Coming to Grips with the Objects We Grasp: Detecting Interactions with Efficient Wrist-Worn Sensors

Eugen Berlin **Jun Liu** **Kristof van Laerhoven** **Bernt Schiele**
{berlin | kristof | schiele}@cs.tu-darmstadt.de
Department of Computer Science, Technische Universität Darmstadt, Darmstadt, Germany

ABSTRACT
The use of a wrist-worn sensor that is able to read nearby RFID tags and the wearer's gestures has been suggested frequently as a way to both detect the objects we interact with and to identify the interaction. Making such a prototype feasible for longer-term deployments is far from solved however, as plenty of challenges remain in the hardware, embedded algorithms, and the overall design of such a bracelet-like device. This paper presents several of the challenges that emerged during the development of a functioning prototype that is able to sense interaction data for several days. We focus in particular on RFID tag reading range optimization, efficient data logging methods, meaningful evaluation techniques, and long-term deployments.

Author Keywords
wrist-worn RFID, gesture detection, wearable interaction

ACM Classification Keywords
H.5.2 Information Interfaces and Presentation: Miscellaneous

General Terms
Design, Experimentation

INTRODUCTION
With the introduction of small and inconspicuous RFID tags, wearable tag readers have been proposed as early as 2000 for the detection of objects we interact with [13]. Other work evolved designs from clunky glove-based prototypes (e.g., [12, 1, 11, 5]) to sleek bracelet-like designs such as the Intel iBracelet [4, 9]. The combination of sensed RFID tags and inertial data for detection of what gestures are performed with the objects, has been mentioned and explored in a large body of work in years after. In one of the earlier articles [16], the authors showed that the characteristic motion patterns in the inertial data, combined with the knowledge of which tools or objects were grabbed by the user, gave in many cases very good results in the recognition of various daily activities such as *"brushing teeth"* or *"making tea"*. Other work such as MIT's ReachMedia [2, 3] combines inertial sensors and an RFID reader to detect particular on-the-move interactions from mobile users.

Applications suggested for such a wrist-worn sensor range from medical applications, e.g., detecting Activities of Daily Living to follow the routines of elderly patients that are living independently [4], to generic mobile interfaces that are designed to be more natural than existing mobile interfaces, e.g., an input device for wearable computers [3]. One of the aspects of this work remains underdeveloped, however, even though it is a main requirement for the acceptance of wrist-worn RFID-accelerometer-sensors in these applications: a battery-driven device at that location needs to be both light and power-efficient enough.

Figure 1. The wrist-worn prototype identifies grabbed objects and the physical interactions with them, using a combination of RFID and inertial sensors. This paper's aim is to explore light-weight and power-effienct solutions in particular, to facilitate long-term deployment.

This paper specifically presents the challenges that were encountered during the development of a bracelet that detects interactions performed with detected objects, while focusing on a low-power solution that is deployable over long stretches of time. The main contributions of this paper are threefold: First, we mention the technical procedure in optimizing a wrist-worn RFID antenna. Second, a benchmark is presented that allows evaluation of different antenna configurations. Third, an approximation algorithm is demonstrated which makes recognition of short gestures possible. A final section reports on experiments where this prototype was deployed in a domestic setting for several days.

TEI 2010, January 24–27, 2010, Cambridge, Massachusetts, USA.

COMBINING RFID READING AND INERTIAL SENSING

The basic principle of RFID-tag reading is that a reader is able to power nearby tags by induction, using relatively large antennas. Having the tag reader at the wrist means that grabbed objects and tools can be detected using just these tags, assuming the reader's range is large enough to power and communicate with the object's tag. Since the reader needs to be mounted at the wrist, this type of wrist-mounted RFID sensing comes with a harsh energy constraint: the antenna and its circuitry need to be strong enough to detect hand-held tags, yet power-efficient enough to not drain the battery after a short while.

Inertial sensors such as accelerometers have been suggested and applied for the recognition of physical actions that have characteristic motions or postures. Work such as [18] propose the characteristic wrist positions and motions to detect short actions in a car scenario such as *"pulling the handbrake"* or *"opening the oil tank"*. In our prototype, after detection of a *claw hammer* one might for instance expect the actions *"hammering"* or *"pulling out nails"*, depending on the gestures detected with the wrist-mounted accelerometer. Accelerometers are known to be power-efficient but their data is also not as rich, compared to IMUs containing also gyroscopes and magnetometers. The knowledge of what object the user has grasped can help here to distinguish between a limited set of interactions.

