A Novel Micro-Vibration Sensor for Activity Recognition: Potential and Limitations

Dawud Gordon, Hedda Rahel Schmidtke, Michael Beigl
Karlsruhe Institute of Technology
[firstname.lastname]@kit.edu

Georg von Zengen
Technische Universität Braunschweig
g.vonzengen@tu-bs.de

Abstract

This paper researches the potential of a novel ball switch as a wearable vibration sensor for activity recognition. The ball switch is available as a commercial, off-the-shelf sensor and is unique among such sensors due to its miniaturized design and the low mass of the ball. We present a detailed analysis of the physical properties of the sensor as well as a recommendation for circuit design, sampling method and a feature generation algorithm for activity recognition. The analysis reveals that it is sensitive to vibrations between 1.5 kHz and 8 kHz, where the acceleration sensor is responsive below 1.6 kHz. Furthermore, the ball switch is substantially cheaper (3x), smaller (2x) and uses less power (50x) than an accelerometer based system, but delivers less information. We also present the results of a case study in activity recognition done in parallel with an acceleration sensor using 5 subjects and 8 different activities. It shows that the ball switch can increase recognition rates when added to an accelerometer-based system, demonstrating that it can sample activity-pertinent information which an accelerometer can not. We conclude that this ball switch can be used to recognize high-frequency activity components and effectively improve recognition rates while representing a very low cost sensor in terms of price, device size and power consumption.

1 Introduction and related work

Intelligent devices are increasingly expected to recognize their environment and situations. The most common method of fulfilling these expectations is by using acceleration sensors which are rapidly becoming ubiquitous in modern day technology. They are embedded in devices from cell phones and laptops, to every-day items such as tennis shoes and TV remote controls [3]. Their effects range from smart phones which are capable of adjusting themselves based on their orientation to devices that can recognize individual users and situations [3][9][16].

Several applications have already been developed using multiple acceleration sensors worn at different body locations to recognize different activities, i.e. [2][7][14][15]. Other examples use one single sensor location but multiple sensor modalities to recognize a variety of activities such as daily routines in [16], or a broad spectrum of activities in [3][9][18]. The resulting systems can automatically recognize and adjust to certain situations and activities without the user having to explicitly input anything after a training phase. These applications are usually wearable or mobile and must therefore be energy aware in order to avoid maintenance activity such as battery replacement or charging.

In this paper, a new approach to sensory feature creation for activity recognition in wearable computing is presented. This approach is based on a novel, low-power vibration sensor system which is used to recognize certain activities and situations while consuming significantly less power than an acceleration sensor. In [10, 11, 15], other novel sensors have also been introduced to the activity recognition and wearable community in much the same way. In this paper we will present this ball switch as a tool for context recognition. Along with a method for feature generation and information extraction specifically designed for this type of ball switch, we will also present the strengths and weaknesses of the ball switch in the context of wearable activity recognition.

The vibration sensor is a miniaturized ball switch (Fig. 1), referred to as a micro-vibrational sensor (MVS) by the manufacturer, available as a commercial, off-the-shelf device (COTS). A conductive sphere rolls between two charged plates, closing the circuit in a certain position. With a diameter of $800\,\mu\mathrm{m}$, the sphere's physical properties are different than those in traditional ball switches, especially in terms of sensitivity even to extremely low-intensity vibrations, as well as sensitivity in all three dimensions [13].

Other work done with this sensor in [6] indicated that some types of activities generate vibrations on the human body, and that the novel sensor is especially useful for detecting these vibrations. This motivated the hypothesis that the better time resolution of the vibration sensor may outweigh the better data resolution of acceleration sensors (3

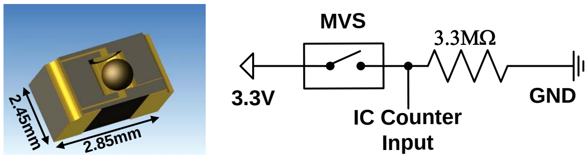

Figure 1. The MVS from [13] and schematic

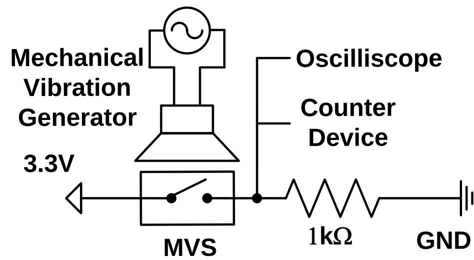

Figure 2. The vibrational analysis setup

analog acceleration vs. 1 binary vibration value) in some situations. The intention of this work is to gain a deeper understanding of the characteristics of the sensor and its suitability for specific types of activity recognition.

