Empirical Study of Design Choices in Multi-Sensor Context Recognition Systems

Nagendra B. Bharatula[1], Mathias Stäger[1], Paul Lukowicz[2], Gerhard Tröster[1]

[1] Wearable Computing Lab, ETH Zürich, Switzerland
[2] Institute for Computer Systems and Networks, UMIT Hall, Austria
{bharatula, staeger, troester}@ife.ee.ethz.ch, paul.lukowicz@umit.at
http://www.wearable.ethz.ch

Abstract. This paper deals with the design and implementation of a highly miniaturized, multi-sensor context recognition system. It represents an optimal trade-off between power consumption and recognition performance rather than straightforward maximization of the recognition rate.
We present a thumb-sized, 8 gram platform that combines sound, acceleration and light sensing with processing power, wireless communication and a battery. Based on this platform we make an experimental evaluation of design choices present in such multi-sensor context recognition systems. We introduce a design method to achieve an optimal power consumption vs. recognition rate trade-off through variations of the sampling rate, feature selection and choice of classifiers. Power consumption analysis indicates that our system can operate for 300 hours without having to recharge the battery.
An important and somewhat surprising result of our analysis is that the addition of a sensor may be a power efficient way to improve the overall system performance.

1 Introduction

Wearable Computing encompasses a wide variety of systems ranging from smart badges, through intelligent textiles to backpack worn high end computers. One important part of the wearable computing vision is that of miniaturized sensor networks seamlessly integrated in different parts of the user's outfit including parts of the clothing and accessories such as jewelry, watches, keychains, glasses etc. It has been shown by many researches that such body worn sensor networks can provide valuable information about user context and activity. This information can enhance the usefulness and usability of a variety of mobile devices [1]. Examples include assessing the user interruptability [2], improving user interfaces, and providing a proactive services such as automatic delivery of appropriate manuals to a maintenance worker. Thus body worn sensor networks can be considered as essential infrastructure needed to enhance the functionality of a whole range of mobile and wearable devices and applications.

Seamless integration in clothing and accessories invariably means that the sensor nodes have to be small and unobtrusive. Placing a sensor inside of a ring effectively means that only a couple of cubic millimeters of space is available for the system. Since in most cases a wired connection to an external power source is not desirable (cables running from a ring are unacceptable for most situations), this space has to accommodate not only the electronics, but also the battery and/or a power generation system. Thus reducing the power consumption of the system is a major design objective. In fact without the ability to run the sensor nodes at sub-milliwatt power levels the whole concept of body worn sensor networks is not feasible.

1.1 Paper Scope and Contributions

The design and implementation of low power, miniaturized sensor systems is one of the main research areas of our group. Previous work included the implementation of experimental sensor systems [3, 4], analysis and design of the electronic packaging concepts [5] and the development of systematic design methodologies for such systems [6]. Investigations on the influence of different system parameters (e.g. sampling rate, resolution) on context recognition accuracy [7, 8] have also been made. A key idea developed through our work was to include power consumption concerns at all layers of system design: starting with the hardware and going up to the choice of features and classification methods. For the latter this means that the purpose is not just to achieve the best possible recognition rate, but rather the best trade-off between power consumption and recognition rate. In most cases, the trade-off has to be determined empirically for a specific application domain as part of system training.

In [7, 9] we have applied such a power consumption oriented system design to a sound based context sensor. In this paper we extend this approach to a highly miniaturized multi-sensor activity recognition node. First, we present a thumb-sized, 8.2 gram hardware platform that contains two accelerometers, a visible light sensor and a microphone, processing logic, a low power wireless transceiver and a rechargeable battery (Sect. 2). With all components constantly running at full power it consumes just about 8.5 mW and can run off the integrated battery for 57 hours. Secondly, we present a detailed experimental investigation of opportunities for power savings in embedded context recognition systems. Based on an office/household scenario we show how an optimal trade-off between recognition rate and power consumption can be found (Sect. 3.2 – 3.5). The analysis includes the optimization of the sampling rate, the computed features and the classifiers. The final system has an average power consumption of 1.6 mW leading to a battery life time of 300 hours, as shown in Sect. 4.

An interesting and important result of our study is to show that with appropriate feature choice, an additional sensor (in our case the light sensor) can improve the recognition rate with nearly no cost in terms of power. This emphasizes the point that designing for low power needs an integrated approach ranging from hardware specification to the tuning of classification algorithms.

1.2 Related Work

Many research groups have built their own sensor boards that contains one or more sensors. Examples include the Smart-Its [10], Smart Badge [11], Ubisensor

[12] or the TEA Device [13]. However the focus is usually not on small and low-power hardware.

