Towards a monitoring smart home for the elderly: One experience in retrofitting a sensor network into an existing home

Giovanni Moretti, Stephen Marsland, Debraj Basu, Gourab Sen Gupta, the MUSE group *,
School of Engineering and Advanced Technology, Massey University, Palmerston North, New Zealand
{G.Moretti, S.R.Marsland, D.Basu,G.SenGupta}@massey.ac.nz

Abstract. There has been a lot of research interest in research smart homes for behaviour recognition and related tasks supporting the elderly living alone. Amongst the many challenges of such research are the selection of sensors and the secure storage of data. However, there are other important issues such as reliable data collection, acceptance of sensor systems into the inhabitant's life, and ensuring that the sensor system does not create hazards in the house.

In this paper we report on our experiences retrofitting a set of wireless sensors into the house of an elderly person living alone and an exploration of some of the data integrity issues faced.

Keywords: Smart home, elderly monitoring, wireless sensor network, data collection

1. Introduction

The impending 'grey tsunami' is well-recorded [2]: in 2005 the number of people over the age of 60 was 10% of the overall world population, and this number is projected to increase to more than 21% by the year 2050 [12].

As they age, people prefer to stay living in their own homes, even where the ageing process is accompanied by decreasing physical and cognitive abilities, with the consequent reduction in quality of life. There is also evidence that moving out of a familiar environment can lead to a sudden decrease in cognitive abilities; many of the 'activities of daily living' (ADLs) such as eating, dressing, and grooming [28] are over-learned and automated processes that become difficult when they are performed in an unfamiliar place [3].

A possible solution to the challenges of ageing in-place is the monitoring smart home, which has been of significant research interest over recent years, for a recent overview see e.g., [16]. In order to be completely relevant, the smart home needs to be the person's existing home, and the sensors need to be retrospectively fitted. The aim of the home is to observe the behaviour of the elderly inhabitant in order to detect when problems occur. In these events, the system may issue an alert to the inhabitant, or call a relative or caregiver.

A monitoring smart home consists of a sensory system, data storage facility, and an ambient intelligence that interprets the sensor observations. To offer effective assistance, the ambient intelligence must analyse the activities performed in the home, and infer the behaviours from those actions, as they are observed through the sensors. We use the term 'behaviour' to mean a particular set of activities performed with a particular aim in mind, such as doing the laundry, cooking dinner, making breakfast, or watching TV. These are sometimes known as Activities of Daily Living (ADLs). One of the main challenges in behaviour recognition is that the exact activities are not directly observed. The only information provided are the sensor observations, which could be that the kitchen light is on, the oven is turned on and the burner is on; the

*http://muse.massey.ac.nz

inference that therefore somebody is cooking is left to the intelligent part of the system.

In this paper we describe a system that we have designed and implemented for the collection of data about an elderly person living in their own home. The system was installed into the house in April 2012, and data collection has been performed ever since.

Unlike other smart home datasets (see section 1.1 for a brief overview), our system is installed in the existing home of an elderly person. While this has advantages in terms of realism, it does mean that we have chosen not to collect any ground truth in terms of behaviours identified. We have previously described algorithms that can perform unsupervised learning of behaviours [5] and will use this, together with discussions with the house inhabitant, to produce an annotated dataset in the future.

1.1. Relevant Literature

There are several smart home datasets that are concerned with activity or behaviour recognition available; see `http://boxlab.wikispaces.com/List+of+Home+Datasets` for a partial list and [9] for a review.

Of particular relevance for our work are the CASAS datasets from Washington State University [6,7] and the MIT PlaceLab [30,22]. Both of these come from testbed facilities that have a very large number of sensors (over 900 in the case of the MIT PlaceLab) throughout the living spaces. People, often students or professionals, live in these facilities for some period of time, and as well as their activity data being collected, annotations marking their actual behaviours, either from on-the-spot labelling or through observation are attached to the data.

While these are clearly very valuable datasets for activity recognition, they do not tell us much about the actual experiences of installing sensors retrospectively into a person's home, nor how such data from an elderly person living their 'normal' life differs from younger people moving temporarily into prepared accommodation. It is these two things that are of primary interest to us in this paper.

The most directly relevant work has generally looked at multiple site installations, either within an elderly care facility (the TigerGate project [27]) or distributed across many homes, such as the ORCATECH Living Laboratory [20]. Another approach of great relevance is the so-called 'Smart Home in a Box' that has very similar aims to our own work [11].

There has been some work on user perceptions of smart homes, e.g., [8,25], but these are primarily based on anticipated benefits rather than real experience, and so our study provides the first opportunity to investigate real experiences from a monitoring smart home, an area that we will be following up in future work.

2. Project Aims

The Massey University Smart Environments (MUSE) group are studying many aspects of smart home research, from sensor development and networking [14] to behaviour recognition and context awareness [4,15], through to emotion recognition [17]. We have also presented possible use cases for smart home research [23]. This paper describes the design and implementation of our first foray into data collection for the project.

The philosophy behind our approach is that privacy and unobtrusiveness are vital to acceptance of smart homes. Through informal conversations with many people we have identified that while many elderly people are happy to be monitored (particularly when they feel that it reduces the perceived burden on their family or carers) they are also wary of losing their privacy. We have therefore chosen to consider only relatively uninformative state-change sensors, and cameras and microphones are completely ruled-out.

There are some disadvantages to avoiding cameras and microphones. If the system believes that something has gone wrong, data from these sensors would enable the carer to view images from inside the home, which would quickly clarify whether there was any cause for concern or not. However, not using cameras and microphones, as well as obviating privacy concerns, has the added advantage that the computational analysis of sound and image data is not required.

With regard to unobtrusiveness, there are two primary considerations. One is that people modify their actions according to how much they feel that they are being watched, and the second is that the sensors should be relatively easy to install, and avoid damaging the home. These considerations are important for any real smart home system deployed in the future, but it is particularly important for finding unpaid research subjects. The guideline that we have actually applied is the slightly more relaxed "no visible damage", meaning that the odd small screw hole where it cannot seen is not generally considered to be a problem (although this is checked with the owner before proceeding). It also led us to focus on wireless networking systems;

while battery life needs to be considered, the benefit of avoiding wires far outweighed the costs. This criteria also precluded modifications to the existing electrical system in the home, or to the plumbing. Being monitored at all is somewhat invasive, but if these tenets are adhered to, they maximise the acceptability of having sensors in the home and minimise the chance that the inhabitant will alter their normal behaviour (at least, after enough time has passed that they forget about the installation).

In common with many other researchers, our behaviour recognition algorithms assume that data from the sensor stream are available in the form of a sequence of 'tokens', i.e., there has been some preprocessing of the sensor readings (and possibly the fusion of different sensor data) into a sequence of observations that can be used by some form of ambient intelligence system. For our purposes, we consider that the task of the smart home is to segment this token stream into different behaviours, and to identify whether or not this behaviour is typical of the user, and if not, whether it is sufficiently abnormal to warrant an action (such as calling a carer or interacting with the inhabitant, for example in the form of reminders: 'did you remember to turn the gas off?').

The data that is collected by a smart home system needs to be sufficient for the ambient intelligence to do its job. The minimal requirements that we have identified for that task are as follows:

behaviour recognition and segmentation The data collected by the sensors provide only a partial snapshot of the person's activities, and it is likely to be noisy. The smart home needs to be able to deal with such data, segment it into separate behaviours, and reliably identify those behaviours.

novelty detection Once the inhabitants' current behaviour has been recognised, the system needs to decide whether or not the behaviour is novel (i.e., they are doing something that they have not done before, so that the classification failed) or if they are doing it in an odd way (either the activities they are performing are not quite correct, or the time or place are unusual). For a monitoring smart home, this is the part of the system that is likely to produce warnings for carers, or some form of interaction with the inhabitant.

the ability to utilise temporal and contextual data Using this data can improve the results of both the behaviour recognition and novelty detection, indeed it may be essential to detecting novelty in many cases.

life-long learning Things change within houses and in people's behaviours all the time. The system needs to be able to learn about these things so that it does not annoy the inhabitants and carers. In addition, not all behaviours may appear during a training phase, and so the system has to be able to receive feedback and update itself.

