Empath2: A Flexible Web and Cloud-based Home Health Care Monitoring System

Robert F. Dickerson
Dept. of Computer Science
University of Texas at Austin
Austin, TX USA
rfd@cs.utexas.edu

Enamul Hoque
Samsung Research America
Mountain View, CA USA
eh6p@virginia.edu

Ifat Afrin Emi and John A. Stankovic
Dept. of Computer Science
University of Virginia
Charlottesville, VA USA
iae4qb@virginia.edu,
stankovic@cs.virginia.edu

ABSTRACT

Home health care sensing systems are projected to streamline the efficiency of the practice of medicine by decreasing the costs of senior care and by providing preventative care to keep people out of hospitals and nursing homes. Many current sensing systems are not yet flexible enough to easily handle widely different medical applications. Empath2 provides a flexible three layer architecture that uses the Cloud and can easily be instantiated for different home health care applications. To demonstrate the flexibility of the architecture, Empath2 was instantiated for three widely different purposes. We present the design of Empath2 stressing properties that support flexibility and discuss its differences from other flexible home monitoring architectures. Evaluations for three sets of real home deployments (two of which with actual patients, and one with healthy people) are presented showing the short deployment times, short software development times, and its effectiveness for the applications at hand. Lessons learned are also presented.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Wireless communication; D.2.2 [**Network Protocols**]: Protocol architecture

1. INTRODUCTION

In many countries of the world we are seeing rapidly aging populations that are overwhelming the health care system [3]. Home health care monitoring is being touted as one solution. Such systems have been in development for many years and there are many companies selling various related products. Most of these systems to date either monitor activities of daily living (ADLs) or are focused on a single medical issue (e.g. Parkinson's disease). In order for these systems to be able to handle wide varieties of home monitoring, various flexible architectures and middleware solutions have been developed [14, 18, 19, 22]. However, what they lack is that they fail to demonstrate their flexibility in actual deployments with real patients that study multiple medical issues. We find that the realities of actual deployments for real patients with real medical problems give rise to many new requirements, including: the need to use off-the-shelf devices, support for non-expert users, easy installation and maintenance, ease of developing the specific software required for a particular instantiation, remote monitoring, extensible monitoring over time, handling the large scale of data, dealing with complex humans and environments, and support for the evolution of human behaviors.

In this paper we present a new architecture called Empath2 and evaluate it for multiple home health care purposes. The main contributions of this paper are:

- A demonstration of the flexibility of the Empath2 architecture by adapting it for three case studies of deployments with real patients. These are: examining the relationship of sleep and stress on the number of seizures that people with epilepsy experience, the relationships between nighttime agitation and incontinence events for those with Alzheimer's disease, and depression monitoring.
- Deployment details showing short installation times and few number of intermediate visits needed. Most deployments require a half-hour time and can be done by non-technical users. The system configuration times for different instantiations of Empath2 are also shown to be short.
- Introduction of a web-service based stream abstraction designed for heterogeneous and multidimensional stream data, These streams are merged or forked to form inference trees. When data is requested, inference is lazily evaluated. We also show examples of storing this stream data in a document-based backend (MongoDB).

2. EMPATH2 SYSTEM DESIGN

The basic architecture of Empath2 is similar to many home health care sensing architectures and consists of three main layers: the sensing layer, the basestation, and a cloud-based web server and database shown in Figure 1. However, in contrast to single purpose architectures, e.g., those that

PETRA '15, July 01 - 03 2015, Island of Corfu, Greece.

DOI: http://dx.doi.org/10.1145/2769493.2769508ACM. 2015.

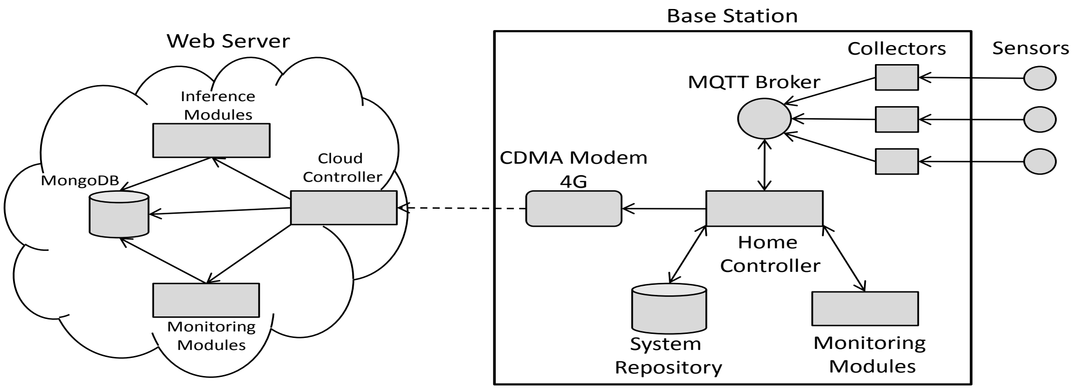

Figure 1: Home area network framework.

monitor ADLs, because of the specific design decisions made for Empath2, new instantiations for different home health care applications can be built quickly and simply. Also, in contrast to other flexible architectures [14, 19], Empath2 was instantiated and evaluated with real patients on three widely different monitoring applications.