Accelerometers are most commonly used for activity recognition. Widely studied activities are walking patterns like walking, standing, sitting or climbing/descending stairs [4, 14–16] More than 20 daily life human activities were recognized using features from multiple accelerometers heterogeneously distributed over the wearer's body by L. Bao et.al. [17]. They compared the performance of several algorithms with respect to recognition rates without major emphasis on power consumption in the system. Combining the data from acceleration sensors with an additional sensor such as a microphone was studied in [18], where major emphasis was on the recognizing the tasks with improved accuracy.

Energy and performance considerations at different hardware layers in miniaturized sensor nodes were investigated in the Smart Dust project [19]. However, algorithms for context recognition tasks and their complexity were not considered. System design approaches to power aware mobile computers were analyzed in [20]. An other systematic high level approach for designing distributed wearable systems can be found in [6]. First ideas for a more hardware oriented approach, that tries to minimize the complexity of the algorithms in respect of the hardware resources, were reported in earlier papers by our group [8, 9].

This paper presents a more advanced stage of the methodology. It describes the steps required to obtain an optimized algorithm which respects the recognition rate vs. power consumption trade-off. We explicitly analyze the use of a cheap sensor (in terms of power consumption) to improve the overall recognition rate without increasing the power consumption considerably – something that to the best of our knowledge has not been analyzed before.

2 Low Power Hardware Design

As described in the introduction, the main goals of our design were lowest possible power consumption and miniaturization of the platform. Together with our experience on the usefulness of different sensor types for context recognition, this has determined the choice of hardware. The system layout is show in Fig. 1. The two accelerometers and the microphone were included since in the past they had proven to be useful for user activity recognition [18, 9]. The light sensor was added as a power-wise very cheap sensor on the assumption that it might allow us to reduce the amount of data and feature complexity needed from the other more power hungry devices. Moreover, the system is also an intermediate step towards the implementation of a future ultra low-power sensor button, presented by our group in [5]. Therefore the same sensors as in this design study were used.

Hardware description: We have compared different candidates before considering the final specification in Table 1. The analog signals from the microphone and the accelerometers are low pass filtered with a 3 dB cut-off frequency of $f_c = 1.4\,\text{kHz}$ and $f_c = 50\,\text{Hz}$, respectively and fed into a 12 bit 8-channel analog to digital converter (ADC) which is controlled by the microcontroller. The

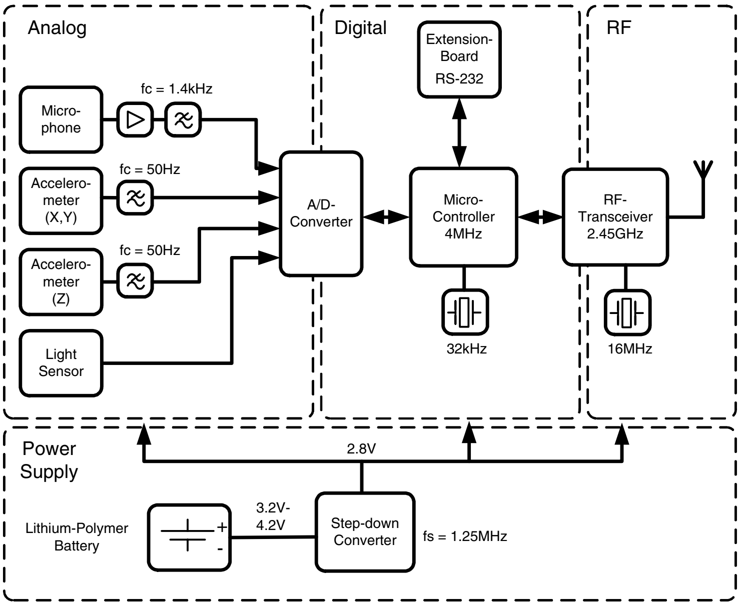

Fig. 1. Overall system architecture of the sensor button

microcontroller does some local preprocessing and forwards the data to the wireless transceiver. In a second version of the board we replaced the microcontroller with the slightly larger, but more powerful model MSP430F1611. The required space can be compensated, since it already includes an 12 bit ADC. The clock for the microcontroller (max. 4.5 MHz) is generated by an internal digital controlled oscillator (DCO). The DCO is adjusted and stabilized by an external 32 kHz clock crystal. This allows very energy efficient low-power modes.

The nRF2401 transceiver was chosen for it's small size and because it's one of the most energy efficient on the market: due to a special burst transmit mode it requires 26 nJ/bit for transmitting and 57 nJ/bit for receiving [12]. Moreover, the required antenna space for a 2.4 GHz system is rather small ($9 \times 15\,\mathrm{mm}^2$ in our case). The whole device is powered by a lithium-polymer battery which has

Table 1. Components of the sensor board

Component	Final Specifications
Accelerometers	ADXL311(Analog Devices)
Microphone	SP0103 (Knowles Acoustics)
Visible light sensor	SFH3410(Osram)
RF Transceiver	nRF2401(Nordic Semiconductors)
Microcontroller	MSP430F123(Texas Instruments)
Analog-Digital-Converter	AD7888 (Analog Devices), 12bit
Battery	LPP402025(Varta Microbattery)

a high capacity of 130 mAh and a good form factor ($20 \times 20 \times 3.8\,\text{mm}^3$) and fits under our PCB board. To get good power efficiency, a step-down converter is used instead of a linear regulator. It provides a voltage between 2.7 to 2.8 Volts.

