

The 2nd International Conference on Ambient Systems, Networks and Technologies
(ANT)

Energy-Efficient Assessment of Physical Activity Level Using Duty-Cycled Accelerometer Data

Stephan Bosch^a, Raluca Marin-Perianu^a, Paul Havinga^a, Mihai Marin-Perianu^b

^aUniversity of Twente, Enschede, The Netherlands

^bInertia Technology B.V., Enschede, The Netherlands

Abstract

This paper describes an energy efficiency improvement of the IMA accelerometer-based method for estimating the level of physical activity of a person. The sensor sampling and data processing requirements are significantly reduced by duty-cycling sensor sampling, thus making implementation and long-lasting operation possible on resource-constrained devices as sensor nodes. By duty-cycling, the system maintains adequate bandwidth, while still reducing the effective number of samples taken per unit of time. We analyze in detail the impact of duty-cycling on the accuracy of the method and show that we can reduce the duty-cycle to as little as 10 %, incurring a mean error of only about 4 %. This translates into energy saving of up to 60% on the sensor node.

© 2011 Published by Elsevier Ltd. Selection and/or peer-review under responsibility of Prof. Elhadi Shakshuki and Prof. Muhammad Younas.

Keywords: physical activity; energy expenditure; accelerometer

1. Introduction

In recent years, monitoring the level of daily human activity has gained interest for various medical and wellbeing applications, mainly because 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), cardiovascular disease, obesity, osteoporosis, and diabetes [2, 3]. The reason is that persons suffering from chronic conditions can enter a vicious circle: being active causes discomfort for these persons, making them reduce their level of activity, thus making them progressively more sedentary, which in turn deteriorates their health even further. This vicious circle can be broken by monitoring the level of daily activity and by providing feedback and assistance to better manage their physical condition, e.g. by stimulating them to perform exercises. Apart from these medical applications, measuring the level of daily activity also has applications in (professional) sports [4] and personal fitness [5].

The activity level of a person is best assessed in terms of energy expenditure [6]. This is a relatively complex and intrusive process, because the rate of human metabolism needs to be measured in a lab. As

Email addresses: s.bosch@utwente.nl (Stephan Bosch), mihai@inertia-technology.com (Mihai Marin-Perianu)

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 well-being 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 is to use a triaxial accelerometer to record the bodily motion, which can provide a measure that usually correlates well with the actual energy expenditure [6]. Most accelerometer-based solutions produce a value that correlates well with the amount of motion, i.e. the motion signal energy, sensed within a certain period of time. The produced value is dimensionless in most cases and calculated directly from the accelerometer output, e.g. by summing the absolute values of all sensor axis measurements within the specified time period.

An important consideration is the sample frequency of the accelerometer: to obtain a reliable result, the sample frequency should account for the bandwidth of (voluntary) human motion [6]. This dictates a certain minimum sample frequency, which is in most cases specified between 20 to 40 Hz [6]. It is tempting to choose a sufficiently high sample frequency with a considerable margin, to guarantee maximum performance. However, especially when a person's level of physical activity needs to be monitored continuously throughout the day, the energy consumption of the device needs to be reduced as much as possible in order to maximize battery life. The chosen sample frequency directly influences the energy consumption of the device, since higher sample frequencies increase the sensor's power consumption and require more CPU processing. Therefore, for choosing a proper sample frequency, a trade-off exists between the energy consumption and the reliability of the measurement.

Because a certain minimum sample frequency is necessary to maintain good performance, another approach is necessary to reduce the energy consumption further. We observe that a person's level of activity is unlikely to change very quickly. This suggests that it is sufficient to measure the level of activity only intermittently, avoiding the need to sample continuously and allowing the system to enter sleep mode on a regular basis, thereby conserving energy. Therefore, we propose a duty-cycling approach in which the system measures the level of activity only at a fraction of the time.

In this paper, we evaluate this idea by assessing the impact of duty cycling on the performance of a representative accelerometry-based algorithm for estimating a person's level of physical activity. We simulate the performance at various algorithm configurations using data collected from an experiment involving real daily activities. We also compare duty cycling with the impact of reducing the sampling frequency directly. We first briefly survey the available methods for measuring a person's level of physical activity and we describe in detail how this is achieved with an accelerometer. Then, we describe our duty-cycling method, which we subsequently evaluate in a series of simulations.

