Design of a Home Sensing System for Monitoring Agitation in People with Dementia with Real World Deployment Considerations

A Thesis

Presented to

the faculty of the School of Engineering and Applied Science

University of Virginia

in partial fulfillment

of the requirements for the degree

Master of Science

by

Joshua Dugan

August

2016

APPROVAL SHEET

The thesis

is submitted in partial fulfillment of the requirements

for the degree of

Master of Science

AUTHOR

The thesis has been read and approved by the examining committee:

John Lach

Advisor

Kamin Whitehouse

Steve Wilson

Accepted for the School of Engineering and Applied Science:

Craig H. Benson, Dean, School of Engineering and Applied Science

August

2016

Abstract

Alzheimer's disease affects a large number of elderly people throughout the world. Caring for a person with Alzheimer's dementia who lives at home rather than in a nursing facility can be extremely taxing on the family of the person with dementia (PWD). Caregivers report high levels of stress and depression, and that episodes of agitation are the most difficult part of caring for a PWD. An inability to adequately deal with agitation is the most important reason caregivers cite for moving the PWD to a nursing facility. The Behavioral and Environmental Sensing and Intervention (BESI) project aims to increase the caregivers' self-perceived ability to deal with agitation by predicting agitation events and providing the caregivers with real-time notifications and targeted interventions. Agitation is predicted based on the environmental conditions around the PWD measured continuously using a sensor system deployed in the home. Physical agitation and the behaviors that lead up agitation events are measured using a wearable inertial sensor on the PWD. The inertial and environmental data will be fed into a model personalized to each PWD that predicts agitation and suggests targeted interventions to the caregivers.

This thesis presents the design of a sensor system for use in the BESI project that measures the environmental conditions in a house occupied by a caregiver and a PWD and detects agitation using a wearable accelerometer on the PWD. The system continuously collects sensor data from multiple rooms, and aggregates it in a central location for analysis. Designing home monitoring systems like this one, presents a number of significant engineering challenges that need to be addressed to collect, transmit, and synchronize data from many disparate sensing modalities. These challenges along with considerations for in-home deployment, such as system reliability, monitoring, repair, and appearance are discussed. The sensor system has been deployed in the homes' of two caregiver-PWD dyads with each deployment lasting seven days, and the results from these deployments are presented in terms of the amount of data collected from the environmental and body-worn sensors.

Acknowledgements

The work presented in this thesis would not have been possible without the support of my professors, colleagues, friends, and family. Thank you to everyone who has helped me and a special thank you to the following people:

I would like to thanks my advisor, Professor John Lach, for his valuable advice and guidance during my time studying at the University of Virginia.

The progress made so far on the BESI project and the work presented in this thesis is only possible because of the efforts of the other students at UVA who have worked on the project: Ridwan Alam, Nutta Homdee, Ben Ghaemmaghami, Neeraj Ghandi, and Sriram Dandu.

This thesis would also not be possible without the inputs of our collaborators throughout the BESI project: The researchers at North Carolina A&T led by Dr. Tonya Smith-Jackson and the Virginia Tech Carilion clinical team: Temple Newbold, Dr. Kemi Bankole, and Dr. Martha Anderson.

I appreciate the feedback and advice from all of the members of the INERTIA group who provided valuable guidance on my research throughout my time at UVA.

To the members of my Master's thesis committee, Professors Steve Wilson and Kamin Whitehouse, thank you for time and advice to improve to quality of this thesis.

Finally, I would like to thank my family for their love and support.

This work was funded in part by the National Science Foundation under grant #IIS-1418622.

Table of Contents

I. Introduction

An estimated 5.3 million Americans, most over the age of 65, have Alzheimer's disease [1]. Alzheimer's is the most common form of dementia with about 70% of dementia caused by Alzheimer's disease [2]. Common symptoms of dementia are memory loss, cognitive, visual, and vocal impairment. Dementia gets progressively worse over time, and many people with dementia rely on caregivers to accomplish daily tasks. These caregivers are often untrained family members living with the person with dementia (PWD), or trained nurses in assisted care facilities. In 2014 caregivers provided an estimated $17.9 billion in unpaid care for people with Alzheimer's disease. Caring for a family member with dementia can be incredibly taxing on the caregiver, and about 40% of caregivers suffer from depression, while 60% report high rates of emotional stress [2]. As Alzheimer's disease progresses, the PWD requires more attention and becomes more difficult to care for. Family caregivers often feel that they can no longer effectively care for their loved one with dementia, and feel forced to move them to an assisted care facility. This can place a significant financial and emotional burden on the family of the PWD and can be detrimental to the mental and emotional health of the PWD because of the change in living environment [3].

Caregivers of people with dementia report that the most important factor leading to moving a PWD from home to an assisted living facility is dementia related agitation [4]. Agitation is a set of behaviors that are generally either repetitious behavior without a purpose, like pacing, or socially inappropriate or aggressive behavior, such as verbal outbursts or hitting [3]. Agitation is generally assessed based on the frequency and severity of a variety of behaviors based on observation of the PWD or caregiver interviews. Agitation is thought to be caused by discomfort or pain experienced by the PWD that they are unable to relieve on their own and cannot properly vocalize to others. Three common sources of agitation are physical maladies, medication, and the environment. Physical pain is often a part of normal aging, and side effects from medication can exacerbate this and lead to increased agitation. Dementia patients often have trouble dealing with uncertainly and change, so deviations from their normal routine or unusual environmental conditions, such as temperature deviations, dim lighting or excess noise, can increase the frequency and severity of agitation [3].

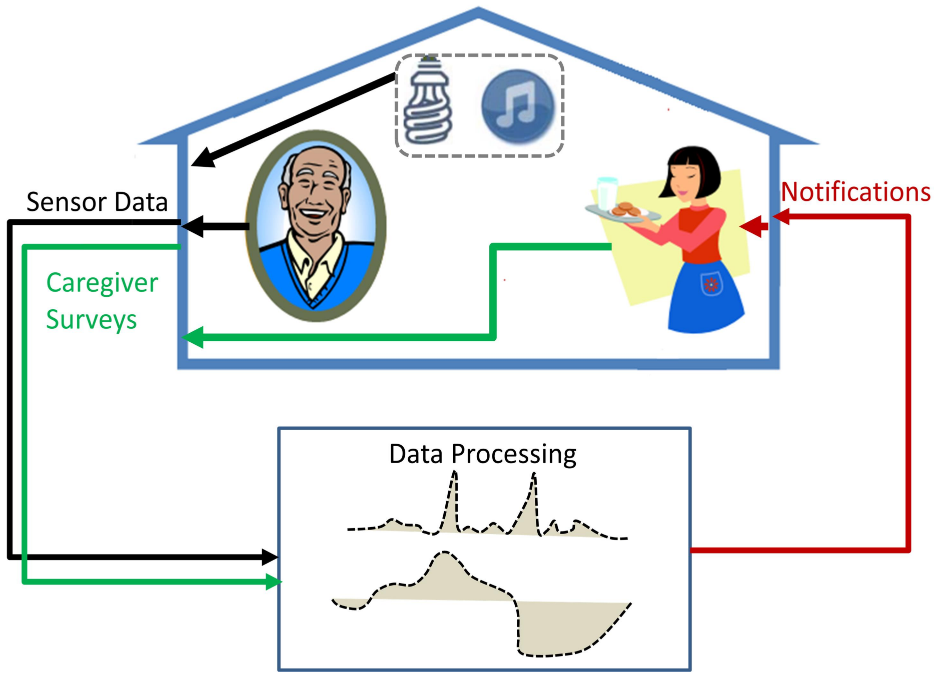

Figure 1. Conceptual representation of the BESI project. Environmental data, wearable sensor data, and caregiver surveys are combined to detect and predict agitation and to notify the caregiver. (Figure is a modified version of one originally made by Dr. Philip Asare)

This document describes the design and implementation of a sensor system that, when placed in the home of a caregiver-PWD dyad, detects agitation in a PWD and measures the environmental conditions that they are experiencing leading up to and at the time of agitation. The sensor system is part of a larger project, Behavioral and Environmental Sensing and Intervention (BESI), which has the goal of predicting or detecting agitation in its early stages and notifying the caregiver of the PWD with personalized interventions to directly interact with the PWD or to change the PWD's environment. The project is attempting to show that these interventions can empower the caregivers of people with dementia and increase their perceived ability to care for their loved ones. To supplement the information obtained from the sensor system, the caregiver regularly fills out surveys on a tablet computer on the activities and mood of the PWD and to document periods of agitation. The caregiver also reports his or her own mood through surveys. The information from the surveys and the sensor data will be used in models personalized to each deployment that predict and detect agitation and suggest interventions to the caregiver. Figure 1 shows a conceptual representation of the BESI project. The project is a collaboration between electrical and computer engineers from the University of Virginia, industrial engineers specializing in human-computer interface (HCI) from North Carolina A&T, and medical experts in geriatric psychiatry and nursing from the Virginia Tech Carilion School of Medicine. The researchers at UVa are primarily responsible for developing the sensor system and analyzing the data. The researchers at NCAT have designed the interface of the caregiver tablet application, and will assess issues of usability with tablet and the sensor system and develop models of the living space of the caregiver-PWD dyad. The clinical team from Carilion interfaces with the participants and potential participants for recruitment, deployment planning, and providing any assistance that the dyad requires during a deployment. They also conduct assessments of the caregiver and the PWD before and after the deployments and provide input to ensure that everything placed in the home during a deployment is as safe and unobtrusive as possible.

This BESI system is deployed in single-family homes that are occupied by a PWD and their caretaker. The study is specifically targeted at patients that are living at home with a family member or friend as their caregiver rather than those that are living in assisted-care facilities. The sensor system monitors the environmental conditions in every room of the house in which the PWD spends any significant amount of time and collects motion data from the wearable sensor while the PWD is in the house. The environmental conditions monitored are: light level, temperature, audio noise level, and the number of people in the room with the PWD. Physical agitation is assessed using a wrist-worn accelerometer on the PWD that measures acceleration in three dimensions and is used assess the PWD's level of motion. Verbal agitation is measured using the same microphones that are used to monitor the environment.

The deployments are organized into three phases: system validation, model development, and intervention. The system validation phase consists of two of deployments lasting one week each in which the main goals are to validate that the data is correctly collected from the sensor system and tablet surveys and to collect feedback from the dyad on having the BESI system in their house. The model development phase has eight deployments of one month each, where the data from the sensor system and tablet are used to develop models personalized to each dyad to detect and predict agitation. This phase is still just data collection and does not provide any suggested interventions to the caregiver. The third phase will be twelve deployments each lasting 60 days. The first half of each deployment in this phase will be used to develop personalized models for each dyad as in phase two, and in the second half, the BESI system will use the models to recommend interventions developed by the clinical team to the caregivers in real-time to assist them in preventing or lessening agitation. Ideally these interventions will reduce the severity and frequency of agitation and improve the caregiver's mood as reported in the tablet surveys. For all three phases, environmental data is collected 24 hours per day and accelerometer data is collected whenever the PWD is in the house.

We have conducted the two phase one deployments of the BESI system in participants' homes to validate the functionality of the sensor system and tablet application. The focus was on our ability to measure things like physical agitation, verbal agitation, and environmental context, and to get feedback on the usability of the tablet and sensor system. In this phase and in phase two, the sensor system collects and stores the sensor data for post-deployment analysis, but in phase three the system will need to be capable processing the sensor data in real time to generate notifications that will be sent to the caregiver. These notifications need to be sent in a timely manner so that the caregiver can quickly intervene, although the relationship between notification timing and the effectiveness of the interventions is not currently known.

Sensor System Requirements

The sensor system for BESI is required to detect physical and verbal agitation in the PWD and to measure the indoor environmental conditions that the PWD experiences prior to, during and after periods of agitation. Physical agitation can be measured using a wrist-worn 3-axis accelerometer as demonstrated in [8] and discussed in the related work section. Verbal agitation requires a microphone close enough to the PWD to detect speech or other verbal outbursts. The environmental conditions that need to be measured are light, noise and temperature, which can be measured by an ambient light sensor, a microphone, and a temperature sensor respectively. The environmental sensing modalities were selected based on input from the medical team. Because these conditions are not constant throughout a building, the system requires multiple instances of each environmental sensor placed throughout the area of interest to get an accurate picture of the conditions that the PWD is experiencing at any given time. The logical way to deploy the sensors is with one of each kind of sensor per room because conditions within a room tend to be fairly constant compared to conditions between rooms. The system also needs to measure motion to determine the number of people in the same room as the PWD as a proxy for the level of social interaction that the PWD is experiencing.

Because the sensor system is deployed in actual patient's homes for extended periods of time, the design needs to consider the constraints that this imposes. The system is deployed in a single family home occupied by a PWD and their caregiver. Every room in the home in which the PWD generally spends time needs to be monitored with the environmental sensors. Because the system needs to be usable in a large variety of potential deployment homes and the specific floor plan and rooms to monitor might not be known in advance of a deployment, the number of rooms and the total space that the system can monitor need to be quickly configurable. The system also needs to be deployed for up to two months at a time. This establishes a requirement for the amount of data the system needs to store, but it also creates the need to check for problems in the data collection while a deployment is in progress and if possible to address issues without requiring a visit to the dyad's home. While the system does not need to look like a commercial product, it does have to look acceptable enough for people to be willing to have it in their homes. The system needs to present minimal risk of physical or mental harm or discomfort to the subjects, and it needs to be resilient to bumps or drops that may occur during a deployment.

The sensor data generated over long deployments needs to be stored for post-deployment analysis and it needs to be time synchronized. The sensor data generated across multiple rooms for long deployments can exceed 100 GB for a single deployment. All of this data needs to be tagged with the room, sensor, and deployment from which it was obtained and stored either at the deployment site or on a remote cloud server so that researchers and clinicians can access it post deployment. The timestamps associated with the data from different rooms and different sensors also needs to be synchronized so that the correlations between sensors can be determined. The synchronization only needs to be to within a few seconds, but because of the length of deployments, even a small drift in timestamps between two data streams can lead to large synchronization errors.

The deployments in the first two phases of the BESI project are only for data collection with the data analysis occurring after the deployment is complete. The important metric for evaluating the performance of the sensor system is the amount of time when data is collected during a deployment. The results section of this thesis evaluates the system using this metric. While the initial deployments of the system are just for data collection and not intervention, the system eventually needs to be able to do the data transfer and processing required to send real-time notifications to the caregiver. In the first deployments that do not involve real-time agitation notifications and personalized intervention suggestions to the caregiver; however, the system must be designed in a way that it can eventually be extended to do real-time analysis and notification.

In addition to the application requirements listed above, the sensor system is also designed to be modular and flexible. While the required sensor modalities were determined prior to the design of the system, the results from the first deployments may demonstrate the need for additional types of sensors or show that a particular sensor that is included provides no useful information. The hardware and software for each sensor is designed to be modular to facilitate the easy addition or removal of sensors. The system is also modular in that the number of rooms to be instrumented can be easily changed. Because the primary goal of the sensor system is to conduct research it is designed to allow a variety of architectures in terms of where data is stored and processed to facilitate exploration.

