

Keep on Moving! Activity Monitoring and Stimulation Using Wireless Sensor Networks

Stephan Bosch¹, Mihai Marin-Perianu¹, Raluca Marin-Perianu²,
Paul Havinga¹, and Hermie Hermens³

¹ University of Twente, The Netherlands

{s.bosch,m.marinperianu,p.j.m.havinga}@utwente.nl

² Inertia Technology, The Netherlands

raluca@inertia-technology.com

³ Roessingh Research and Development, The Netherlands

h.hermens@rrd.nl

Abstract. Because health condition and quality of life are directly influenced by the amount and intensity of daily physical activity, monitoring the level of activity has gained interest in recent years for various medical and wellbeing applications. In this paper we describe our experience with implementing and evaluating physical activity monitoring and stimulation using wireless sensor networks and motion sensors. Our prototype provides feedback on the activity level of users using a simple colored light. We conduct experiments on multiple test subjects, performing multiple normal daily activities. The results from our experiments represent the motivation for and a first step towards robust complex physical activity monitoring with multiple sensors distributed over a person's body. The results show that using a single sensor on the body is inadequate in certain situations. Results also indicate that feedback provided on a person's activity level can stimulate the person to do more exercise. Using multiple sensor nodes and sensor modalities per subject would improve the activity estimation performance, provided that the sensor nodes are small and inconspicuous.

1 Introduction

In recent years, monitoring the level of daily human activity has gained interest for various medical and wellbeing applications. It has been shown that health condition and quality of life are directly influenced by the amount and intensity of daily physical activity [1]. This is particularly relevant to persons with chronic conditions, such as Chronic Obstructive Pulmonary Disease (COPD), asthma and diabetes. The reason is that persons suffering from chronic conditions enter a vicious circle, in which being active causes discomfort, making them progressively more sedentary, and deteriorating their health. Monitoring the daily activity can stimulate people to perform exercises and to be more active in general by providing feedback and assistance to better manage the physical condition. Apart from medical applications, daily activity monitoring can also

be useful for healthy users that want to assess and improve their overall fitness level. Activity monitoring systems can remind, stimulate and motivate people to be more active, especially when used within groups with a competitive nature.

The activity level of a person is best assessed in terms of energy expenditure [2]. This is a relatively complex and intrusive process, because the rate of human metabolism needs to be measured. As an alternative, we can estimate the energy expenditure rather than measuring it explicitly. Estimation can provide an unobtrusive alternative, which assesses the level of activity without affecting the subject's wellbeing and freedom of movement. The energy expenditure is commonly estimated by measuring the amount of motion a person performs during daily life. A proven method [2] is to use an accelerometer to record the bodily motion, which can provide a measure that usually correlates well with the actual energy expenditure. This solution is limited to a single acceleration sensor mounted near the center of mass of the subject, e.g. worn at the waist of the person.

However, the quality of this accelerometer-based activity measurement method is not always optimal. People are engaged in many more activities than those usually considered in current activity monitoring and recognition research. In this paper we show that some activities, such as skiing, do not cause much net acceleration, although the amount of effort involved in these activities can be significant. This causes a discrepancy between the estimation and the actual energy expenditure. Using more than a single sensor node for the estimation could improve the quality of the estimation. However, there exists a trade-off between the number of sensors spread across the user's body and the system's unobtrusiveness. Therefore, activity monitoring applications are usually limited to only one sensor. The experiments outlined in this paper represent the motivation for and a first step towards robust complex activity monitoring with multiple sensors distributed in a wireless sensor network.

Concretely, the contributions of this paper are as follows:

- We show that it is feasible to implement activity monitoring on resource-constrained wireless sensor nodes, which provides the basis for having activity monitoring on distributed nodes, making the system better suited for a large diversity of activities.
- Using practical experiments, we show that the estimation of energy expenditure using only one sensor node with an accelerometer is not always adequate and is sensitive to context dependencies.
- We investigate the feasibility of simple and intuitive feedback devices, such as ambient light, in order to make the system accessible to non-technical users.
- We conduct a usability study to evaluate whether:
 1. Users are more willing to be active when they get feedback on their activity level and when they compete with other people wearing sensors.
 2. Users are willing to wear multiple sensor nodes to monitor their activity levels.
 3. Users are prepared to share activity data with others or have privacy concerns.