2. Background

The problem of measuring a person's level of activity has gained much attention in recent years, mostly because there is a direct relation between a person's level of physical activity and his or her health condition [1, 3]. Regular physical activity can prevent or delay the onset or the progress of certain chronic diseases, such as Chronic Obstructive Pulmonary Disease (COPD), for which low levels of physical activity are related to a higher risk of hospital readmission and shorter survival [1]. Other examples include cardiovascular disease, obesity, osteoporosis, and diabetes [2, 3]. Apart from medical applications, measurement of physical activity also has applications in sports [4] and personal fitness [5].

Doubly-labeled water (DLW) and calorimetry methods are often used in research as a 'gold standard' reference [1, 6, 2]. These methods measure the body's energy expenditure by measuring the rate of metabolism. These are very expensive and obtrusive measurement processes, which can only be performed in a lab. Another important disadvantage is that such methods do not provide any information about the duration, frequency and intensity of the performed physical activity [1].

Pitta et al. [1] explore the available methods for assessing the level of physical activity of COPD patients. The survey explores subjective methods, like questionnaires and diaries, and motion sensor-based methods using pedometers or accelerometers. The survey concludes that questionnaires and diaries can be unreliable, are influenced by the persons involved, are unsuitable to measure low-intensity activities and are affected

by the subject's memory limitations. Questionnaires mainly have value determine the level of physical activity for a group rather than individuals. As an alternative, the survey considers motion sensors such as pedometers to provide an objective measurement. 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], multi-axial accelerometers can provide more detailed information on activity patterns, time and intensity of activities.

In a more broad survey, Warren et al [3] include the possibility of using heart rate monitoring to measure physical activity. Although such measurements are relatively cheap to obtain, have a good association with the body's energy expenditure and also provide information about the intensity of the activity, this technique is not suitable for light activity and is influenced by the fact that a person's heart rate can vary due to causes other than physical activity. Warren et al. also describe an option to combine accelerometry and heart rate monitoring, thereby combining the advantages of both techniques and negating some of the disadvantages, with the penalty of additional complexity and cost.

Probably due to its ease of operation, unobtrusiveness and low cost, accelerometry is a very popular technique. Much research is conducted with commercially available devices such as the ActiGraph, the ActiCal, the RT3, the IDEEA, and the SenseWear accelerometers and others [7, 8, 9]. An important problem to solve is how to extract a useful measure for the level of physical activity from the accelerometer data. Also, the produced value needs to be calibrated using reference experiments, involving a treadmill for example, to find the mapping between the accelerometer activity value and the actual level of physical activity or the bodily energy expenditure. Common algorithms are surveyed by Twomey et al. [10] and Godfrey et al. [9].

A good example of such an algorithm is provided by Bouten et al. [6]. Their extensive experiments show that a person's energy expenditure can be accurately estimated using a triaxial accelerometer and their algorithm. The estimate is based on an activity value that is calculated from the accelerometer signal in the period T starting at t_0 as follows:

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

This IMA ('integral of the modulus of the accelerometer output') activity value correlates with the signal energy of all the accelerometer axes over the period T ¹. It therefore corresponds well with the amount of motion the sensor experienced within the period T . Before integration, the signal is limited to a bandwidth of 0.11 to 20 Hz, to remove the DC (mostly gravity) part and vibration components that are not part of voluntary human motion [11]. Therefore, apart from the remaining noise, the calculated IMA value is very close to zero when there is no activity. In the activity recognition field, a very similar integral is known as the Signal Magnitude Area (SMA) and is used to distinguish between static and dynamic activities [12]. This feature value does not, however, necessarily include any filtering beforehand. The activity value period T for the IMA integral is chosen depending on application requirements. For the work by Bouten et al. it is set to one minute.