Sensing Layer: Empath2 permits easy installation of various wired and wireless sensors that generate either continuous or event-based data. Some examples incorporated into our framework include weight scales, contact reed switches, PIR motion detectors, accelerometers, sound sensors, and user surveys. The framework allows integrating sensing devices that use different communication protocols (e.g., Bluetooth, Zwave, X-10, etc.) by having the associated communication radio adaptor at the basestation. Because some sensors generate high data-rate streams, we typically employ reducers which convert the high data rate or highly multidimensional data to lower data rate and fewer dimensions. For instance, this can be achieved by buffering data in the time period of an epoch (for instance 1 minute), and then calculating the statistics such as the mean, variance, min, max, count as the reported as output. This epoch is instantiated, and forwarded to a basestation over the network.

To support ease of development and flexibility, each sensor program can be built as an independent module implemented in the programmer's choice of language such as Java, C, or Python and be launched from the command line. The modules that we have built, available in a library, are meant to be cross-platform and have been tested on a variety of platforms such as the Raspberry Pi, Beaglebone, and laptops. Besides modules in the library, there are others created by third parties such as the Fitbit and Withings scale that have their own mechanisms to send data to their own server in the cloud. Currently, we have incorporated both the Withings weight scale and the Fitbit, which publish the data to the Withings and the Fitbit web servers, respectively. Because many commodity components release an API for their data, they can be easily incorporated into Empath2's system along with the sensor data collected internally. Consequently, third party devices can be integrated easily with Empath2.

Empath2 also supports adding new sensors to the architecture. When a developer integrates a new sensor device, they must extend the class `AbstractDataCollector` which implements basic functionality such as maintaining a connection to the message broker and handling the serialization and publishing messages. In addition, they must implement the virtual functions for initialization, starting, stopping, and shutting down the device. For each new type of sensor, in addition to the above class, a new data collection class has to be created to read the raw sensor data and publish it to the MQTT broker. Finally, the home controller module has to subscribe to this new sensor.

One interesting feature of our sensor layer is that an entire Zigbee-based mesh network can be treated as a single node attachment to Empath2. Each mote in a Zigbee network shares a common Personal Area Network (PAN) assigned during configuration. Each mote has a pre-configured 64-bit value as its ID. One mote must be configured as the Coordinator and is attached to the basestation, and the other devices are set as routers or end-devices (depending upon duty cycling requirements). When a packet is received by the coordinator, the data is converted to a stream ID for the particular device that it came from.

The Basestation: When the basestation receives an epoch from the sensing layer, it is stored temporarily and staged for syncing to the web service. Empath2 also supports a higher level middleware communication layer on the basestation to simplify the networking between the components. We defined a common protocol for sending data to other devices. In the home, one device is designated as a basestation, and the message broker runs on this device and listens for incoming connections. We use the MQ Telemetry Transport (MQTT) protocol [16] and the Mosquitto server for our message broker. MQTT implements a publish/subscribe message pattern to provide a one-to-many message distribution so that we can adequately decouple the modules from one another. In each of the packets, the payload is a series of Epoch data serialized as a JSON string. MQTT provides this basic network connectivity with small transport overhead, the fixed-length header is just 2 bytes.

When a sensor is installed, the IP address for the message broker must be specified and a name for that sensor. For example, a broker can be located on a machine at `10.0.0.1` and the bed sensor is on `10.0.0.8` and the name is **bed1**. When the bed sensor collects 50Hz data for one second, the mean and standard deviation are computed for each of axis of each accelerometer, resulting in a 12 dimensional feature vector. A new *Epoch* is created using the current timestamp, duration, and feature vector. This epoch is serialized to a JSON string and published to the topic: `sensors/bed1`. The Empath2 controller subscribes to sensor data by subscribing to the topic group `sensors/+`. When a message arrives, the message is stored in the Controller's local database for later syncing with the Cloud. This local database was implemented using SQLite3 which is a self-contained, serverless, transactional SQL database engine. Once data is synced to the cloud it is deleted from the basestation.

Cloud: The third layer is the cloud layer. Empath2's cloud-based server is implemented by a Java web application that can be installed into any Java servlet container (Tomcat, Jetty, JBoss, etc). We used the Spring3 framework for handling the Model View Controller (MVC) pattern for handling requests from the clients and serialization of JSON messages. The Spring Security extension was used to implement authentication and access control. Each user is given a set of roles such as Patient, Clinician, Technician, Administrator, Coordinator, Researcher for controlling access to requests to resources. For a more robust access-control mechanism, these roles can be predicated, so that a user can be a Clinician for X, and X is a Patient of study Y and Researcher is a member of Study Y. Before any resources are served, a user must sign in, and a session ID is created and stored as a cookie in the HTTP client, and the communication provided through an HTTPS tunnel.