Traditional ball or tilt switch sensors have been used before to successfully classify activities in [15] and [16] based on the evaluation of snapshots from multiple ball switches (tilt switches) to infer limb position and attitude. In [15] an approach similar to the one presented here was attempted with multiple switch inputs to a spiking neural network with mixed results. In contrast to the above work our system uses a single but more sensitive sensor to recognize activity information directly extracted from sensing the *vibrations* on the body of the subject wearing the sensor. The approach in [15] effectively discards information generated by the ball switch between snapshots (samples). The novel methods for feature generation and information extraction presented in this paper allow us to perform continuous recognition with high resolution but with very low power consumption. In this way the dynamics in vibrations can be taken into account over a period of time with a very fine resolution even at low sampling rates, rather than relying on snapshots of the system state to recognize activities.

2 The Micro-Vibration Sensor

As indicated in [4] it is necessary to examine the physical properties of a novel sensor before being able to make a statement about what can and cannot be recognized using the output of that sensor, ergo what can and cannot be classified by a recognition algorithm. A vibration is defined as a periodic, back-and-forth motion of a body and is either a free vibration, meaning the object is disturbed and then left to its own devises until the disturbance reaches an equilibrium, or a forced vibration, where the object is maintained in motion by external forces [1].

In order to analyze the behavior of the MVS, an experiment was conducted which subjected the MVS to various different forced vibrations. Although not directly tested, this experiment will also provide insight into the behavior of the sensor under free vibrations as will be discussed later. The experimental setup involved connecting the output of the MVS to a digital counter and applying a constant voltage as input. The counter increments memory each time that a transition from 0 to 1 (positive flank) is sensed on the input line, recording the number of ball switch events which have occurred. The MVS was then placed on the membrane of a loud-speaker which was connected to a sinusoidal signal generator as indicated in Fig. 2. The speaker used was a VISATON® W 250 9067, 8 Ω 10" speaker with a frequency response range up to 6000 Hz [17]. Modulating the voltage level and wave frequency output of the signal generator creates a vibration on the ball switch with a frequency equal to that of the wave being generated and an amplitude which is a function of the frequency and voltage level output of the generator. Although some of the frequencies generated in the experiment surpassed 6000 Hz, this did not affect the results (see Section 6).

The MVS was exposed to a constant vibration in terms of frequency and amplitude for a period of 10s, during which the amount of ball switch events was recorded. Frequency and modulation were then changed and the experiment was repeated. The output of this experiment can be seen in Fig. 3, where the horizontal axis is the frequency of the wave applied to the speaker and the vertical axis is the number of events which occurred per second on average. Each data set on the graph indicates a measurement campaign with a different speaker input signal amplitude. The displacement Δ value was measured at rest for the input sine wave maximum and minimum voltages (1 Volt through 5 Volts, in steps of 1 Volt). Fig. 3 indicates an increase in sensitivity above 1.5 kHz, with another dramatic increase above 3 kHz. Below 1.5 kHz the response of the MVS is unpredictable, as vibrational frequencies may or may not illicit a response from the ball switch depending on amplitude.

These results are measured under constant forced vibrations, and would therefore not be directly valid for free vibrations. In that case the results of the experiment indicate that if a free vibration, e.g. an impulse, is partially in a frequency spectrum in the sensitive area above the approximate 1.5 kHz margin it will generate a response from the sensor whose intensity is dependent on both the amplitude of the vibration and the frequency. If the vibration is above

[1] Encyclopedia Brittanica: http://www.eb.com

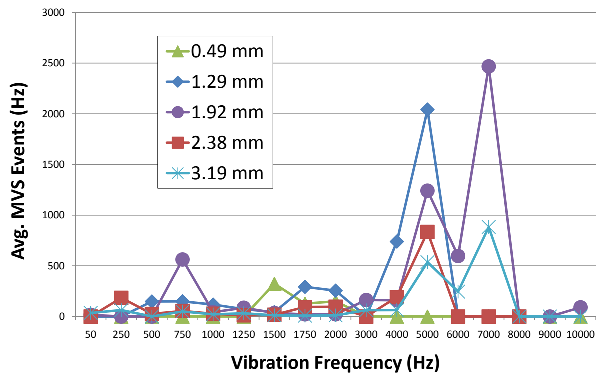

Figure 3. Frequency analysis of the MVS

or below the sensitive range, it will generate indeterministic output, meaning the sensor may or may not generate a response to the vibration depending on the amplitude and frequency of the vibration. In terms of amplitude, $.49\,\mathrm{mm}$ appears to be below the threshold of measurable displacement for the ball switch, where the optimal displacement is dependent on the vibrational frequency.

3 Data analysis and feature generation

In order to evaluate the new activity recognition techniques using the vibration sensor, sampling hardware was used which simultaneously gathered sensory data from the MVS and an accelerometer. The experiment utilizes the Akiba wireless sensor node which conducted measurements using an on-board MVS micro-vibration sensor (MVS) from Sensolute [13] and an external ADXL335 3D accelerometer (referred to as the ADXL) board from Analog Devices [1]. Each axis of the ADXL is directly connected to one of the 10bit-wide A/D ports of the processor (Microchip PIC18F14K22 [8]), and the MVS output is connected to the 16bit timer1 input as seen in Fig. 1.