2.1. Data: privacy and security

The subject has agreed for their data to be used for this research, and this project has been evaluated by peer review and judged to be low risk. It was therefore exempted from review by the Massey University Human Ethics Committee. It is important to note that the data being collected relates to the activities of the participant going about their daily lives. It is not data related to the individual's health, and therefore does not fall under the requirements of various regulatory requirements relating to health information, such as the New Zealand Health Information Privacy Code [29]. Nevertheless, in keeping with good practice, the subject's personal details (name, address, etc.) are not stored with their activity data, and access to this data is restricted to those needing it for research purposes.

3. Preparation

3.1. Subject Selection

As was discussed earlier, there is very little data available describing how an elderly person lives in their own home, despite the fact that this is the primary use case of the monitoring smart home. This then, was our aim: to find a suitable subject living alone in their own home, who was prepared to act as a test case. The research was funded by the Order of St John, a New Zealand charity that provides healthcare services, including an elderly care system in the form of monitored alarms based on pendants. We asked a local representative of St John to select some potential subjects, and to ask if they would be interested in participating. Following this we chose an elderly person living in a fairly typical New Zealand home.

We believed that it was important to make clear that there is little of no perceived benefit to our house inhabitant or their family, since this is only a prototype system for data collection at this stage. However, as well as discussing what would be useful, and ensuring that our plans did not inconvenience them, we also identified a few things that would be useful, such as a warning system if the kitchen tap was left turned on accidentally or the fridge left open.

3.2. Sensor Selection

As has already been stated, our aim was to use binary state-change sensors as far as possible. We also wanted to err on the side of redundancy amongst sensors rather than missing information, since statistics and information theory can be used to identify useful sensors from this data to discover which sensors were truly informative. This would allow fewer sensors to be used in later installations, without significant data loss.

We began by considering the potential activities that could be monitored using this type of sensor, arriving at the following list:

- room door open/close
- cupboard open/close
- fridge open/close
- seat and bed occupancy
- movement within a room
- appliance use
- lights on/off
- stove/cooker in use
- laundry use
- tap use (waterflow)
- toilet use
- shower use

We also considered a few different types of environmental sensor in order to provide some contextual information, such as:

- temperature
- humidity
- luminance level

Depending upon the particular types of sensor used, it would also be possible to identify a person's location [19], and vital signs [24] or the use of particular cups, etc. [18]. However, in order to make a simple system we chose not to consider these options.

3.3. Networking and Sensor Availability

The desire to avoid modifications to the house and to keep the visual impact of the system as low as possible meant that battery powered wireless sensors were preferred over wired versions. They allow more freedom in their placement than do wired sensors. Even if the sensor requires AC power, there is still an advantage in not requiring cables for communication as this would make deployment in existing homes much more difficult and the occupant is unlikely to be happy with visible cabling.

While many companies advertised "proof-of-concept" sensors or notes on intended future products, there were not many options of sensors and data collection infrastructure for immediate purchase. After investigation of the options available we chose to use a system provided by the Monnit Corporation in the USA (`http://www.monnit.com`). They had a range of sensor modules intended for commercial applications along with the infrastructure to support them.

3.3.1. The Monnit Sensors

Physically, the Monnit sensors are all based around the same 45x25x20 mm package, with a short (12cm) external wire antenna. The sensors are powered by a 20mm disc battery (no. CR2032) that slots into the sensor. Depending on the sensor's setting, the battery can last for more than a year. For longer battery life, a larger alternative packing that uses AA cells is available, but these were not used for this project.

The range of these sensors is stated as 80-100m, which is more than sufficient for use within a home. The sensors communicate with a local hub in the home using proprietary wireless (not WiFi) protocol on either 433MHz, 868MHz or 900MHz. The hub also acts as a gateway onto the Internet. The differing frequencies allow customers to choose whichever conforms to their country's radio transmission regulatory requirements.

The range of Monnit sensors is still in development. At the time that we purchased the sensors (mid-late 2011) the ones that appeared to be the most useful were:

- the Open/Close sensor: this is a reed-relay-based sensor that can detect the presence of a magnet. Either the sensor or the magnet can be mounted on the door and the other on the frame.
- the Dry Contact sensor: this has a pair of unconnected wires that the customer can connect to some form of switch. The sensor will send a notification when the switch is opened or closed. These sensors provide a way to integrate custom-designed sensors into the Monnit data collection infrastructure.
- the Movement sensor: this is a passive infra-red (PIR) sensor that can detect the movement of peo-

ple within a 110° angle horizontally and about 90^o vertically. It has a range of around 5m.
- the Lux sensor: this can measure the light level within a room, and optionally send an alert when the level crosses a threshold.
- the vibration sensor: depending on its sensitivity, this could be used to detect the vibration from some appliances, for example a washing machine.
- temperature sensors
- humidity sensors

Notably absent were sensors for detecting the use of mains-powered appliances or waterflow. In order to overcome this we designed our own mains and waterflow sensors, which were then interfaced to the Monnit *Dry Contact* sensor. These designs are described in section 3.4.

There were also no pressure sensors available that were suitable for monitoring seat occupancy, and after investigating potential solutions to this, together with their costs, it was decided not to go ahead with this sensor for the first round of implementation. Recently, Monnit have released a seat occupancy sensor that consists of a small pressure pad with a cable to the sensor module. The pressure pad is intended to be placed on the seat of a chair. Obviously, it is preferable if the pad is concealed with a cushion or covering, but for many styles of chair there is no way to hide their presence. If the occupant has a favourite chair or two, these could be monitored, but there are too many chairs in most homes for monitoring them all to be practical.

The final set of sensors installed is shown in Table 1.

Figure 1 shows a schematic of the house together with the locations of the sensors and a description of them for the initial installation. Most Monnit sensor are around US$50-60 each with the exception of the PIR sensors which are US$119 each. The installation has nine PIR sensors and sixteen others giving an approximate cost for the Monnit sensors of US$1951. Allowing US180 for all of the locally constructed sensor front-ends (four appliance-use and two waterflow sensors - to be described shortly) gives a sensor cost of US$2131. Adding an Ethernet gateway (US$229) brings the total to US$2360. The Ethernet Gateway requires access to an Internet LAN. If this were not available (or not convenient), the cost would increase. An alternative (now available) would be to use a gateway that connects via the cellphone network.

Sensor type	**Location**
Movement	Kitchen
Light detection	Pantry
Appliance use	Kettle
Appliance use	Toaster
Water flow	Kitchen basin
Open/Close	Kit. cupboard, lower Left
Open/Close	Kit. cupboard, lower Right
Dry Contact/microswitch	Kit. cupboard, upper Left
Dry Contact/microswitch	Kit. cupboard, upper Right
Dry Contact/microswitch	Fridge door
Temperature	Stove
Movement	Dining room
Movement	Main bedroom
Movement	Lounge
Movement	Bathroom
Movement	Toilet
Movement	Front bedroom
Movement	Hallway
Movement	Hobby room
Appliance use	Hobby room
Appliance use	Washing machine
Water flow	Bathroom basin
Temperature	Bathroom
Light detection	Back door
Humidity	Bathroom

Table 1

The sensors currently installed provide greater coverage in the more commonly used parts of the house. Half of the sensors (upper section) are located in the kitchen/dining rooms.

3.4. Custom made sensors

As Monnit did not have sensors capable of monitoring appliance use or waterflow, these were designed locally and interfaced to the 'Dry Contact' Monnit sensor. This provides a way to send the sensor data from a custom sensor via the Monnit infrastructure. Monnit sensors upload their current battery state along with the sensor data on every update, but without modifying the Monnit sensor to use the same batteries, there is no way without using an additional Monnit sensor to upload the state of any batteries powering the custom sensor. Therefore the battery-powered waterflow sensor described below has to have its battery status monitored independently.