To further demonstrate flexibility, we implemented the system on two Cloud platforms: the Amazon Web Services and the Google App Engine. For the Amazon Web Service, we launched a two small EC2 instances with 2 GB of memory. One EC2 instance ran a Jetty9 instance for our application, and another ran only the MongoDB database. The database node mounted a RAID10 array with 8 GB of Elastic Block Storage formatted with the XFS filesystem. For the Google App Engine, we used their native components, such as the Jetty web server and the Google Data Store for the database. Changing the deployments only required changing the Spring's XML configuration file describing what DAO implementation was to be used for the local stream data.

A key component of Empath2 is an abstraction we developed that meets the special needs of sensor streams called a `StreamFeed` [5]. This new steam abstraction is similar to a regular web resource such as HTML or an image that it can be referenced by an URL. As such, they can be both the source and the target of hyperlinks. All data in the Empath2 system are exposed simply by a URL, where the UUID is a 128-bit UUID: `http://www.XYZ.com/stream/{UUID}`

Likewise, all Epochs inside of a stream also can be retrieved using a URL: `stream/{UUID}/epoch/{ID}`

Each of these URLs point to a StreamFeed that can be fused, processed and filtered to create new StreamFeed. The result can be repeated to eventually produce an inference tree. For example, data from multiple sensors in the home can be combined to form better estimates of the occupant's behaviors. Because the returned stream is often quite large, a series of time range parameters are encouraged. A particular attribute can be used as a filter parameter. Consider the following HTTP GET request to:

`sensor/{UUID}/epochs/filter/?min=100`

This URL-based interface follows the RESTful (Representational State Transfer) principles [6]. The advantage of using a RESTful interface is that there is an inherent standardization placed on the operations that can be applied to the resources, without needing to explicitly define descriptions of the methods. The URL contains all the information that is needed to return to a particular state of a web service.

There are three basic types of streams in Empath2: First, there are persistent streams that are stored in a database. Second, and very common, are memory streams that do not have persistence and are populated upon request. This is useful for streams are only needed to produce some report to a caregiver. This allows the `Evaluator` objects to store the results temporarily like a scratchpad, so that clients can quickly query for the information without requiring the entire inference chain to be recomputed. Third, there are web streams which are data sources that are not stored locally in the Empath2 system, but rather through another webserver on the Internet.

We used a document-based database for serveral reasons. Most relational databases (such as MySQL) enforce ACID guarantees for atomicity, consistency, isolation, and durability. However, our data being generated rarely changes once being committed, also even if it is changed, the data does not need to be consistent across the database replicas, just eventually consistent. Document-based databases achieve much higher performance in the write once, read many times use case.

When a deployment is setup, many stream `Processors` are created and their operating parameters are set. Afterward, streams are "wired" to the input and output ports for these `Processors`. Consider a simple example with scoring questionnaires. A `PHQ9Evaluator` processor is created and the input port "PHQ-9 Responses" is wired to a persistent stream A holding the item responses, and to the output port "PHQ-9 Score", a memory stream B is wired to the port. When the PHQ-9 stream is queried for the first time, the stream holds no epochs and is marked 'dirty'. Because of this, the StreamService invokes the PHQ-9 Evalator's evaluate() function which will query A for the all epochs in that time range. Because A is persistent and not marked dirty, all the values are available and evaluation does not need to be taken to another level. Next, for each of the epochs, a score is produced and the result added to stream B. This evaluation method uses a lazy evaluation approach for fetching stream information because the rate of querying for higher level data is much less frequent than the production of lower level data. When a lower level stream gets appended to, the streams above it are marked 'dirty' for reevaluation. There are some requirements to this structure, most importantly, there cannot be any cycles in the inference chain or evaluation will never halt.

The monitoring modules in the cloud run each morning to check if data have been uploaded from the deployment site and if there is any inconsistent data that represents one / more sensors may be broken. The status is automatically emailed to a monitoring team so that they can react in case of any sensor / device failure. The monitoring modules in the basestation continuously logs the system memory consumption, Internet connectivity status, and battery level of the laptop along with checking consistency in the generated data. These logs are also automatically uploaded in the cloud. This multi-level monitoring ensures that we can detect problems as soon as possible and also often determine the exact problem.

Table 1: **Stream Metadata Examples**

Key	Description
Creator	UUID of the user who created the stream
DeploymentID	UUID of the deployment the data came from
TargetID	UUID of a user the stream might relate to
PreferredRenderer	Bar plot, time series, table, etc
Device	Specification of the device make and model

2.1 Time to Instantiate

The time it takes to instantiate an instance of Empath2 for the use of an new application depends on the time it takes to implement inference and sensing modules. The inference modules are application-dependent and vary considerably depending on the complexity involved in the application. For example, the development times for the inference logic presented in this paper took the following amount of time: incontinence (7 days), depression (15 days), and epilepsy (2 days). These inference modules are then simply wired into the framework. The time it takes to implement sensing modules also varies depending on their complexity (e.g., we developed a bed sensor module that took a total of 300 lines of Java and C code).