Overall, the system has a size of $27 \times 32\,\text{mm}^2$, a thickness of 9 mm and weights 8.2 grams, including battery (see Fig. 2).

Additional Hardware: The back side of the hardware provides a 10 pin connector which allows to stack the sensor board with an extension board. So far, we have implemented a programming board, a second wireless board with a RFM DR3001 transceiver and a board which contains a RS-232 to USB converter. This last board allows a transfer rate of up to 921.6 kBaud to a PC.

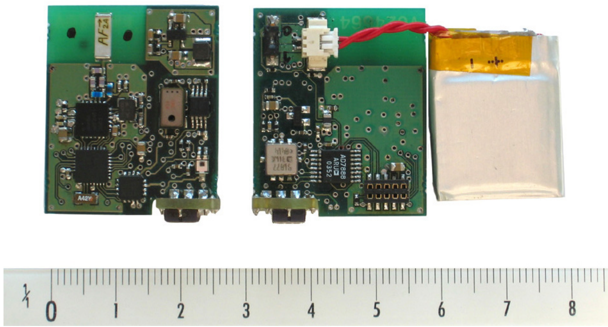

Fig. 2. Picture of the hardware. Left: top view; Right: bottom view with battery

3 Power-Optimized Recognition-Method Design

Given a fixed hardware platform, the next stage in the design of a power optimized system is the choice of the recognition method. The parameters that can be modified to reduce power consumption include the choice of sensors, their sampling rate, the type of features extracted from the signal and choice of classifier. The aim of the design process is to find out how these parameters affect the recognition performance and then select a point where acceptable performance is achieved with minimal power consumption.

As a consequence the first step in our investigation is to inspect the sensor data and to check whether a reduction in bit resolution or sampling frequency is possible. Next, different combinations of features and classifiers are evaluated for their recognition performance. In a third step, the computational complexity of the features are compared to each other. Finally, by combining the two metrics – recognition rate vs. computational complexity – the optimal recognition method is chosen.

3.1 Experimental Scenario and Setup

Since we had studied sound based recognition in much detail in previous research [9] we have concentrated on acceleration and light for this work. In particular, we wanted to find out if and how a cheap (in terms of power consumption) sensor like a photo diode can contribute in improving the recognition performance. Thus we have chosen a scenario where there was reason to believe that both acceleration and light information would be useful. This consisted of the following set of 6 daily office or household activities:

- fast typing on a keyboard
- moving a computer mouse
- writing on a whiteboard
- lifting a cup and drinking from it
- opening a cupboard
- opening a drawer

Three test subjects with the sensor board mounted on their right wrist were asked to perform each of the activities at least 20 times. The experiment was then repeated on another day with the same 3 test subjects under similar conditions. The last activity 'opening a drawer' was intentionally excluded from the feature and classifier optimization process in Sect. 3.3. This activity was used to performed a cross check on the selected features in the end. Using smaller data sets reduces the significance of the results (e.g. which features are important). Since the goal is to give insight into different design options rather than to present globally valid results, our small set is justified.

Accelerometers were sampled at 100 Hz and the light sensor at 5 Hz. The unprocessed data was recorded on a laptop using the RS-232 to USB module. The collected data was pre-segmented and labeled by hand. All further post-processing and classification was done on a PC using Matlab and Weka Toolbox [21].

3.2 Reducing the Amount of Input Data

Although this step is very obvious, it is often forgotten and recognition algorithms unnecessarily need to deal with high resolution, high sampling-frequency signals from multiple (often redundant) data sources.

Sampling Frequency and Bit Resolution: In our case a first visual inspection on the acceleration data showed that no significant information is contained in frequency above 15 to 20 Hz. This complies with the result in [8]. Following the Nyquist criteria, first trials were made with resampled signals at $f_{s\_acc} =$ 40 Hz. The sampling rate of the light sensor was kept at 5 Hz. Similarly, we tried different bit resolutions, but in our case 12 bit proved to be useful, since amplitude of the acceleration signals are quite different for different classes (e.g. for typing on a keyboard and for opening a cupboard).

Number of Data Streams: To further simplify the algorithms, only one axis of the accelerometers was used. Namely, the axis pointing in the direction of the thumb – perpendicular to the arm, in the plain of the hand.

3.3 Finding an Optimal Set of Features and Classifiers

The next step deals with finding features and classifiers that give a high recognition accuracy. Although very complex features are ruled out, the complexity is in general not taken into account yet.