2 Related Work

The state-of-the-art electronic method to assess physical activity is to use an integral of three-dimensional acceleration measurements. For instance, Bouten [2] provides a solution for activity monitoring called the Tremor, describing a method that uses a three-dimensional accelerometer to estimate metabolic energy expenditure. It uses the integral of the modulus of bodily acceleration, referred to as the IMA value, as measure for physical activity. Bouten provides validation for this method by establishing the correlation between the true metabolic energy expenditure and the IMA result using two laboratory experiments and one field experiment. In a later publication [3], Bouten et al. explore the effects of orientation and placement of the sensor on the resulting IMA value. This yields the conclusion that the sensor is best placed near the waist at the lower back. Bouten also recognizes that static exercise can cause discrepancies between the IMA value and the actual energy expenditure. Later work by Steele et al. [4] and Fujiki [5] et al. use a different integration algorithm than IMA, but in essence the methods are similar. Fujiki [5] et al. also provide further insight in the effect of sensor placement on the activity estimation.

Alternatively, pedometers can be used to estimate physical energy expenditure. However, pedometers are only capable of detecting step counts and not the true intensity of the activity. This is particularly a problem for slow-walking patients. According to Pitta et al. [1], multiaxial accelerometers can provide more detailed information on activity patterns, time and intensity of activities. In comparison to pedometers, accelerometers have good sensitivity to low-intensity activities and the small activity variations seen in relatively inactive, i.e. sedentary, populations.

Also, subjective methods like questionnaires can be used to quantify physical activity. The work by Pitta et al. [1] provides a survey and a comparison of the subjective methods and the methods involving motion sensors. It concludes that subjective methods do not provide an accurate means to assess physical activity quantitatively, showing large variation among different subjects.

Recently, physical activity monitoring has gained interest also for the consumer market. FitBit [6] is a commercial product to monitor and review one's fitness and quality of sleep. Data is collected continuously and uploaded to a website to provide feedback. ActiSmile [7] is another commercial wearable device to estimate the physical activity level. It first classifies the type of activity among sitting/lying, walking and running, then it uses this result to assess the level of activity and it provides feedback on the display using a smiley face.

With respect to the related work, we focus in this paper on three important aspects. First, we claim that users are involved in more types of physical activities than typically considered and that one sensor node may not be enough to assess energy expenditure in all situations. Second, we implement and test the IMA algorithm on resource-constrained wireless sensor nodes, which creates the basis for having activity monitoring on distributed nodes, and thus being better suited for a large diversity of activities. Third, we implement and evaluate the usability

of light-based feedback, taking therefore the computer out of the loop and making the system accessible to also non-technical users.

3 System Description

In this section we provide an overview of all system components, describing how our system operates, the hardware involved, the user interface and the algorithms we use to assess energy expenditure.

3.1 System Overview

The core of our system is represented by wireless sensors nodes that monitor the users' daily activities. The nodes include an accelerometer and are capable of processing the accelerometer samples into a measure of activity. This activity level is stored internally and can be communicated to a local user interface when the user wants to review his or her current daily or hourly activity. For the evaluation presented in this paper, the sensors also log the last day of activity measurements with a granularity of one hour. This enables us to compare the daily activities reported by our test subjects to the logged activity levels.

The subjects wear the sensors on their pant belt. Earlier research [3] has shown that the center of mass, i.e. the waist, is a sensible position for wearing such a sensor, being less prone to bias towards particular activities. For example, when the sensor were located on an arm it would show high levels of activity when the person is sitting somewhere and waving his arms, which does not correspond well with the amount of activity actually involved.

The user interface of our system is a spherical LED lamp, also referred to as Orb, that lights up in different colors. Using a color in between red and green, the user can get a rough idea of the level of activity in the last 60 minutes and




Fig. 1. An overview of the activity monitoring prototype




Fig. 2. ProMove inertial sensor board

the last 24 hours. The user can send commands to the user interface by tapping the sensor node.

Figure 1 shows an overview of the system. The sensor node is worn on the user's belt. When it is tapped by the user, it sends a message to the gateway containing the daily and hourly activity statistics. The gateway interprets the data and controls the orb to change its color accordingly.