3. EVALUATION

The evaluation of Empath2 consists of instantiating and deploying in homes with real patients and collecting data from three very different home health care applications. Our main goals are to demonstrate the flexibility of the framework and ease of instantiation. We also show examples of the collected data which was produced as specified by our collaborating medical professionals. This data is currently being used by them to obtain preliminary information on researching these medical conditions. It is not our intent to show how the system solves medical problems. Such a result would require large scale clinical trials which are beyond the scope of this work.

3.1 Incontinence and Dementia

Alzheimer's disease is the most prevalent type of dementia in the US, affecting 4-6 million people, and is estimated to expand to 17 million by 2050. Alzheimer's frequently presents with episodes of nighttime agitation, which are highly burdensome and costly for caregivers and for the health care system. Additionally, people with dementia such as Alzheimer's are also much more prone to urinary incontinence than others. There is anecdotal evidence that there is a relationship among incontinence, sleep disturbance, and agitation in these patients. Urinary incontinence is thought to trigger awakening from sleep, with subsequent agitation, although there is a lack of systematic evidence that these phenomena are related. Based on previous work with urinary incontinence and sleep in persons with Alzheimer's disease, this study with real patients aims to describe the relationships among the times of occurrence of nighttime agitation, sleep continuity and duration, and urinary incontinence in persons with Alzheimer's disease by using innovative, non-invasive sensing technology.

Instantiation: As shown in Figure 2, to instantiate a system for this incontinence study, we simply create Empath2 with instances of an accelerometer-based sleep monitoring module, an activity monitoring module that consists

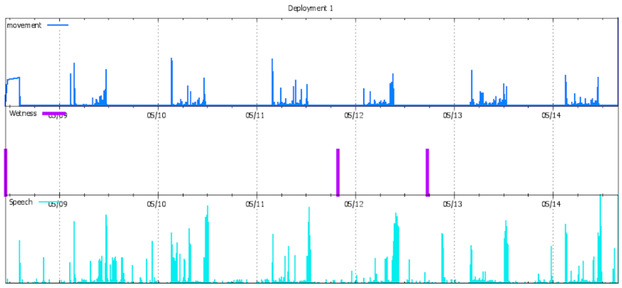

Figure 3: Incontinence study over 6 nights, investigating relationship among movements (top), wetness episodes (middle), and vocal outbursts (bottom).

of a wetness sensor only, and an acoustic based speech module.

The primary objective of this study is the ability to detect when a wetness episode occurs. The DryBuddy is a small, off-the-shelf, lightweight, and wireless sensor that uses a magnetic locking system to keep the sensor in place on the outside of the undergarment or pajama. It uses electrical conductivity within an incontinence pad to determine if a wetness event has occurred. We monitor audible speech outbursts by way of a microphone. The mechanisms presented in the previous section permit easy linking of these sensors into the system. In total, sensing aspects of the instantiation took only a few hours and the specific inference and higher level processing (which can not be general) took 7 days.

To date, this on-going project has produced 12 deployments with actual ill patients. Figure 3 shows a summary of the data for our first deployment. The first wetness event shown in the figure was a test before deploying the system to ensure that the X10 receiver was placed within the range of the DryBuddy. During the deployment, two wetness events were detected. From the data, the analysis to find the relationship among incontinence, sleep disturbance, and agitation is ongoing and done by another research group consisting of nurses and statisticians and that analysis is not within the scope of this paper.

The results of all the deployments are summarized in Table 2. For lack of space, we do not show the details of bed and audio sensor data from the deployments. Table 2 shows the age and sex of the patients of each deployment. The table shows that in initial deployments graduate students visited the patients' homes to set up all the devices. The process takes about 20 minutes. However, in four of the deployments, two nurses set up the devices themselves with the help of the instructions provided to them. Deployment by nurses takes extra time because they have to be in phone contact with the graduate students remotely monitoring the system to check that all the sensors are connected and reporting data. This shows the ease of deployment of Empath2.

3.2 Depression Monitoring

Depression is a major health issue that affects over 21 million American men and women each year. Depression often goes unrecognized and untreated, and even once treatment begins it is often difficult to monitor its effectiveness.

Table 2: Summary of incontinence study deployments

ID	Age	Sex	No. of Nights	No. of Nights Wetness Detected	No. of Nights Wetness Reported	Who Deployed	Deployment Time (Minutes)	No. of Intermediate Visits
1	65	Female	6	2	5	Grad. Student	20	0
2	87	Male	8	5	5	Grad. Student	20	0
3	71	Male	5	4	5	Grad. Student	20	0
4	73	Female	5	5	5	Nurse	30	0
5	79	Male	5	0	0	Nurse	30	1
6	88	Male	5	4	5	Nurse	30	1
7	89	Female	5	5	5	Nurse	30	0
8	95	Female	5	3	5	Grad. Student	20	1
9	102	Female	7	7	7	Grad. Student	15	1
10	76	Female	6	3	4	Grad. Student	25	0
11	78	Female	7	3	4	Grad. Student	20	2
12	79	Male	7	7	7	Grad. Student	15	1