The algorithm by Bouten et al. is simple and feasible for implementation on hardware with few processing and memory capabilities, such as sensor nodes. That is why our research builds upon this work and uses the same algorithm. We implement this formula almost exactly, except for the fact that we employ a simple high-pass filter in stead of the prescribed bandpass filter.

3. Energy-Efficient IMA

We aim to provide an efficiency improvement on the IMA accelerometer-based energy expenditure estimation method by Bouten et al. [6], as described by the IMA formula of Equation 1. We focus on reducing the sampling frequency f_s , thereby lowering the system's processing requirements and energy consumption.

¹The 'tot' subscript of IMA_{tot} in the original publication [6] refers to the *total* of all axes, but we simply refer to it here as *IMA*.




Fig. 1. Overview of duty cycled sampling scheme.

However, a certain minimum sampling frequency is necessary to obtain a reliable result. This means that at some point, reducing the sample frequency further would impair the reliability and accuracy of the IMA output too much. This is mostly due to the fact that a certain minimum bandwidth is required to capture the most significant spectral components of the movement signal, which according to Bouten et al. are located roughly between 1 Hz and 20 Hz.

However, we could still improve efficiency in the time domain: it is probably not strictly necessary to keep sampling continuously. If we duty cycle the accelerometer sampling, we could achieve a higher level of efficiency, while maintaining sufficient bandwidth. Assuming that the level of physical activity does not change much in the inactive periods of the duty cycle, i.e. when the system is not sampling, the impact on performance would be minimal. Also, during the inactive periods, the system could enter a sleep mode to conserve energy. For this optimization, the IMA algorithm itself is not altered in any way and its band-pass filter is not reset at any point in time.

The defining parameters of a duty cycling scheme are the duty cycle D , i.e. what fraction of time the system is active, and the duty cycle period (T_D) or frequency ($f_D = 1/T_D$), which dictates how often the system switches between activity and inactivity per unit of time. Figure 1 schematically shows the proposed duty cycled sampling procedure for one activity value period T . Two alternatives are displayed. The top plot shows a duty cycle period that matches the period T , meaning that all samples are collected at the beginning of that period. The duty cycle is set to 20%. The bottom plot shows a similar scheme, but the period T is divided in multiple duty cycle periods (three in this case). The bottom alternative has the advantage that the samples are spread more evenly over the period T , which makes the chance of missing short but important significant activity smaller. The IMA output is probably going to be less reliable when the duty cycle period is very long, simply because the chance that the activity level changes significantly in the inactive period is higher. However, choosing a very short duty cycle period is also not beneficial, since the act of switching to and from sleep mode will claim system resources as well. The length of the duty cycle period is therefore an important design choice.

Note that when f_D is very high relative to the sampling frequency at moderate duty cycle, the net effect will be very similar to just reducing the sampling frequency. As an extreme example, when f_D is equal to half the sample frequency and the duty cycle is 50%, every other sample is dropped, which means that the *effective sample frequency* is equal to half the sample frequency. Even when f_D is set to a more useful value, we can still calculate the effective sample frequency $f_{s,eff}$, which is equal to the actual amount of samples collected per second ($f_{s,eff} = Df_s$). We use this as a measure for the amount of effort involved in the IMA calculation. Obviously, when the duty cycle D is 100%, f_s equals $f_{s,eff}$.

Summarizing, the following parameters are important for our implementation of the IMA algorithm:

- *Activity value period (T)*: the interval between successive IMA activity values.
- *Sample frequency (f_s)*: the sampling rate of the accelerometer at the active periods of the duty cycle.

This is what the hardware is configured to and the true rate at which samples enter the system when active.

- *Duty cycle (D)*: The fraction of time the system is sampling. The *effective sample frequency*, which indicates how many samples are actually collected per unit of time, depends on this parameter and is equal to $f_{s,eff} = Df_s$.
- *Duty cycle frequency (f_D)*: The frequency at which the system switches between active and inactive within the duty cycle scheme. The activity value period (*T*) should be an integer multiple of the duty cycle period ($T_D = 1/f_D$).