To summarize, the two sensing technologies have proven their worth in preliminary studies where prototypes were used in feasibility studies. Several questions remain regarding their operation in applications which require functioning longer than a few hours. In particular, the following questions emerge and are tackled in this paper:

- How can the working range of an RFID antenna be increased to reach from wrist to the abject in the hand?
- Which antenna performs better? And how do we test this?
- How dense should RFID samples be to detect grabbed objects, yet save energy?
- How can acceleration patterns be approximated efficiently?

The basis for our wrist-worn sensor design is the Porcupine sensor [7], an accelerometer-based module that allows power-efficient capturing of inertial data, enriched by light sensors, a temperature sensor, and a real-time clock and calendar chip. It is able to log up to 4 GB worth of sensor data on a small microSD card, or wirelessly transmit chunks of sensor data to a nearby station. For the reading of RFID tags, the M1-mini from SkyeTek was chosen. This exact module is used by most other research [12, 16, 3] as it is one of the smallest and power-efficient modules on the market – this module comes with a small on-board antenna which has a reach of about 3 centimeters. To have a reading range larger than that, one has to attach an external antenna and a matching circuit. Figure 2 shows the current version of the bracelet design, with the oval-shaped PCB antenna and matching circuit, the M1-mini RFID reader, and the Porcupine modules on top of each other (without the battery).

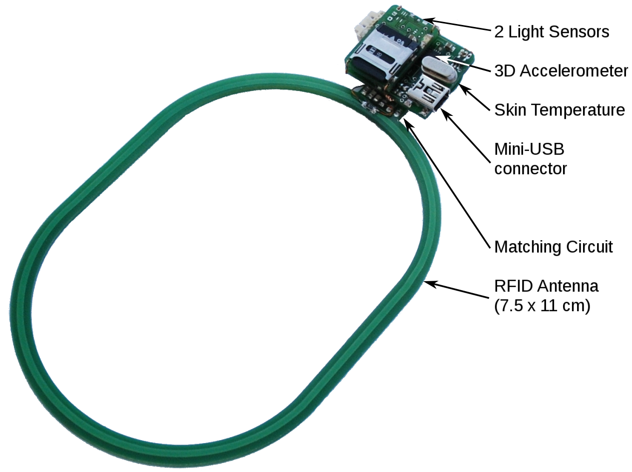

Figure 2. The Porcupine with PCB antenna and SkyeTec RFID reader. The module measures 36x26x20 millimeters and weighs about 32 grams including the straps and battery.

OPTIMIZING RFID READING RANGE

One of the most time-demanding efforts in building a wrist-worn RFID sensor is the optimization of the bracelet's built-in antenna. Although [17] gives an extremely good introduction in the use of RFID tags and readers in pervasive computing applications, it requires a deeper knowledge to maximize the potential of a given antenna. In [4], a customized bracelet-like antenna is described, but the authors provide no details about their engineering and tuning processes of their antenna. As the reading range for RFID tags is crucial in this type of work, this section is dedicated to a detailed description of the antenna design and tuning process. The antenna design presented in this paper is additionally freely accessible on `http://porcupine2.sf.net`.

The Q-value is a measure of "quality" for the antenna and is directly related to the reading range. Generaly, increasing Q will result in a higher power output of the particular antenna, allowing higher reading distances. On the other hand, a too high Q will conflict with the band-pass characteristics of the RFID reader. Since a sufficient bandwidth for the wireless communication between the reader and the RFID tags has to be assured, the relation between the quality and the bandwidth is thus reciprocal, resulting in a trade-off between the quality of the antenna and the bandwidth.

The M1-mini SkyeTek reader can be equipped with an external antenna, which in our case is a requirement as the built-in antenna does not reach further than a few centimeters. The driver impedance of the reader is typically set to 50Ω. Every antenna has its own impedance value, and to be able to use our customized single-loop coil with the M1-mini reader, the antenna impedance has to be matched to the reader's impedance front end. There exist different matching approaches such as gamma, transformer or capacitance matching, whereby the latter was used in our work.