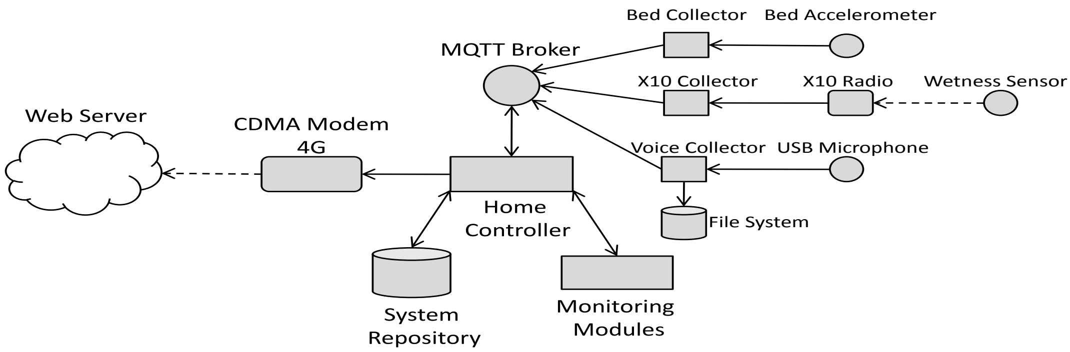

Figure 2: The incontinence study uses wetness sensor, bed accelerometer, and bedside microphone.

This poses particular challenges for the diagnosis and treatment of depression, particularly for those who avoid visiting a doctor or therapist due to social stigmas or a lack of energy. Currently, depression diagnosis is based on subjective screening questionnaires or structured clinical interviews that rely on timely in-person visits as well as accurate recollections by the patient, these have been shown to be inaccurate since symptom reports are often exaggerated or left incomplete. Yet early detection and treatment of this debilitating disorder has been shown to improve patient outcomes considerably [7, 15]. However, people are quite complicated and exhibit depression is different ways. This requires a multi-modal sensing approach.

As shown in Figure 4, the depression detection system we implemented monitors sleeping patterns, changes in behaviors, weight gain or loss, social interaction, and feedback from the person via a medically approved subjective questionnaires called the PHQ-9. The architecture of the system is the same as Figure 2 except the wetness sensor and with additional X10 sensors (as described below) reporting to the X10 receiver.

Instantiation: For the depression study, we instantiate Empath2 with an instance of an accelerometer-based sleep monitoring module, an instance of an activity monitoring module that consists of multiple X10 sensors, an instance of acoustic based speech module, an instance of the weight monitoring module, and an instance of the subjective mood scoring module. Again, the instantiation of the sensing modules took a few hours, and due to the complexity of this application the specific inference streams and high level logic took 15 days.

We implemented and deployed Empath2 for depression in two homes supporting all the modalities listed above. The first deployment was in an apartment over a period of 4 weeks. Although these results are not meant to investigate any medical hypotheses, it however shows an example of the system in operation and how it is able to collect useful data about a depressive episode continuously in the home. Adding these new modalities involved registering an X10 receiver to the broker to add activity recognition and a touch-screen client that supports the questionnaire. In addition, on the server, the weightscale API was configured to pull information from the webservice.

It took less than one hour to install Empath2 in the subject's home. X10 devices were attached to the stove, freezer, refrigerator, kitchen sink, microwave, spice cabinet, plate cabinet, glasses and cups cabinet, front door, cleaning closet, medicinal closet, bathroom sink, trash can, wardrobe closet, and shower. The weight scale was placed on the floor of the bathroom. A computer with the client software was placed in the living room. The total cost of the system excluding the laptop and phone is less than $500.

Each morning, the subject reported his subjective rating of the previous night's rest as being good or poor (this is

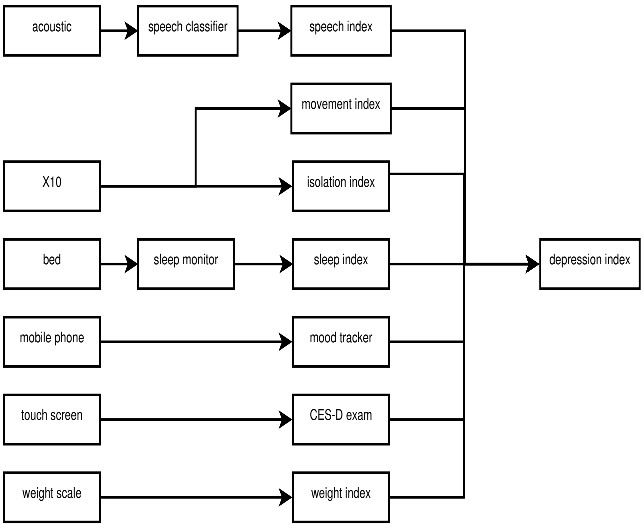

Figure 4: The inference system for depression risk index.