Features: The features that were considered are shown in Table 2. While there are more complex features around that can be used for context recognition [16, 22, 23], only simple features were considered in order to keep the overall complexity of the algorithms at a minimum. As it will be shown, still good recognition rates can be achieved with a selection of these features. The features were applied to a 4 second window of the data-stream without looking for a segmentation or a start of an activity. Since our test subjects have repeated the activities continuously (with short breaks in between), this seamed a valid approach.

Table 2. List of simple features considered in time and frequency domain

	Feature	Symbols used	
time	Mean	mean_{acc},	mean_{li}
	Standard deviation	std_{acc},	std_{li}
	Variance	var_{acc},	var_{li}
	Fluctuation of amplitude	fluc_{acc},	fluc_{li}
	Zero Crossing Rate	zcr_{acc},	zcr_{li}
	Mean Crossing Rate	mcr_{acc},	mcr_{li}
	Gradient	grad_{acc},	grad_{li}
	Mean of gradient	mgrad_{acc},	mgrad_{li}
	Short time average energy	energ_{acc},	energ_{li}
frequency	Bandwidth	BW_{acc},	BW_{li}
	Frequency Centroid	FC_{acc},	FC_{li}
	Spectral flux	FLUC-S_{acc},	FLUC-S_{li}
	Spectral Rolloff Frequency	SRF_{acc},	SRF_{li}
	Band Energy Ratio	BER_{acc},	BER_{li}

Optimal Feature Set: An initial test showed poor recognition results (around 50%) with the frequency domain features alone. Subsequently, it was decided to use only time domain features. A mixture between time and frequency domain features would have been possible but for complexity reasons (Sect. 3.4) this idea was abandoned.

Different feature sets were defined using the Weka toolbox [21]. Starting with a full set that contained all features, subsets were calculated that are highly correlated with the class and show a low intercorrelation between the individual features. Table 3 shows the resulting feature sets.

Classifiers and Metrics for Performance Evaluation: It was assumed that each activity is equally probable and therefore an overall recognition rate is a useful

Table 3. Optimized Feature sets

Feature Set	No. of Features	Features
F1	9	mean$_{acc}$, fluc$_{acc}$, mcr$_{acc}$, std$_{acc}$, mean$_{li}$ mgrad$_{acc}$, fluc$_{li}$, mcr$_{li}$, std$_{li}$
F2	5	mean$_{acc}$, fluc$_{acc}$, mcr$_{acc}$, std$_{acc}$, mean$_{li}$
F3	4	mean$_{acc}$, fluc$_{acc}$, mcr$_{acc}$, mean$_{li}$
F4	3	mean$_{acc}$, fluc$_{acc}$, mean$_{li}$
F5	3	mean$_{acc}$, fluc$_{acc}$, mcr$_{acc}$
F6	2	mean$_{acc}$, fluc$_{acc}$,

and meaningful metrics for the performance comparison of classifiers. Otherwise, other metrics like recall, precision or false positives need to be considered [24]. The optimized feature sets were tested with different classifiers: the C4.5 decision tree [25], a k-Nearest Neighbor (k-NN) classifier (with $k = 5$), Bayes Net and Naive Bayes classifier [21] with 10 fold cross-validation. The results are shown in Fig. 3. C4.5 and k-NN provide the best results which is consistent with work from other groups [17].

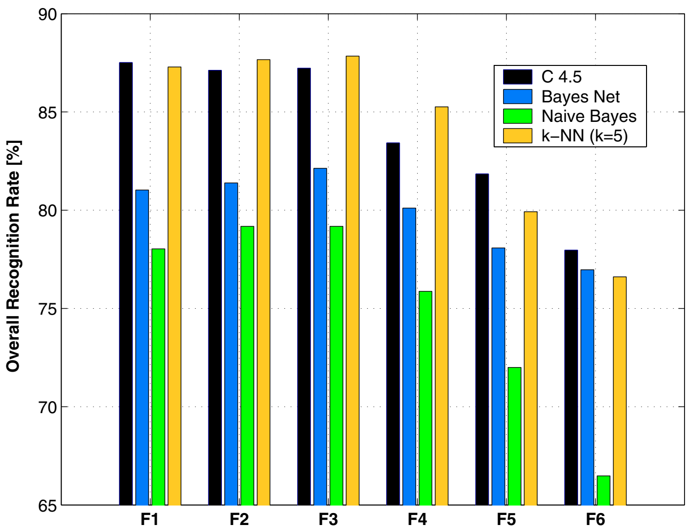

Fig. 3. Classifier performance for time domain feature sets at $f_{s\_acc} = 40\,\mathrm{Hz}$

Validation of the Original Assumptions: To validate the original assumption that 40 Hz sampling frequency of the acceleration signal is high enough to give good results, we calculated the recognition rate for feature sets F1 to F6 for sampling frequencies f_{s_acc} from 5 to 100 Hz with the C4.5 classifier. The results are depicted in Fig. 4. First, it can be seen that even lower sampling frequencies than 40 Hz would have been possible without huge loss in recognition rate. On the other hand, a maxima at 40 Hz can be observed for feature sets F2 and F3.

