Sensing for Obesity: KNOWME Implementation and Lessons for an Architect

Sangwon Lee[1], Murali Annavaram[1], Gautam Thatte[1], Vikor Rozgic[1], Ming Li[1], Urbashi Mitra[1], Shri Narayanan[1], Adar Kravitz[2] and Donna Spruijt-Metz[2]
[1]Ming Hsieh Department of Electrical Engineering, [2]Keck School of Medicine
University of Southern California*

Abstract

The use of bio sensors for monitoring an individual's health continuously and in real time promises to revolutionize health care in near future. This paper focuses on using a biometric sensor platform called KNOWME to assess pediatric obesity and its interdependence on social, environmental and physiological parameters. The KNOWME platform integrates off-the-shelf sensors with a Nokia N95 mobile phone to continuously monitor the vital signs of a subject and correlate these signs with geo-spatial information. Based on the development and in field deployment studies this paper enumerates several hardware design impediments that need to be overcome for the broader applicability of KNOWME. Using some preliminary data, this paper motivates the need for energy efficient sensor management and hardware supported data privacy for future KNOWME implementations.

1. Role of Mobile Sensing In Obesity

The alarming rise in pediatric obesity rates is an issue that needs to be dealt with using a multi-pronged approach. While it should be obvious that physical activity (PA) plays an important role in reducing obesity rates, it has been observed that PA declines precipitously during early adolescence. While research on the medical front is pushing the boundaries of our understanding of the genetic and biological reasons for obesity, the social and environmental impact on PA are not well understood [5][4]. Currently the primary impediment for understanding these concerns is the lack of mechanisms to non-intrusively measure PA in real-time and real-space. It is in this context, we believe, that a combination of health sensors and mobile phones will play a significant role in improving our understanding. The ability to record and interpret PA continuously with minimal intervention from the subject is the first step in understanding the complex interplay of environmental, genetic and socio-economic reasons for weight gain[2][3].

At USC, we have been working on a mobile health platform, called KNOWME [1], to develop low cost technology centric tools for pediatric obesity prevention and treatment. KNOWME interfaces wireless sensors with mobile phones to precisely monitor behavior, aspects of metabolism and geo-spatial information. In the current implementation, KNOWME uses off-the-shelf sensors to measure physical activity, blood pressure, sleep, heart rate, galvanic skin response and blood glucose levels and communicate the measured information to a mobile phone using Bluetooth interface. We combine the external sensor data with the mobile phone's in-built sensors, in particular Geographic Positioning Systems (GPS), accelerometer, audio and video tags. We then exploit the communication capabilities of mobile phone to transmit the combined health record to a data server in real time. KNOWME also exploits the increasing compute capabilities of mobile phones to process the data collected from sensors to perform complex signal processing operations to detect anomalies in user state. In addition, to reduce battery consumption the mobile phone runs context-based sensor scheduling algorithms to activate only the minimum set of sensors to collect only the necessary sensor samples to detect user physical activity [8]. KNOWME is currently being deployed in a user study with about 20 subjects in the Los Angeles area.

In this paper, we present the KNOWME system overview and describe current implementation. We will describe several challenges and lessons learned during our current implementation. The primary intent of this paper is to present research challenges for hardware designers in the area of body sensor networks based on observations from our current implementation.

The rest of this paper is organized as follows. Section 2 describes the KNOWME system and how various sensors application modules interact. Based on the observations from development and deployment efforts Section 3 discusses three important impediments that need to be addressed by hardware designers to make KNOWME broadly applicable. We then conclude on Section 4.

2. KNOWME System Overview

KNOWME is built using three-tier architecture as shown in Figure 1. The first tier is the sensor tier integrated with a mobile device that forms the primary data collection layer. The second tier is a web server tier that receives the data from sensor tier and performs authentication and data scrubbing. The third tier is a data storage layer that stores encrypted data received from the web server. Since most of the hardware complexity in KNOWME is in the sensor tier in this section we describe this tier in detail.

* This research is supported by NIH, Nokia and Qualcomm.

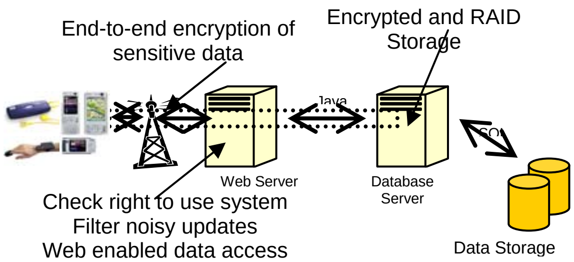

Figure 1 Three-Tier KNOWME Architecture

2.1 Sensor Tier and Mobile Application

Each data collection node in KNOWME consists of a Nokia N95 mobile phone and four health and ambient sensors, namely accelerometer (ACC), electrocardiograph (ECG) monitor, Pulse Oximeter (OXI), and GPS. We use the built-in GPS and accelerometer sensors on N95. The ECG and OXI sensors are Bluetooth enabled off the shelf sensors from Alive Technologies. The ECG and OXI sensors are capable of sampling either at 300 Hz or 75Hz and the data can be either transmitted over the Bluetooth channel or can be stored locally on flash memory.