3.4.1. Waterflow sensors

The monitoring of water use is problematic as the sensors are being retrofitted into the subject's home, and the desire to have as little impact as possible pre-

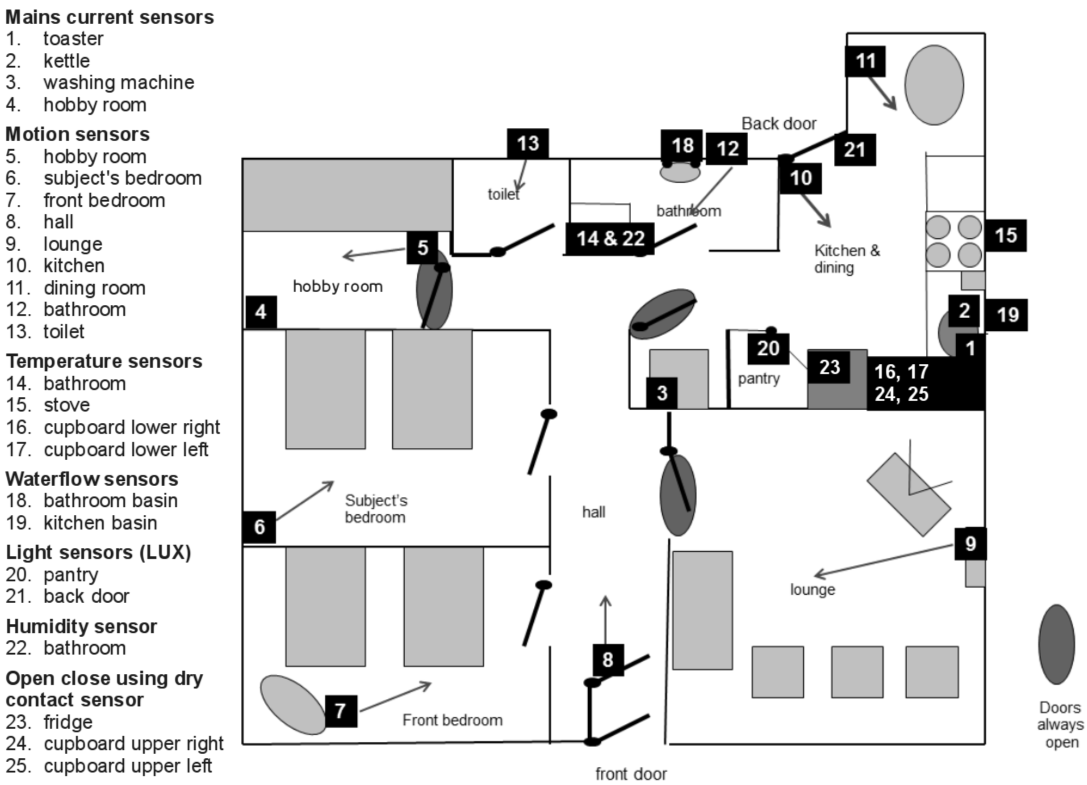

Fig. 1. Schematic of the house showing sensor locations for the initial installation

cludes the use of a conductivity sensor (since this would need trailing wires above the sink), and also the use of an impeller sensor, which would require modifications to the plumbing. The solution finally arrived at was the use of an acoustic water-flow sensor. The turbulence produced when water flows through a constriction (such as a partially open tap) can produce sufficient noise for it to be used as a proxy for water-flow. Care has to be taken to attach the microphone as close to the tap as possible. It could be argued that this violates our 'no microphone' policy, but since the actual audio data from the (very short-range) microphone is not available, there is no conflict with the the privacy argument.

Two measures were taken to ensure that kitchen noises do not trigger a water-flow indication. Firstly, a high-pass filter excludes frequencies below 10 kHz. This effectively removes voices and most music. Secondly, the noise has to be present without interruption for at least two seconds. Almost all uses of a tap will satisfy this condition and a shower too, whereas voices, the banging of pots on the kitchen bench and most other noises will not. They will be above the detection threshold, but will usually drop below it within the two second limit, and therefore will not pass the time threshold. This method of water-flow detection is not ideal, but it does work and satisfies the "minimal impact" requirement. The block diagram of the water flow sensor module is shown in Figure 2.

Larson [21] describes a much more sophisticated system that, from a single point, can distinguish between the use of different waterflow outlets, based on the differences in pressure waves as the valves open and close. This would offer an all-in-one unit for getting detailed waterflow data from multiple outlets. However, the system described appears to be a research prototype. If this system were available commercially, it would be a more elegant solution than that based on individual sensors described here, although our approach is simple, inexpensive, easy to build and easy to integrate with the Monnit sensor network.

3.4.2. *Electrical Current Usage Sensors*

Another piece of data that we felt was important was concerned with the use of electrical appliances. In accord with our desire to avoid modification of the house,

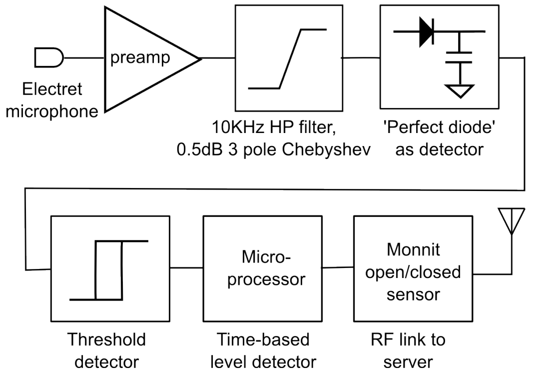

Fig. 2. Block diagram of acoustic waterflow detector

Fig. 3. The electrical current sensor attached to a Monnit Open/Close module

we finally chose to use an intermediary between the appliance plug and the wall socket, as is shown in Figure 3. While this is visible to the users, as the plug projects further from the wall socket, it does mean that the system can be installed without recourse to electricians.

This sensor is placed between the plug of an electrical appliance and the wall socket. It senses the current when the appliance is switched on and shorts the wires from the Monnit sensor using a relay. The block diagram of this sensor module is shown in Figure 4; for more details see [1]. An LED is also used to indicate that the electric appliance is currently in use. The sensor circuitry for the appliance sensor is mains-powered so there are no batteries that need monitoring. The sensor was built and tested by a qualified electrician before use in the house.

Using individual sensors for appliances has some appeal as each directly senses the current taken by a single appliance, giving a definitive indication of whether the appliance is in use. Patel describes a single point connection to monitor all power outlets based on their impact on the power line, learning the effect each has on the power line using effects such as electrical noise [26]. The system is similar in concept to that of Larson mentioned earlier for detecting waterflow. Patel's system also requires training. We have opted for simpler alternative of the direct sensing the current used by selected appliances and integrating the output of those sensor with the Monnit network.

3.5. Accessing the Data

It was important that data was collected as often as possible by the researchers, but without inconveniencing the house inhabitant. We therefore decided that all of the data from the sensors should be collected by, and available from, a server at the University. This also reduces the equipment needed at the subject's home. By far the simplest method of transporting the data is to use the Internet. This was another benefit of the solution provided by Monnit, which includes an Internet gateway hub. In New Zealand, virtually all home Internet access is via broadband using ADSL over the home telephone line. In cities, some Internet service providers (ISPs) offer Internet using dedicated radio links, and in the last year, the national program of "fibre to the home" has commenced, but neither are widely available yet.