Design Challenges

The rest of this document presents the development of the sensor system for detecting agitation in people with dementia and for monitoring the conditions they are experiencing before, during and after episodes of agitation. It also discuses some of the general issues and potential solutions associated with developing this type of sensor system for in-home use. The design of a home monitoring sensor system presents several significant challenges that will be addressed throughout the remainder of this document:

1. **Multi-room Data Environmental Collection (Section III.B) -** Collecting environmental data from multiple rooms requires a distributed architecture that either has the capacity to store sensor data locally in the room where it is collected, stream it to a central location, or both.
2. **Temporal Synchronization (Section III.C) -** A distributed architecture makes time synchronization of sensor data more difficult because individual components have their own

sense of time based on a local clock. The nodes in the sensor system need to communicate to agree on a global time, and this communication needs to be wireless to make the system easy to install and as unobtrusive as possible.

3. **Body-Worn Wireless Sensor Handoffs (Section III.D) -** Wearable sensor platforms with a form factor that is acceptable for long-term use do not have the wireless communication range to cover an entire house. This means that the system needs to have multiple access points for the wearable to connect to throughout the house and a mechanism for reliably switching between access points.
4. **Remote Monitoring (Section IV.A) -** Monitoring the state of the system during a deployment to detect errors quickly and prevent extended periods of time when sensor data is lost is challenging because it requires a node in the sensor system can communicate with the outside world. This node needs to have enough information about every other component in the sensor system to determine if anything is wrong, which requires communication for data aggregation within the system. Remote monitoring is important to prevent long periods of time where no useful data is collected
5. **Unreliable Wireless Communication (Section IV.B) -** Wireless communication is prone to dropped packets, variable latencies, and lost connections, which makes designing a reliable system difficult. Because the sensor nodes need to reliably communicate wirelessly for time synchronization, data aggregation, and remote monitoring, the sensor system needs to have mechanisms to detect communication problems and automatically correct them if possible.
6. **Remote Maintenance (Section IV.A) -** Fixing problems that occur during a deployment is especially difficult because visits to the deployment site or even contact with the dyad are limited by the availability of the caregiver, and even when the caregiver is available, home visits or even calls to the dyad are a significant imposition on them. This requires that as many errors

as possible be fixed remotely during a deployment without disturbing the dyad. Common problems need to be identified before deploying the system and fixes developed that can be done without physical access to the system.

7. **Real-World Deployment Constraints (Section IV.C) -** Designing for in-home deployment places constraints on a sensor system that are not present when testing the system in the lab. Size, appearance, and safety to the dyad need to be considered along with ease of installation and disassembly. The system also needs to be quickly configurable for different houses.

This thesis is organized as follows: section II presents the background and related work in agitation detection and home monitoring. Section III contains a description of how the sensor data is collected, synchronized and aggregated by addressing challenges 1, 2, and 3. Section IV describes issues and problems that needed to be addressed to make the sensor system suitable for real world deployment (i.e., how challenges 4, 5, 6, and 7 were addressed). Section V contains the results from these deployments, and section VI presents the reasoning behind the significant design decisions made in developing the sensor system. Section VII concludes and presents some directions for future work.

II. Background

A. Introduction to Sensor Systems for Health Applications

With recent advances in sensor and wireless communication technology and the development of small, low-power sensor platforms, ubiquitous sensing has become a heavily researched topic in a variety of application domains. One area of particular interest is the use of sensors for health applications to provide patients and doctors with information that was previously unavailable. Body-worn or environmental sensors can be used to enhance or automate assessments to measure patient's levels of pathology that doctors currently perform in clinical setting. Using sensors to automate performing tests that are currently done can increase accuracy and reduce doctors' workloads. Sensor data can also be used in conjunction with new tests or to enhance existing ones by providing medical experts with information that was previously unavailable. Sensor systems can also be used monitor patients continuously to provide clinicians data while the patients are going about their daily lives. One potentially rich source of new information to medical professionals is data about the patients and their environments when they are at home. Collecting this data requires systems that incorporate many sensors spread throughout the home and presents challenges related to power, size, data storage, and synchronization.

Because different health applications often rely on data from different sets of sensors, specific applications often require some level of customization in developing a sensor system for continuous monitoring. To successfully collect and analyze data and to provide useful information, these systems need to address common issues and the best solutions for each issue often depend on the specific application for which the system is being deployed. The following list identifies some of these challenges and briefly discusses the tradeoffs associated with addressing each one.

Sensing Modalities – Different health applications require measuring different physical phenomena, and often rely on data from a collection of diverse sensors. This is the main reason that custom systems are required because a single system that contains every sensing modality that might be required would be too complex to be practical. For most health applications, sensors are used to measure physiological characteristics of the subject. This usually requires some type of body worn sensor platform, such as an accelerometer to record motion, or requires recording video or audio of the subject. Body-worn sensors present less of a privacy concern to subjects, but they can be uncomfortable and annoying, especially if worn for extended periods of time. If the subject's environment is also of interest then the system needs to include sensors to monitor it. In some cases these can be placed in fixed positions in a specific area if the location of the subject is restricted to this area, or if only the data when the subject is in this area is of interest. In other cases, the environmental sensors need to also be worn on the body to capture relevant measurements as the subject moves between locations throughout their day. Systems with multiple types of sensors often need to use different sampling rates for different sensors, and the required sampling rate for a given sensor can also vary depending on the application.

Data Storage and Communication – Sensor systems that are deployed for several days or weeks can generate huge amounts of data which needs to be stored in the system or streamed to a remote storage. The data collected by a particular sensor node can be stored on the node, communicated to and stored in a central location at the deployment or sent to a remote cloud data store. On-node storage avoids the need to communicate data, which can introduce additional points of failure to the system, but it requires significant storage space on each node and does not allow analysis of data from multiple nodes during a deployment. Because of this it is not feasible in applications that require real-time notification. Central or remote data storage both require that sensor nodes have some way of transmitting data, and in most cases require wireless communication. This adds additional issues to the system because wireless communication is subject to limited range and unreliable performance. The choice of storing data at a central point at the deployment or on a cloud data store is dependent on whether the sensor system has the necessary computation and storage resources and if the system has internet access to send to a remote location. Sensor data can be stored at multiple points within a system for redundancy at the cost of adding extra storage at each layer of the system.

Power and Battery Life – Providing power to all of the components of a sensor system is a significant challenge because it limits where sensor can be placed, adds to the size of sensor nodes, or restricts the length of time a system can be deployed. Using sensor platforms with wall plugs provides an effectively infinite and reliable source of power, but restricts where the sensor can be placed because they need to be within range of a plug. It also adds wires that can get in the way of the subjects. Sensors with batteries can be more freely placed and can be worn on the body, but the batteries need to be periodically changed or charged for longer deployments. In the near future ultra-low power sensor nodes with energy harvesting from solar, thermo-electric, or other sources have the potential to provide benefits of battery-powered sensors without the drawbacks.

Time Synchronization – Most algorithms applied to sensor data for health applications depend on temporal synchronization between different sensor nodes. Clock drift produces error between nodes even if they are initially synchronized, and these errors can become significant without periodic resynchronization. This requires communication between nodes in the sensor system to agree on a global clock for each node to have the correct time, or requires that the data be streamed to a central node that can time stamp it based on arrival time. This second approach does not account for variations in transmission time between nodes, but the timing errors this introduces are insignificant for many applications.

Safety, Reliability, and Security – A basic requirement of any system deployed near human subjects is that it does not pose a significant risk to their health or well-being. Systems that provide notifications or interventions on which the subject's health depends need to reliably deliver these notifications and interventions. A system that collects health information needs to be secure to maintain the confidentiality of the subject. The exact requirements for safety, reliability, and security depend on the application, but they should be considered in the design and testing of any sensor system.

In-Home Deployment – Any sensor system that is deployed many different environments, such subjects' homes, needs to be easy to install and remove, and needs to be configurable to work in homes with varying floor plans. Many sensor systems are designed to allow nodes to be added or removed without affecting the performance of the system, which allows for deployments in diverse homes. Because installing and removing the systems requires visiting the subjects' homes it can be a significant disruption to them, so the installation and removal process should be as fast as possible and not damage the home in any way. Systems that require on-site calibration or collecting training data for machine learning algorithms at installation time can add significantly to the disruption caused by installation, so these requirements should be minimized and the installation procedure should be optimized to run quickly.

B. Related Work

Detecting Agitation

Using sensor systems to detect agitation in people with dementia has the potential for relieving burden on caregivers and clinicians, but it presents a number of challenges. In clinical setting, the Cohen-Mansfield Agitation Inventory (CMAI), or other similar behavioral questionnaires are generally used as the standard for evaluating the extent of agitation in people with dementia, and attempts to use sensors to measure agitation usually either try to find sensor data that correlates with CMAI scores or to detect specific behaviors that are part of the CMAI. The CMAI comes in several forms that ask slightly different sets of question, but all of the versions ask the caregiver to recall the frequency of behaviors over the preceding two weeks. The behaviors fall into four categories: physically aggressive, physically nonaggressive, verbally aggressive, and verbally nonaggressive. Examples of each of these categories are hitting, hoarding objects, complaining, and yelling respectively. The choices for frequency for each behavior are:

- **Never**
- **Less than once a week**
- **Once or twice a week**
- **Several times a week**
- **Once or twice a day**
- **Several times a day**
- **Several times an hour**

Each question is assigned a score from 1 (Never) to 7 (Several times an hour) based on its frequency of occurrence. The scores can either be left as a measure of specific behaviors, or combined into a measure of agitation in each of the three categories. The CMAI relies on the recall of caregiver of the PWD about the PWD's behavior over the previous two weeks, which can potentially make the results unreliable especially if the caregivers are family members who are not trained in assessing agitation. The CMAI also focuses on behaviors that are more common among people with advanced dementia and may be less applicable to the subjects of BESI with more mild Alzheimer's who are still living at home and exhibit lower levels and frequency of agitation [6].

The authors in [7] present a framework for detecting agitation using video cameras to automate detection of specific physically agitated behavior in an assisted-care facility. This approach is validated against a clinical assessment that uses real-time observations made by nurses in institutions that care for people with dementia For this application privacy is not a large concern because the expectation is that the people with dementia will be visually monitored by the nursing staff regardless of the presence of cameras; however, for dementia patients living at home, privacy is a major concern that generally precludes the use of cameras or microphones.

In [17], the authors showed that actigraphic measures from a wrist-worn sensor correlate with CMAI scores. This method is similar to what is used in the BESI system, but they only compared correlations between CMAI scores and the level of activity over a day. The authors did not examine the value of actigraphic data over smaller periods of time to detect specific instances of agitation, which is required for the BESI system.

The pilot study that was the precursor to the BESI project used accelerometers worn on the wrists, ankles, and waist of a PWD to assess agitation [8]. The study showed some correlation between CMAI scores and motion as measured using the Teager energy of the accelerometer data. The study showed a stronger relationship between specific instances of agitation and the Teager energy during those times, which suggests that an accelerometer can be used to detect physical agitation in people with dementia. The accelerometer located on the dominant wrist of the subject produced data most strongly correlated with agitation. The sample size was only six patients, and they were selected based on a high risk of agitated behavior, so it is not clear how well the accelerometer will work in detecting agitation in a population that is less prone to agitation, like the participants in BESI.

Some researchers have argued that using body-worn sensors is too much of a burden on people with dementia and that systems with just environmental sensors will be more widely accepted. One such system is the KinVocal, which uses a Microsoft Kinect to detect verbal agitation [9]. It uses a combination of speech processing and speech-to-text to provide features for machine learning algorithms that detect instances of the verbal agitation behaviors asked about in the CMAI, such as cursing, negativism, and repeated calls for help. While the results of KinVocal are promising, recording raw audio data presents a significant privacy concern that may limit its acceptance by potential participants, and it only addresses verbal agitation.

Home monitoring Sensor Systems

Pervasive, long-term home sensing systems can provide beneficial information in a number of applications from home security to smoke alarms to systems that assist in the care of children or the elderly. These systems continuously collect data from a variety of sensor, with the specific sensing modalities dependent on the application, spread throughout the house and analyze the data for anomalies that require automated or human intervention. Someone in the home or a third party expert can then be notified to take appropriate action in the event that the system detects something requiring human interaction. Designing these systems presents a number of significant engineering challenges that need to be addressed to collect, synchronize, and analyze data from many disparate sensing modalities.

Several universities have launched labs in part to examine the challenges and opportunities of home monitoring systems. The Georgia Tech Aware Home is a 5,040 sq. ft. home that is used for designing and testing research systems for in-home deployment [10]. The research focuses are on sustainability, infrastructure innovations, digital content delivery and health and well-being. The MIT PlaceLab is a one bedroom condominium equipped with a variety of sensors that monitor the activities of the inhabitants and the environment [11]. It is used to research how the environment in the home affects people's behavior. These facilities are used for a wide range of research projects, but the ones most closely related to the BESI system are for the assistance of the elderly to live independently. They usually focus on allowing the elderly to live alone using things like robotic assistance. The BESI project is aimed at a group of elderly that are no longer able to live alone with the aim of providing better information to their caregivers.

Empath is a home monitoring system that uses several sensing modalities to monitor people at home who suffer from depression. The system collects data from a variety of wireless sensors, such as microphones, motion sensors, and bed-mounted accelerometers [12]. The system includes support for sensor platforms that use different wireless protocols, but all of the sensors need to be within wireless communication range of a base station laptop. Empath2 generalizes the Empath system to provide a flexible framework that can be used to form sensor systems for varying applications using commercial wireless sensors [13]. The system includes data collection and transmission to a cloud server with automatic detection of some errors in the system, such as a loss of internet access for the in-home base station. The Empath2 system relies on each sensor having a wireless communication link to a base station in the house. This is not the case for the wearable accelerometer platforms used for agitation detection, because currently available devices with suitable form factors and times between recharge have limited communication ranges. Thus the Empath2 system would require significant modifications to be used in a system like the BESI sensor system.

A common application of home sensing systems for the elderly is in measuring activities of daily life (ADL) to allow people with declining physical and cognitive health to live alone with a system that detects events that require help from others. In [14], the authors use many small state change sensors taped to various appliances to collect data for activity detection. The sensors were battery powered and did not have any form of communication, which precludes any real-time processing and notification. The data was collected from the individual sensors and analyzed post-deployment. The authors of [15] and [16] both present sensor systems for monitoring the ADL in the elderly using ZigBee-based wireless sensor systems. Unlike Bluetooth, ZigBee has a transmission range long enough to cover an entire house, but it is not included in most smartphones so it is not widely used. The system in [15] uses only environmental sensors, such as contact sensors on doors, monitors on appliances, and light, temperature, and humidity sensors. In [16] the system uses three body-worn accelerometers and PIR motion sensors to detect anomalies in the behavior of elderly people, such as lying down in the bathroom or not eating at mealtime. Using a ZigBee-based accelerometer platform has the advantage of eliminating the need for multiple access points because of its range, but ZigBee achieves this increased range with greater transmission power, which limits the battery life of ZigBee devices and makes them impractical for continuous sensing and streaming applications like BESI.