A data collection and visualization application is developed using J2ME and PyS60 that runs on N95. The goal of this mobile application is to gather data from multiple sensors with minimal intervention from the user and with no interruption to the mobile device functionality. Note that since the data collection task is anticipated to run for 12-14 hours per day for multiple weeks in our field study, the application must be robust and continue to operate even under unanticipated conditions. Hence, we divided the mobile application into two components: a server component for data collection and a client interface application for configuring the sensors and data visualization. Although the N95 supports several SDKs, we chose the S60 Python for the server side and the J2ME for the client side. Since these languages are not restricted by platform, it will be easy to port the application for other mobile phones in future.

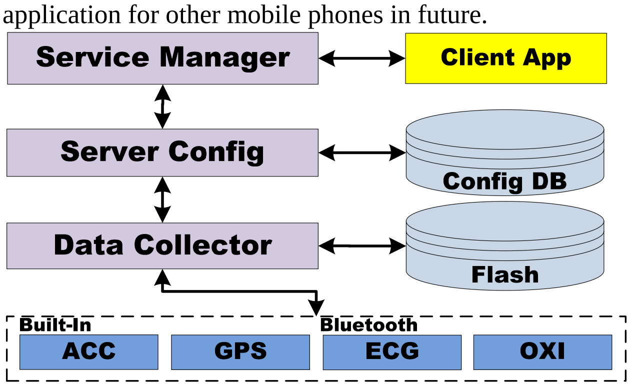

Figure 2 Mobile Application Interactions

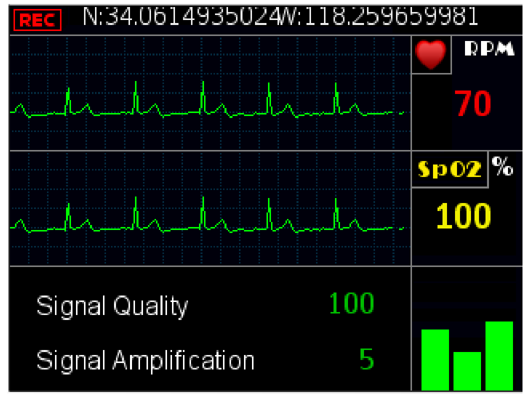

Figure 3 Mobile Visualization Client Screenshot

The server application comprises of four tasks: server configuration manager, data collector, device manager and service manager. Figure 2 shows the various components in the server application interact with the client. During the initialization phase of KNOWME the server reads the sensor configuration database. This database stores user-specific information such as user id, data encryption keys, and sensor information such as a list of available sensors, the MAC addresses of sensors, mechanisms for controlling the sampling rate. After reading the configuration database, the server creates a collector thread for managing the sensor data collection and storage. It is the responsibility of the collector thread to receive data from multiple sensor threads and store the data in the mobile phone flash storage as well as provide the data to any other requestor. The collector thread creates one device manager thread per each sensor. The device manager thread reads data from its corresponding sensor continually and in real time. The device manager thread deals with the vagaries of medical sensors, such as lost connection, noisy data. After receiving the data from the attached sensors, it buffers the data locally and sends the buffered data to the collector thread at regular intervals. Since the collector thread receives data from multiple device manager threads it is imperative that the each device manager buffer and send a coarse data packet infrequently to the collector. The last task of the server application is to service data requests from clients. The server creates a single service thread that listens to client connections on a local communication socket. On receiving a client data request it gets the data from the collector thread and sends it to the client. In addition it also receives sensor configuration data from the client whenever a new sensor is connected to the KNOWME platform.

The client component is designed to provide a user interface because the server cannot be user interactive due to real time constraints on sensor data collection. It communicates with the server using local sockets. Every new sensor that is added to KNOWME is first configured through the client component. For instance, during the addition of a new sensor to KNOWME the client application is used to scan for a new Bluetooth device and find the Bluetooth MAC address. These values are then communicated to the service thread within the server manager which will then store these values in the device database which will be retrieved during the server initialization and configuration, as described earlier. Figure 3 shows the screen shot of the client visualization software.