3.2 Hardware

We use the ProMove [8] wireless inertial sensor nodes for our activity monitoring application. The ProMove board shown in Figure 2 features a three-axial accelerometer, a two-axial gyroscope and a three-axial digital compass. For this application, we only use data from the accelerometer sampled at a frequency of 100Hz. These inertial sensors are sampled through a low-power MSP430 microcontroller [9], which communicates the data wirelessly through the CC2430 radio [10]. The radio is IEEE 802.15.4-compatible and transmits at 2.4 GHz. A useful feature of the CC2430 is that it combines an 8051 CPU with a radio peripheral. This allows us to implement our system in a two-tiered fashion: performing data processing on the MSP430 and wireless networking on the CC2430.

As shown in Figure 1, we use a CC2430EB evaluation board from Texas Instruments [10] as gateway for our nodes. These boards only have the CC2430 radio processor and are equipped with an RS232 serial interface, that we use to connect to the orb (see Figure 1).

3.3 Assessing Energy Expenditure

We use the accelerometer of the ProMove sensor nodes to assess a person's energy expenditure. We implement the IMA algorithm described by Bouten [2] on the ProMove nodes. The accelerometer signal is first high-pass filtered to remove the gravity component of the accelerometer signal. The resulting acceleration signal is then integrated according to Equation 1 to obtain the IMA value. The

integral produces a value that correlates with the signal energy over the three axes, which is a good measure for the intensity of the measured motion. The integration time T is set to one minute. To calculate the IMA value for one hour or one day, the IMA minute values are averaged over the period involved.

$$IMA = \int_{t=t_0}^{t_0+T} |a_x(t)| dt + \int_{t=t_0}^{t_0+T} |a_y(t)| dt + \int_{t=t_0}^{t_0+T} |a_z(t)| dt \quad (1)$$

3.4 User Interface

We use the Ambient Orb [11] as our user interface. This is a LED-based light housed in a diffuse glass orb. It contains the necessary control electronics to make the orb light up in many different colors. The device has a serial RS232 interface through which it accepts simple commands to change the color of the light. The color can either be selected from a table of pre-defined colors or custom colors can be composed by specifying the required RGB color channel values.

As shown in Figure 1, we connect the Orb to one of our CC2430 evaluation boards using its serial interface. This evaluation board acts as an IEEE 802.15.4 coordinator and listens for incoming commands from the user nodes. Currently, two commands are distinguished: display the activity level of the last 24 hours and display the activity level of the last 60 minutes. The user can issue these commands by tapping the sensor node, by tapping it twice or three times respectively. When the hourly activity level is requested, the orb first becomes white and then glows towards the activity level represented as a color between red and green. To distinguish the daily activity level display, it is indicated by a bright blue initial color instead of white.

Multiple users can interact with the orb simultaneously. Each sensor node has an internally stored identifier code that is displayed as a two-color sequence on the orb when a user first interacts with it. Each time a different user interacts with the node, his identifier is shown on the node first, followed by the activity level display after a short pause.

The tap commands are picked up by the accelerometer of the sensor node. A tap causes high-frequency vibration, which is picked up as a high frequency acceleration pulse by the accelerometer. The actual tap instances are detected by high-pass filtering the signal with a cut-off frequency of 35 Hz, then rectifying the signal and finally low-pass filtering the result to obtain a signal that correlates with the amount of high-frequency vibration. The actual tap detection uses thresholds to mark the onset and end of a tap. The main advantage of this technique is that the sensor nodes do not need additional hardware to receive user commands.

3.5 Energy Saving

The sensor nodes that are attached to the users need to monitor activity for several days. Energy conservation is therefore an important aspect. The main

power drain on the sensor nodes is the activity of the radio. The sampling process on the MSP430 draws about ten times less current, thus it is imperative to turn off the radio when it is not needed. Since our sensor nodes do not need to communicate unless the user tries to interact with the orb, the radio is turned on only when the user taps the sensor. If the user does not issue any commands for 30 seconds, the radio is turned off by putting the CC2430 chip in deep sleep mode. In this situation the whole ProMove board consumes only 7.2 mW, mainly for sampling the accelerometer.