Ideally, the occupant would already be connected to the Internet and the sensor network could route data via their network. The sensor data traffic would be minimal and not cause any significant increase in traffic that would count against the monthly broadband data allowance. If broadband is not available, as long as a telephone line is available, the researchers could fund the installation of ADSL broadband and provide complementary access as a side effect. There is however a possible complication to this plan. In New Zealand, a healthcare provider (St John, who funded part of this research) offers an emergency monitoring service for the elderly. They carry a pendant with an alarm button that when pressed will alert a monitoring service who will first telephone the home and if there's no response, send an ambulance. The signals from the pendant base station are also routed over the telephone lines. A local ISP (Inspire Net NZ Ltd.) kindly agreed to sponsor the project by providing broadband access. They chose to avoid possible interactions with the telephone-based emergency monitoring equipment by not using ADSL, but providing Internet access via their radio-based broadband service. It also has the advan-

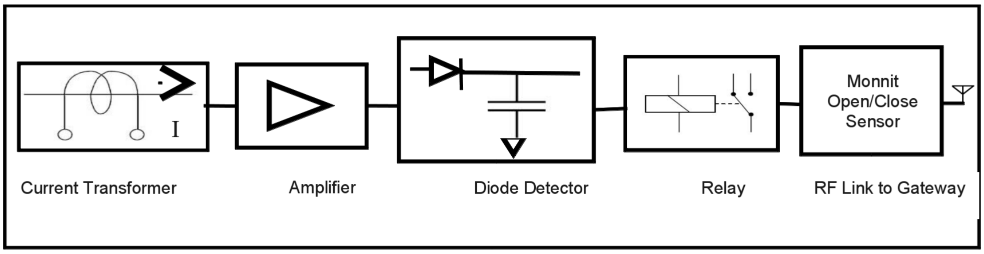

Fig. 4. Block diagram of the mains current sensor, used to detect when appliances are being used

tage that the local hub need not necessarily be placed in the house, as the RF link from the sensors to the hub will is not usually affected when passing through the walls of the house. If the ADSL-based access to broadband is infeasible, the sensor data could be routed as data traffic via the cellphone network.

4. Deploying the sensors

4.1. Construction of Houses in New Zealand

Houses in New Zealand are primarily built using wooden framing with an outer cladding, either of painted wood, brick or composite materials. There is very little metal in either the internal or external walls. Internal walls are almost always plasterboard which is either painted or wallpapered. This form of construction is fortunate as none of the commonly-used materials will significantly attenuate or reflect the signals from the wireless sensors. As the range of the wireless sensors greatly exceeds that required for communication within the home, the sensor hub could also be in a nearby building such as in a car garage, which in older New Zealand homes is often a separate building with AC mains connected.

As the plasterboard wall cladding is normally painted or wallpapered and neither is easily repaired, this precludes mounting sensors directly on to the walls. The only options are to mount sensors on existing items of furniture or on the tops of tall cabinets or the wooden door frames. The Monnit sensors are small and light and can be mounted using the supplied double-sided adhesive tape. If this method of mounting is used, care must be taken to ensure they are not accidentally dislodged as they then pose a potential trip hazard if stood on, which can have serious repercussions, especially for the elderly.

4.2. Installing sensors: The pragmatic reality

4.2.1. Sensor Mounting

While double-sided tape can be used to install many of the sensors, there are other considerations. One is that finding ways to attach sensors that are directional and need to be oriented correctly (such as infra-red movement sensors) is tricky, and the other is that the natural location of many sensors can inconvenience the user. For example, door open/close sensors for kitchen cupboards cannot be placed where they will obstruct the placing of things into the cupboards.

For the infra-red sensors we were also concerned that they could fall onto the floor, where a person with imperfect vision could miss them and trip over the sensor. We therefore chose a more secure method of mounting, which was to screw the sensors to the top or sides of the wooden internal door frames. As the top of the door frame is not visible to the home occupants, there is usually no objection to drilling small holes for mounting screws. An L-shaped bracket screwed into the top of the door frame provides a secure way to mount infra-red movement sensors (Figure 5), but one that allows it to be pivoted in the direction of interest. In this case the steel L-bracket was bent slightly to orient the sensor downwards, as well as to allow rotation.

The monitoring of a particular electrical appliance is straightforward if the appliance stays in a fixed location, but some appliances such as a portable electric heater are often moved. The electrical current sensor is attached to the power plug and while moving it with the heater is possible, it would be awkward and and could probably annoy the occupant. Monitoring undedicated power outlets can show when they're in use, but without any way of finding out by which appliance, and some appliances (like heaters) are turned on and then left on for hours with the occupant moving from room to room. It is not directly useful for indicating where the occupant is spending their time, which is quite a difficult problem. IBM's smart floor[13] un-

Fig. 5. Infra-red movement sensor mounted on the top of a door-frame using a L-bracket, which allows the sensor to be pointed in the direction of interest

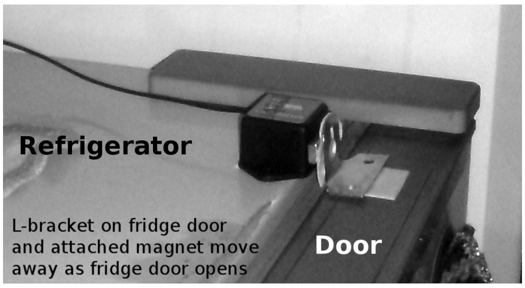

Fig. 6. Assembly of an a magnetic 'Open-Close' sensor on the fridge door

ambiguously indicates the occupants location, but, barring cameras, all other methods are indirect. PIR sensor indicate movement, but not location. For example, the sensor will activate when the occupant enters a room and sits down. If they're reading and not moving, the movement detector will return to its quiescent state. These detectors are good for detecting initial arrival but are limited beyond that.

Another challenge was the refrigerator door. The magnet-based 'Open-Close' sensor initially seemed ideal, but this sensor has to be carefully aligned with the magnet within 2mm or so for the sensor to be detected. This alignment of the magnet/switch assembly was difficult to arrange and maintain without using screws. Initially, the sensor itself was attached to the fridge with double-sided tape, and the magnet attached to an L-bracket that was wide enough to extend over the rubber door seal as can been seen in Figure 6. This proved to be an unstable arrangement, as the alignment of the magnet with the sensor on the fridge drifted. Attaching parts with double-sided tape can work well, but is not ideal when the parts require precise alignment. This arrangement was replaced with a microswitch attached to the refrigerator. The arm of the microswitch extended diagonally down into the gap between the refrigerator and the door. When the door is closed, the rubber seal pushes the microswitch arm against the refrigerator, closing the microswitch. A Monnit Dry Contact Sensor was used to forward on the microswitch contact status to the server.

The last place where our constraint on modifications to the house limited our sensing ability was with the kitchen stove. While monitoring this is potentially very useful, electrical monitoring of the stove requires direct access to the built-in wiring, and we decided that this would be unsafe without a qualified electrician present. While there would be other options (thermal IR imaging, for example) we chose to miss this information for this initial investigation.

4.2.2. Installation Process

While it would be ideal for the researchers to have free and open access to the home for one or two days, and be free to move their equipment – ladders, electric drills, boxes of sensor, tools, documentation, laptops for checking the sensors and so on – in our experience, the reality is quite different.

Instead, it is necessary to work with the occupant at home, usually in two to three hour slots whilst trying to not place items where they might be tripped over (the occupant's eyesight may be poor), and working with the constraint of not damaging the existing home in any significant way. This makes the installation of the sensors significantly more time-consuming than was initially expected. Sensors often cannot usually be placed in their 'ideal' location, and this can preclude the use of sensors that require precise alignment. Unfortunately, this includes the magnet-based Open/Close sensor which meant that these could not be used for their intended purpose. As can be seen in Figure 5, on the facing surface the door frames are not right-angled, but are chamfered and rounded for aesthetic reasons. This means that it is impossible to mount the sensor on the door frame without either a custom-shaped wedge to make up the missing material or having the sensor on a little arm screwed into the top of the door frame.

The kitchen cupboards likewise presented considerable difficulty as there were no suitable mounting points for the sensors within the cupboard that would not have required drilling into the visible parts of the cupboard interior. This was eventually worked around by using microswitches with extender arms screwed into the top of the cabinet, which fortunately did not extend to the ceiling. The practical difficulties of mounting sensors or switches on cupboard doors meant that only the most commonly-used kitchen cupboards were monitored.

