisation systems," *Comput. Commun.*, 2012.
- [4] A. K. Bourke, S. Prescher, F. Koehler, V. Cionca, C. Tavares, S. Gomis, V. Garcia, and J. Nelson, "Embedded fall and activity monitoring for a wearable ambient assisted living solution for older adults," in *Eng. in Medicine and Biology Soc.*, 2012, pp. 248–251.
- [5] A. Kealy, K. McDaid, J. Loane, L. Walsh, and J. Doyle, "Derivation of night time behaviour metrics using ambient sensors," in *Proc. 7th Int. Conf. Pervasive Computing Technologies for Healthcare (PervasiveHealth)*, 2013, pp. 33–40.
- [6] N. Bal and M. Schwarz, "Ambient assisted living for type 2 diabetic patients," in *Proc. 4th Int. Conf. Pervasive Computing Technologies for Healthcare (PervasiveHealth)*. IEEE, 2010, Conference Proceedings, pp. 1–3.
- [7] Y. Charlon, N. Fourty, W. Bourennane, and E. Campo, "Design and evaluation of a device worn for fall detection and localization: Application for the continuous monitoring of risks incurred by dependents in an alzheimer's care unit," *Expert Sys. with Applicat.*, vol. 40, no. 18, pp. 7316–7330, 2013.
- [8] S. Mowafey, S. Gardner, and R. Patz, "Development of an ambient intelligent environment to facilitate the modelling of "well-being"," in *IET Seminar on Assisted Living*. IET, 2011, Conference Proceedings, pp. 1–6.
- [9] A. A. Chaaoui, P. Climent-Pérez, and F. Flórez-Revuelta, "A review on vision techniques applied to human behaviour analysis for ambient-assisted living," *Expert Sys. with Applicat.*, vol. 39, no. 12, pp. 10873–10888, 2012.
- [10] Y.-M. Kuo, J.-S. Lee, and P.-C. Chung, "A visual context-awareness-based sleeping-respiration measurement system," *IEEE Trans. Inf. Technol. Biomed.*, vol. 14, no. 2, pp. 255–265, 2010.
- [11] R. Wang, G. Medioni, C. J. Winstein, and C. Blanco, "Home monitoring musculo-skeletal disorders with a single 3d sensor," in *Proc. IEEE Conf. on Comput. Vision and Pattern Recognition Workshop*. IEEE, 2013, pp. 521–528.
- [12] A. Paiement, L. Tao, M. Camplani, S. Hannuna, D. Damen, and M. Mirmehdi, "Online quality assessment of human motion from skeleton data," in *Proc. British Machine Vision Conf.*, 2014.
- [13] M. Mubashir, L. Shao, and L. Seed, "A survey on fall detection: Principles and approaches," *Neurocomputing*, vol. 100, pp. 144–152, 2013.
- [14] J. Winkley, P. Jiang, and W. Jiang, "Verity: an ambient assisted living platform," *IEEE Trans. Consum. Electron.*, vol. 58, no. 2, pp. 364–373, 2012.
- [15] P. Gupta and T. Dallas, "Feature selection and activity recognition system using a single triaxial accelerometer," *IEEE Trans. Biomed. Eng.*, vol. 61, no. 6, pp. 1780–1786, 2014.
- [16] C. Filippaki, G. Antoniou, and I. Tsamardinos, "Using constraint optimization for conflict resolution and detail control in activity recognition," in *Ambient Intell.* Springer, 2011, pp. 51–60.
- [17] D. Grammenos, X. Zabulis, A. A. Argyros, and C. Stephanidis, "Forthics internal rtd programme ambient intelligence and smart environments," in *3rd European Conf. on Ambient Intell.*, 2009.
- [18] D. Cook, M. Schmitter-Edgecombe, A. Crandall, C. Sanders, and B. Thomas, "Collecting and disseminating smart home sensor data in the casas project," in *Proc. CHI Workshop on Developing Shared Home Behavior Datasets to Advance HCI and Ubiquitous Computing Research*, 2009.
- [19] A. Fleury, M. Vacher, F. Portet, P. Chahua, and N. Noury, "A french corpus of audio and multimodal interactions in a health smart home," *J. on Multimodal User Interfaces*, vol. 7, no. 1-2, pp. 93–109, 2013.
- [20] S. S. Intille, K. Larson, E. M. Tapia, J. S. Beaudin, P. Kaushik, J. Nawyn, and R. Rockinson, "Using a live-in laboratory for ubiquitous computing research," in *Pervasive Computing*. Springer, 2006, pp. 349–365.
- [21] N. Zouba, F. Brémond, M. Thonnat, A. Anfosso, E. Pascual, P. Mallea, V. Mailland, and O. Guerin, "A computer system to monitor older adults at home: Preliminary results," *Gerontechnology*, vol. 8, no. 3, pp. 129–139, 2009.
- [22] T. Van Kasteren, A. Noulas, G. Englebienne, and B. Kröse, "Accurate activity recognition in a home setting," in *Proc. 10th Int. Conf. on Ubiquitous Computing*. ACM, 2008, pp. 1–9.
- [23] L. Chen and I. Khalil, "Activity recognition: Approaches, practices and trends," in *Activity Recognition in Pervasive Intell. Environments*. Springer, 2011, pp. 1–31.
- [24] I. Gorton and J. Klein, "Distribution, data, deployment: Software architecture convergence in big data systems," *IEEE Software*, vol. PP, no. 99, pp. 1–1, 2014.
- [25] S. Soro and W. Heinzelman, "A survey of visual sensor networks," *Advances in Multimedia*, 2009. [Online]. Available: <http://dx.doi.org/10.1155/2009/640386>
- [26] M. Munaro and E. Menegatti, "Fast RGB-D people tracking for service robots," *Autonomous Robots*, pp. 1–16, 2014.
- [27] R. R. Coifman and S. Lafon, "Diffusion maps," *Applied and Computational Harmonic Analysis*, vol. 21, no. 1, pp. 5–30, 2006.
- [28] S. Gerber, T. Tasdizen, and R. Whitaker, "Robust non-linear dimensionality reduction using successive 1-dimensional laplacian eigenmaps," in *Proc. 24th Int. Conf. on Machine Learning*. ACM, 2007, pp. 281–288.
- [29] Bluetooth SIG, "Specification of the Bluetooth System - Covered Core Package version: 4.0," [Online]. Available at: <https://www.bluetooth.org/en-us/specification/adopted-specifications>, 2010.
- [30] M. Siekkinen, M. Hiienkari, J. Nurminen, and J. Nieminen, "How low energy is bluetooth low energy? Comparative measurements with Zig-Bee/802.15.4," in *Proc. IEEE Wireless Communications and Networking Conference Workshops (WCNCW)*, 2012, pp. 232–237.
- [31] E. Mellios, A. Goulianos, S. Dumanli, G. Hilton, R. Piechocki, and I. Craddock, "Off-body channel measurements at 2.4 ghz and 868 mhz in an indoor environment," in *Proc. 9th Int. Conf. on Body Area Networks (BODYNETS)*, 2014.
- [32] N. Y. Hammerla, R. Kirkham, P. Andras, and T. Ploetz, "On preserving statistical characteristics of accelerometry data using their empirical cumulative distribution," in *Proc. Int. Symp. on Wearable Computers (ISWC)*, 2013, pp. 65–68.
- [33] M. Haghghi, M. Bocian, O. Oddbjornsson, J. H. Macdonald, and J. F. Burn, "Synchronous data acquisition from large-scale clustered wireless sensor networks," in *10th IEEE Vehicular Technology Society APWCS*, August 2013.