

A Long-Term Sensory Logging Device for Subject Monitoring

Dawud Gordon, Florian Witt, Hedda Schmidtke and Michael Beigl

Distributed and Ubiquitous Systems
Technische Universität Braunschweig,
Braunschweig, Germany

Email: {gordon, fwitt, schmidtke, beigl}@ibr.cs.tu-bs.de

Abstract—In this paper we introduce a device for monitoring subjects by logging sensory data over long periods of time. The system consists of a sensory measurement unit, a memory unit and an application for data preprocessing tasks, such as converting sensory measurements to desired units or calculating averages. In order to demonstrate usage, an activity level monitoring system inspired by medical applications is implemented using the device. A rudimentary threshold-based embedded classifier is trained using several different activities, yielding an activity level indicator. 2 subjects are monitored to train the classifier, and the system is then evaluated on new data using those two subjects plus a third not involved in the training process. The results indicate an activity classification of 74% into three levels using 2 simple data thresholds, with a system lifetime of 26 days on 2 AAA batteries.

I. INTRODUCTION

In health-based medical monitoring applications, having to involve medical personnel in the monitoring process is costly and intrusive, hindering normal user activity. In order to reduce costs and improve quality of life of the subjects, a process for unsupervised monitoring of patients in their everyday environments is needed. Inspired by this scenario, we have developed a low-power, high sample-rate sensory device which can be worn by subjects over long periods of time.

The system consists of a sensory measurement unit, a memory unit and a monitoring application. The sensor board has 3 sensors, a light sensor, a temperature sensor and a vibrational sensor, with an interface for an external acceleration sensor if required. The sensor unit is also equipped with an RF communication module which has been disabled for the purpose of this application. The MMU provides a FAT32 interface between the sensor board and a microSD card used as a data sink. The software application periodically samples the sensor, performs any preprocessing necessary and passes the data to the MMU to be logged.

A simple activity level monitoring application was constructed in order to evaluate the system. The application consists of a rudimentary threshold-based classifier that provides an activity level given a vibrational intensity. Low activity levels have been shown to be a risk factor for many illnesses, and higher levels have been indicated to prevent the onset of sickness [4], [5], [9]. Activity levels have been correlated with overall disease risks as indicated in [4], spiraling obesity in [5], as well as osteoporosis and muscle atrophy in [9], demonstrating

the need for monitoring the activity levels of individuals over longer periods of time. Inspired by these phenomena, we decided to test this device as a platform to allow medical experts to develop activity level monitoring applications.

The experiment is done as follows: three mutually exclusive activity levels are defined and example activities are selected in order to evaluate the classification. 2 subjects are monitored to gather an initial dataset which is then used to train the classifier by adjusting the threshold levels. The classifier is adjusted to differentiate between the different activity levels, yielding a quantized (as opposed to analog) activity level indicator. The system is then evaluated using 3 subjects in terms of recognition and power consumption levels. Finally the results and their implications are discussed.

II. RELATED WORK

Activity level is a term used to describe the amount of energy expended per time period (e.g. Joules per week) as is indicated in [4]. In the medical literature, this value is often measured using the doubly-labeled water technique to ascertain total energy expenditure of a test subject [5]. Some other approaches to ascertaining a subject's activity level are personal estimation as in [4] or using accelerometers as in [6].

Devices for logging acceleration data, such as in [1], do not provide customizable data preprocessing, but rely on external processes to interpret the data. The approach in [6] is based on the use of accelerometers and generates accurate measurements of energy expenditure, but also requires a separate system to interpret the data. Activity recognition applications based on small low-power nodes using accelerometers such as in [8] also use a back end classification system to analyze the data, placing unnecessary location limitations on the subject. Other embedded classification approaches such as [3] and [7] are concentrated on context recognition and transmission as opposed to the logging aspect as presented here.

In this paper, we opt for a long-term approach to patient monitoring, detailed with an example using a vibrational sensor to gather data on the activity level of the subject being monitored. We attempt to combine the features in the aforementioned publications by taking advantage of the low-power properties of all of the components in the system.




Fig. 1. The Activity Level Logging Device

III. METHODS

The monitoring device consists of two separate modules: the sensor sampling and preprocessing system and the data storage system. The sensor board is responsible for sensor sampling and preprocessing and a separate board acts as an MMU providing access to a microSD card. The sensor board is based on the PIC18F14K22 microprocessor and the memory board is implemented using a PIC32MX, both of which are available from Microchip¹. The device is pictured in 1, the sensor board is on top, followed by the memory board below, with the battery case at the bottom.