The use of 3M double-sided mounting tape intended for clean removal was sometimes useful, but as this requires the sensors be mounted flat against the wall, it was only be useful for sensors that aren't directional, such as room temperature, humidity, or light on/off. Any extension (such as a hinge) would add weight and leverage, increasing the risk that the assembly would become detached.

All of this installation is usually done with the occupant still at home, leaving those installing the sensors acutely aware of his/her presence. While it may be possible to ask the occupant to go out for the morning or afternoon, this was not felt to be appropriate as the installation team were guests, not close friends of the occupant.

The underlying message is that the retrofitting of sensors into an existing home is fraught with pragmatic difficulties, takes far longer than is initially envisaged, and can require multiple visits. Some of these difficulties could be ameliorated by the creation of custom-made brackets (possibly 3D printed) but again this would require multiple visits as it would be necessary to photograph the site (e.g. the door frame), creating the design files, printing the parts, then visiting to check they fit, and possibly iterating. Unfortunately, as there is little uniformity in interior design, the next house would likely require different custom fittings. Installing the sensors is a protracted process that needs good technical support, preferably with multiple people, each covering particular rooms with sufficient tools for them to work independently.

4.2.3. Data Upload

As was previously explained, we chose to use the Monnit sensor system, which includes uploading of data via a hub that acts as a gateway to the Internet.

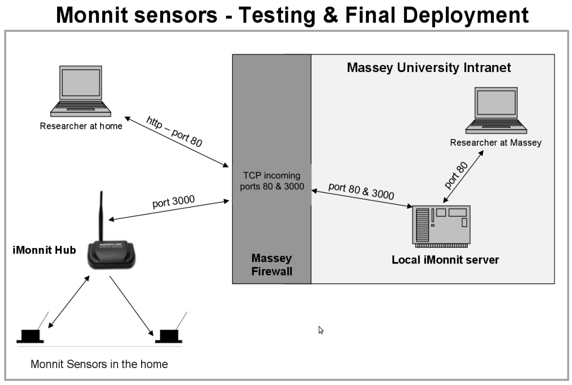

Fig. 7. Monnit Sensor network using Massey-based Monnit Enterprise server

The Monnit hub itself is quite small, around the size of a home Wifi router, so it is not difficult to place it inconspicuously. It requires access to mains power supply and a wired LAN with DHCP. If the sensor-to-hub range of 80-100m is insufficient, a repeater is available from Monnit. Although released too late for use with our initial installation, Monnit have now released a hub that uses the cellphone network instead of a wired LAN for Internet connectivity. This offers the possibility of placing the hub at any convenient power point within 100m of the sensors without requiring on-site Internet access. Once the sensor data is at the hub, it is then forwarded on to a central web site for collection. The architecture of our deployed site is shown in Figure 7.

Other than several status LEDs and a reset button, the hub has no controls. It is preconfigured to download its configuration settings and upload its data to a Monnit-hosted server. This server has pages for viewing or downloading the received data and for configuring the sensor and hub settings. By default, the hub communicates with `https://www.iMonnit.com` or if a license for Monnit Enterprise server has been purchased, the server can be hosted locally. In this case, the central Monnit server can be set to redirect the hub to the local server for subsequent uploads and downloads of hub/sensor configuration updates.

We did purchase a Monnit Enterprise license as this enables the collected data to be held 'in-house' which is preferable for privacy reasons. This also allows us to directly access to the tables the Enterprise server uses to store the sensor data.

4.2.4. The Heartbeat Timer: Updating Hub and Sensor Settings

Many net-connected devices, such as home routers and many printers, use a web server on the device as a way of changing configuration settings. This approach is not used by Monnit. Instead, all configuration updates for both hubs and sensors are downloaded from a central site. It is a 'pull' update, initiated by the hub or sensor, not the Enterprise server, although the server will queue updates from the user. The reason for this is to extend the battery life of the sensors, as it enables them to remain asleep most of the time. Although not necessary, the same mechanism is used to update the setting of the hub.

The sensor settings are updated whenever the sensor sends data or when the 'heartbeat' timer expires. As part of the acknowledgement cycle, any changes in the sensor setting will be downloaded to the sensor. As the sensors spend most of their time asleep, without an update caused by the 'heartbeat' timer, there is no way to be certain that the sensor is still active, or to change its configuration settings. The heartbeat for the hub is initially set to five minutes but can be reduced down to one minute. The means the alterations might not be downloaded to the hub immediately, but the delay is sufficiently small that it is not a problem.

The fact that sensors are mostly asleep can be problematic as it may be some time before they next send data or wake on a heartbeat timeout. During this time it is not possible to contact the sensor, therefore exploring the large number of configuration settings can be a protracted process. When the settings are changed on the server, it will queue them for download when the sensor next checks in. For easy experimentation this should be very frequently, but setting the timeout interval to low values will quickly drain the sensor's battery. For example, the 'dry contact' sensor consumes 50,000 times more current (35mA) when transmitting than the 0.7μA when it is asleep. For long battery life, transmissions must be kept to a minimum. The setting of the heartbeat interval is therefore a compromise between operator convenience and battery life.

4.3. Testing Sensors for Installation

Before installation we tested the sensors. In theory, installation simply requires finding one of the appropriate type, recording its number and location, inserting a battery and then mounting it. However, when we tested each sensor, we found that a few sensors would not power up, meaning delays. A secondary part of this was that the communication of each sensor with the hub was also required. As this installation was for a single home, we only purchased a single internet hub. Unfortunately, this meant that we could not test any further sensors—such as those purchased after deployment—at the University, but only during deployment in the home, since that was where the internet hub was. We have since purchased a second hub to enable testing and further development.

5. Points found from the initial deployment

5.1. Sensor and network failure

As the system is designed to collect data without operator intervention, it is possible for events to be missed, and for this to go unnoticed. These events can range from a single sensor failing to report in, through to a complete lack of data from the home. Complete lack of communication will be due either to a failure in Internet connectivity or a failure of the hub. However there is no straightforward way of finding out which.

The hub will usually be behind a NAT (network address translation) layer of a broadband modem/router so there is no way to ping the hub (it is also unclear whether the hub would respond). While it is technically feasible to pinhole the router and forward data to the hub, this is somewhat invasive. The occupant is unlikely to understand what is being asked, it may compromise the network (albeit in a very minor way), and it requires that the occupant have the login details for the router, something many domestic users either have lost (*"were those papers important?"*) or have never known. When all the sensors stop reporting, the only recourse is to ask the occupant to "turn-it-off-and-on-again" or visit the site.

While the Monnit sensors report battery state within their heartbeat information, there are other potential fail points within a sensor, and it is not always possible to know if a sensor is not working, or has not been activated.

5.2. Sensor proxies, noise and errors

On occasion, there is redundancy in sensor collection possibilities, for example because there is a useful proxy for sensing an event. In the current installation, mounting 'open-close' sensors on the doors of an internal pantry in the kitchen was causing difficulty, but the pantry (being internal and dark) had an internal light that activated when the door was opened. Rather than installing another switch, a light sensor placed inside the pantry was used to detect when the door was open.

Sensors may also sense more than is anticipated. A PIR movement sensor was mounted on one side of a room pointing at the favourite armchair opposite, however the chair was beside the door to the hallway. The 5m range of the PIR sensor meant that the sensor could ‘see’ into the hallway and could trigger when movement in the hall was detected, and this would be reported as movement in the room. The wide-angle coverage of the PIR sensor is usually, but not always advantageous. Where it is not desirable, it is, of course, possible to get finer resolution by point the sensors downwards from the ceiling, and possible by restricting their angular range using contact paper. This approach is used in the CASAS smart homes, which have a very large number of PIR sensors [10].

The waterflow sensor can also be subject to false readings. For example, if an electric blender or waste disposal in the kitchen can both produce sustained noise. Fortunately, neither were present in the home used for the trial deployment.

6. Managing the collected data

The web interface provided by Monnit Enterprise provides many facilities including the pages for configuring both the parameters of the hub and individual sensors, pages for viewing (Figure 8) and graphing the data. It can also send emails if newly arriving data falls outside user-defined limits for that sensor, and allows the last 5000 updates to be downloaded in CSV form, either individually or for all sensors. There is no facility to download all sensor readings. This meant that we needed to download data on a regular basis for analysis to check for possible errors. Initial data analysis is described in section 8.