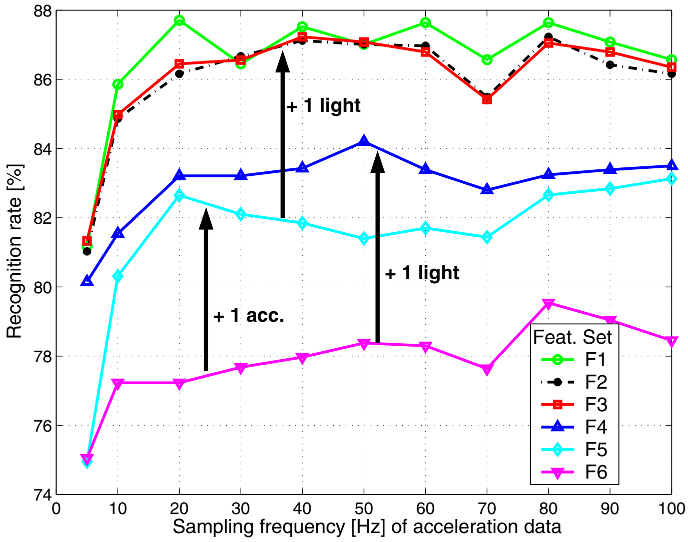

Fig. 4. Recognition rate for different feature sets vs. sampling frequency

Fig. 4 shows that at one point the recognition rate saturates and even taking larger feature sets does not increase the recognition rate considerably (F1, F2, F3). Furthermore, the figure confirms our hypothesis, that by adding a feature from a different sensor (in our case mean_{li}) the recognition rate can be greatly improved (indicated with arrows). Comparing the plots for F4, F5 and F6 it can be seen that the feature from the light sensor even adds more information than an additional feature from the accelerometer. Further simulations showed that no maximal recognition rate can be achieved with only acceleration features. Higher recognition values always contained light sensor features.

3.4 Reduction of Computational Complexity

Metrics for Complexity and Power: Although it is clear how to select features and classifier based on it's recognition performance, it's difficult to choose a specific algorithm that has low complexity with an emphasis on power consumption. We address this problem by calculating the number of instructions for each feature. Table 4 shows the computational complexity associated with some of the features as a function of number of sampling points N, which is of course proportional to the sampling frequency for a given time window. The instruction count for the frequency domain features does not include the complexity to calculate an FFT.

Table 4. Computational complexity based on instruction count for different features

	Feature	ADD	MUL	BRANCH	other
time	mean	$N-1$	1	0	0
	fluc	$2N-1$	$N+4$	0	200
	std	$2N-1$	$N+4$	0	0
	mcr	$\frac{5}{4}N$	N	$N-1$	0
	mgrad	$2N+1$	0	0	0
frequency	FC	$2N-2$	N	0	36
	BW	$5N-4$	$5N$	0	72
	FLUC-S	$2N-2$	$N+4$	0	144
	SRF	$\frac{7}{4}N-1$	$\frac{7}{4}N+1$	$\frac{3}{4}N$	0

As far as classifiers are concerned, C4.5 is favored over k-NN since it is less complex to compute [9]. In our case, the C4.5 trees resulted in only about 30 leaves which can be translated to roughly 120 instructions.

If the sensory data originates from different sensors, like in our case, the computational complexity is not the only metrics that needs to be taken into account. Size and power consumption of the sensors need to be considered as well. However, in case of an accelerometer and a light sensor, the accelerometer clearly dominates in both terms and we therefore neglect this metrics.

Overall Instruction Count: Certain assumptions have to be made, to get an overall instruction count which can be considered proportional to the consumed energy. Those assumptions depend very much on the chosen hardware architecture. Energy per operation ranges from 10 pJ for basic processors [19] and ASICs to 1 nJ for low to medium performance CPUs [6]. Furthermore, a ratio between energy per operation for additions, multiplications, etc. needs to be found. We assumed that a multiplication needs 4 times more energy than an addition or a branch instruction. This can be motivated by internal hardware design of simple array multipliers which are composed of half-adders and full-adders [26]. Simulations show a ratio of 1:3.6 for energy per operation between an adder and a multiplier [27], therefore our assumption seems justified.

Based on Table 4 the computational complexity of the feature sets were calculated. The fact, that certain features can be reused to calculate other ones (e.g. mean for mean crossing rate) was taken into account here. Fig. 5 shows the total number of instructions as a function of the sampling frequency of the accelerometer. Since the light sensor is only sampled with 5 Hz, processing it's data adds very little to the overhead complexity, so that the difference between F6 and F4 or between F5 and F3 is not visible.