CareWatch is a system designed to assist the caregivers of people with dementia by monitoring their location. It focuses specifically on monitoring at night to improve caregiver sleep [18]. It uses a commercial security system to detect when the PWD changes rooms, leaves the house, or gets out of bed at night. While the caregivers reported that the system was of assistance, there was no significant difference in sleep time or quality between the treatment and control groups [19]. This system is limited to only measuring occupancy and movement, and lacks the necessary sensing modalities for the BESI application.

III. Data Collection and Synchronization

This section discusses the distributed system architecture required to collect sensor data from multiple rooms and the communication between the system components to stream the data to a central node in real time. Because the wearable sensor platform moves between rooms as the PWD moves, the system needs some way of automatically switching which access point the wearable is connected to, which is also described in this section. First each of the components in the sensor system is described along with the communication links between components. Next is a description of how the components are used to form a system to collect data from multiple rooms and to communicate it to a central location. The procedure for sampling the body-worn sensor after it connects to a communication hub within a room is essentially the same as the procedure for sampling the environmental sensors in that room, but the need to hand off the wireless connection as the PWD moves throughout the house makes interfacing with the wearable unique among the sensors in the system. The procedure for performing these handoffs is presented in this section. Finally, the process for determining timestamps that are relative to the same global time for all of the sensor data is described.

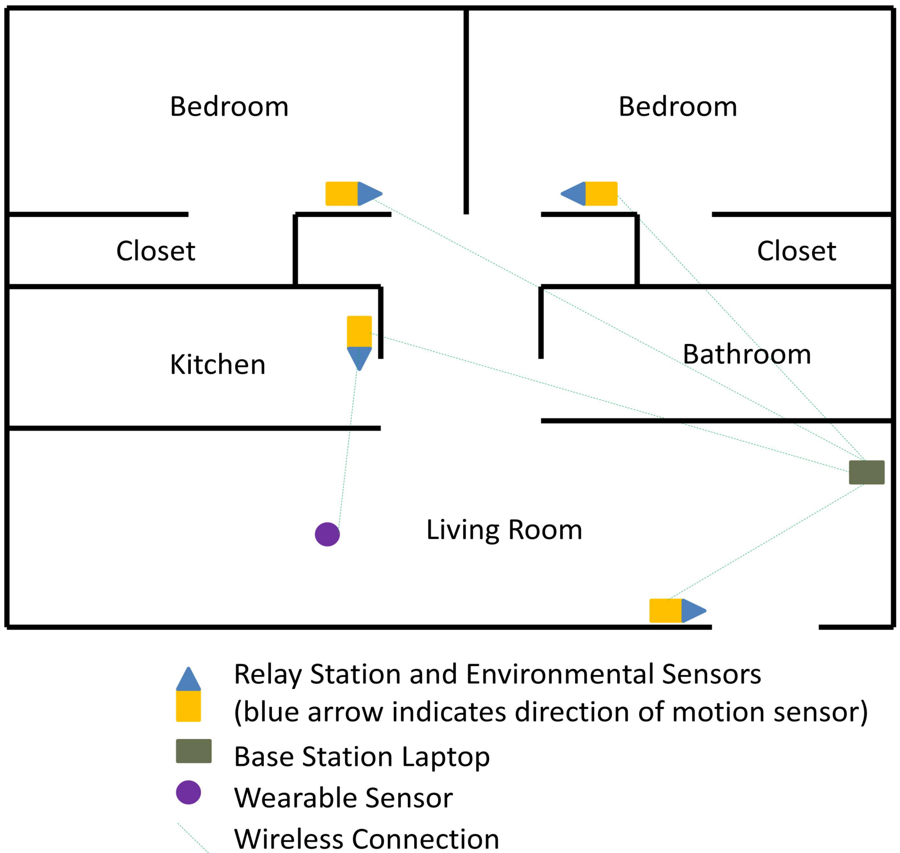

Figure 2. A typical deployment of the BESI sensor system in a single family home. Each relay station collects sensor data from a single room and streams it to the base station.

The BESI system consists of three types of hardware components: sensors, relay stations, and a base station. Each room contains a set of the environmental sensors, and a single relay station. The environmental sensors used are a microphone, ambient light sensor, temperature sensor, and motion sensor. The relay station samples each of the environmental sensors at periodic intervals, and transmits the sensor data to a base station within the home. The data from the inertial wearable sensor platform on the PWD is sent to one of the relay stations within its transmission range. The data is then either stored on the base station or transmitted to a cloud data storage. Figure 2 shows the setup for a typical deployment of the BESI system. One relay station is placed in each room on the house with the exception of closets, storage spaces and bathrooms, which are not instrumented because they are rarely occupied for extended periods of time. The base station is placed near an Ethernet port, and the wearable sensor is placed o the subject.

A. Sensor System Components

Shimmer3

The Shimmer3 is a wearable inertial sensor platform from Shimmer Research Inc. that contains an accelerometer, gyroscope, magnetometer, and an analog-to-digital converter (ADC) expansion port built around an MSP430 microcontroller [20]. Figure 3 shows the Shimmer and the strap used to attach it to the wrist. The Shimmer is slightly bigger than a typical watch, and a watch strap is used to allow the Shimmer to be comfortably worn. The hooks that come with shimmer to attach it to the straps do not reliably hold the Shimmer in place, so tape is used to secure the Shimmer to the hooks. The Shimmer has a Bluetooth 2 radio to stream data and an SD-card for on-node storage. The firmware for the Shimmer is open source and highly customizable. For the BESI system, the Shimmer is used to stream accelerometer data, so the firmware used is a modified version of the default BtStream v0.7.0 from Shimmer Research with only the 3-axis low noise accelerometer turned on. The accelerometer is sampled at 256 Hz with 12 bits used to represent the value for each axis. The raw values reported by the Shimmer are integers in the range 0-4095 for each axis that can be converted into acceleration using offset and sensitivity values obtained from calibration. The Shimmer provides a timestamp with each data sample, which is the number of ticks since the Shimmer began streaming of a rolling 16-bit counter. A single data packet with the timestamp, each axis, and four bits of padding is 72 bits, so the total Bluetooth bandwidth used when sampled at 256 Hz is 18.4 kbps.

Figure 3. Shimmer3 wrist worn accelerometer platform with the wrist strap. The tape holds the Shimmer in place on the band and covers the LEDs.

When the Shimmer is first turned on, it begins broadcasting its Bluetooth discovery message until another device connects to it. After connecting, the Shimmer accepts a set of commands, one of which is the start streaming command. When the Shimmer receives this command, it begins sampling its sensors and streaming the data to the connected device. The Shimmer runs off of a rechargeable 450 mAh battery. The battery life of the Shimmer was determined by running several fully charged Shimmers until they shut off due to low battery with the configuration used for the BESI system (streaming only the low noise accelerometer sampled at 256 Hz). The battery life of the Shimmer with these settings is about 19 hours. For multi-day deployments, the PWD will wear one Shimmer on their wrist during the day while another charges and swap them at night.

Microphone

The microphone used in the BESI sensor system is an omni-directional foil Electret microphone [21]. This microphone is combined with a 100x gain amplifier on a breakout board available from Sparkfun [22].

The microphone has a sensitivity of -44 dBV, which our testing indicates is more than enough to cover a typical room. The breakout board is powered using a 1.8 Vdc supply, and the output is an analog signal with maximum amplitude of 0.9 V centered around an offset of 0.9 V.

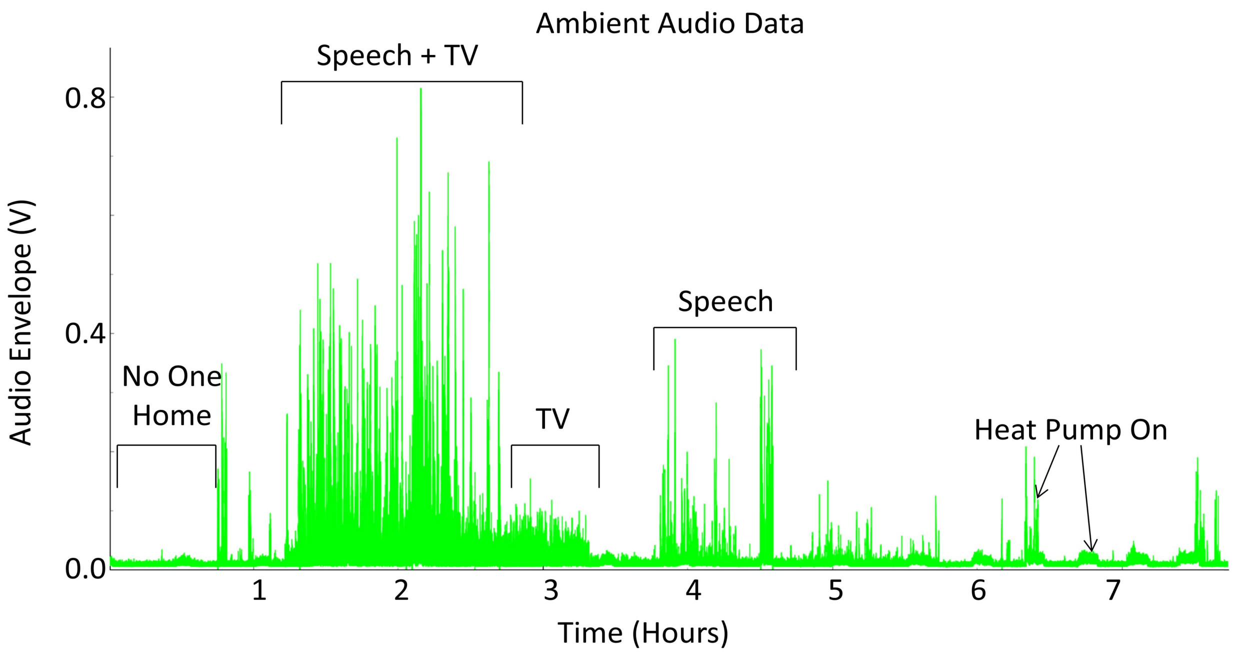

Figure 4. An example of audio data captured by the BESI sensor system. This data was recorded in a researcher's home starting at approximately 4pm and running to 12am.

Because of the higher sampling rate required to accurately capture audio, the data from the microphone is preprocessed on the relay station to reduce the bandwidth requirements for transmission to the base station. The microphone is connected to one of the ADC channels on the relay station microcontroller and sampled at 10 kHz with 12 bits of representation. Human speech is generally contained in the frequencies below 4 kHz, so sampling at 10 kHz is sufficient to accurately capture most of the audio signal of interest. To eliminate very low frequency vibrations from sources like people walking that are outside of the audible range but tend to dominate the signal produced by the microphone, the signal is high-pass filtered after sampling using the first order infinite-impulse response (IIR) digital filter specified by the following difference equation

$$y[n] = 0.9y[n-1] + x[n] - x[n-1]$$

where $y[n]$ and $x[n]$ are the nth samples of the filter output and sampled audio data respectively. The -3 dB cutoff frequency for this filter is approximately 165 Hz for a sampling rate of 10 kHz. The output of the microphone is an analog signal centered around 0.9 V, so the signal is normalized to be centered around 0 V by subtracting 0.9 from it before it is filtered. The filter output is rectified, low-pass filtered using a 100-point moving average, and down sampled to 100 Hz. This produces an envelope of the audio signal sampled at 100 Hz, which is then sent to the base station. Figure 4 shows what this envelope data looks like for a test deployment done in a researcher's home. The events marked in figure 4 are estimations based on the researcher's recollection, but speech typically has higher peaks and more variance compared to artificial noise sources like the TV.

Temperature Sensor

The temperature sensor used is the Texas Instruments LM60 integrated-circuit temperature sensor [23]. It has an accuracy of ±2°C at room temperature, which is enough for this application because errors tend to be a constant offset for a given sensor, and we are more interested in changes in temperature rather than absolute values. The sensor works by using transistors whose characteristics change based on the ambient temperature. These transistors are used in an amplifier that has a gain that depends on temperature to produce an output voltage that changes linearly with changing temperature. The LM60 produces an analog voltage proportional to the ambient temperature using the equation

$$T = \frac{(Vmv - 424)}{6.25}$$

where T is the temperature in °C and Vmv is the output of the temperature sensor in millivolts. The LM60 is powered using a 3.3 Vdc supply from the relay station, and at room temperature it produces an output of about 600 mV. The temperature sensor is connected to one channel of the relay station's ADC and is sampled at 1 Hz. Figure 5 shows 36 hours of temperature data measured in the lab. From this data it is possible to determine when the air conditioning turns on and off in the building and to see the day-night climate control pattern.

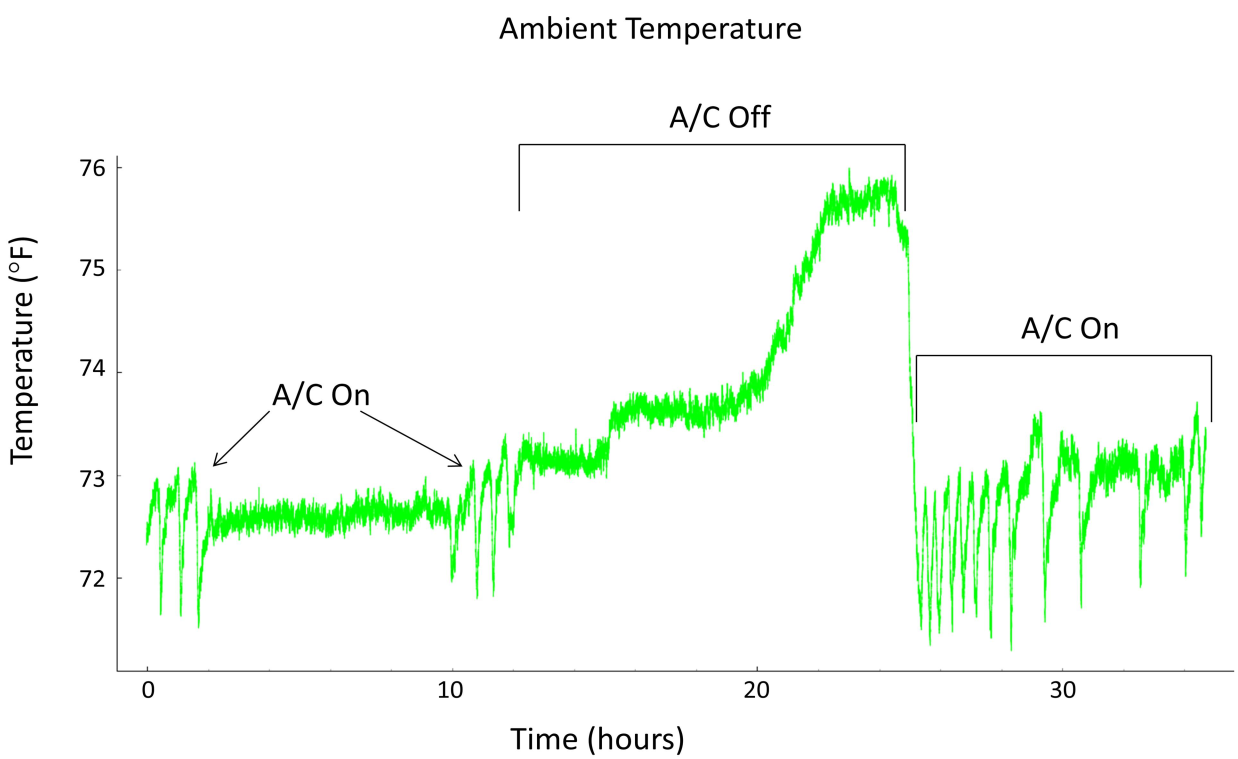

Figure 5. Temperature data in a research lab as measured using the BESI sensor system. The data begins at approximately 7am and ends around 7pm the next day.