The effect that these parameters have on the system's performance and efficiency are evaluated in Section 4. Note that when the bare output from integral formula of Equation 1 is used, its magnitude does not depend only on the level of activity. Parameters like the sample frequency f_s , the activity value period *T* and the duty cycle *D* also have a significant influence on the resulting value. Therefore, we always normalize the activity values before any comparison is made, scaling it such that the theoretical (and practically unattainable) maximum IMA activity level will just fit into a 16 bit unsigned integer ($= 2^{16} - 1$).

4. Experiments and Results

To assess the impact of our efficiency improvement on the accuracy and reliability of the IMA activity value, we perform a series of experiments. We compare the output of the original unoptimized system at the highest sample frequency with the output of the system at various different optimized configurations. To be able to simulate different algorithm configurations using exactly the same conditions, we need to collect raw accelerometer data. This means that the actual experiments need to be performed just once and that the algorithm itself is executed offline in the simulation. Because the output is only compared between different algorithm configurations, and not between different activities or users, the absolute values of the system output are of little importance. This means that we can suffice with only one user, provided that a sufficiently broad set of activities is performed.

For this experiment, we use the ProMove inertial sensor node platform [13]. We collect raw three-dimensional accelerometer data at 200 Hz from one person performing daily activities. The sensor range is configured at $\pm 6g$. We collect approximately one hour of data involving activities like cycling, walking, standing and sitting. The sensor node is mounted at the user's waist on his belt. The data is logged to the node's on-board flash memory and downloaded wirelessly at the end of the experiment.

The accelerometer data from the experiment is shown in the top plot of Figure 2. The data is logged from a bicycle errand to the local city. The first 2.5 minutes involve sitting at a desk, walking down some stairs and out of the building. After that time, a short cycling trip is started, which ends at the person's home. At round the 6 minute mark, the trip ends and the person's activities involve walking around his home and standing still to talk to people. At minute 10, the actual bicycle trip to the city is started, which ends at the minute 27. At minute 24, the person needs to stop at a traffic light. Starting at minute 27, the person walks a significant distance which ends at the 31 minute mark at the counter of a shop. Up until minute 46, the person waits, which mostly involved standing and short walking activity. After that time the person walks back to his bike. Close to the 50 minute mark, the person reaches his bike and cycles back to the office. This trip involved a long stop at a traffic light starting at minute 52 and the person arrives at the office building at the 62 minute mark. The log ends with the person walking back into the building and taking a seat behind his desk.

The middle plot of Figure 2 shows the resulting IMA activity values for three different simulated sample frequencies. The IMA values are produced at an activity value interval (*T*) of 10 s and normalized as explained in Section 3. The sample frequency is reduced from 200 Hz to lower frequencies by down-sampling the raw accelerometer data; this also involves low-pass filtering to prevent aliasing effects. At the full 200 Hz, the IMA values are the most accurate. When the sample frequency is reduced to 100 Hz or even 50 Hz, the IMA values do not differ very much from the 200 Hz case (not shown in the plot). When




Fig. 2. IMA calculated from accelerometer data at various configurations




Fig. 3. Frequency and duty cycle performance statistics

the sample frequency is reduced further to 10 Hz, the difference first becomes significant, especially for the more intense activities, such as walking. Notice that the produced IMA value is consistently lower than the equivalent at 200 Hz, which is to be expected since some spectral components that contribute to the signal energy are lost at lower sample frequencies. However, even though the IMA values differ significantly from the 200 Hz case, the overall trend is mostly maintained, which means that the activity intensities retain similar relative magnitudes. Finally, for a sample frequency of 2 Hz, this trend is lost for the most part. This is evident from the fact that walking (e.g. at minute 30) and riding a bicycle (e.g. at minute 25) become indistinguishable in terms of activity level. Apparently, some important frequency components for distinguishing physical activity levels are located above 1 Hz.

Figure 3(a) shows error statistics for a whole range of sample frequencies. The plot shows the mean error, expressed as the difference with the ‘ideal’ IMA output at 200 Hz. The error magnitude is shown as a fraction of the maximum IMA value encountered in the 200 Hz simulation. We notice that the error is relatively small when the sample frequency is above 40 Hz and becomes very significant when it drops below 10 Hz.