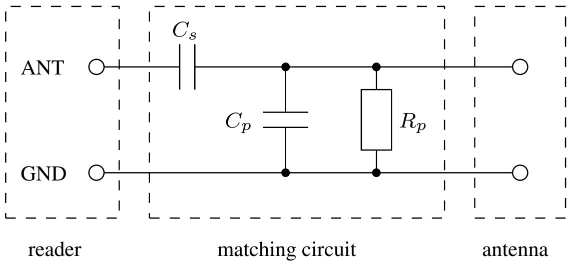

Figure 3. The matching circuit for capacitance matching.

The matching circuit (Figure 3) consists of a serial capacitor C_s, a parallel capacitor C_p and a parallel resistor R_p. The values chosen for those capacitors and the resistor have a great impact on the quality and the bandwith of the antenna. Based on [14], the following approach has been taken to compute the values for the capacitors and the resistor. Thereby, the following variables are: the measured antenna inductance L, the target impedance $Z_0 = 50\Omega$, the Q-value of $Q = 30$, and the frequency $f = 13.56MHz$.

First, we compute the parallel resistor value:

$$R_p = Q * \omega * L \quad \text{with} \quad \omega = 2 * \pi * f$$

Then, using the equations

$$G_L = \frac{1}{R_p} \quad \text{and} \quad B_L = \frac{-1}{\omega * L},$$

we compute the parallel capacitor value

$$C_p = -B_L - \sqrt{\frac{G_L}{Z_0} - G_L^2}$$

and the serial capacitor value

$$C_s = \frac{1}{\omega * Z_0 * \sqrt{\frac{1}{Z_0 * G_L} - 1}}.$$

Using a network analyzer, for our PCB antenna prototype we measured an antenna inductance of $L = 322nH$. With the given approach the optimal values have been found to be $R_p = 820.288\Omega$, $C_p = 371.663pF$ and $C_s = 59.807pF$. After utilizing the off-the-shelf available 820Ω resistor and assembling a combination of fixed and variable capacitors to meet the computed values, we used the network analyzer to fine-tune the antenna to match 50Ω. Finally, we resulted in a reading distance of up to 14 cm, thus slightly exceeding the reading ranges of the antennas mentioned in [4] and [3].

Surprisingly, evaluating the antennas reading range by increasing or decreasing the distance from tag to antenna, is not the best way to benchmark the accuracy in detecting tags. By starting close by the antenna and increasing the distance gradually, will lead to more optimistic reading ranges, since the tags that are charged up by the reader are able to bridge some additional distance. The circumstances when wearing the bracelet and grabbing tagged objects tends to differ substantially from these lab tests. We propose a novel benchmark that will be presented in a later section in this work.

OPTIMIZING RFID READING FREQUENCY

Every reading cycle, where the RFID reader module seeks for the nearest tags, tends to last from 20 in best case up to 68.4 milliseconds in worst case[1], depending on the tag type and whether the reading was successfull. The reader draws approximately $60mA$ in current while in active, $15mA$ in idle and $60\mu A$ in sleep state. This means that a trade-off exists between detection speed and power consumption: More frequent searching for nearby tags means that objects are likely to be found faster yet the reader will demand more from the wearable battery, less frequent reading will make the battery last longer but might result in missing tags or a slower detection.

Assuming that we want to search for nearby tags 16 times per second, we will end up with a time slot of maximally 62.5 milliseconds for one reading. Since the reader also needs time to go to sleep as well as wake up, and additionally the worst case time for one seek even exceeds this time slot, no idle or sleep states are possible in this scenario. This means that the RFID reader will constantly consume power in active state, i.e., a draw of $60mA$. With a light-weight rechargeable battery[2] with a capacity of $600mAh$, and assuming the rest of the bracelet consuming approximately $10mA$ (in worst case), this will result in a runtime of about 8.5 hours.

Reducing the reading rate to 1 reading per second allows the RFID reader to go into the sleep mode to preserve power. In worst case, it will take approximately 70 milliseconds to search for a tag and up to 100 milliseconds for switching modes, giving the RFID reader rougly 730 milliseconds for staying in the sleep state. Since in our approximated calculation we are able to neglect the power drain of $60\mu A$ during sleep, the runtime for the same battery and same configuration mentioned above will account to about 28 hours. This is more than thrice the runtime as in the high frequency reading case where the RFID reader's power efficient sleep mode can not be utilized.