This constellation allows A/D conversion and counting to run independently of the processing tasks. The sensor node conducted readings from both A/D (ADXL) and timer1 (MVS) registers at a frequency of $60\,\mathrm{Hz}$ and outputted the measurements to an external memory management unit which logged the data on a microSD card for further analysis. An in-depth analysis of the sensory device and memory management unit in terms of energy consumption and sampling methods can be found in [6].

Unlike the signal produced by the analog acceleration sensor, the output of the MVS is an asynchronous, digital, binary vector as shown in Fig. 4(1). The relevant information in these signals are the unary transitions between the two states of the signal. The vibrational data is a time-series of sequential events whose only important unit is their time

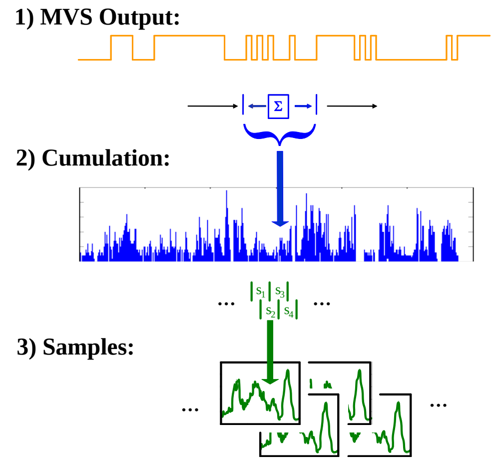

Figure 4. The MVS preprocessing algorithm

stamp, or position on the time line. These events are signaled by a voltage change on the output pin of the vibration sensor, from zero to a logical one or one to zero.

In order to be able to recognize a specific pattern within this system, namely a pattern generated by a certain activity, this signal must be converted into a form which can be analyzed using pattern matching and recognition algorithms. To create such a signal from the time-series, a cumulation method was developed which creates a wave form from the individual events. This function uses a history window to construct a wave based on the number of events in that window. The window is passed over the time line creating a new signal as depicted in Fig. 4(2). This wave, although digital in nature, can now be treated as a digital representation of an analog signal, namely the vibration levels measured by the MVS. In Fig. 4(3), this wave is cut into separate samples to be classified by the recognition algorithm as to the activity being performed.

4 Sensor hardware comparison

As mentioned in the introduction, our goal is embedded activity and context recognition in an ultra-low power wireless sensor node based on the MVS. For this reason the following power consumption analysis was conducted using the PIC18F14K22 [8] microprocessor from Microchip® based on the circuit in Fig. 1.

Power consumption The ADXL335 3D acceleration sensor was chosen because of its ease of use as well as its typical power consumption signature. In the data sheet the current drawn by the sensor is indicated to be close to $425\,\mu\mathrm{A}$ at an operating voltage of $3.3\,\mathrm{V}$. At that voltage the rate of consumption of the ADXL is $P_{ADXL} = 1.4\,\mathrm{mW}$. The schematic for the integration of the MVS 0608.02 shown in figure 1 implements a $3.3\,\mathrm{M}\Omega$ pull-down resistor and therefore pulls a total current of $1\,\mu\mathrm{A}$ at $3.3\,\mathrm{V}$. This yields a calculated consumption of $P_{MVS-calc} = 3.3\,\mu\mathrm{W}$.

The different resistive values for the MVS pull-down resistor ($3.3\,\mathrm{M}\Omega$ in Fig. 1 versus $1\,\mathrm{k}\Omega$ in Fig. 2) is due to the internal capacitance of the measurement equipment (counter device and oscilloscope for visualization) used in the experiment in Section 2. This capacitance introduced a higher rise time τ which essentially reduces the sensitivity of the ball switch, acting as a low-pass filter. For this reason a lower resistance was selected for the sensitivity analysis.

The MVS has two states as with any switch: `ON` and `OFF`. In the `ON` state the consumption is $P_{MVS-calc} = 3.3\,\mu\mathrm{W}$, but in the `OFF` state the consumption is virtually zero, since no current flows over the sensor. Due to the construction of the MVS, the sensor is in either state at any given time with a probability of 50%, meaning that the actual consumption is only half of the calculated consumption, or $P_{MVS} = P_{MVS-calc}/2 = 1.45\,\mu\mathrm{W}$. This is approximatley one full order of magnitude less than that of the acceleration sensor.