In the next simulations, we explore what happens when the duty cycle is varied. We fix the simulated




Fig. 4. Duty cycle performance statistics at varying duty cycle periods

sample frequency at an appropriate level of 50 Hz. The activity values are again produced with an interval T of 10 s, which is now also the duty cycle period (T_D). The normalization of the IMA value also accounts for the effect of the duty cycle. The bottom plot of Figure 2 shows the resulting IMA activity values for three different simulated duty cycles. As the plot indicates, the main effect of duty cycling is that the IMA output is less stable. This is to be expected, since the system is sampling only at a fraction of the time, and at those instances the amount of motion may be significantly higher or lower than the motion that is skipped and thus not included in the result. The largest deviations happen at those instances where a brief dip or peak in the activity level happens just in the sample period of the sensor. Obviously, this effect is reduced when the duty cycle is higher and less samples are skipped. This effect is visible in Figure 2 for example at minute 40. An important difference with the reduction of the sampling frequency, as explored in the earlier simulations, is that the overall magnitude of the IMA output does not change, indicating that the system retains sufficient bandwidth. Figure 3(b) shows error statistics for a whole range of duty cycles 0.01 % and 100 %. As shown, between 10 % and 100 %, the impact of duty cycle on performance seems close to linear. However, when the duty cycle drops below 10 %, the difference grows much faster.

For the results in Figure 3(b), the duty cycle period T_D was set equal to the activity value period T of 10 s. As explained in Section 3, this may not be optimal since 10 s is a long time. Therefore, we investigate what value for T_D would be optimal. To extend the range of possible values a little, we increase T to one minute. We investigate duty cycle periods T_D that are integer divisors of T . The results are shown in Figure 4. The sample frequency f_s is still set to 50 Hz. The horizontal axis of the plot shows the effective sample frequency, which directly maps to the chosen duty cycle ($f_{s,eff} = Df_s$), as explained in Section 3. Clearly, reducing T_D is beneficial, as the error drops consistently for almost the whole spectrum of effective sample frequencies. However, reducing T_D below 2 s shows no more clear improvement (not shown in the plot), which suggests that this is the optimum value in this case.

It is now interesting to check how the reduction of the duty cycle - and thereby the reduction of the effective sample frequency - compares to the reduction of the real sample frequency shown in Figure 3(a). Figure 4 also includes a plot for the performance of varying the sample frequency directly, while keeping the duty cycle at 100 %. The comparative performance impact of reducing only the duty cycle versus reducing only the real sample frequency is clearly visible in the plot. Particularly to attain a low effective sample frequency (well below 25 Hz) to minimize processing effort, the use of duty cycling can be beneficial.

For example, we could aim to achieve an effective sample frequency of 5 Hz. We can do this by setting a duty cycle of 10 % at a real sample frequency of 50 Hz, which incurs an error of only about 3-5 % at a duty cycle period below 15 s. In contrast, reducing the sample frequency directly to 5 Hz incurs an error of more than 10 %, which is up to three times worse. Conversely, approximately the same 3 % error is incurred

at a real sample frequency of 15 Hz, which means that, at that same level of error, the amount of samples that needs to be taken and processed can be decreased by about three times (i.e. 66 % less) using a proper duty cycling scheme. Also, at a duty cycle of 10 % the system can enter sleep mode up to 90 % of the time and, at a moderate duty cycle period of a few seconds, it can sleep for long consecutive periods, avoiding the burden of frequent sleep-wakeup cycles. This amounts to a very significant efficiency improvement, yielding estimated energy savings of around 60 % (including some additional overhead) when compared to reducing the sample frequency directly.

5. Conclusion

In this paper we investigated improving the IMA algorithm in terms of its efficiency. We achieve this by effectively reducing the amount of accelerometer samples taken per unit of time. This is done by duty-cycling the sensor sampling in stead of reducing the (hardware) sampling frequency directly. This preserves the system's bandwidth, while still reducing the sampling and processing needs. We perform simulations with the IMA algorithm for various possible configurations using data collected from a real-life experiment. The simulations show that a reduction in the energy consumption of up to 60 % is feasible using the duty-cycling method, with only a minor (up to 5 %) difference from the results produced at very high sample frequencies with no duty cycling.