In addition to the idle and sleep modes between readings, it is also possible to have the RFID reader change into sleep state and suspend the reading if no movement is registered by the bracelet's accelerometer. E.g., no movement will occur for longer periods during the night when the subject wearing the bracelet is sleeping, or when the bracelet has been temporarily taken off the wrist. Combining the sleep states of the RFID reader and the Porcupine in such situations will save battery power and result in a longer total runtime. Assuming the night phase to last about 7 hours and neglecting the power consumption during this period, the total runtime with the same battery will account to 35 hours. Utilizing a battery with slightly higher capacity will allow to deploy the bracelet for two full days without the need for recharging.

[1] Figures from the SkyeTek M1-mini RFID reader datasheet, http://www.skyetek.com

[2] KLIC-7002: a Li-Ion Rechargeable Digital Camera Battery, http://www.kodak.com

Figure 4. Some of the objects from the box test, chosen to have a high variety in shape, weight and material, tagged with RFID tags.

EVALUATION OF RFID READING: THE BOX TEST

To obtain the optimal RFID parameters, both reading range and frequency, we designed a benchmark in which subjects wear the bracelet prototype and load a variety of tagged objects in and out a box, which is also tagged with several RFID tags. The advantage of this test over straightforward measuring increasing or decreasing distances between a test tag and antenna, is that the whole system is immediately tested under realistic circumstances. This test has the user's wrist and hand present in the middle of the antenna, and the amount of objects, as well as the speed at which objects are grabbed and released tends to be challenging enough. The subjects loaded and unloaded the box three times, and closed it each time it was loaded with all objects.

Furthermore, this test can employ a variety of objects with a wide range of properties that might impact later usage. They can be chosen to fit target applications, or they can be selected according to a variety of shapes and materials. Some of the objects are illustrated in Figure 4. Tags were mostly placed on the areas where people tend to grab the objects, 8 tags were also placed on the flaps of the box to detect the closing and opening of the box. Video footage of the test was also taken to be able to interpret sequences where things went wrong – Figure 5 shows some frames from this.

To evaluate how well an object was recognized in the box test, we count the number of hits, events when the RFID reader correctly found an object's tag as it was taken in the test subject's hand, and divide this over all occurrences when the tag should have been detected (excluding the tags affixed to the box); This measure will in the remainder of this paper be referred to as *hit rate*. Test were done for the shape of the antenna, for the number of loops in the antenna, the Q value, and reading frequency.

Antenna Shape

Figure 6 shows the hit rates between two types of antennas we tested in the early design phase: one PCB antenna which is round, as mentioned in [16], and one new design that has the antenna slightly tilted at a downward angle and is more oval in shape. The latter is only slightly larger, but makes it easier to put on the antenna as a bracelet. A third type of antenna, which uses flexible coils and snap-on metal buttons

Figure 5. Various subjects who participated in the box test had to load and unload a box with tagged objects.

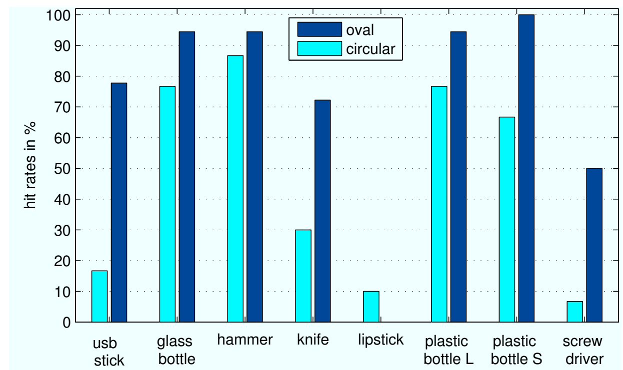

Figure 6. Hit rates for several objects during the box test, using two antennas: one circular as proposed in [16], and one oval-shaped that was tuned to $Q = 30$ and worn at an angle.

to close the antenna loops as in [3], is more promising for comfort reasons but harder to implement and unpredictable to design a matching circuit for, as its shape tends to change.