6.1. Event driven sensors and their heartbeat data

Broadly speaking, there are two types of sensors, those the report in at the intervals specified by the heartbeat (e.g. temperature every 60 minutes) and those that report in when some event occurs (e.g. the fridge has been opened, or closed).

The heartbeat on the sensors cannot be disabled, as it is required to ensure timely configuration updates. Therefore, a cupboard sensor (e.g. #14567) set to have a 60 minute timeout could send data such as (where

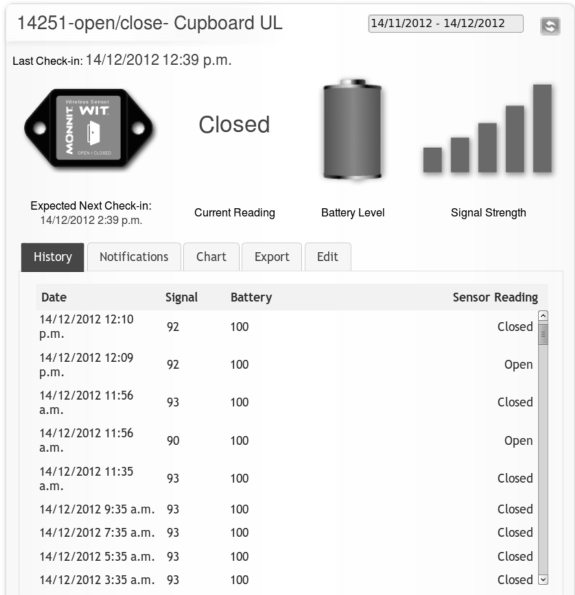

Fig. 8. The web page interface for viewing data from the Monnit sensors

the first entry is the time, the second the sensor ID, and the third the current status):

```
0100, 14567, CLOSED
0201, 14567, CLOSED
0301, 14567, CLOSED
0400, 14567, CLOSED
```

However, because of the heartbeat, it is possible to get multiple “OPEN” readings:

```
1100, 14567, CLOSED
1158, 14567, OPEN    # Cupboard opened
1201, 14567, OPEN    # Heartbeat timeout
1205, 14567, CLOSED  # Cupboard closed
1300, 14567, CLOSED
```

These duplicate “active” readings need to be filtered out by ignoring all following data that has the same state:

– if the previous state was ‘low’, drop ‘low’ events
– if the previous state was ‘high’, drop ‘high’ events

Note that there is some potential variation between the precise minute when data are collected, which means that identification of the heartbeat is not as easy as might be anticipated (and, of course, a sensor event could also be detected at the same time). This will yield only the *change-of-state* data, which is what is wanted (for event-driven sensors). For those sensors that report in at a fixed interval, the heartbeat check-in contains the actual data reading.

6.1.1. Sensor / Sensor Reading Database

Each sensor has an unalterable factory-defined numeric ID number, and can also have a user-defined comment, which initially includes the ID number. The downloaded CSV file includes the sensor ID, the reading timestamp, the user-defined comment along with the state or reading, battery status, and received signal strength. Initially, it was thought that loading the data from the CSV files into a database would be sufficient, as the comment could be used to record where each sensor is located, however it quickly became apparent that this is far too simplistic.

The most obvious difficulty is that the location of a sensor is recorded in the comments field, but this freeform text may be arbitrarily changed. For example, if the sensor in the bedroom was labelled as *"15413-Movement:Master-bedroom"* and it was replaced with another sensor and the comment entered was*"15997-Movement:Main-bedroom"*, if the analysis used the part of the comment as a key, these would appear to be two different rooms (*Master-bedroom* → *Main-bedroom*). As it is far too easy to make these types of errors, they must be prevented by design, with all names picked from a central register. This a classic case of the need for normalising the data, to ensure that there is only one name for each location. The same need applies to the different sites for a multiple house installation, as the design must allow for a main bedroom at each multiple sites, without changing the data access methods.

6.2. Data Organisation

The data available via the web interface includes only recent readings. This may suit the commercial/industrial uses that appear to be the envisioned market for the Monnit sensors, but it is not ideal when scriptable access to all the data is required, nor is direct access to the database of the Monnit Enterprise server database, as the data reported by the web interface is a blend of data from several tables.

An alternative method could be use the mechanism provided for the Monnit Enterprise server to forward newly arriving sensor data on to a generic URL using HTTP. This effectively means that we can collect the data in any form we like once the server end has been created. A wide variety of fields can be used when creating the HTTP URL. Some of these parameters are:

Parameter	Meaning
SensorID	Unique ID for this sensor
Timestamp	Date message was delivered
Data	Transmitted data
DisplayData	UI Display value for data
Battery level	
Signal Strength	

However, as there is no facility for applying one forwarding configuration to all sensors, this configuration must be done on a sensor-by-sensor basis. Our initial purchase was for 69 sensors, and the configuration of these would be tedious. It would also have been error-prone as there is no way to see which sensors have been configured for forwarding, other than by individually inspecting each sensor via its web-based configuration page.

This facility, if augmented, would provide a clean way of collection incoming data without having to manually export tables from either the web interface (with its 5000 row limit) or directly dumping the database tables to CSV files.

Once the data has been collected, whether via processing exported CSV files or from the data forwarded using the HTTP request, it is important that changes at the sensor level, such as replacing a sensor with another of the same type in the same location, be transparent to whatever is using the data. All that is important is that these are the readings from use of the toaster in the kitchen, not which sensor collected the data or even which sort of sensor it was.

6.3. Design of Database and Application API

An API to provide access to the data for analysis and hide implementation details such as which sensor is in which location (or when it was there) requires the following tables:

siteList: <List of people whose houses are being monitored>

locationList: # what room locations are there?

- Bedroom i
- Bathroom
- DiningRoom
- Kitchen
- Hallway
- None (if sensor is alive and reporting but not in use)

Sensors: # A registry of all sensors (active or not)

- sensorID
- sensorTypeID
- StartUseDate/None # None if purchased but not yet activated
- DecommissionedDate/None

SensorType # A register of sensors by class, with possible translation between sensor state and meaning (e.g. a Dry Contact sensor with its contacts shorted (Closed) might mean a cupboard is Open)

- type (O/C, Temp, Lux, Light, DryContact ...)
- readingRange
- class: (event or timeoutInterval)

Sensor Location # Records which sensor is at a particular location within a site, and during which time interval it was there.

- start-time
- end-time (or None)
- sensorID
- locationsListID
- siteListID

It is also necessary to be able to specify periods in which the sensor network is 'down for maintenance'. These will happen for example when the researchers visit the site and are replacing batteries, testing and possibly replacing sensors. During this period the hub needs to be online so it can be seen that newly installed sensors are active, but while these readings will be recorded on the server, they should be excluded from any analysis of the occupant's activities.

Once the data has been regularised, it could then be made available in different forms: via CSV files or via a REST web interface returning XML or JSON. However as the data is about research subjects, access may need to be restricted to those with appropriate rights, and possibly recorded in access logs.

7. Communication port and protocol details are essential

As part of the initial purchase, Monnit upload initial configuration for all purchased sensors onto a free account at `http://www.iMonnit.com`. As the sensor hub is preconfigured to upload data to this server, as soon as batteries are inserted into the sensors, they will start uploading data. It is then possible to login and view the data uploaded from the sensors, alter the sensor parameters and parameters relating to the hub configuration. However, firewall issues may cause updates to the sensors or Hub configurations to be blocked. If this happens, the system silently fails: no warnings are given. A 'Hub Communication Test' panel has been requested from Monnit.

The hub uploads data using a TCP connection to the server on port 3000 and receives any reconfiguration updates as part of the reply. The sensors get their configuration updates the same way, as part of the *upload-data/get acknowledgement* cycle.