A. Sensory Data Acquisition

The sensor board firmware can be described as a multi-rate sampling and low-power data processing application, which can be programmed for adjustable-frequency sampling of the sensors. Data preprocessing is then limited by the power of the microprocessor as well as the effort incurred by the sampling rate. Both preprocessing tasks and sensor sampling share the same processing resources, therefore they both affect power consumption rates as more work consumes more energy. Sampling rates of sensors can be implemented well over 100Hz depending on the amount of processing required.

For vibration measurements, we use a miniature ball switch sensor, in which a conductive sphere rolls between two charged plates, closing a circuit in a certain position. With a diameter of 0.8mm, the sphere's physical properties are different from those in traditional ball switches, especially in terms of sensitivity even at low-intensity vibrations, as well as in three dimensions [2]. A diagram of the sensor as well as the circuit used to sample it can be found in Figure 2.




Fig. 2. The Micro Vibration Sensor from [2] and Schematic

When agitated, the ball within the sensor rolls back and forth, generating digital output which is sampled by the micro controller. Although the output is in digital form (see Figure 3), the valuable information which is used is only unary, meaning only the transitions in the output is noted, and not the state in which the ball switch is in. The output of the sensor is connected to the counter input of the controller, using dedicated hardware to keep track of the number of transitions, also referred to as events, of the vibration sensor in a given period of time. A periodic sampling of the counter yields the number of events which occurred since the last measurement.




Fig. 3. Sensory Output for Typing (Top), Walking (Middle) and Jumping Rope (Bottom)

For light level measurements, the APDS 9003-020 photo sensor from Avago Technologies Ltd.² was used. The analog sensor delivers an indicator of light intensity levels across the entire spectrum recognizable by the human eye. For temperature measurements the TC1047A temperature-to-voltage converter from Microchip Technologies¹ was selected, with a range from -40°C to $+125^{\circ}\text{C}$ and an accuracy of $\pm 0.5^{\circ}\text{C}$.

The sensor board software periodically samples the sensors, performs preprocessing (e.g. converting temperature sensor readings into $^{\circ}\text{C}$), and then passes the information to the memory board for logging. The frequency of the measurements, as well as the preprocessing methods are application specific and should therefore be customized to fit the scenario. Both the light and temperature sensors are analog, meaning analog-to-digital conversions for each of the sensors are necessary.

IV. APPLICATION EXAMPLE: ACTIVITY LEVEL CLASSIFICATION

In order to create a basis for evaluation, a system to classify three different activity levels based on vibrational intensities was constructed. The implementation described here is extremely simplified and is not meant for actual monitoring applications, but rather is only used to demonstrate the limitations and capabilities of the monitoring device. The scenario is motivated by the medical literature in related work and our experiment illustrates principle technical feasibility, however an evaluation as a medical tool is clearly beyond the focus of this paper.

¹Data sheet available at www.microchip.com

²Data sheet available at www.avagotech.com

A. Activity Level Recognition

The activities used for testing and training were selected such that they approximately fell into three different categories of activity levels as can be seen in table I. Two subjects performed each of the three activities over the course of 30 seconds in order to generate a basis on which to train the classifier. The subjects were wearing the device attached at the hip, programmed with a sample rate of 1Hz. Table II displays 10 seconds worth of raw sensory data as read from the counter, where each row of the table constitutes one second of the specified activity.

TABLE I
ACTIVITY LEVELS, REPRESENTATIVE ACTIVITIES, IDENTIFIERS AND THRESHOLDS

ID	Level	Repr. Activity	Levels
1	Low	Typing	< 550
2	Medium	Walking	551 - 1508
3	High	Jump-Rope	≥ 1509

Using this data, thresholds were calculated to separate the data into three classes. This was done by calculating the mean number of ball switch events per user per event and selecting the midpoint between the upper value of one class and the lower value of the class above it. The resulting threshold values which were used to create the rule-based classifier are indicated in the third column of table I. Under real conditions and depending on medical indication and treatment plan, these thresholds would need to be determined by a medical professional. They are strongly application and subject dependent as will be discussed later.