3. New Hardware Challenges for KNOWME

While the current KNOWME platform uses off-the-shelf components there are several circuit and architecture research challenges that need to be addressed to improve the broader usability of KNOWME. In this section we will elaborate on some of these challenges.

3.1 Sensor Energy Management

The first and foremost challenge in body area sensing is improving energy efficiency of sensors. Figure 4 shows the battery level of the mobile phone as we turn on various sensors used in KNOWME. The X-axis on the figure is length of time measured in minutes. The curve labeled ALL_UnBuffer shows the battery level on the mobile phone as we collect data from all the sensors (ECG, OXI, ACC, GPS) and write them to the local flash drive without any buffering. The curve labeled ALL_Buffer shows the battery level if we buffer the writes to the flash drive and send large packets to write to the flash. As can be seen buffering improves the battery life from 240 minutes to 299 minutes, which is a 25% improvement. The curve labeled Only_BT shows battery life when using only the Bluetooth enabled sensors (ECG, OXI). This graph shows the energy cost of using Bluetooth alone for continuous communication. Finally, the graph labeled Only_Int shows the battery drain when using N95 internal sensors only (ACC, GPS). Most of the energy consumption in the internal sensors is due to GPS. A single GPS reading typically consumes 6.616 joules whereas an accelerometer reading costs only 0.359 joules. Note that the manufactured rated battery life of N95 in standby mode (i.e. no communication or computation) is ~200 hours. Using the mobile phone for health sensing dramatically reduces the battery life in all cases.

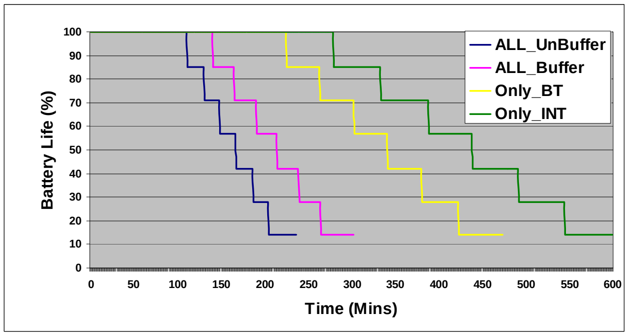

Figure 4 Battery Usage of KNOWME

These results show that the traditional focus in computer architecture to reduce the computational energy cost will have little impact on the usability of mobile devices in body area sensor networks. Rather there is a pressing need to reduce the communication energy cost. New wireless communication standards such as Zigbee and MiWi are focusing on low power implementation. There is a significant cost for supporting a plethora of existing wireless standards and hence supporting new standards on a mobile device will be an undesirable solution. Rather, reusing an existing radio to dynamically configure a wireless standard will be an interesting research direction. As such exploring efficient implementation of software defined radios on a mobile CPU will provide viable solutions.

3.2 Data Privacy and Security Hardware

As the popularity of mobile devices in medical information collection and management increases there is a growing realization that information collected about an individual user can compromise one's privacy and potentially security. Thus data privacy must be provided by design rather than as a policy measure. Existing mechanisms for data privacy are based on the notion of k-anonymity in databases [6]. Traditionally k-anonymity is implemented on databases where the data changes relatively slowly. However, in body sensor networks data is collected continuously and rapidly. Hence, providing k-anonymity in real time for continuously changing data with limited computational resources is a challenging task. For instance, the best known runtime complexity of computing k-anonymity is exponential in k, assuming no data is modified. We believe that hardware assisted low power data sorting algorithms and sensitive data suppression are areas of future research that will transform the usability of body sensor networks.

On the data security front, there are two levels of security that are needed for KNOWME. First, data being transmitted from the mobile device to the data server need to be encrypted during transmission. For the data transmission, we believe there is a need to tradeoff data vulnerability during transmission with encryption costs. Traditional approaches to data security such as PGP or AES encryption are power intensive. Note that once the data is received by the server, strong encryption mechanisms such as PGP or AES can be implemented on the server side. Hence, hardware supported light weight encryption mechanisms will be more ideal for data transmission security. The second level of data security must be provided for the data that is stored on the local flash drive. Strong data security must be provided for the local flash drive storage. However, the energy cost of such strong security may be acceptable since the encryption process can be opportunistically scheduled to a time when the mobile device is being charged. Note that encryption of flash data does not need to be done in real time as the vulnerability factor is much lower for locally stored data.