7.0.1. Corporate firewall issues

The hub communicates with the server using TCP port 3000, but which direction was unclear from the documentation. If the Monnit Enterprise server is behind a corporate firewall, data arriving from the outside must be forwarded through the firewall to the server (a pinhole is needed). As it is now common for servers to also have their own firewall, an additional pinhole will be needed on the server. It was only with the hub inside the firewall on the same intranet as the server, but with no data being uploaded that it was realised what was happening. Microsoft Server 2008 was blocking the incoming connections. Once a pinhole was added to the server firewall, everything functioned as expected.

8. Data Analysis

In this paper we describe some initial software that we have developed for data analysis to ensure that sensors are working and that sensible data is being observed. It is primarily a set of data visualisations for human observers rather than an intelligent system. However, it is also capable of producing sets of tokens in the form of a data stream that is suitable input for behaviour recognition systems.

Two primary methods of data visualisation were chosen, and examples of them are shown in figure 9 and 10. In the first, a schematic of the house is shown, with sensor locations marked as circles. A period of time is selected, and a movie is shown of the sensors being activated and inactivating, through changing the colour of the sensor markings, with one hour of real time being compressed into 20 seconds. This is a useful way to see the actual observations as they occur. Figure 10 instead shows an overview of an entire 24 hour period, with the times when each sensor was activated clearly delineated. Together, these two visualisations can help to identify sensors that are not working, and highlight other areas for investigation. They also enable us to pick out potential patterns in the data.

The illustrative data shown in figure 10 provides a lot of information about our system and also the inhabitant (for this reason, the data has been slightly modified to aid privacy). It can be seen that there is good correlation between the motion sensor in the kitchen and the various sensors in the kitchen (e.g., cupboards) and that neither the hobby room nor the washing machine were used during that day. The house inhabitant uses the kitchen/dining room as a principal recreation area. There appear to be some issues with the motion sensor in the bathroom, which were fixed at a later time.

9. Conclusions

In this paper we have described our planning process and experience in retrospectively installing a set of wireless sensors for a monitoring smart home into the house of an elderly person. While this is only one case study (and we are currently preparing to install a second set of sensors into a different house) we hope that it will highlight some important issues for researchers planning to undertake a similar investigation. The currently-available datasets, which reflect very detailed observations of people living in prepared spaces for relatively short periods of time are very useful for developing behaviour recognition algorithms, but we need complementary studies such as this one to ensure that the specific issues of elderly people and retrofitting are considered.

A few of the lessons that we have learnt along the way that we did not expect are that the actual installation process is far more awkward and time-consuming than we expected (working in somebody else's home while they are present, and without causing any damage), that dealing with sensors that are asleep makes testing and changing settings time-consuming and problematic, and that it is critical that all aspects of the sensors and network are well tested before as well as during and after installation. The second of these is particularly interesting, since if a sensor is asleep, updates to it do not happen immediately. You can change settings on the server but these will not be downloaded to the hub until its next timeout (which can be quite short e.g. 5 minutes), and once at the hub, the update will not get to the sensor until either a sensor heartbeat timeout occurs, or its event happens (e.g. an open-close sensor activates). This timeout can be much longer, for example up to 59 minutes for an sensor with an hour heartbeat. As the 'pull update on heartbeat' is essential to getting a good battery life from the sensors, it is unlikely that much can be done about this, but it does make experimenting with sensors a protracted process.

Sensor networks are complex mixtures of hardware and software, and time needs to be put in to ensuring that each part works. Software tools to assist in this validation are an important part of the system that we have developed. We have been modifying a variety of parts of our system over the past few months, trying out different sensor configurations and finding out what works and what does not. Over the next few months we will prepare to release a dataset of human activities over a reasonable time period (such a three months) with a static sensor configuration. Unlike many readily-available datasets we do not have a ground-truth annotation, but we will provide some interpretation of the data with it.

Over time we also hope to learn the difference between anticipated and actual benefits and disadvantages during the preparation, installation, and actual usage of a smart home monitoring system through interviews with our current subject and future participants.

Acknowledgements

The MUSE group, listed as an additional author on this paper, include Hans Guesgen, Ayesha Hakim, Paul Lyons, Amal Punchihewa, An Cong Tran; see `http://muse.massey.ac.nz` for further details. Financial support for this project was provided by the Priory in New Zealand of the Most Venerable Order of the Hospital of St John of Jerusalem and InspireNET (Palmerston North, New Zealand). We are also very grateful to Jillene Durham of St John's for useful discussions. The work was granted low-risk ethics notification by the Massey University Human Ethics Panel. Finally, many thanks are due to our anonymous house inhabitant and their family for allowing us into their home and lives.

References

[1] D. Basu, G.S. Gupta, G. Moretti, and S. Marsland. Wireless sensor network based smart home: Deployment and monitoring. In *IEEE Sensors Applications Symposium*, 2013.

[2] J. Beard, S. Biggs, D.E. Bloom, L.P. Fried, P.R. Hogan, A. Kalache, and S.J. Olshansky. Global population ageing: Peril or promise? PGDA Working Papers 8912, Program on the Global Demography of Aging, January 2012.

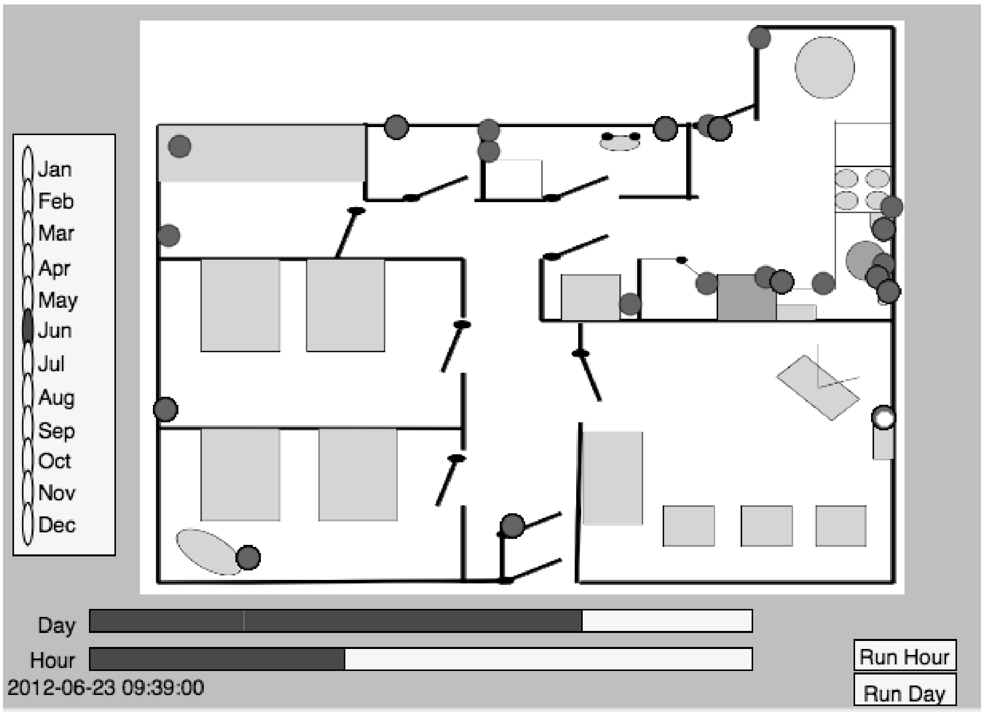

Fig. 9. A snapshot from a movie showing a speeded-up visualisation of the house observations. Grey circles mark the locations of sensors that are not currently active, while those with a white dot in show the location of currently active sensors.

[3] R.S. Bucks, D. Ashworth, G. Wilcock, and K. Siegfried. Assessment of activities of daily living in dementia: Development of the Bristol activites of daily living scale. *Age and Ageing*, 25:113–120, 1996.

[4] S.-L. Chua, S. Marsland, and H.W. Guesgen. Behaviour recognition from sensory streams in smart environments. In *Proceedings of the Australasian Joint Conference on Artificial Intelligence*, 2009.