3 ADC operations are necessary to convert the measured acceleration for each ADXL axis represented in voltage to a digital value, each costing $1.2\,\mathrm{ms}$ giving a total of $3.6\,\mathrm{ms}$ when the PIC18LF14K22 is in low power mode, e.g. is clocked at 31.25 kHz. Each ADC read requires 2 `MOV` commands to transfer the 10 bit values from the SFR to memory, each costing 1 processor cycle, yields 12 processor cycles. Each processor cycle requires 4 clock cycles yields a total $1.536\,\mathrm{ms}$ per ADC read. Together, converting an analog value to a digital one and transferring it to specific location costs $T_{ADXL} = 1.536\,\mathrm{ms} + 3,6\,\mathrm{ms} = 4.368\,\mathrm{ms}$. Vibration readings and cumulation are directly carried out by a hardware component of the processor, the timer/counter. This is a low power module which operates independently from the rest of the embedded processor [8]. Reading this value, checking and accounting for overflow and subtracting the previously read value incurs on average 64 clock cycles which requires $T_{MVS} = 8,192\,\mathrm{ms}$ at $31.25\,\mathrm{kHz}$.

As the processor pulls $15.5\,\mu\mathrm{A}$, its power consumption is $P_{proc} = 51.15\,\mu\mathrm{W}$ at $3.3\,\mathrm{V}$. One accelerometer measurement lasts $T_{ADXL} = 4.368\,\mathrm{ms}$ with a consumption rate of $P_{proc} + P_{ADXL} = 1.45115\,\mathrm{mW}$. For the vibration sensor, one reading uses a total of $P_{proc} + P_{MVS} = 54.45\,\mu\mathrm{W}$. This indicates that the energy required to sample the MVS is approximately 14 times less than that necessary to sample the acceleration sensor. The validity of these calculations will be confirmed later in section 5.

It is important to note that these values will not scale indefinitely for higher clock rates of the processor, as there is a ceiling on minimum A/D conversion time due to capacitor load time, where the MVS wave construction only consists of processor register operations. This implies that for higher clock rates the ratio of power consumption between the two sensors will tip even farther in favor of the MVS, though overall system consumption will increase.

Size, cost and responsiveness The physical size of both sensors is also comparable; the MVS has a footprint of $2.45\,\mathrm{mm}$ x $2.85\,\mathrm{mm}$ where the ADXL sensor is significantly larger at $4\,\mathrm{mm}$ x $4\,\mathrm{mm}$. Both sensors require external circuitry in order to operate properly; the MVS requires one resister where the ADXL requires 4 capacitors, one for each axis and one for power stabilization.

One very large difference between the two sensors are the frequency response ranges. The ADXL has a measurement bandwidth of $0.5\,\mathrm{Hz}$ to $1.6\,\mathrm{kHz}$ for X and Y axes, and $0.5\,\mathrm{Hz}$ to $550\,\mathrm{Hz}$ for the Z axis [1]. The analysis of the MVS on the other hand revealed that the sensitive frequency ranges of the ball switch begin above the upper limits of the acceleration sensor, namely upwards of $1.5\,\mathrm{kHz}$, above a certain amplitude threshold (see Section 2).

The ADXL335 is one of the more costly acceleration sensors at about 5.50 USD with other comparable models priced as low as 3.00 USD. The MVS on the other hand is a far simpler sensor and is therefore less expensive. The current cost of an MVS (version MVS0608.02) sensor is approximately 1.75 USD, so the sensor is quite competitive, even at the lower end of the acceleration sensor pricing. The costs of the MVS can also be expected to fall as it is a relatively new device and increased production run length and volume would further reduce costs. On a side note, the MVS requires a counter input pin from the processor while the ADXL uses 3 A/D processor inputs.

5 Evaluation and results: a case study in activity recognition

In order to evaluate the MVS as an activity recognition tool, a case study was performed involving 8 different everyday activities. The data which was gathered during the course of the case study and was used for this evaluation are available on the Internet [5].

Experimental settings and parameters The measurement and logging device described in [6] was used to gather the data for this case study along with an external acceleration sensor. The measurement logging device was powered by a plastic battery pack containing two AAA batteries. The device itself was fixed at the subject's hip between the belt and the subject's pants and the belt was fastened firmly to hold the sensor in place (see Fig. 5).

In total, 5 subjects were used to create a basis for the evaluation. 8 activities were selected consisting of riding

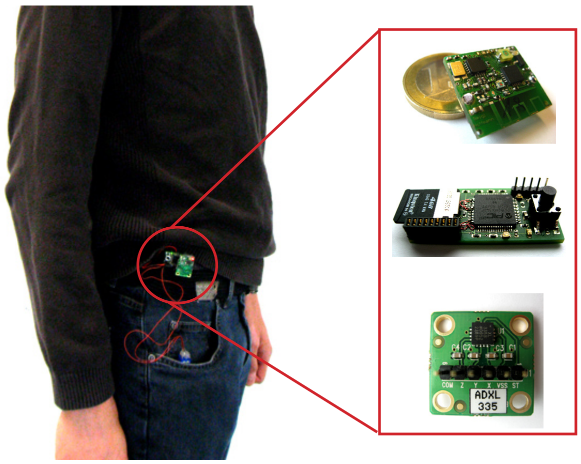

Figure 5. Subject wearing the Akiba node (top), memory extension and ADXL335 board

the bus, riding a bike, walking, jogging, riding the elevator (lift), typing while seated, climbing the stairs and standing at rest. The subjects performed the selected activities, switching the device on to record and using a button to delimit activities if necessary, creating a method for annotation after the fact. During periods where no relevant activity was being performed the device was turned off, effectively limiting the data the selected activities. Three acceleration axes, the ball switch counter, as well as light and temperature sensors were all sampled synchronously.