4 Experiments

4.1 Method

The normal mode of operation of the system is to report the activity level in the last 24 hours or the last 60 minutes to the person wearing the hardware. However, for our experiments the system needs to log the activity during a longer period of time at multiple instances. To achieve this, the software is amended with a storage subsystem that logs the hourly activity levels into the controller's flash memory. It can store the activity of the last 36 hours. The radio protocol is adjusted as well, making it send the full data log each time the user issues one of the tap commands. A separate CC2430EB evaluation board listens for the transmitted data and forwards it to a computer through its serial link.

The hourly recording yields a relatively low data granularity. Consequently, the boundaries between successive activities can fall in the middle of a recorded hour, which makes it difficult to infer to what extent the recorded activity value corresponds to which of the two activities. To mitigate this problem, our current approach is to only consider those hourly data points that pertain to only a single activity.

4.2 Tests

To obtain a view on how well the IMA value maps to the energy expenditure of a particular person, we conduct tests with multiple subjects performing various activities of interest. The tests are performed in the course of several experiments.

Office. The main and longest experiment is performed by nine members of our department. This includes the daily office and home activities. The sensors are worn the whole day, except during sleep. The experiment lasts three days.

As the nodes can only record 36 hours, the logs are collected each day during a collective meeting in the afternoon. The subjects keep a coarse log of what they have been doing all day. These logs are evaluated during the meeting and compared to the logged hourly activity levels.

The recorded activities include the normal office activities, home activities and transportation.




Fig. 3. Subject wearing the ProMove node and receiving feedback on the Orb

Ski. Two sensor nodes are taken on a skiing holiday in the mountains. The two subjects keep a coarse log of their daily activities, including skiing, resting and driving the car.

Sports. Four separate experiments are performed in which the subjects are engaged in sports: tennis, badminton, volleyball and cycling. Tennis is performed by two persons, volleyball is performed by four persons and cycling is performed by three persons at the same time. The badminton experiment is spread over multiple parallel matches involving four subjects. Each experiment lasts more than an hour to record only the activity of interest.

5 Results

Table 1 shows the average IMA value for various activities. Because the sensor nodes log the activity level on a per-hour basis, the presented values are the hourly averages of the same activity performed by multiple persons at multiple time instances. The table also shows how many hours of data are involved in computing the average and how many subjects were involved in a particular activity. Since activities do not necessarily change on the hour boundaries, the IMA averages are solely based on those hours during which only one activity was performed with certainty.

The recorded activities are as follows:

- *Not worn*: the sensor node is not worn by anyone and just lies stationary somewhere.

Table 1. Activity recording results

Activity	Recorded Hours	Persons Involved	IMA
Not Worn	129	9	373
Home Activities	33	6	1581
Office Activities	58	9	1144
Sitting	25	6	1120
Walking	4	2	7767
Cycling	3	3	6466
Tennis	2	2	11411
Badminton	6	4	10874
Volleyball	4	4	9464
Skiing	40	2	1496
Driving a Car	31	2	827
Presenting	1	1	2655

- *Home activities*: the subject is engaged in the typical home activities, such as cooking, washing, childcare etc.
- *Office activities*: the subject is engaged in the typical office activities, which mainly consist of sitting, but also intermittent walks between offices and fetching coffee.
- *Sitting, walking*: the subject is continuously sitting or walking respectively.
- *Tennis, Badminton, Volleyball, Cycling*: the subject is playing tennis, playing badminton, playing volleyball or riding a bicycle respectively.
- *Skiing*: the subject is skiing downhill.
- *Driving in a car*: the subject is driving in a car.
- *Presenting*: the subject is giving a presentation during a meeting.

Figure 4 graphically illustrates the difference between the activities in terms of the measured IMA value. We notice that the results show that the IMA value always greater than zero, even when the sensor is not worn, due to the integration of accelerometer noise.

The results show that home activities produce higher activity values than office activities. The reason is that in the office people barely leave their desk for hours. At home, activities such as doing laundry, cleaning etc. contribute to a higher activity level. It is interesting to see that what people describe as sitting matches well with the office activities.

The activity value for cycling is the average of the IMA value from three persons making a bicycle trip. For this test the activity results vary wildly, with values ranging from 5635 to 7542. These differences are probably caused by differences in the bicycles, the exact path taken on the road and the way the subjects ride their bicycle.