Light Sensor

The ambient light sensor is the TAOS TSL2561 luminosity sensor [24] on a breakout board available from Adafruit [25]. The sensor measure ambient light using two photosensitive diodes, one that responds to the full spectrum of light (including infrared) and one that only responds to infrared. These diodes are both incorporated into circuits that produce voltages that vary with the amount of ambient light and the output of each is sampled by an ADC on the breakout board. The voltage measured from each diode is transmitted using an Inter-Integrated Circuit (I2C) bus to the relay station where the lux (lumens per square meter) is calculated with an equation that removes the effect of infrared light on the full spectrum diode to produce an output that matches the visible light spectrum [24]. The relay station samples the light sensor at 1 Hz. Figure 6 shows light data from five day-night cycles measured in three rooms of a researcher's home. During the day the light level measured is due almost entirely to the sun with large variation between relay stations based on their proximity to windows. After sunset the relay stations can detect if the lights in a room are turned on or off.

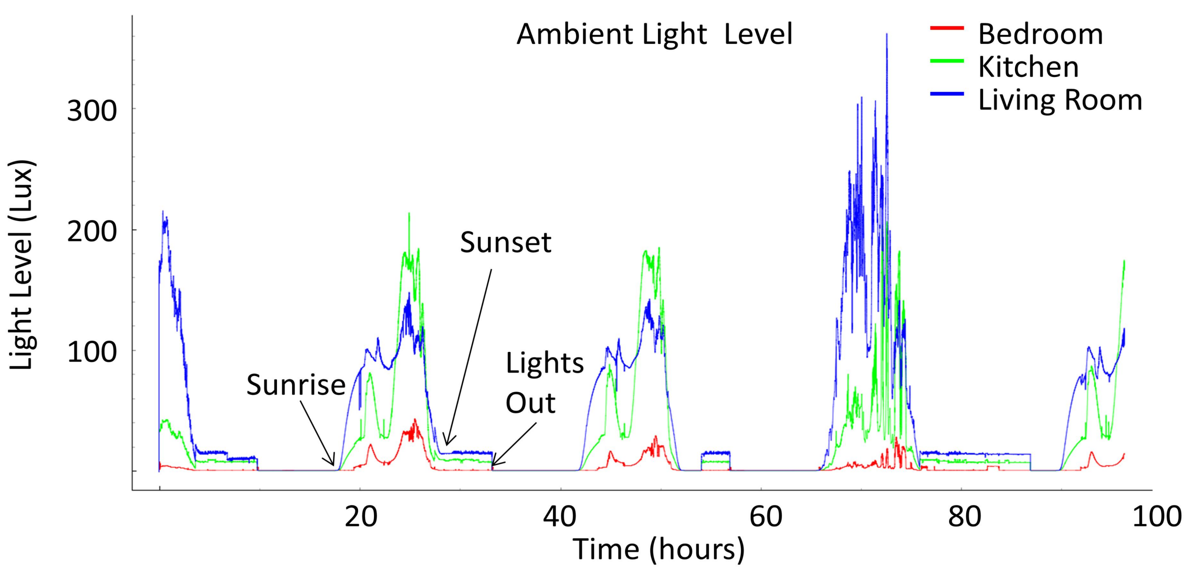

Figure 6. Ambient light data from three rooms during a five day deployment in a researcher's home. During the day, sunlight produces a higher lux level with high variance compared to artificial light in the evenings.

Motion Sensor

The goal of the motion sensor in BESI is to estimate the level of social interaction of the PWD. As a proxy for social interaction, we track the number of people in the room with the PWD. To achieve this, the motion sensors used have two passive infrared (PIR) motion sensors with narrow, slightly offset fields of view to measure the direction of movement. Figure 7 shows the door motion sensor case and board. When someone moves in front of the sensors their output voltages spike based on the change in heat in their field of view. Because the apertures in the case are offset, one sensor is triggered before the other, and the order in which they are triggered indicates the direction that the person is moving. Figure 8 shows the door sensor output when one person repeatedly walks back and forth past the sensor set up facing across a door.

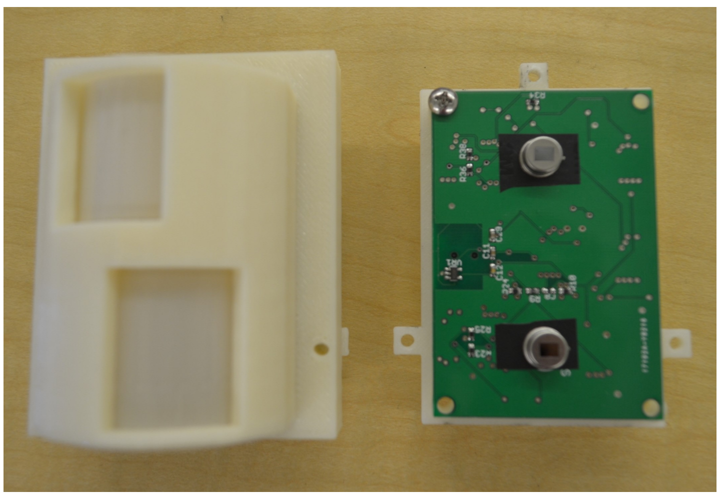

Figure 7. Door motion sensor cover and board. The offset apertures in the case lead to a different channel sensing motion first based on the direction of the motion.

In a deployment, the door sensors are set up to monitor the movement in and out of each room. Each door sensor will be positioned next to a door pointing parallel to the plane of the door as shown in figure 2. Based on the number of entries into and exits out of a room as determined from the door sensor data, the occupancy of each room in the house can be determined at any given time as long the initial occupancy of each room is known.

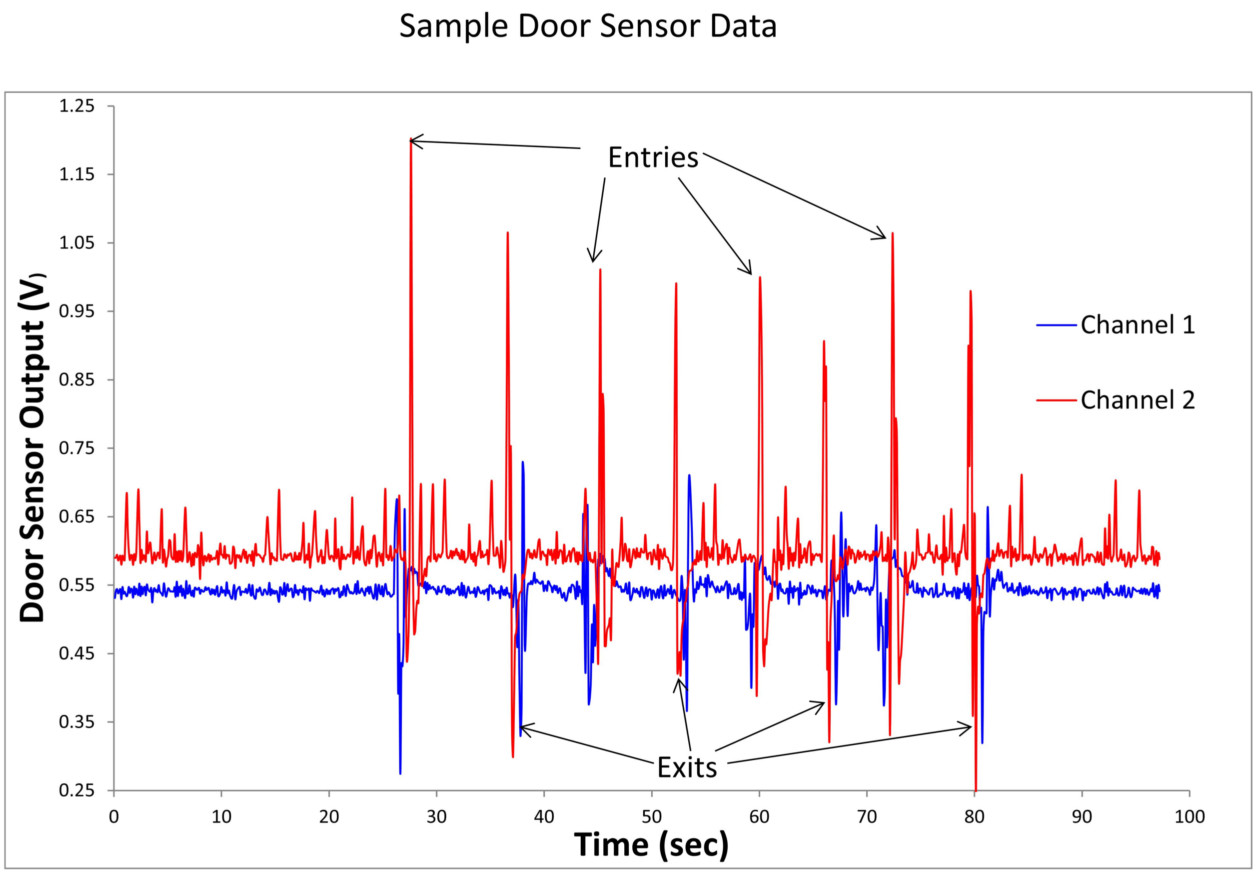

Figure 8. Door motion data from a single person repeatedly walking in and out of a room. The direction is determined by which channel spikes first.

Relay Station

The relay stations collect the data from the sensors and transmit it to the base station with one relay station per room. At the core of each relay station is a BeagleBone Black computer board that provides processing power and storage and interfaces to the various sensors [26]. The BeagleBone is a low-cost ($55) hobbyist board that has a 1 GHz ARM® Cortex-A8 processor, 512MB of RAM, 3D graphics and NEON floating point accelerators, and 4GB of flash memory. It also has pin headers that allow easy access to the BeagleBone's I/O ports, such as GPIO, ADC, I2C, and SPI [26], and in addition to the pin headers, the BeagleBone has Ethernet, USB and HDMI ports to support higher-level communication interfaces. The BeagleBones in the relay stations are running Debian, an embedded Linux-based operating system. The relay stations are powered using 5Vdc 2A wall adaptors, and the BeagleBone provides, 5Vdc, 3.3Vdc, and 1.8Vdc pins to power external components.

The relay stations use wired and wireless connections to communicate with the sensors and base station. The temperature sensor, microphone and the two channels of the motion sensor are each connected to one channel of the BeagleBone's ADC. The light sensor is connected to one of the I2C interfaces on the BeagleBone. To provide wireless communication, the BeagleBone has an 802.11bgn Wi-Fi USB dongle [27] and a CSR Bluetooth 4.0 USB dongle [28]. The Bluetooth is used to communicate with the Shimmer wearable sensor, and the Wi-Fi is used to send the sensor data to the base station. The wearable is the only sensor that is not physically located near the relay station.

The relay stations all run an application that samples the sensors and streams the data to the base station. All of the code running on the relay stations is written in Python 2.7 except for the code to interface with the ADC, which is written in C to achieve the sampling rates required for the microphone. The wireless connection between the relay station and base station for streaming data uses the Python Socket library to stream, which creates bi-directional TCP/IP connections between remote devices [29].

After connecting to the base station, the relay station program spawns threads to handle interfacing with the light sensor, the ADC, and the Shimmer and sending the data to the base station. Each thread is passed a reference to the socket used for its sensor which is used to send data to the base station. The data from each sensor is transmitted to the base station once every second.

The light sensor uses the Adafruit I2C library to read the light value measured by each of the two photodiodes and calculates the ambient light level in Lux using the formula described in the light sensor section above and detailed in [24]. The sensor is sampled at 1 Hz, and the exact timestamps are determined using the clock on the BeagleBone to measure the time between the start of data collection and the current time. Every second the most recent lux measurement and timestamp is transmitted to the base station.

All of the sensors that use the ADC: the temperature sensor, microphone, and the motion sensor are controlled using a single thread because multiple threads cannot concurrently access the ADC. The python ADC code runs a C program using the Python Subprocess module that reads one second of data from each sensor, timestamps the data, and sends it to the python program using a pipe. The temperature sensor is sampled at 1 Hz, each motion sensor channel is sampled at 1 kHz, and the microphone is sampled at 10 kHz. The microphone is high-pass filtered and the envelope of the signal is sampled at 100 Hz as described in the microphone hardware section. The envelope of the motion sensor is sampled at 10 Hz. The C code calculates the timestamp for each sample relative to the start of its execution (which resets about every second), and the python code converts this into a time relative to the start of the deployment. The new ADC data is transmitted to the base station every second using three sockets, one for each sensor.

The data that the accelerometer sends to the relay station when the Shimmer is connected is sampled using the Bluetooth connection. The Bluetooth connection and data transfer is done using the Python Lightblue Bluetooth library. The data is sent by the Shimmer as byte data with 2 bytes for each accelerometer channel, 2 bytes for the timestamp and 1 byte of padding for a total of 9 bytes per sample. On the relay station the accelerometer data is parsed into the timestamps and the measurements for each of the three axes for each packet. The byte of padding between each data packet is used to delineate consecutive packets. Every second the data from the Shimmer is transmitted to base station.

The BeagleBone, environmental sensors, and USB dongles are contained in a custom 3D printed case to protect the subjects from the system and vice versa. Figure 9 shows a relay station in its case. The case consists of a rectangular box with a removable lid to hold everything except the motion sensor. The lid has holes to allow the temperature sensor, light sensor, and microphone access to the environment. One side of the box is open to allow easy access to the BeagleBone power adaptor port, power and reset buttons and the USB ports. The motion sensor is contained in a separate 3D printed case that hangs off of one side of the box. It has two offset apertures that each constrains the field of view of one of the PIR sensors. The case has hooks on one side that allow it to be hung from temporary hooks on the wall, or it can be placed on a small table several feet from the ground. We have two versions of the case that are mirror images of each other, which allows it to be hung with the door sensor facing either right or left to allow us to be as flexible as possible in selecting locations while deploying the system.