3.5 Trade-off: Recognition Rate vs. Instruction Count

Combining the results from Sect. 3.3 and 3.4 – especially Fig. 4 and Fig. 5 – the recognition rate can be plotted against the overall instruction count, which

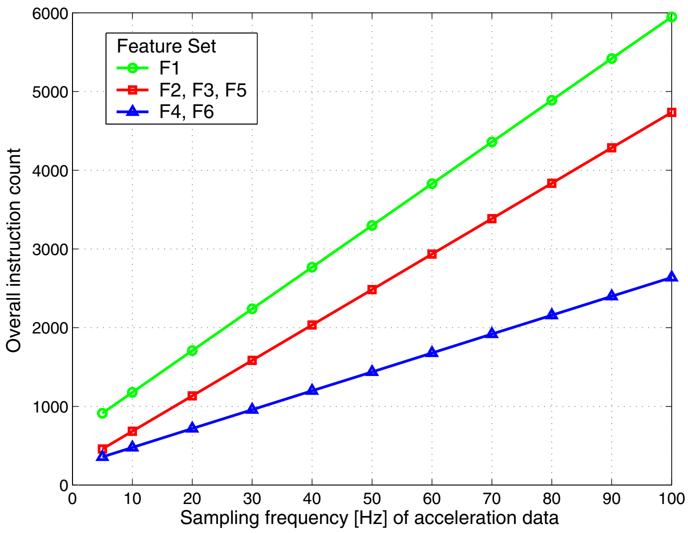

Fig. 5. Instruction count for different feature sets vs. sampling frequency

this time also includes the complexity for the classifier – hence the change in nomenclature from 'F..' to 'Algo..' in Fig. 6. With a large diversity of features and classifiers, a pareto-front could be plotted (indicated by dotted line) and the optimal point chosen. In this case it would be feature set F3 with a C4.5 classifier.

Here again, we see that an algorithm with includes information from the light sensor wins in terms of recognition rate AND complexity. Clearly, a multi-sensor platform such as ours is justified in context recognition tasks.

Confusion Matrix and Validation: With the algorithms fixed, the individual recognition rates of the 5 activities can be computed from the confusion matrix, given in Table 5 (with $f_{s\_acc} = 40\,\mathrm{Hz}$).

Table 5. Confusion matrix for Algo3 with 5 classes

a	b	c	d	e	← classified as	Accuracy
345	5	0	2	27	a=mouse	91.03%
3	547	2	92	4	b=whiteboard	84.41%
3	4	532	31	4	c=drinking	92.68%
3	73	22	500	12	d=cupboard	88.20%
25	2	19	13	441	e=keyboard	87.65%

In order to verify that the feature set, that was chosen based on 5 classes, works with more classes as well, the confusion matrix for all 6 activities is given

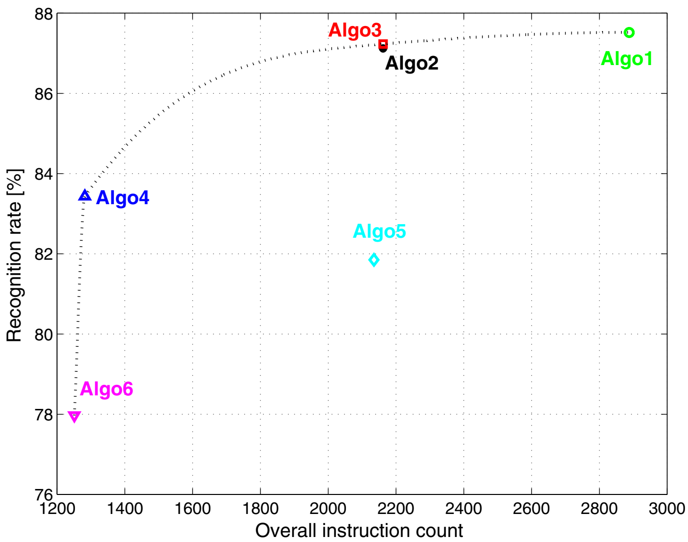

Fig. 6. Recognition Rate vs. Instruction Count for a C4.5 classifier and $f_{s\_acc} = 40\,\text{Hz}$

in Table 6. Compared to the 5 classes the overall recognition rate drops from 87.65% to only 85.77%.

Table 6. Confusion matrix for Algo3 with 6 classes

a	b	c	d	e	f	← classified as	Accuracy
343	4	0	4	27	1	a=mouse	90.50%
1	530	3	82	4	28	b=whiteboard	81.79%
2	1	525	31	5	10	c=drinking	91.46%
2	67	19	486	4	32	d=cupboard	79.67%
34	2	20	9	430	5	e=keyboard	86.00%
0	23	23	38	2	497	f= drawer	85.25%

4 Power Consumption Analysis

The power consumption of the system depends on the specification of the individual components and the modes in which they are operated, in particular the speeds and duty cycles. The former is partly determined by our platform design. The latter is a result of the recognition method designed in the previous section. However it also depends on two other factors: the partitioning of the computation between the sensor node and a remote node and the number of classification steps that need to be performed per each second.