The results for tennis and badminton are very similar, since these two sports and the movements involved therein are very similar as well. Volleyball produces lower activity values. This is due to two main reasons: more people are involved in the ball game, spreading the activity over more persons, and the game is interrupted for longer periods of time.




Fig. 4. IMA values for various activities

The most interesting data results from the skiing activity. Although skiing is certainly not an activity with a low energy expenditure, the sensors show a low activity value. This value is lower than walking, even though skiing is much more exhausting. The reason is that during skiing the amount of variation in the acceleration of the person is relatively limited in comparison to walking. The regular motion of the gait of walking is not present when sliding down a hill on skis. The energy expenditure from skiing is more the result of the act of leaning in the bends and keeping ones legs steady.

According to the results displayed in Figure 4, walking is more active than riding a bicycle, but the spread in the activity values is relatively large. This is unexpected, since the walking experiments were conducted with a normal walking pace, whereas the cycling experiments were performed at speed. Therefore, we would expect to see a higher activity result for the cycling experiments. We attribute this discrepancy to the same effect as for the skiing experiments.

The ski experiment also shows that the orientation and the way sensor is mounted influences the result. After in the middle of the experiment, the pouches in which the sensors are worn were exchanged between the subjects. The persons kept their own sensor. During the week, the recorded activity values were similar, with the trend that one was reported to be more active than the other most of the time. After the exchange of the pouches, this trend inverted. The main observable difference between these pouches is the fact that one is worn vertically and the other horizontally.

Figure 5 shows an example of an activity log for two subjects. For most of the day, these subjects were together and engaged in the same activities. The graph shows a clear correlation between the two subjects. At 19:00 they played a tennis match for one hour, which is much more active than anything else recorded during the experiment. Between 23:00 and 7:00 the sensors were not worn. As shown, the IMA values are constant in stationary conditions.

Results indicate that the use of a single accelerometer is not always adequate to accurately and consistently estimate the level of physical activity. Also, the orientation and way the sensor is mounted on the subjects have an influence on




Fig. 5. Example activity values recorded during 27 hours from two subjects

the measurements. However, the activity measurements do correlate well between subjects that are engaged in identical activities.

5.1 Usability

To evaluate usability aspects of the system, the nine test subjects filled in a questionnaire. The questions target the wearability and unobtrusiveness of the system, the feedback quality of the user interface, the effectiveness of the activity stimulation that the users perceive, and the privacy concerns users may have.

Wearability and Unobtrusiveness. The first question pertains to how subjects feel while wearing the sensor node. Most subjects (78 %) are not constantly aware that they are wearing a sensor node or they are not bothered by its presence. One subject indicates that it is unpleasant to wear, mainly due to its size but also due to its looks. Another subject notices that a relatively large belt-mounted device can be annoying when sitting in a chair, lying somewhere or when doing sports. All subjects agree that the sensor node should be smaller than it is now.

We also ask how many sensor nodes subjects are prepared to wear. Most subjects (67 %) answered that they would accept wearing two or three sensors. Two subjects are prepared to wear more than three sensors and one user accepts only one sensor. However, the subjects indicate that the limit on the number of nodes is inversely proportional to its size. The intended mounting positions and the type of sensors have an influence on the acceptable number of sensors, such that the integration of very small sensors into clothing would be preferred.

Feedback Quality. To assess the feedback quality of our visual user interface, the users are asked about their opinion on the Orb light interface. Most subjects

(67 %) find the interface appealing, but not too intuitive. One subject likes it unconditionally. Two other subjects have a very hard time interpreting the colors, especially when multiple subjects are interacting with the Orb. In fact, most subjects are initially confused by the display of the user identifier and struggle to distinguish between the hourly and daily activity display. It should be clear without much explanation what the color sequences mean, but that turns out to be more difficult than expected.

Activity Encouragement. The first question regarding activity encouragement is whether subjects feel motivated to be more active by wearing the sensor node. Most subjects (67 %) indicate that wearing the sensor stimulates them to be more active. Three subjects did not really feel motivated to do more exercise and regarded the sensor to be just a monitoring device.

The second question extended this assessment to include the possibility that the competitive nature of a group could further stimulate users to be more active. We ask whether subjects feel motivated to be more active when their results are compared to others. Most subjects (67 %) do feel that this competitive property of the system further stimulates their activity. The same three subjects that were not motivated as much by the sensor alone were also not impressed by the competitive stimulation.