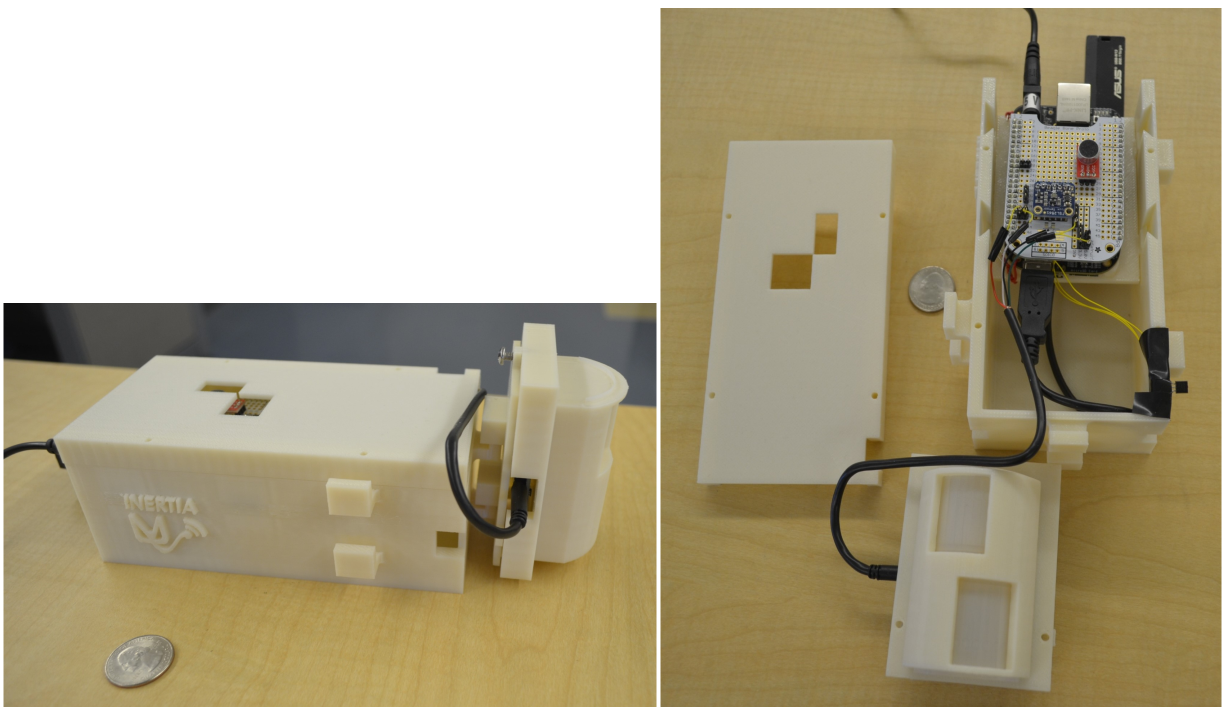

Figure 9. Side view of a relay station in its case (left). Top view of a relay station with the top of the case removed (right).

Base Station

The base station is a laptop with a Wi-Fi adaptor and an Ethernet wired connection. For all of the tests and deployments, we use a Lenovo ThinkPad laptop running Microsoft Windows 7. It has an Intel® i3 2.3 GHZ processor, 4 GB of RAM, and a 300 GB hard drive. The base station uses the Wi-Fi adaptor to connect a local area network (LAN) to communicate with the relay stations, and the Ethernet is used to connect to the internet. The LAN is established by a router without an internet connection that the base station and the relay stations connect to. The router currently used is a TP-LINK TL-WR41HP, a high power 802.11n wireless router. The router is configured with static IP addresses for each of the relay stations and the base station to allow communication between components using these IP addresses.

Al of the code for BESI running on the base station is written in python 2.7, and it also uses the Python Socket library to communicate over Wi-Fi. The base station stores all of the raw sensor data it receives from the relay stations in text files and plots the data as its received using the Python library PyQtGraph [30].

B. System Architecture

This section describes how the sensor system components are organized to form a hierarchical structure that collects environmental data from multiple rooms and aggregates it in a single location and how this addresses challenge 1 from the list of challenges in the introduction. Each room contains a relay station that samples each of the sensors in that room, and the relay station transmits this data to the base station. The Shimmer communicates with one of the relay stations within Bluetooth range, which is usually the relay station in the same room as the PWD. Figure 10 shows the structure of the sensor system and the communication links between components.

At the beginning of a deployment when the base station and relay station applications are first run, each relay station attempts to establish a Wi-Fi connection to the base station. After a connection is established, the relay station begins sampling each sensor using the hardware interfaces described in the previous section. For all of the sensors, the relay stations collect data, determine the timestamps, and send it to the base station. The base station stores the data in temporary files and plots the last few seconds of data for each sensor for real-time monitoring and debugging. Periodically, the base station processes the temporary files to synchronize the timestamps and saves the results in new files that are analyzed post-deployment.

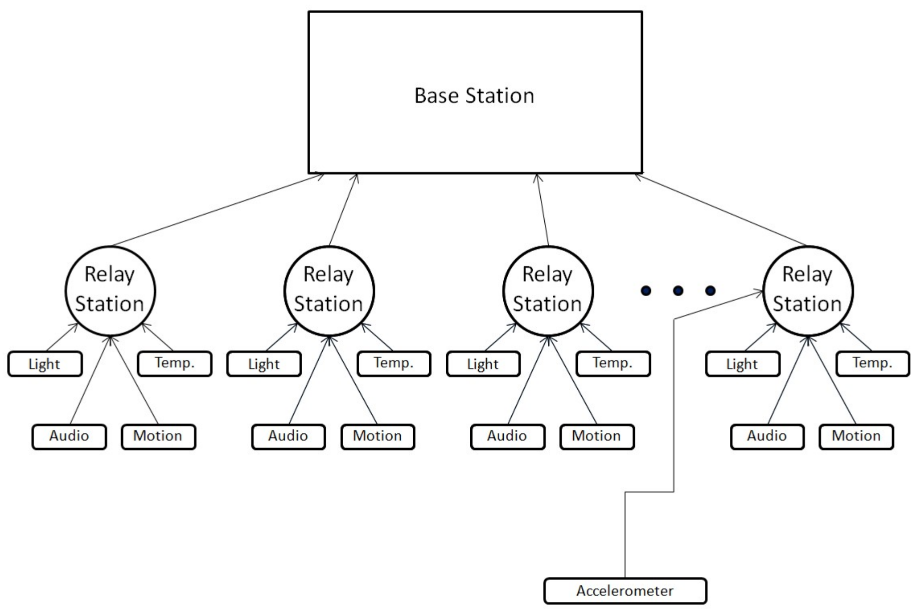

Figure 10. Components of the BESI sensor system and communication links between components. For the environmental sensors (light, temperature, audio, and motion) the communication is wired, and for the other links it is wireless. The arrows represent the flow of sensor data through the system with the direction of the arrow indicating the sender and receiver.

Relay Station Application

The relay station attempts to connect to connect to the base station over Wi-Fi until successful, and then samples and transmits data from each sensor. After connecting, the relay station application spawns one thread for each sensor that reads and transmits the data from that sensor. If the connection to the base station is lost, the relay station attempts to reconnect to the base station, and begins transmitting data again when successful. No data is read during the period when the connection to the base station is lost. Figure 11 shows the state machine for each relay station. The details on the connection to the Shimmer are discussed in **Connecting to the Body-Worn Sensor** below.

.

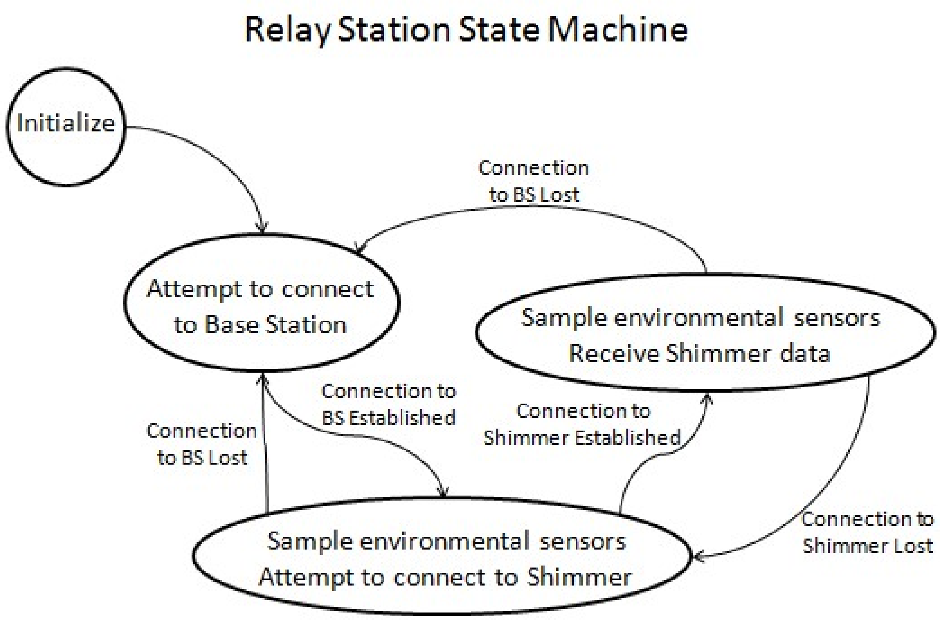

Figure 11. State machine for the firmware running on each relay station. BS refers to the base station.

The Wi-Fi connection between the relay stations and the base station is initially used to send configuration information to the relay stations and then used to send sensor data from the relay stations to the base station. The base station starts listening on a specific network port, and the relay station creates a socket connection using the IPv4 address of the base station and the port number. The relay stations and base station are on the same local network, which is required to connect. When a relay station first establishes a connection to the base station, the base station sends the Bluetooth addresses of the Shimmers to which the relay stations will try to connect. The relay stations then create a separate connection to send the data from each sensor to the base station. The port numbers used for each relay station to connect to the base station are static and are set prior to deployment on the base station and relay stations. The lowest port number for each relay station is used as a unique identifier for that relay station.

Because typical deployments will use two Shimmers to allow one to charge while the other is worn, the relay station will alternate trying to connect to two Bluetooth addresses, which are initially sent from the base station. While not connected to a Shimmer, the relay station attempts to connect to one of the Shimmers every five seconds. If the connection is successful, the relay station sends the command to start streaming and begins receiving the accelerometer data. If the connection between the relay station and Shimmer is lost (usually due to the subject moving out of the Bluetooth range), the relay station will go back to trying to connect to either Shimmer, and will begin sampling accelerometer data again when a new connection is established.

Base Station Application

When the base station application starts, it waits until the relay stations establish connections, then stores and plots the sensor data as it is received from the relay. The base station begins by starting a new process for each relay station to handle the data streaming and processing. This is done to allow a separate window for plotting the data from each relay station and to make the application easily scalable to setups that require different numbers of relay stations. Each process begins listening for a connection on the port corresponding to the ID of its relay station. After a connection is established, each base station process sends the Bluetooth addresses of the Shimmers used for that deployment to its relay stations and begins listening at the ports that correspond to each individual sensor on its relay station. After the connections to the relay stations have been established, the base station processes check every 50 milliseconds for data from any of the sensors on their relay stations, and if new data is available, it is saved to a temporary file without any processing. New data is also parsed into sensor readings, and added to a list for plotting. These plots show the data from the last few minutes and are used for quickly assessing the state of the system. When the base station receives the first piece of data from a sensor, the time is marked in the file to indicate the start of the data collection, and the timestamps for all of the data in that file is relative to this start time. Separate files are used for each sensor on each relay station. The data transmitted from the relay stations from the light and temperature sensors is sent and stored as strings, while the data for the microphones, door sensors and accelerometers is sent and stored as byte data. Sending the data as strings has the benefit of producing human-readable files and requiring less processing, but the bandwidth required is about double that of byte data. Because of this, strings are used for the low sampling rate sensors, while bytes are used for the data from the sensors with higher sampling rates. Figure 12 shows the state machine for each process on the base station.

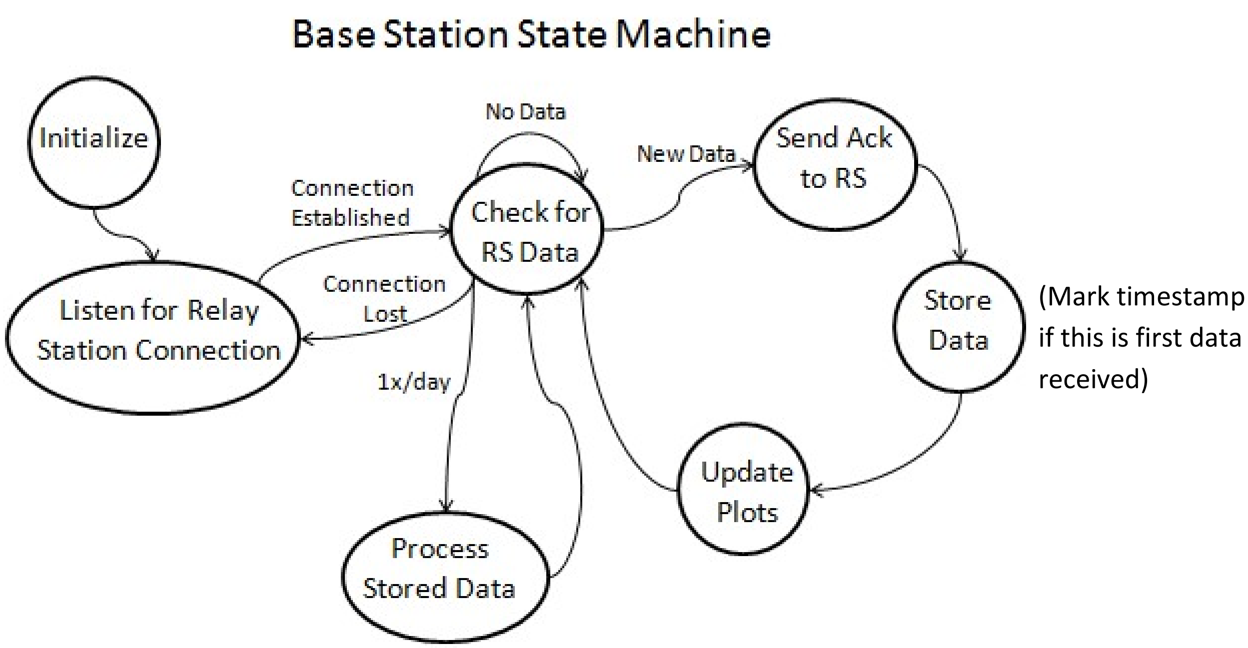

Figure 12. State machine for the software running on the base station. This diagram represents the interactions between the base station and a single relay station and is duplicated for each relay station.

The data is plotted to enable us to easily see if all of the relay stations are connected to the base station and successfully collecting data during a deployment. When the base station receives a packet of one second of sensor data from a relay station, it parses the packet to determine the value of each sensor reading and adds these to a first-in first-out list for each sensor. The lists for the microphone and door sensor have 1000 items, and the lists for the accelerometer, light sensor, and temperature sensor have 200 elements. The elements of each list are plotted versus their index, and the graphs are displayed on the base station. The plots are updated whenever new data is received. The values plotted are: raw ADC values for the accelerometer and door sensor, the envelope for the audio data, degrees Fahrenheit for temperature, and lux for the light sensor. Figure 13 shows an example of the plot window for one relay station.