In Fig. 7 the following three scenarios with respect to those factors have been analyzed for their power consumption (all values at 2.7 V supply voltage)

- *A: All sensors (no local processing):* If all sensors are active and sampled, the microcontroller needs to run at approximately 4 MHz. The data rate of 3.3 kbit/s requires the nRF transceiver to be on for 4.2% of the time and therefore it requires only 0.46 mA. Otherwise, i.e. in continuous operation, the transmitter would consume 10.5 mA at -5 dBm output power. We measured a current consumption of 2.6 mA which complies with the calculated value. Considering a 90% efficiency of the step-down converter, the total power consumption is 8.5 mW.
- *B: Without microphone (no local processing):* If the microphone is not used, the power consumption is reduced due to the reduced clock frequency of the microcontroller (1 MHz) and the lower data rate at the transmitter (0.3 kbit/s). The overall power consumption is estimated to be 3.0 mW.
- *C: One accelerometer + light sensor with Algo3:* If algorithm 3 as described in Sect. 3.5 is implemented, the power consumption can be reduced even further. Only one accelerometer is needed and the transmitter needs only to transmit the classification results. It therefore runs with 0.01% duty cycle. The microcontroller runs at 350 kHz to calculate one classification per second and consumes just 92 μA. The overall power consumption in this case is computed to 1.6 mW. To calculate 5 classifications per second a microcontroller frequency of about 1.5 MHz is needed and the total power consumption would be 2.5 mW.

With the 3.7 V, 130 mAh battery that our board is equipped with, the sensor board can be run for 57, 160 or 300 hours for scenarios A, B or C, respectively.

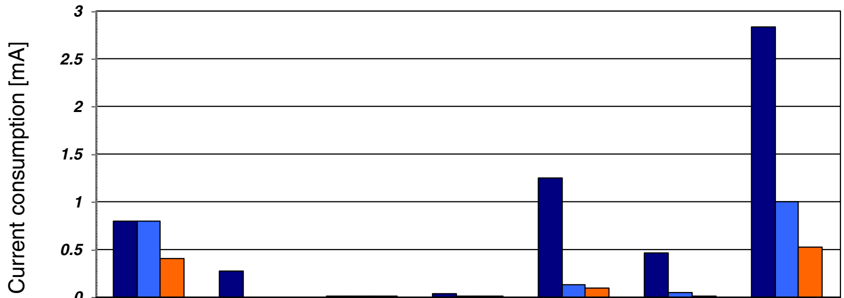

	Accelerometers	Microphone	Lightsensor	ADC	MSP	nRF	Total
A: direct - all sensors	0.8	0.27	0.01	0.04	1.25	0.46	2.83
B: direct - without mic.	0.8	0	0.01	0.01	0.13	0.05	1.00
C: online - Algo3	0.4	0	0.01	0.01	0.09	0.01	0.52

Fig. 7. Calculated current consumption [mA] for different operating modes

5 Conclusion and Future Work

We have demonstrated the feasibility of multi-sensor context recognition a on highly miniaturized (thumb-sized), wireless sensor platforms which consumes 2 to 3 mW of power including wireless transmission. This is a first important step in our work towards fully autonomous sensors nodes seamlessly integrated in the user's outfit. A key contribution of the paper was including power consumption concerns in the recognition-method design and optimizing the system for the best energy vs. recognition rate trade-off rather than maximizing the recognition rate only. We have shown that adding a simple sensor can be a good way of improving performance without increasing power consumption and might be preferable to increasing classifier and/or features complexity.

Currently we are working on an improved version of the hardware platform that contains a more powerful microcontroller. We then intend to conduct more extensive experiments that use all three sensors on the board. Besides reviewing features and classifiers that are suited for other applications as well, we will look into methods to segment the incoming data stream to perform continuous recognition, e.g. we are thinking of a cheap sensor waking up the more expensive sensors. Furthermore, we will attempt to refine the power consumption metrics to be able to compare more complex system. We are also looking into the networking aspect and will carry out experiments with sensor boards distributed over the wearer's body. For this purpose we will extend the power-aware system-level design methodology to a distributed sensor network.