TABLE II
EXAMPLE SENSOR DATA FOR THREE ACTIVITIES

Typing	Walking	Jump-Rope
0	648	1966
0	228	1266
0	594	2040
0	1000	2734
0	1188	1628
0	1444	1898
34	2172	1284
14	1506	1236
0	1484	1972
0	1262	1986

B. Data Logging

The memory unit of the device consists of a microprocessor and a microSD card which is connected to the activity recognition module via UART. The sensor node sends the ID of the activity level recognized to the memory board with a period of one second to be logged in a file on the SD card. The processor implements a FAT file system on the microSD card, allowing the user (in this case the medical care giver) to remove the card at any time to analyze the activity level log.

One issue which comes up when attempting stand-alone data logging with time stamps is accounting for clock drift. Fortunately, since we are attempting long term data logging,

this error can be corrected by a comparison between log time and system time when the microSD card is inserted into the physician's machine, allowing for a general estimate of clock drift over the monitoring period. This approach assumes that the time period from removal of the card to insertion into the evaluation machine can be estimated, and that the estimation error is small compared to the amount of clock drift. A more accurate method would be to enable the RF communicator located on the sensor board and create a beacon system which synchronizes the system clock when in range. This is outside the scope of this paper and would affect the power measurements presented here.

V. EVALUATION

The evaluation of the sensory monitoring and logging system is divided into two parts. First, the classification methods are evaluated using a short study. Second, the physical attributes of the system are evaluated in order to generate an estimate of the system lifetime.

A. Activity Level Classification

In order to evaluate the accuracy of the activity level classifier, 3 users were asked to perform the activities associated with the activity levels for 30 seconds per activity. These evaluation data are distinct from the training data, gathered under similar conditions at a different point in time. Table III indicates the results of the evaluation, where subject 3 was not included in the classifier training process. We evaluated correct or incorrect recognition of activity levels again with the representative activities for the respective activity levels (see Table I).

TABLE III
CLASSIFICATION RATES IN PERCENT BY USER AND ACTIVITY

	Subject 1	Subject 2	Subject 3
Low (Typing)	81	97	100
Medium (Walking)	72	69	86
High (Jump-Rope)	0	66	97

These results indicate very successful classification rates for typing at the computer, especially when considering that random probability distribution would equate to a classification rate of only 33.3%. Walking was also relatively well recognized as indicated by recognition rates of the medium activity level from 69% to 86%. Jumping-rope on the other hand, shows confusing results with respect to one subject, indicating recognition rates from 0% for subject 1 to 97% for subject 3. Overall, the system was able to classify the three activities correctly 74% of the time.

TABLE IV
CONFUSION MATRIX FOR CLASSIFIER OUTPUT IN PERCENT

	Low	Medium	High
Low (Typing)	95	3	2
Medium (Walking)	6	75	18
High (Jump-Rope)	2	43	55

The activity jumping rope shows the effects of interpersonal variations between the different subjects of the evaluation, as illustrated by the last row of Table III. The inability on the part of the classifier to produce the correct activity level for jumping rope (see Table IV) is due to these interpersonal variations and can be partially attributed to the overly simplified classifier design and training process. For an actual medical application, the activity levels and thresholds would obviously need to be fine-adjusted by the medical expert based on the medical indication and a patient's general health status.

B. System Lifetime Analysis

Two aspects must be observed when evaluating the maximum lifetime of this system: energy consumption and memory consumption. The energy consumption was measured while the activity level recognition and logging process was underway. In order to calculate how much energy was consumed, the amount of current pulled by the device was dynamically measured. The amount of energy consumed over the period of time beginning at t_0 and ending at t_1 is then given by

$$\int_{t_0}^{t_1} V_{supply} \times I_{supply} dt$$

The result of the analysis is an energy consumption of $4.478mJ/s$ or $4.478mW$. An average AAA battery has $5071J$ with an average voltage of $1.225V$ during discharge and an electric charge of $1150mAh$, yielding an estimated system lifetime of 26.1 days of continuous measurement. This is in part due to the low consumption of the sensor which has a rate of consumption of only $.5 \mu W$. These values are strongly influenced by the application, specifically the sample rate (in this case 1Hz) and the amount of preprocessing needed and will therefore vary from application to application. These rates are also dependent on the microSD card used, as the consumption values vary from twice to five times that of the microprocessor per byte written.