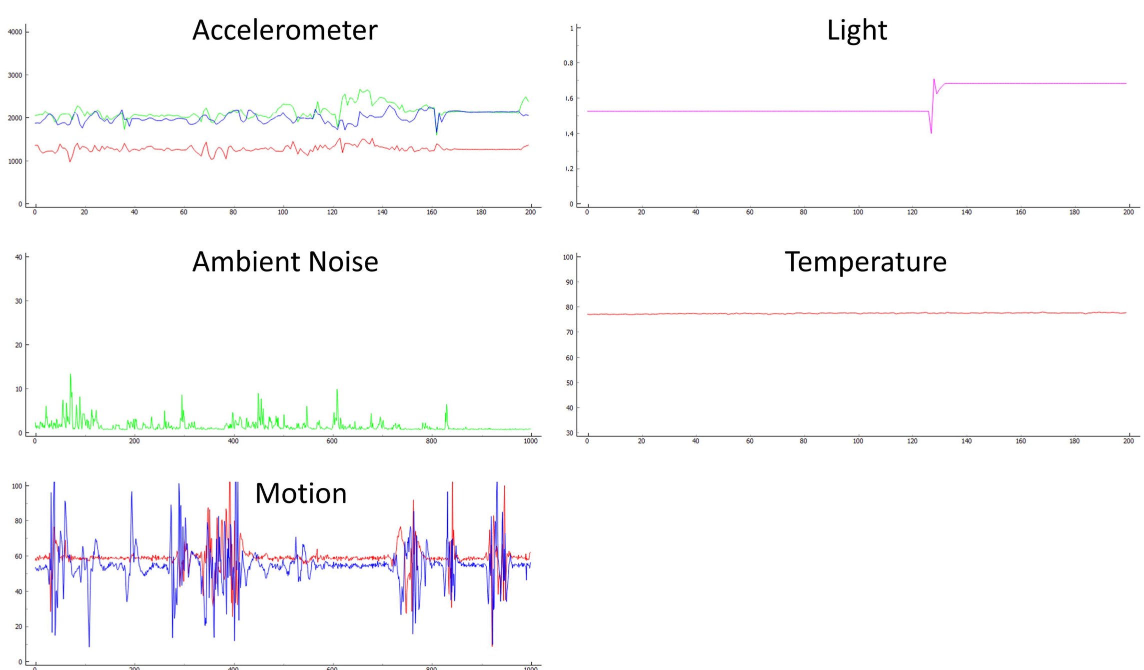

Figure 3. Data plots displayed on the base station showing the most recent data received from a single relay station. These plots are used for remotely monitoring the sensor system during a deployment.

At fixed intervals (every 24 hours for our deployments) the data in the temporary files is processed and stored in permanent files. Each permanent file starts with metadata that lists the start time, deployment and relay station numbers, and the sensor from which the data comes. The format for all the data in the permanent files is <timestamp in seconds>, <sensor value>, with the timestamps relative to the start time at the beginning of the file. The temperature and light data stored in the temporary files is already in this format, so the processing is simply copying the data to the new files. For the door sensor and the microphone, the data is in the appropriate format, but is stored as byte data, so the processing for the data from these sensors involves unpacking the sensor readings and timestamps into strings and writing them to the permanent file. The accelerometer data also has to be processed in this way, but time timestamps need additional processing as described in the **Time Synchronization** section below. If the base station loses connection to a relay station (the mechanism for detecting this is described in the **System Reliability** section below), the temporary files for each sensor are processed before the relay station reconnects, so that they can be overwritten without losing data when the relay station does reconnect.

C. Connecting to the Body-Worn Sensor

The Shimmer is used to collect accelerometer data from the dominant wrist of the PWD and stream the data in real time to a relay station using Bluetooth. The Shimmer needs to connect to several relay stations as the PWD moves through the house because the range of Bluetooth is limited. This section describes how this is done and addresses challenge 2.

Bluetooth is designed to be a one-to-one connection protocol with one connected device acting as the master and the other as the slave. A master device can be connected to seven slave devices simultaneously, but a slave device can only be connected to one master. When not connected a slave device transmits a discovery message that contains its Bluetooth address and the Bluetooth services it supports. This message can be received my any other Bluetooth device within range, and it is the only one-to-many broadcast message in the Bluetooth protocol. A master device can attempt to connect to the slave after receiving its address from the discovery message, or it can attempt to connect without first receiving the discovery message if the address is already known. To maintain the security and privacy of connections, two devices must be paired before they can establish a connection. Pairing usually requires entering the same PIN on both devices. Bluetooth is widely used in wearable sensors, like the Shimmer, because most smartphones and laptops contain a built-in Bluetooth radio, and Bluetooth is significantly lower power than Wi-Fi, which is the other ubiquitous wireless protocol. Battery-powered Bluetooth devices are typically intended for a range of 10 meters [34], although the range is highly dependent on the environment. Obstacles or interference from other wireless devices in the same frequency band can significantly affect the operational range. Tests with Shimmer indicate that it can maintain a connection with a relay station up to about 20m if the line of sight is not blocked. The range is typically 5m to 10m if walls or other obstacles are in the way.

Any home monitoring system that uses a Bluetooth device that can move throughout the house needs to address several challenges to enable constant data streaming while the device is in the house. Because the range of Bluetooth is smaller than most houses, the system needs to have multiple access points to which the wearable can connect and needs a mechanism for automatically handing off the connection as the person moves out of the range of one access point and into the range of another. Similarly, the system needs to be able to reconnect when the person returns after leaving the house. Because Bluetooth slave devices can only be connected to one master at a time, every Bluetooth device configured as a slave needs to reliably detect when the connection to the master is lost and return to the disconnected state to allow another device to connect.

Home monitoring systems like the BESI systems that use Bluetooth to transfer data can generally be configured in two ways: with the wearable as a slaves and the fixed-position access points as masters or with the wearable as a masters and the access points as slaves. Configuring the wearable as a master has the advantage of allowing it to get the discovery messages from every access point and chose the one with the strongest connection. The signal strength from each of the access points can also be used to determine the location of the wearable within the house [35]. If the wearable is the master it can also stream to multiple access points concurrently for data redundancy. The drawback of this setup is that the wearable needs to perform a discovery scan to see which access points are in range before connecting, which can take up to 20 seconds and performing scans frequently can limit the battery life of the device. This can lead to a significant latency between the wearable moving into range of an access point and it establishing a connection to that access point. Along with lower connection latency, configuring the wearable as the slave, allows each access point to connect to multiple devices, which allows multiple wearables to be monitored simultaneously. For the BESI system the Shimmer is configured as the slave because minimizing the connection latency reduces the chance of missing important data and because USB Bluetooth dongles, which are used in the relay stations, are typically only configured to be masters.

In the BESI system the handoff between access points is accomplished by having the relay stations constantly attempt to connect to the Shimmer until successful and relying on the limit on one connection in the Bluetooth protocol to ensure the Shimmer is only connected to one at a time. If the Shimmer is connected, it will simply reject additional attempts to connect. Prior to each deployment, each relay station is paired with the Shimmers used to that deployment. When the system is first turned on, both Shimmers are not connected to any device, and will accept connections from any relay station. All of the relay stations attempt to connect using the Bluetooth addresses that are sent from the base station without first doing a discovery scan. Each Shimmer will connect to one of the relay stations within range, although it is nondeterministic which relay station it will connect to. The other relay stations will continue to attempt to connect, but the attempts will be rejected while the Shimmer is connected to another relay station. When the Shimmer moves out of range of the relay station to which it is connected, it will again accept connections and one of the relay stations within range will connect to it. If the shimmer moves out of range of al of the relay stations, which happens when the PWD leaves the house, the Shimmer will disconnect and stay disconnected until the PWD comes home.

D. Time Synchronization

The data collected from different sensors on the same relay station and the data collected from different relay stations need to be synchronized to perform analysis on the temporal relationship between data streams. For the BESI project, we are interested in analyzing changes in the environment and in the PWD over minutes or hours, so synchronization to within a few seconds is sufficient. This section presents the time synchronization scheme in the BESI system and addresses challenge 3.

When the base station receives the first data packet from each sensor for a relay station, it records the time that this data is received and uses this as the start of the data collection for that sensor. This ignores the Wi-Fi transmission delay, but this is typically only milliseconds. On the relay station, this first data packet is given a timestamp of 0, and the timestamp for each subsequent data point is calculated on the relay station and sent to the base station. On the base station the time of every data point can be calculated by adding the timestamp for that sample from the relay station to the start time recorded on the base station. Whenever a relay station disconnects and reconnects this process restarts with the base station recording the start time and the relay station starting the timestamps again at 0.

The process for generating timestamps for the accelerometer data is slightly different because the data is timestamped on the Shimmer. The Shimmer timestamps data using a rolling 16-bit counter (maximum value of 65535) running off of a 32.768 kHz clock. The counter rolls over from 65535 to 0 every two seconds, and for a sampling rate of 256 Hz, consecutive samples are separated by 128 ticks. When a relay station first connects to the Shimmer and begins streaming data, the base station marks the time as the start of a data collection. The timestamps reported by the Shimmer are sent from the relay station and are converted into a timestamp in seconds on the base station before being written to the permanent file using the formula:

$$T_{sec} = 2C_{rollover} + \frac{(C_{now} - C_{start})}{32768}$$

where T_{sec} is the time in seconds since the start of data collection, $C_{rollover}$ is the number of times the counter on Shimmer has rolled over from 65535 to 0, C_{now} is the counter value for the current timestamp, and C_{start} is the value of the counter for the first sample. This approach is robust to single dropped packets, or packets with incomplete data from the Shimmer to the relay station.

IV. Deployment Considerations

The previous section describes how the sensor data is collected when the system is working as intended, but it does not address many important issues that need to be considered in real world deployments. This section addresses some of these issues. When something goes wrong in the system detecting and addressing the issue quickly is important to get as much useful data as possible, and this section begins by describing the remote monitoring and repair capabilities of the system. Next it discusses the main source of deployment problems, the wireless communication, and the methods used to develop a more reliable system. This section ends by addressing the physical appearance of the system and the things done to make it more acceptable to potential participants.

A. Remote Monitoring and Repair

In actual deployments on-site maintenance imposes a high cost because it is disruptive the dyad and traveling to the deployment site can take a significant amount of time, so the system has remote repair mechanisms for common problems. Because deployments are huge time and effort commitments both for the researchers and the subjects, we need to be able to quickly detect if data is not being collected and either fix the issue remotely or perform a home visit correctly to maximize the amount of useful information obtained from each deployment. This section addresses challenges 4 and 6.

The sensor system is designed as much as possible to be monitored and repaired remotely without any intervention from the dyad. The base station is connected to the internet via an Ethernet cable, which provides a reliable connection, and is running TeamViewer [32], a desktop sharing application that allows us to monitor and control everything running on the base station through a secure connection.

The real-time plots on the base station of the data from each of the relay stations allow a quick assessment of the current status of the system through the remote desktop access to the base station. The base station plots the data as it is received as described in the base station software section. Figure 13 shows an example of the plots for the data from one relay station. These plots are updated every second while the relay station is streaming data. From these plots we can see if the base station is receiving data from every relay station, if the data from a sensor looks unusual, and if the wearable is connected and streaming to one of the relay stations.

Remote terminal connections to the relay stations can be made from the base station, and thus from a remote computer using a Secure Shell (SSH) with a program like PuTTY [33]. This allows remote monitoring and repairs of the relay stations as long as they are connected to the local network that connects them to the base station. Using this remote access, the applications on the relay stations can be restarted if they get into a bad state, and the BeagleBones can also be restarted. The application code can also be changed mid-deployment if necessary. For example one relay station was frequently disconnecting from the Wi-Fi during the first deployment, so this relay station was modified to not use the Bluetooth, which interferes with the Wi-Fi. This significantly improved the Wi-Fi performance for the rest of the deployment. This change was made several days into the deployment without disturbing the dyad in any way.

B. System Reliability

The two main sources of problems with collecting data using the BESI sensor system are the Wi-Fi connections between the relay stations and the base station and the Bluetooth connection between the Shimmer and the relay stations. The relay stations occasionally disconnect from the Wi-Fi, which requires reconnecting and resetting the data streams to the base station. Because the Wi-Fi connections to the relay stations are used for remote repair, maintaining them is crucial. As the PWD moves throughout the house, the Shimmer switches which relay station it is connect to, but several mechanisms are required to ensure that these handoffs are done reliably. This section describes how challenge 5 was overcome.

Wi-Fi Connection

To ensure that the relay stations are always connected to the local Wi-Fi network, each one continuously runs a script that pings the local network router every five seconds, and resets the Wi-Fi connection if the ping is unsuccessful. In testing, this reliably maintains our ability to remotely access the relay stations, and in the rare instances where the Wi-Fi connection is lost, the connection is quickly reestablished. During one deployment the connection to one relay station was lost and not re-established, but this seems to be isolated to that specific relay station and deployment and has not been replicated in lab testing. The base station uses a wired Ethernet connection to the internet, which is less prone to failure than a wireless connection. The combination of a reliable remote connection to the base station and relay stations that automatically detect and repair issues with wireless connectivity allows us to monitor and repair most problems that occur during a deployment without disturbing the participants.

While the relay stations eventually recover from temporary loss of wireless connectivity to the base station, the sockets used in the data streaming do not recover as gracefully, and the connections need to be reset after Wi-Fi connectivity is restored. The base station can detect when the connection to a relay station is lost because not data is received from that relay station for a period of time. Every time the base stations checks for sensor data from a relay station and finds none, it increments a counter and determines the connection to the relay station is lost when the counter has been incremented a number of times corresponding to about 20 seconds without data. When this occurs, the base station closes the connections to that relay station, processes the data already collected, and begins listening again at the ports corresponding to that relay station. The relay stations cannot as easily detect a lost connection because when a relay station loses its connection to the base station, the commands to send data will continue to execute without throwing an exception, but the data will not be delivered to the base station. To remedy this, the base station sends a single acknowledge byte back to the relay stations whenever it checks for data from that relay station. Periodically at intervals of about 20 seconds, the relay station reads and discards all of the acknowledge messages sent since the last time it checked. If no messages were received in the preceding interval, the relay station determines that the connection to the base station was interrupted, and it tries to reconnect to the base station as shown in the state machine figure 11.

Bluetooth Connection to Shimmer

The mechanism for handing off the Shimmer Bluetooth connection between relay stations also handles the cases when the connection is temporary lost without a corresponding movement between rooms by the PWD. The relay station that lost the connection will being trying to reconnect to the Shimmer, and the Shimmer will enter a mode where it accepts new connections from the relay stations. This leads to a rapid reconnect between the Shimmer and a relay station in range. There are, however, cases when the connection to the Shimmer lost, but it does not enter a state where it accepts new connections. The two cases where this occurs with the Shimmer are described below along with their solutions. Also discussed is a mechanism that stops the Shimmer from streaming data while charging without requiring the caregiver to turn it off.