References

1. Abowd, D., Dey, A., Orr, R., Brotherton, J.: Context-awareness in wearable and ubiquitous computing. Virtual Reality **3** (1998) 200–211
2. Kern, N., Antifakos, S., Schiele, B., Schwaninger, A.: A model for human interruptability: Experimental evaluation and automatic estimation from wearable sensors. In: ISWC 2004: Proc. of the 8th IEEE Int'l Symposium on Wearable Computers. (2004) 158 – 165
3. Lukowicz, P., Junker, H., Stäger, M., von Büren, T., Tröster, G.: WearNET: A distributed multi-sensor system for context aware wearables. In: UbiComp 2002: Proc. of the 4th Int'l Conference on Ubiquitous Computing. (2002) 361–370
4. Junker, H., Lukowicz, P., Tröster, G.: PadNET: wearable physical activity detection network. In: ISWC 2003: Proc. of the 7th Int'l Symposium on Wearable Computers. (2003) 244–245
5. Bharatula, N.B., Ossevoort, S., Stäger, M., Tröster, G.: Towards wearable autonomous microsystems. In: Pervasive Computing: Procc of the 2nd Int'l Conference. (2004) 225–237
6. Anliker et al, U.: A systematic approach to the design of distributed wearable systems. IEEE Transactions on Computers **53** (2004) 1017–1033
7. Stäger, M., Lukowicz, P., Perera, N., von Büren, T., Tröster, G., Starner, T.: SoundButton: Design of a Low Power Wearable Audio Classification System. In: ISWC 2003: Proc. of the 7th IEEE Int'l Symposium on Wearable Computers. (2003) 12–17

8. Junker, H., Lukowicz, P., Tröster, G.: Sampling frequency, signal resolution and the accuracy of wearable context recognition systems. In: ISWC 2004: Proc. of the 8th IEEE Int'l Symposium on Wearable Computers. (2004) 176–177
9. Stäger, M., Lukowicz, P., Tröster, G.: Implementation and Evaluation of a Low-Power Sound-Based User Activity Recognition System. In: ISWC 2004: Proc. of the 8th IEEE Int'l Symposium on Wearable Computers. (2004) 138–141
10. Holmquist et at, L.: Building intelligent environments with smart-its. IEEE Computer Graphics and Applications **24** (2004) 56–64
11. Maguire, G.Q., Smith, M.T., Beadle, H.W.P.: Smartbadges: A wearable computer and communication system. In: The 6th Int'l Workshop on Hardware/Software Co-design. (1998)
12. Kohvakka, M., Hannikainen, M., Hamalainen, T.D.: Wireless sensor prototype platform. In: IECON '03: The 29th Annual Conference of the IEEE Industrial Electronics Society. Volume 2. (2003) 1499–1504
13. Homepage of the TEA project. http://www.teco.edu/tea/ (2000)
14. Randell, C., Muller, H.: Context awareness by analysing accelerometer data. In: ISWC 2000: Proc. of the 4th Int'l Symposium on Wearable Computers. (2000) 175–176
15. Sekine, M., Tamura, T., Fujimoto, T., Fukui, Y.: Classification of walking pattern using acceleration waveform in elderly people. In: Proc. of the 22nd Annual Int'l Conference of the IEEE Engineering in Medicine and Biology Society. Volume 2. (2000) 1356–1359
16. van Laerhoven, K., Kern, N., Gellersen, H.W., Schiele, B.: Towards a wearable inertial sensor networks. In: In Proc. of the IEE Eurowearable. (2003) 125–130
17. Bao, L., Intille, S.S.: Activity recognition from user-annotated acceleration data. In: Pervasive Computing: Proc. of the 2nd Int'l Conference. (2004) 1–17
18. Lukowicz, P., Ward, J.A., Junker, H., Stäger, M., Tröster, G., Atrash, A., Starner, T.: Recognizing workshop activity using body worn microphones and accelerometers. In: Pervasive Computing: Proc. of the 2nd Int'l Conference. (2004) 18–22
19. Doherty, L., Warneke, B., Boser, B., Pister, K.: Energy and performance considerations for smart dust. International Journal of Parallel and Distributed Systems and Networks **4** (2001) 121–133
20. Warren, J., Martin, T., Smailagic, A., Siewiorek, D.: System design approach to power aware mobile computers. In: Proc. of the IEEE Computer Society Annual Symposium on VLSI. (2003) 101–106
21. Witten, I.H., Frank, E.: Data Mining: Practical Machine Learning Tools and Techniques with Java Implementations. Morgan Kaufmann (1999)
22. Cakmakei, O., Coutaz, J., van Laerhoven, K., Gellersen, H.W.: Context awareness in systems with limited resources. In: Workshop on Artificial Intelligence in Mobile Systems, Lyon, France. (2002) 21–28
23. Li, D., Sethi, I., Dimitrova, N., McGee, T.: Classification of general audio data for content-based retrieval. Pattern Recognition Letters **22** (2001) 533–544
24. Provost, F., Fawcett, T.: Analysis and visualization of classifier performance: Comparison under imprecise class and cost distributions. In: KDD '97: Proc. of the 3rd Int'l Conference on Knowledge Discovery and Data Mining. (1997) 43–48
25. Quinlan, J.R.: C4.5: Programs for Machine Learning. Morgan Kaufmann (1993)
26. Rabaey, J.M.: Digital Integrated Circuits - A Design Perspective. Prentice Hall (1996)
27. Rabaey, J.M., Pedram, M., eds.: Chapter 10. In: Low Power Design Methodologies. Kluwer Academic Publishers (1996)