3.3 Sensor Data Collection Quality

Many of the current bio-sensors use leads and contact probes to measure vital signs. The quality of data collected from body sensing is highly dependent on the placement of these leads on the body. These leads are both cumbersome and potentially intrusive to person's daily activities. Detecting these vital signs using contact free sensors is an area of immense potential. There are several possible approaches for such contact free sensing. Existing research in using electromagnetic and/or acoustic wireless sensors can improve the usability of bio sensors. However, very little is known about the computational costs and the hardware circuit implementation feasibility of these sensors. In particular, implementing radar/sonar based technologies on CMOS is a difficult challenge and it is only recently researchers have shown the viability of implementing radar on CMOS technology. The architecture of the radar transceiver for sensors is an area of research that is virtually non existent.

Sensors	# packets per hour	Packet loss (%)
ECG & ECG	14999/14998	0/0
Oximeter & ECG	36008/14931	0/0.0045
Oximeter & ECG & ECG	36007/7351/7312	0/50.9/51.2

Table 1 Data Loss on Bluetooth Channel

In our current implementation, we used a Bluetooth interface to communicate between the mobile device and bio sensors. The external sensors continuously stream their data to the mobile phone. Due to limitations in buffering on the mobile phone receiver, the mobile phone occasionally can not keep up with the data reception rate. Table 1 shows the data packets lost per hour due to this limitation. The data loss rate is relatively high in our current implementation. One way to reduce the data loss is to adaptively reduce the sampling rate of the external sensors [7]. Currently the sampling rate of external sensors can not be automatically controlled. Programmability of the external sensor sampling rate is a necessary feature for future platforms to more effectively manage the data loss without compromising accuracy.

Bluetooth allows a maximum of six connections thereby limiting the number of external sensors in KNOWME to a maximum of six. While this limit is acceptable in the near term we plan to add additional bio sensors such as Galvanic Skin Response unit and external accelerometers. These additions will be limited by the Bluetooth restrictions. There is a fundamental tradeoff between the number of concurrent connections that are allowed in a wireless protocol and the complexity and power consumption of that radio. Hence, it is more compelling to merge multiple bio sensors into a single sensor module and send multiple sensed signal values on a single channel. Integration also can benefit from the sharing of resources in hardware such as transmission buffers, error correction and decoding logic.

4. Conclusions and Future Directions

This paper describes KNOWME, which is an integrated bio sensor and mobile phone platform for observing the social and environmental impacts on obesity management. While the current platform is built using off--the-shelf sensor and mobile phone components, there are several impediments in the current design. These impediments can broadly be classified into three categories: sensor energy consumption, data privacy and data quality. There is a pressing need to develop new hardware designs to improve the usability of bio-sensors in medical information technology. The paper presented a subset of these research challenges in the context of the lessons learned from the KNOWME implementation.

5. References

[1] M. Annavaram, N. Medvidovic, U. Mitra, S. Narayanan, G. Sukhatme, Z. Meng, S. Qiu, R. Kumar, G. Thatte, and D. Spruijt-Metz. Multimodal sensing for pediatric obesity applications. In *Proceedings of UrbanSense08*, pages 21-25, November 2008.

[2] P. Leijdekkers, Gay. Personal Heart Monitoring System Using Smart Phones To Detect Life Threatening Arrhythmias. In *the 19th International Symposium on Computer-Based Medical Systems*, pages. 157-164, 2006.

[3] J.A. Levine, S.J. Schleusner M.D. Jensen. Energy expenditure of nonexercise activity. *American Journal of Clinical Nutrition*. 2000;72(6):1451-1454.

[4] T. McKenzie, J. Sallis P. Nader. SOFIT: System for Observing Fitness Instruction Time. *Journal of Teaching in Physical Education*. 1991;11:195-205.

[5] D. Spruijt-Metz, B. Saelens. Behavioral aspects of physical activity in childhood and adolescence . In: M. Goran, M. Southern eds. Handbook of Pediatric Obesity: Etiology, Pathophysiology and Prevention. Boca Raton, FL. Taylor & Francis Books/CRC Press; 2005:227-250.

[6] L. Sweeney. k-anonymity: a model for protecting privacy,. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, v.10 n.5, p.557-570, Oct 2002.

[7] G. Thatte, V. Rozgic, M. Li, S. Ghosh, U. Mitra, S. Narayanan, M. Annavaram, and D. Spruijt-Metz. Optimal allocation of time-resources for multi hypothesis activity-level detection. To appear *in the proceedings of the 5th IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS'09)*, June 2009.

[8] Y. Wang, J. Lin, M. Annavaram, Q. A. Jacobson, J. Hong, B. Krishnamachari and N. Sadeh. A Framework of Energy Efficient Mobile Sensing for Automatic Human State Recognition. To appear in *proceedings of 7th Annual International Conference on Mobile Systems, Applications and Services (MobiSys)*, June 2009.