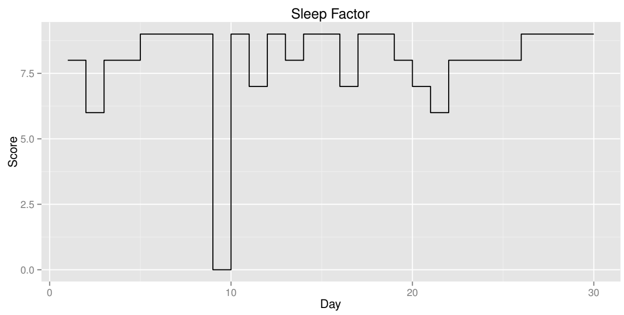

Figure 5: Daily sleep quality factors measured over a month.

used as ground truth). Figure 5 shows the sleep quality score calculated from the bed sensor readings for each night on a scale of 0 to 9, where a higher score represents a better quality of sleep. In the ground truth, the only night when the subject responded that his sleep was poor was on night 10, which appears to correlate with our sleep quality score. The graph suggests that for this subject, Empath2's sleep monitoring score can approximate sleep quality.

The second deployment for depression was in a single resident home; the resident was a 30-year old male. In addition to all the sensors and devices deployed in deployment 1, there were additional sensors to detect when the resident is watching the TV and when he is cooking by using the stove. Once again, it was very easy (less than 1 hour) to add these two new sensing modalities using Empath2 and the higher level inference took minimal time because it had already been created in the previous application.

We collected data from this deployment for six months continuously. All the objective measurements in Figure 4 were calculated and reported. For lack of space, we only show the distribution of the times of the day when resident eats; Figure 6 shows this distribution. Although, eating habits are not part of the inference modules shown in Figure 4, one can easily test if there is any correlation of the resident's eating habits with his mode from the data presented in Figure 6.

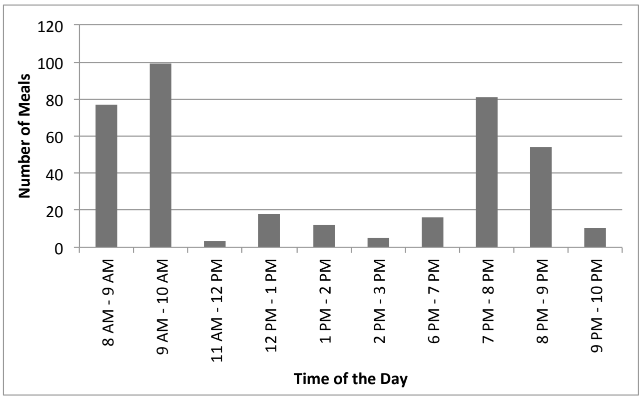

Figure 6: Distribution of the times of day when the resident eats over the six month for Deployment 2.

3.3 Epilepsy and Stress

The most commonly identified precipitant of seizures in people with epilepsy is stress, and the most common comorbidities associated with epilepsy are mood disorders. Research shows that important and complex linkages exist between stress, sleep, and epilepsy. Seizures themselves can also disrupt sleep. Therefore, an understanding of the relationships between stress and sleep is important in seizure control and in improving sleep for patients with epilepsy. In this study, we investigated using an ancient Chinese healing art based in Eastern philosophy, internal qigong, referred to as Reflective Exercise, as an alternative therapy intended to reduce stress, and thus improve the sleep quality.

To measure sleep quality and disturbances, we instrumented the bed with two tri-axis accelerometers (sampling at 50 Hz) by placing them on the right and left sides of the bed. Because the accelerometers are placed just beneath the mattress pad, we are able to detect fine movements while the patient is unaware of their presence. The movement levels for one minute epochs are calculated by computing the mean and standard deviation of the samples. We configured Empath2 for this data collection application. The data then can be accessed through a web interface for the clinical researchers or physicians to investigate the patients condition.

Instantiation: For the epilepsy and stress study, we instantiate XYZ with only an instance of an accelerometer-based sleep monitoring module. Since this sensor modality is already in the library it took almost no time to instantiate the sensing aspects of this application and 2 days for the specialized logic.

We deployed the system in two homes. The first participant was a 19-year-old female who had experienced seizures since adolescence. The second participant was a female of 28-year-old who also started her seizures as a teenager. The system tracked participant's movement on the bed. Figure 7 shows the average movements during sleep for one of the patient for seven days when the patient practiced the internal qigong exercise. Result shows that the therapy improves the sleep quality, and thus the epilepsy condition of the patient.

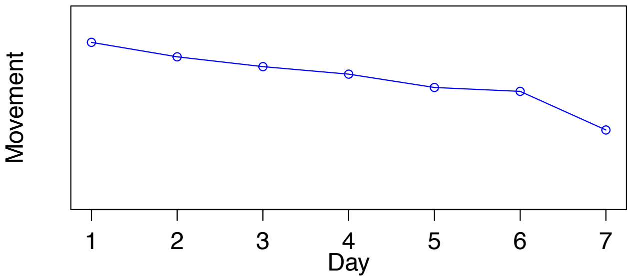

Figure 7: Internal qigong therapy reduces bed movement of epilepsy patient during sleep.