Q-Value

The antennas we evaluated were also configured with Q-values up to 36 (the meaning of the Q was introduced in the 'Optimizing RFID Reading Range' section previously). Figure 7 shows the hit rate results for four different Q values for the same antenna. For the maximization of the reading range, the box test showed that a value of $Q = 30$ gave the best trade-off between the quality and the bandwidth of the antenna prototype.

RFID Reading Frequency

To find out at what frequency the reader should wake up from its sleep state to seek for nearby tags, we conducted an experiment with a reading rate of 16 Hz on various objects. To simulate less frequent readings, the data stored during that experiment was selectively analyzed with an increasing step factor. By doubling the step factor, and this way halving the reading rate, we obtained a monotonic decreasing number of captured tags. The average hit rate started at 100%

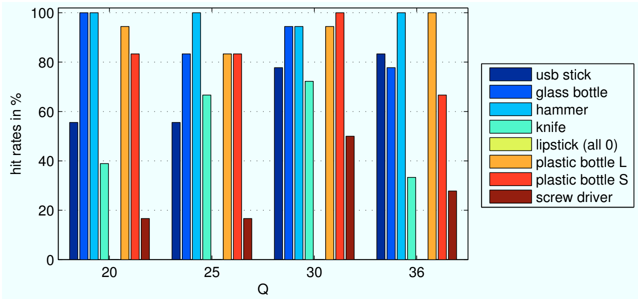

Figure 7. Hit rates for different values of Q, using the oval-shaped antenna. The hit rates for lipstick are 0% for all Q's, and therefore are not visible in the plot.

for the factor 1 (16 Hz), and dropped to an acceptable level of 65% for the step factor values of 4, 8 and 16 (4 Hz, 2 Hz and 1 Hz respectively). Increasing the step factor even more resulted in the hit rate monotously dropping further. From this we conclude that a reading rate of 1 Hz provides a good balance beetween capturing the deployed tags on the one hand and saving a significant amount of battery power on the other. It is interesting to note that this frequency is on par with optimal frequencies found in related work such as [12, 16].

The value of the box test goes beyond a more realistic evaluation of wrist-worn RFID antennas. It offers insights in how well and how fast intended objects are found, while it places the prototype in an environment that is closer to its usage scenario, without demanding a costly experiment setup. For completeness, we also mention the distances reported in previous work as well as our measured maximum range below in Table 1. Note however that not all of these measurements have been taken in the same situation and environment.

project name	max. range	amplified	embedded in
iGlove [12]	3-5 cm	no	glove
SonMicro [10]	4-5 cm	no	glove
Phidgets [5]	10 cm	no	glove
ReachMedia [3]	10 cm	yes	wrist band
iBracelet [16]	10 cm	no	bracelet
Our design	14 cm	no	bracelet

Table 1. Several projects involving glove- and bracelet-based RFID readers, and their reported maximum reading ranges. We stress that a comparison between these distances is hard to do, and offer the box-test as a better estimation of how well a design fits the intended application.

OPTIMIZING INERTIAL SENSING

Work on the detecting of gestures and activities with wearable inertial sensors, in particular accelerometers, has matured more than the wearable RFID sensing. The typical duty cycles of the wearable accelerometer-based sensor are followed by throttling down the microcontroller's operating speed between the sensing, in a way which bears many similarities to that mentioned in [15]. The device is also put into a lower-power mode while reading of the sensor values, compared to the sending of the sensor data buffer. Instead of focusing on these fairly well-known techniques, however, this section mainly explains in detail how the accelerometer data is approximated so that only the essence of gestures is retained.

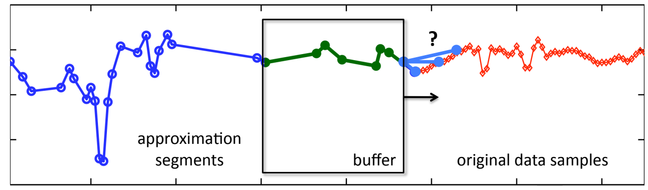

Figure 8. Illustration showing how raw accelerometer data is transformed into a series of segments that approximate the gestures made.