Privacy Concerns. We ask whether subjects would be prepared to share their activity results with others, e.g. on a website. Most subjects (67 %) would be prepared to do so to some extent, but only with friends. Three subjects see no merit in doing so and wonder who would be interested in this information.

Conclusion. Even though our prototype sensor node is not very small (see Figure 3), most subjects are not bothered much by its presence. However, there is the consensus that it needs to be smaller and this is one of our goals for future work. Depending on the form factor, most subjects would be willing to wear more sensors spread across the body.

The light interface needs improvement. The distinction between the hourly and daily display is not intuitively clear and the display of a user's personal color sequence is sometimes confusing.

The opinion about whether the system encourages activity varies among the subjects of the questionnaire. Most subjects are encouraged to be more active, but others only recognize the system as an activity monitor. This is also true for the envisioned competitive encouragement.

6 Conclusions

We described in this paper our experience with implementing and evaluating activity monitoring in wireless sensor networks. The results from our experiments represent the motivation for and a first step towards robust complex activity monitoring with multiple sensors distributed in a wireless sensor network. Wireless sensor networks create a promising basis for achieving distributed activity recognition and monitoring for a large diversity of physical activities.

In our experiments, multiple subjects were monitored during their daily activities (both at home and at the office) and additionally while performing sports such as badminton, tennis, volleyball, cycling and skiing. The results show that the recorded values correspond to the expected intensity for some of the performed activities, but certain activities, such as skiing and cycling, produce lower activity measurements than expected. Using multiple sensor nodes and sensor modalities per subject would improve the activity estimation performance, provided that the sensor nodes are small and inconspicuous, as indicated by the results of our usability study.

The usability study indicates two other directions of future work. First, simple feedback with an ambient light is appreciated by the users, but it is not considered to be intuitive enough. The capabilities of the feedback device should allow for faster user identification and provide more granular information on the activity level. Second, the competitive nature of group interaction should be better exploited to motivate and stimulate users to increase their level of daily physical activity.

References

1. Pitta, F., Troosters, T., Probst, V., Spruit, M., Decramer, M., Gosselink, R.: Quantifying physical activity in daily life with questionnaires and motion sensors in COPD. *European Respiratory Journal* 27(5), 1040–1055 (2006)
2. Bouten, C.: Assessment of daily physical activity by registration of body movement, p. 145. Eindhoven Univ. Technol., Eindhoven (1995)
3. Bouten, C., Sauren, A., Verduin, M., Janssen, J.: Effects of placement and orientation of body-fixed accelerometers on the assessment of energy expenditure during walking. *Medical and Biological Engineering and Computing* 35, 50–56 (1997)
4. Steele, B., Holt, L., Belza, B., Ferris, S., Lakshminaryan, S., Buchner, D.: Quantitating Physical Activity in COPD Using a Triaxial Accelerometer* (2000)
5. Fujiki, Y., Tsiamyrtzis, P., Pavlidis, I.: Making sense of accelerometer measurements in pervasive physical activity applications. In: CHI EA 2009: Proceedings of the 27th international conference extended abstracts on Human factors in computing systems, pp. 3425–3430. ACM, New York (2009)
6. Fitbit, Inc., <http://www.fitbit.com/>
7. Krauss, J., Solà, J., Renevey, P., Maeder, U., Buchholz, H.: ActiSmile: A Portable Biofeedback Device on Physical Activity. In: Proceedings of the Sixth IASTED International Conference on Biomedical Engineering, International Association of Science and Technology for Development, # 80, 4500- 16 Ave NW, Calgary, Alberta, T 3 B 0 M 6, Canada (2008)
8. Inertia Technology: ProMove wireless inertial sensor node, <http://www.inertia-technology.com>
9. Texas Instruments: MSP430 family of ultra-low-power microcontrollers, <http://focus.ti.com/docs/prod/folders/print/msp430f1611.html>
10. Texas Instruments: CC2430 System-on-Chip Solution for 2.4 GHz IEEE 802.15.4 / ZigBee, <http://focus.ti.com/docs/prod/folders/print/cc2430.html>
11. Ambient Devices: Ambient Orb, <http://www.ambientdevices.com/cat/orb/orborder.html>