The capacity of the microSD card is 2GB or 2×10^9 Bytes. Each second requires one byte for storage of the activity level identifier; at 2GB capacity the application would be able to last for 63 years and 5 months before running out of memory. This clearly indicates that energy consumption, not memory, is the limiting factor in system lifetime.

VI. CONCLUSION AND FUTURE WORK

In this paper we presented an adjustable-rate sensory sampling and logging device with light-weight data processing capabilities. The device uses vibration, light and temperature sensors to take a snapshot of its environment at regular intervals. We presented the hardware architecture to support this process consisting of a sensor board for data acquisition and a memory board for logging. A software system was also presented, consisting of an embedded sensor sampling application with processing capabilities running on the sensor board and a FAT file system driver on the memory board. The activity level log is secured to a microSD card on the memory board to allow easy access for data extraction.

Inspired by a medical scenario calling for long term monitoring of patients as to the amount of activity performed, we created an example application which classifies subject activities in 3 different levels based on the amount of vibrations incurred. An evaluation of recognition structures was carried out, using data from two individuals to train the activity level classifier. The evaluation was carried out using separate data from the two previous subjects as well as a third individual not involved in the training process. Overall, an activity level recognition rate of 74% was achieved, illustrating the light-weight processing capabilities of the device. In addition to the recognition rate, the lifetime of the system as a logging device was calculated. Using 2 AAA batteries as a power supply, the system lifetime was calculated to be >26 days.

The results of the evaluation indicate that the device presented here can accurately acquire and log sensory data, as well as perform light-weight data processing operations on the embedded system. They also indicate that these tasks can run over a time frame of approximately a month without interruption, allowing for unsupervised monitoring of subjects.

Enabling the RF communication unit on the sensor board would provide support for the beacon synchronization application as discussed previously. This would solve the clock drift problem, but would create another as RF communication would negatively affect power consumption. For this reason it would then be necessary to adapt the application and re-evaluate the power consumption of the device.

ACKNOWLEDGEMENT

The authors would like to acknowledge funding by the European Commission for the ICT project CHOSen "Cooperative Hybrid Objects Sensor Networks" (Project number 224327, FP7-ICT-2007-2) within the 7th Framework Programme.

REFERENCES

- [1] Hobo pendant g data logger. <http://www.onsetcomp.com/>, 2010.
- [2] Sensolute GmbH. Mvs0608.02 micro-vibration sensor. <http://www.sensolute.com/>, 2010.
- [3] Kugsang Jeong, Jongho Won, and Changseok Bae. User activity recognition and logging in distributed intelligent gadgets. In *Multisensor Fusion and Integration for Intelligent Systems, 2008. MFI 2008. IEEE International Conference on*, pages 683–686, Aug. 2008.
- [4] S Mora, N Cook, JE Buring, PM Ridker, and IM Lee. Physical activity and reduced risk of cardiovascular events: potential mediating mechanisms. *Circulation*, 116(19):2110–8, 2007.
- [5] KH Pietilinen, J Kaprio, P Borg, G Plasqui, H Yki-Jrvinen, UM Kujala, RJ Rose, KR Westerterp, and A Rissanen. Physical inactivity and obesity: a vicious circle. *Obesity*, 16(2):409–14, 2008.
- [6] G Plasqui, AM Joosen, AD Kester, AH Goris, and KR Westerterp. Measuring free-living energy expenditure and physical activity with triaxial accelerometry. *Obes Res*, 13(8):1363–9, 2005.
- [7] M. Stager, P. Lukowicz, N. Perera, T. von Buren, G. Troster, and T. Starner. Soundbutton: design of a low power wearable audio classification system. In *Wearable Computers, 2003. Proceedings. Seventh IEEE International Symposium on*, pages 12–17, Oct. 2003.
- [8] Emmanuel Munguia Tapia, Stephen S. Intille, William Haskell, Kent Larson, Julie Wright, Abby King, and Robert Friedman. Real-time recognition of physical activities and their intensities using wireless accelerometers and a heart rate monitor. *Wearable Computers, IEEE International Symposium*, 0:1–4, 2007.
- [9] Ilkka Vuori. *Physical Inactivity as a Disease Risk and Health Benefits of increased Physical Activity*, pages 29–96. Oxford (UK): Meyer & Meyer Sport Ltd., 2004.