Future work includes building an actual implementation of this optimization and verifying the results at a broad set of activities performed by different people. Also, the actual energy savings can then be measured directly.

References

- [1] F. Pitta, T. Troosters, V. S. Probst, M. A. Spruit, M. Decramer, R. Gosselink, Quantifying physical activity in daily life with questionnaires and motion sensors in COPD, *European Respiratory Journal* 27 (5) (2006) 1040–1055. doi:10.1183/09031936.06.00064105.
- [2] G. Plasqui, K. R. Westerterp, Physical activity assessment with accelerometers: An evaluation against doubly labeled water[ast][ast], *Obesity* 15 (10) (2007) 2371–2379. doi:10.1038/oby.2007.281.
- [3] J. M. Warren, U. Ekelund, H. Besson, A. Mezzani, N. Geladas, L. Vanhees, Assessment of physical activity – a review of methodologies with reference to epidemiological research: a report of the exercise physiology section of the european association of cardiovascular prevention and rehabilitation, *European Journal of Cardiovascular Prevention & Rehabilitation* 17 (2) (2010) 127–139. doi:10.1097/HJR.0b013e32832ed875.
- [4] A. J. Wixted, D. V. Thiel, A. G. Hahn, C. J. Gore, D. B. Pyne, D. A. James, Measurement of energy expenditure in elite athletes using MEMS-Based triaxial accelerometers, *Sensors Journal, IEEE* 7 (4) (2007) 481–488. doi:10.1109/JSEN.2007.891947.
- [5] S. Consolvo, D. W. McDonald, T. Toscos, M. Y. Chen, J. Froehlich, B. Harrison, P. Klasnja, A. LaMarca, L. LeGrand, R. Libby, I. Smith, J. A. Landay, Activity sensing in the wild: a field trial of ubifit garden, *Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems* (2008) 1797–1806 ACM ID: 1357335. doi:10.1145/1357054.1357335.
- [6] C. Bouten, K. Koekkoek, M. Verduin, R. Kodde, J. Janssen, A triaxial accelerometer and portable data processing unit for the assessment of daily physical activity, *Biomedical Engineering, IEEE Transactions on* 44 (3) (1997) 136–147. doi:10.1109/10.554760.
- [7] K. Lyden, S. L. Kozey, J. W. Staudenmeyer, P. S. Freedson, A comprehensive evaluation of commonly used accelerometer energy expenditure and MET prediction equations, *European Journal of Applied Physiology* 111 (2) (2010) 187–201. doi:10.1007/s00421-010-1639-8.
- [8] C. Yang, Y. Hsu, A review of Accelerometry-Based wearable motion detectors for physical activity monitoring, *Sensors* 10 (8) (2010) 7772–7788. doi:10.3390/s100807772.
- [9] A. Godfrey, R. Conway, D. Meagher, G. ÓLaighin, Direct measurement of human movement by accelerometry, *Medical Engineering & Physics* 30 (10) (2008) 1364–1386. doi:10.1016/j.medengphy.2008.09.005.
- [10] N. Twomey, S. Faul, W. Marnane, Comparison of accelerometer-based energy expenditure estimation algorithms, in: *Pervasive Computing Technologies for Healthcare (PervasiveHealth), 2010 4th International Conference on*, 2010, pp. 1–8. doi:10.4108/ICST.PERVASIVEHEALTH2010.8879.
- [11] C. V. C. Bouten, Assessment of daily physical activity by registration of body movement, Technische Universiteit Eindhoven, Eindhoven, 1995.
URL <http://repository.tue.nl/443690>
- [12] A. Khan, Y. Lee, S. Lee, T. Kim, A triaxial Accelerometer-Based Physical-Activity recognition via Augmented-Signal features and a hierarchical recognizer, *Information Technology in Biomedicine, IEEE Transactions on* 14 (5) (2010) 1166–1172. doi:10.1109/TITB.2010.2051955.
- [13] Inertia Technology, ProMove wireless inertial sensor node.
URL <http://www.inertia-technology.com>