Momentary Bluetooth Connection - When the Shimmer connects to a relay station and quickly disconnects, which can happen when the devices are on the edge of Bluetooth range, the Bluetooth radio used by the Shimmer can stay in the connected state even though the connection is lost. Because Bluetooth is designed to allow a slave device, like the Shimmer, to connect to one master at a time, subsequent attempts to connect to the Shimmer by any relay station fail. During a deployment this error does not usually occur when the PWD moves between rooms in the house because the range of the Bluetooth radios on the relay stations overlap, so the Shimmer is usually well within the range of at least one relay station at all times, but this issue can occur when the PWD leaves and returns to the house and often means that the Shimmer needs to be manually reset before it can continue streaming. To address the problem that the Shimmer is sometimes erroneously in the connected state, the Shimmer firmware was modified to automatically detect this and reset the Shimmer when it occurs. When the Shimmer thinks it is connected it periodically increments a counter. If the counter reaches a certain value, which takes about 30 seconds if it is not reset to 0, the Shimmer performs a software reset by writing a value to the Shimmer's microcontroller watchdog timer without using the proper watchdog password. While the Shimmer is streaming data to a relay station, the relay station sends a command to the Shimmer that resets the counter approximately once every two seconds. Thus while the Shimmer is streaming correctly, its behavior is unchanged, but it quickly resets itself if it is in the connected state but is not connected a relay station.

Data Streaming While Charging - Because the dyad is required to switch Shimmers in the morning and at night and to change the Shimmer not in use, one Shimmer is always on the charger while the other is being worn. By default the Shimmer continues to stream when on the charger. During a deployment, the sensor system records accelerometer data from both Shimmers and post-deployment we can determine which data comes from the charging Shimmer because the Shimmer is held in a known fixed position when charging. However, each relay station only connects to one Shimmer at a time, so if the PWD is in the same room as the charging Shimmer and the charging Shimmer is connected to the relay station in that room, we will lose the date from the worn Shimmer unless it is within range of another relay station. To address this problem and to prevent streaming and storing useless accelerometer date from the charging Shimmer, the Shimmer firmware was modified to disable the Bluetooth radio when the Shimmer is on the charger and re-enable it when it is removed from the charger. This prevents any relay station form connecting to the Shimmer while it's on the charger and allows them to always be ready to connect to the Shimmer worn by the PWD. The Shimmer also performs a software reset whenever it is removed from the charger to prevent errors that have occurred from persisting over multiple days.

Bluetooth Connection While Initializing - Another error with the Shimmer that requires a manual reset occurs when a relay station tries to connect while the Bluetooth radio on the Shimmer while it is being initialized. When the Shimmer is first turned on it configures the Bluetooth radio to put it into slave mode and to disable remote configuration. To perform these operations the microcontroller on the Shimmer sends the commands to the Bluetooth module and waits for a response. If a relay station tries to connect while the microcontroller is waiting for a response, it disrupts the process and the response from the Bluetooth module on the Shimmer is lost. This leads to the firmware on the microcontroller hanging indefinitely and requires a manual reset to fix. This became a larger problem because the solutions to the two problems outlined above both require simulating cycling the power on the Shimmer, thus increasing the number of times the Shimmer initializes its Bluetooth module. The modification to the Shimmer that addresses this problem is to have the microcontroller send the configuration commands but not wait for a response from the Bluetooth module. This prevents the microcontroller firmware from hanging at the cost of possibly missing an unsuccessful attempt to configure the Bluetooth module. For the BESI application this is not an issue because the Shimmer's Bluetooth is in slave mode by default and disabling remote configuration is only relevant if the device to which the Shimmer is connected tries to issue configuration commands over the Bluetooth connection, which the relay stations do not do.

C. Physical Appearance

Because the sensor system is deployed for extended periods of time in dyads' homes, its appearance needs to be as unobtrusive as possible. The system also needs to be quickly installed and removed. This section addresses the appearance of the system components and the installation process and addresses challenge 7.

The relay station cases hide all of the exposed wiring and electronics on the base station other than the small openings for the sensors. The cases are made of a dull, light colored material that blends in with the paint on most walls. Because of the door motion sensor, the relay stations need to be near the door and need to not be blocked when the door is open. While repositioning tables to get the relay stations in the correct position is simple in a lab setting, it is often impossible in home deployments. Additional tables in a home can be tripping hazards for the PWD and caregiver, and even when there are tables in the locations that we would like to place relay stations, furniture in homes is often already in use holding pictures and other fragile objects that cannot be easily moved. To make placing the relay stations simpler, the cases have hooks on the sides that can be attached the wall using temporary double-sided tape. The tape can be removed at the end of the deployment without damaging the walls. The cases come in two versions that are mirror images of one another so that the relay stations can be hung on either side of a door to avoid the side where the door hinges. Figure 14 shows a relay station hanging on a wall in the lab. Hanging the relay stations on the walls generally keeps them out of the way, but in both deployments, the PWD or caregiver has run into one of the relay stations and knocked it off the wall. Determining the best places to put the relay stations is an ongoing issue that varies from house to house and needs to be carefully considered in future deployments.

Figure 14. Relay stations are typically hung on the wall next to the door into the room. This makes them easy to install and remove and usually keeps them out of the way.

While graphical displays and blinking LEDs are useful for debugging in the lab, they can be distracting or annoying, especially at night, and can even cause episodes of paranoia in people with dementia [31], so these light sources are turned off or covered. The Shimmer has two LEDs that indicate its state and battery level, which are covered with electrical tape during the deployments. The relay stations have lights on the BeagleBone, USB hub, and wireless adaptors that are also covered with tape. The base station laptop is kept closed during deployments and can even be placed in a cabinet or drawer as long as the power and Ethernet cables are connected.

V. Deployment Procedure and Results

To date the BESI sensor system has been used in two deployments in the homes of people with dementia. This section describes the procedure for deploying and monitoring system and presents the results from the deployments in terms of the amount of data collected. While the goal of the BESI project is to get useful information about agitation from the sensor data, the processing and analysis to extract to this information is left for future work. The deployment issues discussed in the previous section related to the Bluetooth connection to the Shimmer were identified as problems from the results of these deployments, and were not included in the system for these deployments. The solutions to **Momentary Bluetooth Connection** and **Data Streaming While Charging** were added after the first deployment, and the solution to **Bluetooth Connection While Initializing** was added after the second deployment.

A. Deployment Procedure

The medical experts conduct a home visit to the potential dyad, where they describe the BESI study, perform an initial assessment battery, record the floor plan of the home, and note the rooms that are most important to instrument. From this visit, the likely locations of the relay stations are selected, which dictates the number of relay stations required for the deployment. Once this is known, each relay station and the base station are assigned static IP addresses in the LAN router. The configuration file on each relay station is modified to contain the port number for that relay station and the IP address of the base station, and the configuration file on the base station is modified to contain the port number for each relay station, the correct number of relay stations, and the Bluetooth MAC addresses of the two Shimmers used for the deployment. The relay station is also configured with the name of the LAN network and the IP address of the router so it will automatically connect to the network when powered on and periodically ping the router to check the connection. All of this setup can be done in a short amount of time as it only involves editing a few lines in text files on the Relay Stations and the Base station.

At the deployment site, the first step is to finalize the locations of the relay stations, base station and router with the caregiver and the medical team. While the general floorplan of the home is known prior to deployment, details such as the locations of plugs and furniture are not known until deployment time. This requires some adaptation of the locations of the relay stations that were selected base on the floorplan alone. The relay stations need to be positioned near a power outlet and near the door into the room, but they cannot be placed where they will be in the way of pedestrian traffic. Placing the relay stations often requires using extension cords to reach a plug, but the cords cannot present tripping hazards for the occupants, so they are usually run behind furniture. The base station needs to be close to an Ethernet port, but otherwise can be out of site. For best performance the LAN router should be placed near a power outlet in the middle of the house, but the exact placement depends on the layout of the house.

Once all of the sensor system components are placed, the system is turned on, and each relay station is tested individually to ensure that it is transmitting data correctly and that it connects to each of the Shimmers used in the deployment. The data plots displayed on the base station are used to determine this. Next, the whole sensor system is run for 15 minutes to make sure the connection from each of the relay stations to the base station is stable. The dyad is given two Shimmers and a charger and is instructed on using and charging the Shimmers. The caregiver is also instructed on using the tablet for surveys. The setup and testing takes about two hours.

While the data collection is in progress, the sensor system is checked every four hours during the day using the remote monitoring capabilities. TeamViewer allows the researchers monitoring the deployment to see the data plots on the base station in real time from a remote computer and without requiring anything from the dyad. The data saved on the base station can also be checked to determine how much data has been collected from each relay station. If problems arise, they are fixed remotely if possible, or by contacting the caregiver or going to the deployment site if necessary. The most common problems observed were a temporary loss of connection from a relay station to the base station that requires restarting the application on the relay station and errors with the Shimmer that requires a manual reset. Restarting the relay station application can be done remotely without disturbing the dyad, but resetting the Shimmer does require contacting the dyad and asking them to turn it off and back on again. The issues with the Shimmer that require a manual reset are discussed in the deployment considerations section above.

B. Deployment Results

For the two deployments of the BESI system conducted to date, both were performed in single family homes with one caregiver and one PWD. Both houses were approximately 3,000 sq. ft. with a main floor and a basement, and both deployments used six relay stations. The first had four on the main floor and two in the basement, while the second had five on the main floor and one in the basement. The relay stations were positioned so that the door motion sensors could measure movement between rooms and to monitor the conditions in the main rooms in the houses.

Table 1 shows the results for the first data collection in terms of how much data was collected. Relay station 6 disconnected from the LAN Wi-Fi network twice without reconnecting, which led to it only collecting data for about 40 hours, but this was the only relay station to experience this problem, which suggests that it is specific to this relay station. All of the other relay stations collected environmental data for at least 90% of the total deployment time. For relay stations 2, 3, and 5, several hours of data were lost due to operator error during the deployment. Ignoring that loss, the five relay stations each collected data for at least 96% of the deployment. The remaining time without data was due to losses in the connections between the relay stations and the base station before the relay stations automatically reconnected and resumed streaming. These results suggest that the sensor system can reliably collect environmental data even with occasional losses of connection between the relay stations and the base station. An improvement can be made by collecting data even when the relay stations are not connected and transmitting when a connection is established. This is described in more detail in the Discussion of Design Decisions section below.

Table 1. Results for data collected from the first week long deployment in the home of a caregiver-PWD dyad. Bluetooth was disabled on relay station 1 to improve Wi-Fi performance. Relay station 6 lost connection and did not reconnect shortly into the de

Relay Station	Accelerometer data collected (hours)	Environmental data collected (hours)	Percent of time collecting environmental data (out of 172 hour deployment)
1	0.37*	169.3	98.3%
2	3.5	154.9	90.0%
3	35.6	155.3	90.2%
4	1.9	165.9	96.3%
5	3.15	165.2	95.9%
6	7.34	40.9	23.8%

*Bluetooth on relay station 1 was disabled approximately 24 hours into the data collection

The results for the second deployment are shown in table 2. This deployment had more issues with the wireless connections between the relay stations and the base station than the first deployment. Relay stations 1 and 6 especially frequently disconnected from the base station and occasionally required remote restarts. The other relay stations collected a similar amount of data to the relay stations in deployment 1. Deployment 2 also had issues with the remote monitoring which twice required on-site maintenance on the base station. Once the caregiver was guided through restarting the Teamviewer application over the phone, and later in the deployment, one member of the clinical team visited the house to restart the application. The Teamviewer servers were also down twice for several hours each, which prevented remote monitoring during those times. The dyad reported that the quality of their Wi-Fi was lower during the deployment, which suggests that the local network used to transfer data from the relay stations to the base station was interfering with the existing Wi-Fi network. To address these problems, an option to store the raw data on the relay stations has been added for future deployments. This should address the issues with Wi-Fi interference as well as enabling the collection of data when the connection between a relay station and the base station is lost. This is discussed in more detail in the Design Decisions section below.

Table 2. Results for data collected from the second week long deployment in the home of a caregiver-PWD dyad.

Relay Station	Accelerometer data collected (hours)	Environmental data collected (hours)	Percent of time collecting environmental data (out of 166 hour deployment)
1	3.5	142.5	85.9%
2	35.9	149.2	90.0%
3	27.2	156.7	94.5%
4	0.5	158.0	95.3%
5	2.7	159.9	96.4%
6	0	126.9	76.5%

The collection of accelerometer data was not as successful as the environmental data. During the first deployment, data was streamed from two Shimmers: one worn by the PWD and one on the charger. The Shimmer on the charger still streamed data, but because the position of the Shimmer on the charger is fixed, the accelerometer data from the body-worn Shimmer can be distinguished from the data from the charging Shimmer. Table 1 shows the amount of non-charging accelerometer data collected. In total about 50 hours of meaningful data was collected from the Shimmer. Due to the lack of ground truth data on the locations of the PWD and the Shimmers it is not possible to determine the exact situations that caused the Shimmer to stop streaming, but the following are some causes that seem likely based on the data and feedback from the users.

1. On several days, the PWD left the house for an extended period of time. During this time, the expected behavior is to not collect accelerometer data; however, when the PWD re-entered the house, the Shimmer did not reconnect until it was power cycled by the caregiver. This led collecting very little data during the day on the days when the PWD left the house.
2. At night, the connection to the Shimmer would occasionally be lost, likely due to the PWD sleeping on top of the Shimmer.
3. One of the Shimmers was dropped during the data collection, which damaged the case so that it no longer fit snugly in the wrist holder and frequently fell out. This led to the PWD only wearing a Shimmer during the day for the last few days of the deployment.
4. At some times during the deployment, he PWD chose not to wear the Shimmer.

During deployment 2 more accelerometer data was collected than during deployment 1, but it was still significantly less than desired. For this deployment, the charging Shimmer did not stream data, and we collected about 70 hours of data from the accelerometer as shown in table 2. Issues 1) And 3) in the list above were addressed for this deployment, but issues 2) and 4) were still present. The problem with the Shimmer not reconnecting after the PWD returned to the house was fixed using the mechanism described in the **Momentary Bluetooth Connection** section above. For this deployment, the Shimmer was taped to the hook that attaches it to the wrist strap, which prevents the Shimmer from falling out of the hook. The Shimmer did have to be manually reset several times by the caregiver due to the relay stations attempting to connect to the Shimmer while it was initializing the Bluetooth radio. This problem and its solution are described in detail in the **Bluetooth Connection While Initializing** section above.

VI. Discussion of Design Decisions

Designing the BESI sensor system involved making a number of conscious and unconscious decisions that affect the functionality and reliability of the system. This section describes some of these decisions, and the pros and cons of the chosen implementation.