The subjects were computer science undergraduate students with technical backgrounds although not extensively in the field of activity recognition. Each user performed all activities sequentially, and data collection was conducted one subject at a time. In total, 142 minutes of data was collected on a university campus from 5 subjects over the course of one week for the evaluative case study.

Activity recognition The WEKA data mining toolkit [19] was selected for activity recognition for its simplification of the pattern-matching algorithms as well as its acceptance in the community [2][3][9][11][12][14]. Specifically, the C4.5 decision tree [12] was used due to its prevalence in the activity recognition literature using acceleration sensors [2][9][14] and its suitability for the intended extremely resource-restricted sensor node platform. Additionally, the IBk k-nearest neighbors and Naive Bayes classification algorithms were also evaluated in order to provide a comparison between standard recognition algorithms [4][10].

Using the samples generated by the algorithm in fig. 4(3), a set of features is generated for each sample which is used to identify the activity. The features used are identical for both the MVS and ADXL, except for the fact that the acceleration data generates 3 sets of features, one per acceleration axis. This information is not available when using the vibration sensor as only one sensor is being used and the axis of a specific vibration is, at best, very difficult to isolate and is not a part of this work. The other features generated are mean, standard deviation, entropy, area under the curve and FFT-peaks, since these were often cited as being the most decisive [2][7][9][14][18]. The three selected classification algorithms were trained by the WEKA toolkit using the activity feature sets for the vibration data on the one side and the acceleration data on the other. A sample window size of approximately 1 second with 50% overlap was selected for the case study and is constant over all classifications.

In order to evaluate the case study, 3 different classification phases were conducted. In the first phase, the classifiers were trained and tested on the data gathered from all subjects using a 5-fold approach (80% of the data is used for training and 20% for testing). The intention is to analyze how the classifiers performed if data from all subjects was present at training time. In the second phase, data collected from 4 subjects was used to train the classifiers, and the data from the remaining subject was used for testing to provide an indicator of interpersonal variances in the MVS and ADXL output respectively.

In the final phase, the effect of the MVS as a post-hoc addition to a pre-existing activity recognition system was evaluated. To show this, a classifier was trained using the acceleration, light and temperature data of all subjects 5-fold. The C4.5/J48 classifier was selected for this task because of the advantageous property of not being affected by junk features, meaning that redundant and useless information is automatically discarded at training time [12][19]. Then, the same procedure was conducted again with the addition of the MVS data. The goal of this phase is to assess how much novel information is delivered to a system when the MVS is integrated post-hoc, which would not be otherwise available using conventional sensors.

Classifier performance The results of the three separate classification phases can be seen in Table 1. The acceleration sensor performed far better than the vibration sensor in the personalized classification phase no. 1, with an average classification rate over the 3 algorithms of 84.7% as compared to slightly more than half that value for the vibration sensor. The results of phase no. 2 indicate that the ADXL only slightly outperformed the MVS in this phase with a classification rate of 30.5% on average. In general, the k-nearest neighbors classifier is par with the decision tree, where the Naive Bayes classifier performed poorly compared to the other classifiers. Phase no. 3 indicates a 4% increase in overall system classification rates from 92.8% to 96.6% when the ball switch features were included. An

Phase	Type	IBk	J48	Bayes	Average
No. 1	Personalized MVS	46.2	49.2	34.1	43.2
	Personalized ADXL	91.9	96.6	65.6	84.7
No. 2	Generalized MVS	36.1	34.0	21.4	30.5
	Generalized ADXL	23.0	34.1	53.4	36.8
No. 3	ADXL, Light, Temp.		92.8		
	ADXL, Light, Temp., MVS		96.6		

Table 1. Results of the evaluation in percent

activity per activity comparison between the classification rates of the ADXL and the MVS has been omitted here as the rates for the ADXL were relatively even across all activities and outperformed the MVS.