Most accelerometer-based gesture recognition research (e.g., [18]) starts with segmentation of possible gestures out of the streaming data. This is a common method to save energy and avoid having to process each new piece of incoming data as if it was a candidate gesture. Characteristic values that are easy to calculate from a sensor stream, such as the variance of the accelerometers, are in this first segmentation phase used to retain only those segments that could hold a possible gesture. A following phase then uses more processing-intensive algorithms to produce so-called features (e.g., mean and variance, Fourier coefficients, or wavelet coefficients) which classify the candidate segment as a gesture.

Since the resources on our wrist-worn device are limited, we propose to combine the segmentation phase and classification phase by applying a technique used in time series data mining. The raw data is read at 100 Hertz and is approximated by linear segments using a version of SWAB [6], called mSWAB [8], which reports good performance on human accelerometer data and approximates the data faster then SWAB. This means in practice that instead of needing buffers of hundreds of bytes containing the raw data samples, only a few segment coordinates are needed to describe a gesture pattern. Apart from the reduced footprint, there is also a processing advantage of using approximation: matching the approximation of known gestures with that of the current accelerometer data is possible in hardware using simple comparisons using Euclidean distance measures.

Figure 8 shows the basic concept of the approximation algorithm; The information in the sensor data is compacted by representing it as a series of linear segments which are then matched against known series of segments, instead of reading the accelerometers and carrying their raw data further to recognition algorithms. A small buffer is used to convert a small sequence of data into segments (using the Bottom-Up [6] segmentation heuristic), after which the left-most segment is added to the final set of segments. The buffer is then shifted to the right, to the point where the local slope between data samples changes sign (i.e., at peaks within the data), and the buffer starts its segmentation again.

Thus, a stream of original data samples is transformed in a typically much smaller set of segments. These can then simply be matched with those of known gestures, after which the closest match is put out as the most likely classification.

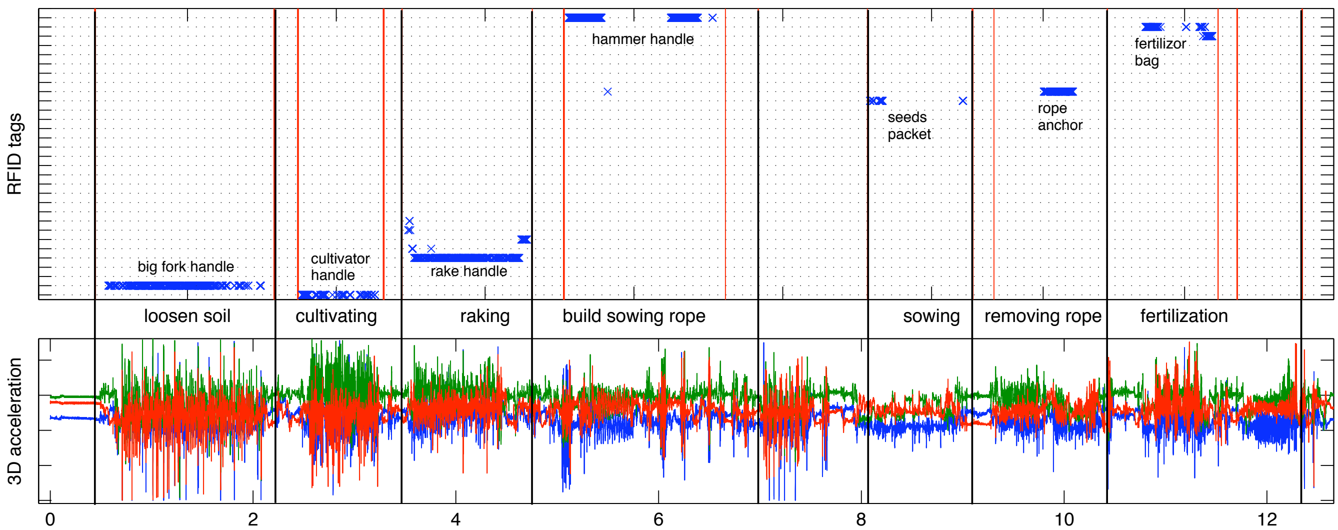

Figure 9. Raw data from both the RFID reader and accelerometer embedded into the bracelet during the sowing scenario from the gardening data. The top plot shows blue crosses whenever a tag was detected at a specific time (x axis) for a specific tag (y axis), while the bottom plot displays the accelerometer time series (approximately lasting 12 minutes). Note that one of the main object's tags was found for most interactions, and that at several occasions (e.g., raking) unrelated tags were detected.