4. LESSONS LEARNED

The Empath2 architecture was refined through multiple real home deployments with real patients for three very different home medical applications. In that process, we discovered many practicalities that must be addressed in any home health care architecture. The lessons learned include the following:

- Deployment time must be as short as possible. The architecture must have the capability to quickly discover the emplaced sensor nodes, activate the system, and importantly test that the installation is fully operational in an end-to-end manner, i.e., from sensors to cloud to users. In our case studies, Empath2 was deployable usually within an hour.
- Deployments must be installable by non-technical experts, e.g., contracted system installers or home caregivers. In our case studies, nurse aides were able to perform the installations.
- Mobile broadband Internet connections are not always available when needed, and the connection can drop while in use. This occurred more often than we expected, most especially in homes in more remote locations. Consequently, the system allows for a local backup (such as a SQLite3 DB on the base station) and a background daemon to reliably synchronize and upload new data when a connection is made.
- The architecture must support frequent monitoring of the correct operation of the system, to ensure that the sensors are working, recording and transmitting properly, and that the data itself appears reasonable. Empath2 provides self-monitoring support at points throughout the architecture (in both the cloud and base station levels).
- The system must be resilient and cognizant to actions of patients and caregivers over the lifetime of the system. Minimizing assumptions required for the system to operate properly is critical. Will the system still work if a sensor or device is turned off or moved from its expected location? What happens when there is a power outage in the area? At a minimum, redundancy of critical components is required and any affected or missing data should be tagged with meta data explaining the situation.

5. STATE OF THE ART

In the first decade of this century, there were various home healthcare systems that are summarized well in a survey paper [12]. Significant research in this area has occurred in university research testbeds. Georgia Tech's *AwareHome* [11] combined context-aware and ubiquitous sensing, computer vision-based monitoring, and acoustic tracking of people. The University of Rochester built their *Smart Medical Home* which is a five-person house outfitted with infrared sensors, biosensors, and video cameras for use by research teams. Other work includes PlaceLab [10], research at Univ. of Florida [9], Univ. of Texas at Arlington[1, 20], *SmartAssist* [2] and HealthOS [13].

Further, to date, there have been an increasing number of companies that have begun to sell their systems directly to people who want to monitor either themselves or their loved ones. Philips provides Lifeline with Auto Alert for elderly people. Also, Philips provided Telehealth solutions where a patient takes their vital signs and answers personalized, clinician-directed surveys and their results are send to a website via a landline or cellular signal. Intel-GE has developed QuietCare that uses advanced motion sensor technology that learns the daily activity patterns of residents and sends alerts to help caregivers respond to potentially urgent situations and major routine changes. Intel also has developed Intel home health Gateway where patients take their vital sign measurements as defined in their care plan and System monitor their health status under the guidance of a healthcare professional. Cisco HealthPresence can improve healthcare between patients, clinicians, and specialists located in distant places. As far as we know all these above systems have not explicitly focused on flexible architectures and evaluated that flexiblity across a widely different medical monitoring applications.

Related work also appears in web-based sensor networks. A recent survey describes web-based sensor networks [17]. The SAPHE health care system for the Sensor Web [4] is able to sense different physiological attributes such as blood pressure, temperature and send them to the web for a specialist or doctor. Another system was developed especially to monitor Parkinson's disease [21].

In addition to systems focused on a single medical issue or on general monitoring of ADLs, there are a number of works that develop more flexible architectures. In [14] an extensible architecture is created but it only focuses on tracking of patients. In [19] GiraffPlus is created as an infrastructure for home care monitoring and used in real homes. However, this architecture concentrates on middleware and sensing layers and to date no evaluations of it can be instantiated for very different home health care applications have been provided. This systems was extended to mobile middleware in GP-m [18]. The paper [22] also discusses architectures for smart home applications but does not evaluate any of them in the context of instantiations of those architectures for different medical conditions. Other work, such as [8] describes an ecosystem for programming flexible assistive environments. Since this work concentrates on the programming paradigm it is orthogonal to our work.

6. CONCLUSIONS

Most commercial home monitoring systems are either focused on single devices or are single purpose. Some support general monitoring, but are not easily extensible. Research systems have focused on creating a system for a specific purpose and are rarely, if ever, tested with actual patients. The Empath2 architecture serves as a basis for easily instantiating versions of it for home health care for a wide variety of purposes. The claim is substantiated by creating, deploying, and evaluating three instances of Empath2. All deployments are in real homes (a total of 17 homes to date). Importantly, 2 of the 3 deployments are with actual ill patients. It is also important to note that the deployments vary from relatively simple ones such as the epilepsy case study, to complicated ones such as the incontinence and depression case studies. In particular, in the case studies we show that using the framework we can quickly implement and deploy multimodal, largely passive behavioral monitoring systems that are useful to caregivers and medical professionals.

7. ACKNOWLEDGEMENTS

This work was funded, in part, by NSF Grant CNS-1319302.