A. Shimmer

Perhaps the most important decision made in determining the architecture of the sensor system was the selection of a wearable accelerometer platform. There are a variety of commercial sensing platforms that provide body-worn accelerometer data from consumer products like smart bands and watches, to inertial measurement units (IMUs) geared towards research like the Shimmer. Other options are research platforms like the TEMPO system [36] that was used for the initial agitation detection study. The most important requirement for the system is a battery life of at least 16 hours so that the wearable can be worn throughout the day and charged at night. Additionally, The IMU should also be customizable in terms of what sensors are used and the sampling rates of these sensors. The raw accelerometer data needs to be available without the use of a proprietary app, which is not the case for most commercial wearables. Requiring an app limits the design space of the relay stations too much in terms of what hardware, operating system and software can be used. Many commercial products only work with a smartphone or computer app that provides processed motion data. The Shimmer was chosen because the firmware is open source, the shimmer can be easily integrated with a custom relay station to stream raw data, it has the needed battery life, and it has the support typical of a commercial product. This dictated that the system needed relays in each room to receive the accelerometer data over Bluetooth. While the Shimmer does meet the technical requirements of the system, it is bulky and ugly, which limits the acceptance of the system by potential participants. It also has limited debugging and remote monitoring capabilities, so if it fails during a deployment, a call to the dyad or home visit is usually required. Wearing a commercial smartwatch would be less of a burden on the participants, so using a smartwatch that meets the technical requirements is preferable to using Shimmer.

B. BeagleBone

Each relay stations needs to interface with the environmental sensors and the Shimmer and send the data to the base station. The BeagleBone was chosen as the foundation of the relay station because it has a variety of I/O ports that can be easily accessed to connect the wired sensors, it can communicate via Bluetooth and Wi-Fi, and has an embedded OS that simplifies coding but provides low-level I/O access. The BeagleBone meets all of the technical requirements for the relay station, but because it is fragile and uses exposed wires to connect to the sensors, it needs to be enclosed in a sturdy case. This adds to the size and weight of the relay stations. The biggest flaw with the BeagleBone is that it is electronically fragile, so if it is unplugged without being shut down first, the board can be destroyed. This creates a risk of failure if a relay station is accidentally unplugged during a deployment or if power is lost.

C. Sensors

Sensors that use wired connections to the BeagleBone were chosen because wired connections are simpler and more reliable than wireless communication. This allows all of the relay station components to be contained in a single package and all of the sensors to be powered off of the BeagleBone. However, this does limit the options for placing sensors because the door sensor dictates that the relay station be placed by the door, so we cannot select other locations that might be better for getting light or audio data. This decision also creates the potential for sensors and the BeagleBone to interfere with each other. For example, the BeagleBone runs fairly hot, which can affect the temperature reading even though the temperature sensor is placed as far away from the BeagleBone as possible within the case.

Additionally, the wires to the door sensor pick up a significant amount of noise on the analog signal from the wire to the USB hub when the Wi-Fi is in use.

D. Relay Station Data Acquisition and Transmission

While the software to sample the sensors on the relay stations is multithreaded, the process of connecting to the base station, acquiring data, and transmitting it to the base station is handled in a single thread. This means that while the relay station is not connected to the base station, no sensor data is collected. A better solution is to use one thread for sampling the sensors that runs all of the time and writes the data to a buffer, and another thread that maintains the connection to the base station and transmits the data from the buffer when a connection is established. The current multithreaded approach also has the disadvantage that each thread has to detect a disconnect on its own.

During the first deployments, it became clear that only collecting data from a room when the relay station is connected to the base station via Wi-Fi leads to unacceptably long gaps in the collected data that make performing the required post-deployment analysis difficult. Constantly streaming the raw sensor data can also interfere with the existing Wi-Fi network in the dyads home. To address these issues, the system has been modified to include local storage of all of the raw sensor data on the Relay stations and streaming only the results of processing the data, which significantly lowers the bandwidth requirements. Currently the relay stations periodically update the base station with the amount of date collected since the last update to ensure that it is functioning correctly; however, future implementations will include the data processing algorithms for agitation prediction and detection on the relay stations that will process the raw data and send the results to the base station. For example instead of streaming raw accelerometer data, the relay stations could instead send the times when the magnitude of acceleration is above a given threshold

To enable local storage on the Relay Stations, two significant issues needed to be addressed: storage space and time synchronization. The BeagleBone has 4 GB of flash memory with about half of that used for the operating system and relay station code. Each relay station collects approximately 100 MB to 300 MB of data per day depending on the amount of time it is connected to the Shimmer. This limits the local storage to about one week of data, but the BeagleBone has a micro-SD card slot that allows the storage capacity to be greatly expanded. Commercially available 64 GB SD cards can be added to the relay stations that allow deployments to run at least 200 days without needing to offload the data from the relay stations, which is well above the 60 days required for the BESI application. Because the BeagleBones do not have real-time clocks, the system time is reset every time the relay stations are turned on. To get the correct time, the base station sends the start time of the data collection to each relay station when they first begin storing data, and each data point is timestamped with the time since the start as measured by each relay station. When the relay stations update the base station on the amount of data collected, which currently happens once every ten minutes, the base station sends back the current time, which the relay station uses to update its own time. This assures that the data from different relay stations are synchronized to within a few seconds, which is all that is required for this system.

E. Base Station Application

When the base station receives a packet of sensor data, the packet is written to a temporary file without and processing, and then the packet is parsed to add the data to a list for plotting. This has the advantage of storing the data correctly in instances when a packet is split in transit, but the plotted data can still be incorrect in this case. However, this approach does duplicate data and requires extra memory to store the lists. It also requires that the temporary files be processed occasionally and no data is streamed while this processing occurs to avoid concurrency issues with accessing the temporary data storage files. The data plots can be used to quickly ascertain which relay stations are currently connected and if their sensor data is reasonable. The plotting does not give any information about the past state of the system, such as how long the Shimmer has been connected to a relay station. A possible improvement is to record important system events to get a better idea of how the sensor system has performed over an entire deployment.

VII. Conclusion

Within the field of wireless health, one of the fundamental challenges is demonstrating that providing patients and medical experts with information from ubiquitous sensing platforms can lead to improved patient outcomes. The BESI project is attempting to show that providing caregivers with more information about agitation in their loved ones with dementia can increase their perceived ability to deal with the agitation and to care for their loved one in general. This thesis describes the sensor system used to detect agitation and the environmental context surrounding agitation events. It also addresses some of the more general issues with deploying and monitoring sensor systems in participants' homes for extended periods of time.

To collect data from multiple rooms, the sensor system uses a set of environmental sensors in each room that are connected to a relay station in that room. The relay station collects all of the sensor data and sends it to a central base station in the home where it is stored. The primary means for detecting agitation is a wrist worn wireless accelerometer platform on the PWD. Whenever the PWD is in the house, this platform connects with and streams data to one of the relay stations within range. The sensor system can be scaled by adding or removing relay stations to cover houses with different numbers of rooms.

Remote monitoring during a deployment is an important part of any sensor system because deployments are intrusive to the subjects, so it is important to detect and correct problems quickly so prevent wasting subjects' time. Correcting issues remotely ideally without even needing to notify the subjects is also an important feature of systems like the one used for BESI. In BESI these goals are accomplished using a remote desktop application to view the real-time data plots on the base station and to view the files containing collected data. This application also allows access to each of the relay stations to correct problems as they arise.

Deploying sensor systems in the homes of actual patients presents several challenges that are not readily apparent when designing the system in the lab. The sensors need to be safe, unobtrusive, and easily installed and removed, and the lights that appear on many of the components used in the sensor system are covered. The BESI sensor system uses the 3D printed relay station cases to hide and protect the sensors. The hooks on the case allow the relay station to be easily hung on the walls of the home and to be removed quickly without damaging the walls. While the Shimmer is a well designed and packages accelerometer platform, it is not designed to be constantly worn and it can be annoying or ugly to some participants. Because of this a goal for future deployments is to develop a better case for the Shimmer or to use a different accelerometer platform with a more ergonomic design.

The BESI sensor system has been used in two data collection deployments in dyads' homes. Over these two deployments, the relay stations collected environmental data for 86.1% of the deployment with most relay stations collection for more than 90% of the time. With the addition of local data storage to the relay stations this number should be much closer to 100% in future deployments. The amount of accelerometer data collected was significantly lower than environmental data due to a variety of issues with establishing and maintaining a Bluetooth connection to the Shimmer. These issues have been addressed and should allow collecting Shimmer data whenever the PWD is in the house in future deployments.

Future Work

The next step for the BESI system is to ensure that it can collect environmental and accelerometer data more reliably than in the first two deployments. The issues with missing environmental data from some time periods is addressed by storing the data locally on the relay stations, which eliminates the requirement that the relay stations be connected to the base station to collect data. The specific issues with Shimmer that occurred have been corrected, but no deployments have been done since then to assess how well the system works with these changes.

Some participants report that he Shimmer is bulky and ugly and, so we are exploring alternate packaging options or the Shimmer, or other wearable accelerometer platforms that look nicer. Using another platform, like a commercial smartwatch, would require some changes to the relay station firmware and might necessitate on-node processing on the wearable rather than streaming to prolong battery life.

In the near future, the BESI sensor system will be used to collect data on more dyads for longer periods of time, and the data will be used to develop personalized models that predict and detect agitation. These models will eventually be used to send recommendations for interventions to the caregivers that will be designed to allow them to better deal with agitation in their PWD.

LIST OF REFERENCES

[1] “2015 Alzheimer’s Disease Facts and Figures,” Alzheimer’s Association, www.alz.org.

[2] “What is Dementia,” Alzheimer’s Association, alz.org.

[3] “Dementia,” PsychGuides.com, www.psychguides.com

[4] J. Sansoni, K.H. Anderson, L.M. Varona, G. Varela, “Caregivers of Alzheimer's patients and factors influencing institutionalization of loved ones: some considerations on existing literature,” *Ann Ig.*, 25(3):235-246, 2013.

[5] J. Cohen-Mansfield, “Conceptualization of agitation: Results based on the Cohen-Mansfield Agitation Inventory and the Agitation Behavior Mapping Instrument,” *International Psychogeriatrics*, 8(S3):309-315, discussion 351-354, May 1997.

[6] J. Cohen-Mansfield, “Instruction Manual for the Cohen-Mansfield Agitation Inventory (CMAI),” The Research Institute of the Hebrew Home of Greater Washington

[7] V. F. S. Fook, P. V. Thang, T. M. Htwe, Q. Qiang, A. A. P. Wai. Automated recognition of complex agitation behavior of dementia patients using video camera, *9th Int'l Conf. on e-Health Networking*, Application and Services, pp.68-73, 19-22 June 2007.

[8] A. Bankole, M. Anderson, T. Smith-Jackson, A. Knight, K. Oh, J.S. Brantley, A.T. Barth, J. Lach, “Validation of Non-Invasive Body Sensor Network Technology in the Detection of Agitation in Dementia,” *American Journal of Alzheimer’s Disease & Other Dementias*, 27(5):346-354, August 2012.

[9] Salekin, Asif, Hongning Wang, and John Stankovic. "Demo: KinVocal: Detecting Agitated Vocal Events." *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*. ACM, 2015.

[10] “Aware Home Research Initiative,” Georgia Institute of Technology, http://www.awarehome.gatech.edu/

[11] “The PlaceLab,” Massachusetts Institute of Technology, http://web.mit.edu/cron/group/house_n/placelab.html

[12] Dickerson, Robert F., Eugenia I. Gorlin, and John A. Stankovic. "Empath: a continuous remote emotional health monitoring system for depressive illness." *Proceedings of the 2nd Conference on Wireless Health*. ACM, 2011.

[13] Dickerson, Robert F., et al. "Empath2: A flexible web and cloud-based home health care monitoring system." *Proceedings of the 8th ACM International Conference on Pervasive Technologies Related to Assistive Environments*. ACM, 2015.

[14] Tapia, Emmanuel Munguia, Stephen S. Intille, and Kent Larson. "Activity recognition in the home using simple and ubiquitous sensors." *International Conference on Pervasive Computing*. Springer Berlin Heidelberg, 2004.

[15] Suryadevara, Nagender Kumar, and Subhas Chandra Mukhopadhyay. "Wireless sensor network based home monitoring system for wellness determination of elderly." *IEEE Sensors Journal* 12.6 (2012): 1965-1972.

[16] Zhu, Chun, Weihua Sheng, and Meiqin Liu. "Wearable sensor-based behavioral anomaly detection in smart assisted living systems." *IEEE Transactions on Automation Science and Engineering* 12.4 (2015): 1225-1234.

[17] Nagels, G., Engelborghs, S., Vloeberghs, E., Van Dam, D., Pickut, B. A. and De Deyn, P. P. Actigraphic measurement of agitated behaviour in dementia. Int J Geriatr Psychiatry, 21(4):388-93, Apr 2006

[18] Meredeth Rowe, Stephen Lane, and Chad Phipps. "CareWatch: a home monitoring system for use in homes of persons with cognitive impairment."*Topics in geriatric rehabilitation* 23.1 (2007): 3.

[19] Meredeth A. Rowe, John A. Kairalla, and Christina S. McCrae. "Sleep in dementia caregivers and the effect of a nighttime monitoring system."*Journal of Nursing Scholarship* 42.3 (2010): 338-347.

[20] “Shimmer3,” Shimmer Research, http://www.shimmersensing.com/shop/shimmer3

[21] ”Omni-Directional Foil Electel Condenser Microphone” Challenge Electronics, http://cdn.sparkfun.com/datasheets/Sensors/Sound/CEM-C9745JAD462P2.54R.pdf

[22] “SparkFun Electret Microphone Breakout,” SparkFun, https://www.sparkfun.com/products/12758

[23] “LM60 Temperature Sensor,” Texas Instruments, http://www.ti.com/lit/ds/symlink/lm60.pdf

[24] “TSL2560, TSL2561 Light-to-Digital Converter,” TAOS, https://cdn-shop.adafruit.com/datasheets/TSL256x.pdf

[25] “Adafruit TSL2561 Digital Luminosity/Lux/Light Sensor Breakout,” Adafruit, www.adafruit.com/product/439

[26] “BeagleBone Black,” beagleboard.org/black

[27] “USB WiFi (802.11b/g/n) Module with Antenna for Raspberry Pi,” Adafruit, www.adafruit.com/products/1030

[28] “Bluetooth 4.0 USB Module (v2.1 Back-Compatible),” Adafruit, www.adafruit.com/products/1327

[29] “socket — Low-level networking interface,” Python Software Foundation, docs.python.org/2/library/socket.html

[30] “PyQtGraph,” http://www.pyqtgraph.org/

[31] “Suspicion, Delusions and Alzheimer’s,” Alzheimer’s Association, www.alz.org.

[32] TeamViewer, teamviewer.com/en/

[33] http://www.putty.org/

[34] Joshua Wright, “Dispelling Common Bluetooth Misconceptions,” SANS Technology Institute, www.sans.edu/research/security-laboratory/article/bluetooth

[35] Liu, Hui, et al. "Survey of wireless indoor positioning techniques and systems." *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 37.6 (2007): 1067-1080.

[36] Powell HC, Hanson MA, Lach J. “A Wearable Inertial Sensing Technology for Clinical Assessment of Tremor”. *Biomedical Circuits and Systems Conference*. Montreal, QC, Canada; November 27-30, 2007:9-12.