Power measurements In order to confirm the calculations done in section 4, measurements were conducted using a BBC Goerz Metrawatt measurement device in a laboratory setting. These measurements were performed without the data logging unit. Each sensor was connected and sampled individually in an endless loop under heavy agitation to mimic activity, and current flow was measured to quantify power consumption. Processor activities performed for the ADXL and MVS were conducted as described in sections 3 and 4. In one cycle (sensor measurement, subsequent processing), an average current flow of $630\,\mu\mathrm{A}$ for the ADXL, and $12.8\,\mu\mathrm{A}$ for the vibration sensor was measured. At $3.3\,\mathrm{V}$ this yields power consumption rates of ca. $2.08\,\mathrm{mW}$ for the ADXL (172.8 J/day) and $42.24\,\mu\mathrm{W}$ for the MVS (3.5 J/day). The lifetime with a watch-type coin cell (CR1620, $1kJ$) would equate to 6 days using the ADXL and 285 days using the MVS in worst case when assuming 24/7 activity of the user. These results show that the MVS would reduce the total measured consumption of the sensor node system by a factor of almost 50 when compared to the ADXL. The difference between the calculated and measured values (MVS: $2.08\,\mathrm{mW}$ vs. $1.45\,\mathrm{mW}$ and ADXL: $0.04\,\mathrm{mW}$ vs. $0.054\,\mathrm{mW}$) is due to the difference between the consumption rates of the processor, A/D and timer unit in the preliminary data sheet and that which was measured. This disparity can either be attributed to measurement device calibration or a documentation error.

6 Discussion

The MVS sensor The vibrational analysis of the MVS sensor provides an insight into the usefulness of the sensor for activity recognition purposes. The unpredictable behavior of the sensor at low frequencies seems to indicate that the MVS is not appropriate for measuring slow, rounded activities. On the other hand, the high sensitivity of the device from $1.5\,\mathrm{kHz}$ to $8\,\mathrm{kHz}$ indicates that the device is very adept at sensing *impulses*. This is due to the fact that impulses or impacts cause a wider range of vibrational frequencies which decay over time seeking equilibrium, and that this decay passes over a wider band of intensity and frequency.

At this point it is important to realize that certain factors in the vibrational FFT analysis of the ball switch may be affecting the results. Firstly, the frequencies applied to the speaker were partially above the optimal frequency response level of the speaker. According to the manufacturers specification [17] this does not distort the vibrational frequency, though it could alter the amplitude of the signal, although to which point the values given by the manufacturer (in dB in reference to the noise level generated) reflect the intensity of the vibrations of the membrane is unclear. Furthermore, it is unclear exactly how much the fluctuating magnetic field generated by the speaker affects the ball switch on the membrane. An additional experiment was conducted by removing the membrane of the speaker and subjecting the ball switch to the magnetic flux without the physical vibrations. This resulted in no MVS events, even at settings that incurred high activity under the normal settings (i.e. $5\,\mathrm{kHz}$ at $1.92\,\mathrm{mm}$), indicating no great affect, although a minor one could still not be ruled out.

Ramifications for activity recognition Activities which produce vibrations within the sensitive spectrum will elicit a response from the switch for the time period for which the vibration remains in that spectrum (free vibrations will decay over time as they approach an equilibrium). For this reason we propose that the ball switch is capable of providing classifiable information for activities that contain events or impulses which produce frequencies above $1.5\,\mathrm{kHz}$ in the human body. This is confirmed by Table 2 which contains the confusion matrix from a personalized classification using the C4.5 (J48) decision tree classifier over the vibration data. The activity jogging contains a series of periodic concussions (footfalls) which stimulate the MVS. For this reason jogging was recognized by the system 79.1% of the time, walking 57.6% of the time and climbing the stairs 47.1%. Another example is the activity of riding a bike, which when conducted outdoors on an uneven surface (as was the case) consists of a series of impacts or free vibrations as the wheels encounter obstacles on the ground, combined with periodic, forced vibrations from peddling. The high frequency free vibrations allow bike riding to be classified over the ball switch feature generation, yielding a recognition rate of 49.2%.

The results also indicate that the ball switch is not suitable for tasks such as gesture recognition, which often rely on the relatively low frequencies[11]. This is especially true when these gesture do not involve impulses, impacts or collisions, but are rather rounded motions such as waving or swiping. Rounded motions will produce low frequency output where the sensitivity is very much a function of the dis-

a Bus	b Bike	c Walk	d Jog	e Lift	f Type	g Stair	h Stand	
27.1	6.5	3.7	0.4	10.1	40.9	4.5	6.7	**a**
9.1	49.2	12.5	0.9	5.5	2.5	17.2	3.0	**b**
2.1	4.7	57.6	8.4	5.9	0.3	20.8	0.2	**c**
0.6	0.9	9.9	79.1	1.8	0.2	7.3	0.3	**d**
7.2	3.6	11.3	1.5	26.0	35.6	10.8	4.0	**e**
2.6	1.5	0.8	0.4	1.4	90.9	0.8	1.5	**f**
3.5	7.6	21.9	9.8	9.0	0.6	47.1	0.6	**g**
5.8	2.0	1.0	0.5	5.2	77.8	1.1	6.7	**h**