[5] S.-L. Chua, S. Marsland, and H.W. Guesgen. Unsupervised learning of human behaviours. In *Proceedings of the Twenty-Fifth Conference on Artificial Intelligence (AAAI-11)*, 2011.

[6] D. Cook. Health monitoring and assistance to support aging in place. *Journal of Universal Computer Science*, 12(1):15–29, 2006.

[7] D. Cook, M. Schmitter-Edgecombe, A. Crandall, C. Sanders, and B. Thomas. Collecting and disseminating smart home sensor data in the CASAS project. In *CHI 2009 Workshop on Developing Shared Home Behavior Datasets to Advance HCI and Ubiquitous Computing Research*, 2009.

[8] J.F. Coughlin, L.A. D'Ambrosio, B. Reimer, and M.R. Pratt. Older adult perceptions of smart home technologies: Implications for research, policy and market innovations in healthcare. In *Proceedings of the IEEE Engineering in Medicine Biology Society*, pages 1810–1815, 2007.

[9] L. Coyle, J. Ye, S. McKeever, S. Knox, M. Stabeler, S. Dobson, and P. Nixon. Gathering datasets for activity identification. In *CHI 2009 Workshop on Developing Shared Home Behavior Datasets to Advance HCI and Ubiquitous Computing Research*, 2009.

[10] A. Crandall. *Behaviometrics for Multiple Residents in a Smart Environment*. PhD thesis, Computer Science Dept, Washington State Uni, 2011.

[11] A. Crandall and D. Cook. Smart home in a box: A large scale smart home deployment. In *Proceedings of the Workshop on Large Scale Intelligent Environments*, 2012.

[12] Department of Economic and Social Affairs. World population prospects: The 2006 revision. Technical Report ESA/P/WP.202, United Nations, 2006.

[13] L.M. Do, T.M. Grigsby, P.A. Nesbitt, and L.A. Seacat. Securing premises using surfaced-based computing technology, 2012. US Patent 8,138,882.

[14] A. Gaddam, S.C. Mukhopadhyay, and G. Sen Gupta. Elder care based on cognitive sensor network. *IEEE Sensors*, 11(3):574–581, 2011.

[15] H.W. Guesgen and S. Marsland. *Spatio-Temporal Reasoning and Context Awareness*, pages 609–634. Springer, Berlin, Germany, 2010.

[16] H.W. Guesgen and S. Marsland, editors. *Human Behavior Recognition Technologies: Intelligent Applications for Monitoring and Security*. IGI Global, 2013.

[17] A. Hakim, S. Marsland, and H.W. Guesgen. A robust joint face model for human emotion recognition. In *Proceedings of Image and Vision Computing New Zealand*, 2012.

[18] H.-H. Hsu, C.-N. Lee, and Y.-F. Chen. An RFID-based reminder system for smart home. In *IEEE International Conference on Advanced Information Networking and Applications*, 2011.

[19] S. Hussain, S. Schaffner, and D. Moseychuck. Applications of wireless sensor networks and RFID in a smart home environment. In *Seventh Annual Communication Networks and Ser-*

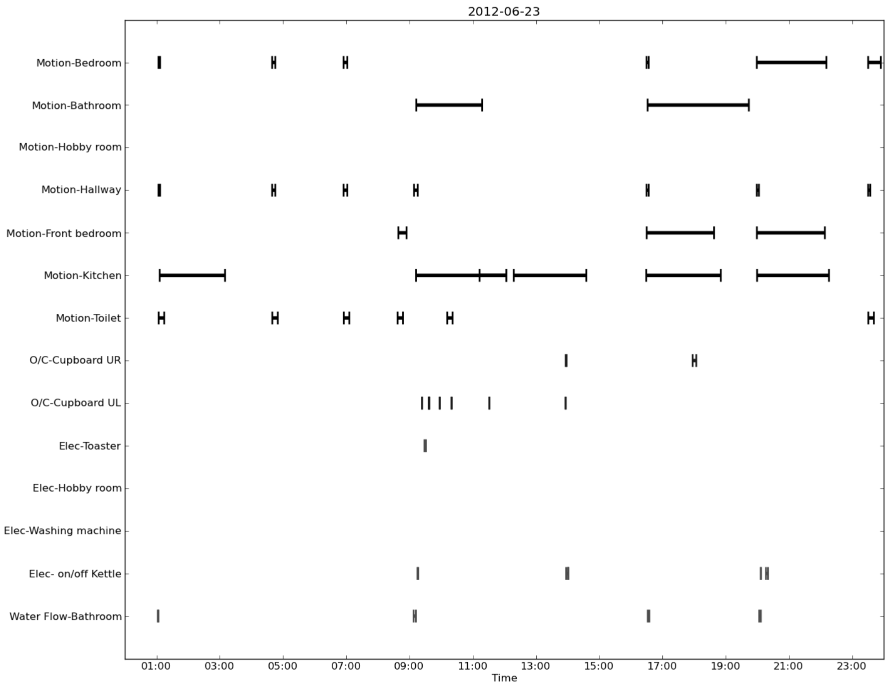

Fig. 10. Illustrative data of possible activations of a subset of the sensors for a particular day.

vices Research Conference, 2009.

[20] J. Kayea, T. Hayes, T. Zitzelberger, J. Yeargers, M. Pavel, H. Jimison, N. Larimer, J. Payne-Murphy, E. Earl, K. Wild, L. Boise, D. Williams, J. Lundell, and E. Dishman. *Deploying wide-scale in-home assessment technology*, pages 19–26. 2008.

[21] E. Larson, J. Froehlich, T. Campbell, C. Haggerty, L. Atlas, J. Fogarty, and S.N. Patel. Disaggregated water sensing from a single, pressure-based sensor: An extended analysis of hydrosense using staged experiments. *Pervasive and Mobile Computing*, 8(1):82 – 102, 2012.

[22] B. Logan, J. Healey, M. Philipose, M.E. Tapia, and S. Intille. A long-term evaluation of sensing modalities for activity recognition. In *Proceedings of Ubicomp*, pages 483–500, 2007.

[23] P. Lyons, A.T. Cong, H.J. Steinhauer, S. Marsland, J. Dietrich, and H.W. Guesgen. Exploring the responsibilities of single-inhabitant smart homes with use cases. *Journal of Ambient Intelligence and Smart Environments*, 2(3):211–232, 2010.

[24] M. Marschollek, O.J. Bott, K.H. Wolf, M. Gietzelt, M. Plischke, M. Madiesh, B. Song, and R. Haux. Home care decision support using an Arden engine–merging smart home and vital signs data. *Studies in Health Technology and Information*, 146:483–487, 2009.

[25] C. McCreadie and A. Tinker. The acceptability of assistive technology to older people. *Ageing and Society*, 25:91–110, 2005.

[26] S.N. Patel, T. Robertson, J.A. Kientz, M.S. Reynolds, and G.D. Abowd. At the flick of a switch: Detecting and classifying unique electrical events on the residential power line. In *UbiComp 2007: ubiquitous computing*, pages 271–288. Springer, 2007.

[27] M. Rantz, M. Aud, G. Alexander, D. Oliver, D. Minner, M. Skubic, J. Keller, Z. He, M. Popescu, G. Demiris, and S. Miller. Tiger place: An innovative educational and research environment. In *AAAI Fall Symposium*, 2008.

[28] N. Robert, K. Taewoon, L. Mitchell, and K. Stephen. Living quarters and unmet need for personal care assistance among adults with disabilities. *Journal of Gerontology: Social Sciences*, 60B(4):S205–S213, 2005.

[29] B.H. Slane. Health information privacy code, 2008. http://www.privacy.org.nz/assets/Files/Codes-of-Practice-materials/Health-Information-Privacy-Code-1994-including-Amendment.pdf.

[30] E. M. Tapia, S. S. Intille, and K. Larson. Activity recognition in the home using simple and ubiquitous sensors. In *Pervasive*, pages 158–175, 2004.