8. REFERENCES

[1] E. Becker, Y. Xu, S. Ledford, and F. Makedon. A wireless sensor network architecture and its application in an assistive environment. In *Proceedings of the 1st international conference on PErvasive Technologies Related to Assistive Environments*, page 25. ACM, 2008.

[2] D. Burmeister, A. Schrader, and D. Carlson. A Modular Framework for Ambient Health Monitoring. *Proceedings of the ICTs for improving Patients Rehabilitation Research Techniques*, pages 3–6, 2013.

[3] M. E. Chernew, R. a. Hirth, and D. M. Cutler. Increased spending on health care: long-term implications for the nation. *Health affairs (Project Hope)*, 28(5):1253–5, 2009.

[4] G. E. Churcher, A. Park, M. Heath, I. Ip, and J. Foley. Applying and Extending Sensor Web Enablement to a Telecare Sensor Network Architecture. In *COMSWARE*, pages 6–11, 2009.

[5] R. Dickerson, J. Lu, J. Lu, and K. Whitehouse. Stream feeds: An abstraction for the world wide sensor web. In *Proceedings of the 1st International Conference on The Internet of Things*, IOT'08, pages 360–375, 2008.

[6] R. T. Fielding. *Architectural Styles and the Design of Network-based Software Architectures*. PhD thesis, University of California Irvine, 2000.

[7] A. Gilbody, Simon and Sheldon, Trevor and House. Screening and Case-finding instruments for Depression: A Meta-Analysis. *Canadian Medical Association*, 2008.

[8] S. Helal, C. Chen, E. Kim, R. Bose, and C. Lee. Toward an ecosystem for developing and programming assistive environments. *Proceedings of the IEEE*, 100(8):2489–2504, 2012.

[9] S. Helal, W. Mann, H. El-Zabadani, J. King, Y. Kaddoura, and E. Jansen. The Gator Tech smart house: a programmable pervasive space. *IEEE Computer*, 2005.

[10] S. S. Intille, K. Larson, J. S. Beaudin, J. Nawyn, E. M. Tapia, and P. Kaushik. A living laboratory for the design and evaluation of ubiquitous computing technologies. *CHI '05 extended abstracts on Human factors in computing systems*, page 1941, 2005.

[11] J. a. Kientz, S. N. Patel, B. Jones, E. Price, E. D. Mynatt, and G. D. Abowd. The Georgia Tech aware home. *Proceeding of the twenty-sixth annual CHI conference extended abstracts on Human factors in computing systems - CHI '08*, page 3675, 2008.

[12] J. Ko, C. Lu, M. Srivastava, J. Stankovic, A. Terzis, and M. Welsh. Wireless sensor networks for healthcare. *Proceedings of the IEEE*, 98(11):1947–1960, 2010.

[13] J. H. Lim, A. Zhan, E. Goldschmidt, J. Ko, M. Chang, and A. Terzis. HealthOS : A Platform for Pervasive Health Applications. In *ACM mHealthSys*, Toronto, ON, Canada, 2012.

[14] C.-H. Lu, C.-L. Wu, and L.-C. Fu. A reciprocal and extensible architecture for multiple-target tracking in a smart home. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 41(1):120–129, 2011.

[15] R. Mojtabai. Does depression screening have an effect on the diagnosis and treatment of mood disorders in general medical settings?: an instrumental variable analysis of the national ambulatory medical care survey. *Medical care research and review : MCRR*, 68(4):462–89, Aug. 2011.

[16] MQTT. MQTT V3.1 Protocol Specification. 2010.

[17] K. A. Nuaimi, M. A. Nuaimi, N. Mohamed, I. Jawhar, and K. Shuaib. Web-based wireless sensor networks : a survey of architectures and applications. In *Ubiquitous Information Management and Communication (ICUIMC)*, Kuala Lumpur, Malaysia, 2012.

[18] F. Palumbo, D. La Rosa, and S. Chessa. Gp-m: Mobile middleware infrastructure for ambient assisted living. In *Computers and Communication (ISCC), 2014 IEEE Symposium on*, volume Workshops, pages 1–6, June 2014.

[19] F. Palumbo, J. Ullberg, A. Štimec, F. Furfari, L. Karlsson, and S. Coradeschi. Sensor network infrastructure for a home care monitoring system. *Sensors*, 14(3):3833–3860, 2014.

[20] A. Papangelis, G. Galatas, and F. Makedon. A recommender system for assistive environments. In *Proceedings of the 4th International Conference on PErvasive Technologies Related to Assistive Environments*, page 6. ACM, 2011.

[21] S. Patel, B.-r. Chen, T. Buckley, R. Rednic, D. Mcclure, D. Tarsy, L. Shih, J. Dy, M. Welsh, and P. Bonato. Home monitoring of patients with Parkinson's disease via wearable technology and a web-based application. In *IEEE EMBS*, volume 02139, pages 4411–4414, 2010.

[22] F. Viani, F. Robol, A. Polo, P. Rocca, G. Oliveri, and A. Massa. Wireless architectures for heterogeneous sensing in smart home applications: Concepts and real implementation. *Proceedings of the IEEE*, 101(11):2381–2396, Nov 2013.