Table 2. Confusion matrix in percent from phase no. 1 for the MVS

	Acceleration	Vibration (MVS)
Power consumption[1]	2 mW	42 uW
optimal recog. freq. range	1-100 Hz[3]	3 kHz-8 kHz
Resolution	3D 10 bit[2]	1 bit
Size	16 sq. mm	7 sq. mm
Suggested use	low-mid. freq. activities personalized detection	high frequency activities unpersonalized detection

[1] at 3.3V
[2] for most small microprocessors
[3] due to slow A/D in small microprocessors. Sensor max is 1.6kHz

Table 3. The combined insight of this paper

placement of the vibration (see Fig. 2). This is evident in the classification rates for activities which generate low-frequency vibrations such as standing (6.7%), riding the elevator (26%) and riding the bus (27.1%).

An interesting phenomenon is noticeable when observing the activity of typing, where a recognition rate of over 90% was achieved. This would appear to indicate that the sensor is well suited to recognize typing as an activity, when actually this is not the case. Indeed, what occured is that when subjects were typing, often no or very few events were generated by the MVS at all, causing an activity to be classified as typing during periods of no activity. This is evident when examining which other activities were confused with typing: standing (77.8%), riding the bus (40.9%) and riding the elevator (35.6%).

This is due to the fact that these activities generate low-frequency vibrations which often produce little or no activity from the MVS. As typing is an activity which consistently produces almost no output, all of the sample features which do not contain any ball switch events are classified as typing, explaining the high confusion rates. The implication is that an activity recognition system which is based on the MVS would benefit by having a "Zero" class into which all sample windows are classified which do not contain any, or only very few events. This would differentiate between activity samples which have been classified and those which simply did not generate enough vibrations to be classified. A possible method for handling such cases would be to increase the sample and cumulation window lengths, which under certain conditions would reduce the number of samples with 0 events, though at the cost of reduced reaction time.

The results of the three-phase classification study demonstrate that the acceleration sensor is capable of delivering quantatively more information of relevance for activity recognition when compared to the ball switch. This can be seen clearly when observing phase no. 1 of the case study where personalized classification using the ADXL was significantly more successful than the MVS for the same activities (84.7% compared to 43.2%).

Phase no. 2 on the other hand, indicates that much of this data is largely subject dependent, making it less useful for a generic monolithic approach to context recognition. In this phase the performance of the classifiers dropped for both the vibration and acceleration data, whereby the reduction in recognition rates on average for the vibration data is significantly less than the acceleration data (29.4% for the MVS versus 56.6% for the ADXL). This would indicate that although the vibration sensor delivers less data than the acceleration sensor, the data is more generic per activity across multiple subjects.

The results from phase no. 3 show an improvement of over 4% in a 3 sensor activity recognition system when the MVS is introduced into the system. This confirms the results of the frequency analysis of the ball switch in Section 2, which indicate that the MVS and the ADXL have complimentary sensitivity ranges in terms of frequency bands and therefore provide activity data which is also complimentary in nature with some overlap.

Lastly, the vibrations which are being measured using the MVS are not usually being generated at that location, but rather these signals must propagate through the human body before arriving at the sensor. This would indicate that sensor location is a crucial aspect when using the MVS for activity recognition as each activity would create a different vibration pattern at a different location, depending on what types of tissue the vibration propagates through. This would suggest that classification rates are only valid for the location where the data was sampled, e.g. are highly location dependent.

7 Conclusion

This paper showed the potential of a novel vibration sensor as a tool for continuous, low-power, wearable activity recognition. Table 3 gives an overview of the characteristics of the vibration sensor system and its use in activity recognition, and presents a comparison with activity recognition based on a 3D acceleration sensor.

We determined that the MVS is capable of sensing activity data pertinent to standard recognition algorithms. On the one side, the MVS does not deliver as much information as the ADXL acceleration sensor, as 3-dimensional information is not differentiable in the vibration data and sensitivity at lower frequencies relatively low. The MVS would therefore be less useful when used for sensing slow movements or activities (under $1.5\,\mathrm{kHz}$), and would not be suitable for gesture recognition of this type. On the other hand, the resolution of the MVS far surpasses that of the ADXL, and vibrational sensitivity in the upper bands (specifically $3\,\mathrm{kHz}$ to $8\,\mathrm{kHz}$) surpasses that of the acceleration sensor. The MVS can therefore be used well to recognize activities which contain concussions and impacts such as jogging, riding a bike on uneven ground, or presumably tapping on a hard surface. Furthermore, the results indicate that the MVS can generate sensory information which can be better generalized over multiple subjects using a generic monolithic classifier approach.

Finally, we evaluated the MVS as an addition to existing activity recognition systems based on standard sensors including acceleration. The work presented in this paper indicates that the MVS can improve recognition rates while costing one third as much as an ADXL acceleration sensor, taking up one half the size, and consuming 50 times less power. All of this makes the MVS a resource-effective, simpler alternative to, or extension of, acceleration sensors for low-power, low-cost wearable activity recognition systems for researchers and developers. As the acceleration and MVS based recognition performs significantly better than just acceleration based recognition, there is strong evidence that high-frequency vibrational signals generated by everyday activities is very useful for activity recognition, and that the MVS is capable of sampling that information.

Acknowledgments The authors would like to acknowledge joint funding by the European Commission under the ICT project “CHOSeN” (Project No. 224327, FP7-ICT-2007-2) and by the Deutsche Forschungsgemeinschaft (DFG) under the project “SenseCast ” (BE4319/1).

References

[1] Analog Devices Inc. The ADXL335 3-Dimensional Acceleration Sensor Specification. http://www.analog.com/, 2010.

[2] L. Bao and S. S. Intille. Activity recognition from user-annotated acceleration data. In *Pervasive*, pages 1–17, 2004.

[3] K. Chang, J. Hightower, and B. Kveton. Inferring identity using accelerometers in television remote controls. In *Proceedings of the Seventh International Conference on Pervasive Computing*, Nara, Japan, May 2009.

[4] D. Gordon, H. Schmidtke, and M. Beigl. Introducing new sensors for activity recognition. In *How To Do Good Research In Activity Recognition: Experimental methodology, performance evaluation and reproducibility. Workshop in conjunction with Pervasive 2010*, Helsinki, Finland, May 2010.

[5] D. Gordon, G. von Zengen, H. R. Schmidtke, and M. Beigl. A data set from wearable accelerometer and vibration sensor. `http://www.teco.edu/~gordon/MVS/`, May 2010.

[6] D. Gordon, F. Witt, H. Schmidtke, and M. Beigl. A long-term sensory logging device for subject monitoring. In *Workshop for Situation Recognition and Medical Data Analysis in Pervasive Health Environments (PervaSense 2010)*, Munich, Germany, March 2010.

[7] J. Lester, T. Choudhury, and G. Borriello. A practical approach to recognizing physical activities. In *Pervasive*, pages 1–16, 2006.

[8] Microchip Technology Inc. The PIC18F14K22 Microprocessor Product Specification. http://www.microchip.com/, 2010.

[9] E. Miluzzo, N. D. Lane, K. Fodor, R. Peterson, H. Lu, M. Musolesi, S. B. Eisenman, X. Zheng, and A. T. Campbell. Sensing meets mobile social networks: the design, implementation and evaluation of the cenceme application. In *SenSys ’08: Proceedings of the 6th ACM conference on Embedded network sensor systems*, pages 337–350, New York, NY, USA, 2008. ACM.

[10] G. Ogris, M. Kreil, and P. Lukowicz. Using fsr based muscule activity monitoring to recognize manipulative arm gestures. In *ISWC ’07: Proceedings of the 2007 11th IEEE International Symposium on Wearable Computers*, pages 1–4, Washington, DC, USA, 2007. IEEE Computer Society.

[11] G. Pirkl, K. Stockinger, K. Kunze, and P. Lukowicz. Adapting magnetic resonant coupling based relative positioning technology for wearable activitiy recogniton. In *Wearable Computers, 2008. ISWC 2008. 12th IEEE International Symposium on*, pages 47–54, 28 2008-Oct. 1 2008.

[12] J. R. Quinlan. *C4.5: Programs for Machine Learning (Morgan Kaufmann Series in Machine Learning)*. Morgan Kaufmann, 1 edition, January 1993.

[13] Sensolute GmbH. MVS 0608.02 Micro-Vibration Sensor Data Sheet. http://www.sensolute.com/, 2010.

[14] E. M. Tapia, S. S. Intille, W. Haskell, K. Larson, J. Wright, A. King, and R. 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.

[15] K. Van Laerhoven and H.-W. Gellersen. Spine versus porcupine: A study in distributed wearable activity recognition. In *ISWC ’04: Proceedings of the Eighth International Symposium on Wearable Computers*, pages 142–149, Washington, DC, USA, 2004. IEEE Computer Society.

[16] K. Van Laerhoven, D. Kilian, and B. Schiele. Using rhythm awareness in long-term activity recognition. *Wearable Computers, IEEE International Symposium*, 0:63–66, 2008.

[17] VISATON. Visaton W-250-9067 Loudspeaker Product Specification. http://www.visaton.com, 2010.

[18] E. Welbourne, J. Lester, A. Lamarca, and G. Borriello. Mobile context inference using low-cost sensors. In *In Proceedings of LoCA*, pages 254–263. Springer-Verlag, 2005.

[19] I. H. Witten and E. Frank. *Data mining : practical machine learning tools and techniques*. Elsevier, Morgan Kaufman, Amsterdam [u.a.], 2. ed. edition, 2005.