EVALUATION OF INERTIAL SENSING: GARDENING

To stress the interaction with tools and objects, a new data set was recorded which utilized both detected RFID tags, at 1 Hz, and raw accelerometer data at 100 Hz. To verify that our approximation technique is indeed able to reduce the accelerometer data to only the bare essentials to detect patterns of interactions, an hour-long gardening scenario was followed to study the detection of interactions per object. The experiment was performed as realistic as possible, taken outdoors in a real garden while planting real flowers and sowing real seeds.

A gardening scenario is sensible to evaluate a range of interactions with objects, since many tools are typically required for the execution of a wide range of tasks. Tools and objects are also used in several ways, depending on the task at hand: the small spade for example was used both to *"dig holes"*, *"create trenches"* and to *"firm the soil around plants"*. A total of 16 different gardening-specific objects were tagged, including spades, rakes, shovels, flower tops, and buckets, with 36 tags being deployed in total. The tools and objects were positioned near each other throughout the recording, resulting in several instances in which a tagged object was moved out of the way and providing our bracelet with a 'false hit' (e.g., a *spade* was detected during sowing). The visualization in Figure 9 shows the raw data during the last gardening scenario.

Three different scenarios were followed: one where a bed of flowers were planted in soil, a second where weed is removed and plants are watered, and a third in which vegetable seeds are sown. Each of the three scenarios contained a series of sub-tasks, such as *"loosening soil"*, *"digging a hole"*, or *"hammering in a sowing line"*. Video footage was taken for post-annotation of the data.

Figure 10 shows the approximation performance of the three previously discussed approximation algorithms: Sliding Windows, SWAB and mSWAB. mSWAB and SWAB show the better approximation to the original data (left plot – residual error[3]) under different approximation costs (threshold). The right plot shows that for all three algorithms, the number of segments produced is in memory footprint far below the raw data. As a result, the approximation with mSWAB is a good candidate for inclusion in the microcontroller's program module: its footprint and processor speed requirements are low enough to be implemented there, while the segments that it produces are still very close to the original data. Current work in this direction includes the incorporation of this code and performing the classification offline (i.e., not on the bracelet, but either afterwards with the logged data, or via a wireless connection on a nearby device).

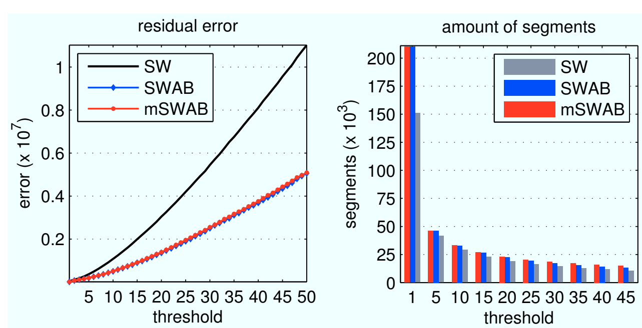

Figure 10. Residual error (on the left) and the amount of produced segments (on the right) of the Sliding Windows, SWAB and mSWAB approximation algorithms using the first scenario of the gardening data set. The other scenarios had very similar results.

[3]Residual error here is used as in [6], i.e., the sum of the distances between the original data and interpolated points on the segments.

Figure 11. Some of the activities conducted during the long-term domestic study, here: watering the flowers, cleaning windows, ironing, and vacuum cleaning.

LONG-TERM DEPLOYMENT

In order to characterize the performance of our system in a relatively long-running experiment, we deployed our sensor bracelet for an entire day, for three consecutive days. During this time, a test subject performed daily a set of 11 household activities such as *"making the bed"*, *"polishing shoes"*, *"vacuuming"*, and *"sweeping the floor"* (see Figure 11 for photos from some more examples). Figure 12 shows some of the objects that were interacted with during these three days. In total, 29 objects were tagged with 43 tags.

Our long-term logging target was met as the prototype was able to capture all data – both accelerometer data and RFID data – for this extended period, while the majority of the objects' tags were detected during each interaction. Although the battery was left charged overnight, the longest continuous log still lasted 18 hours and the bracelet's lightweight battery was never drawn to depletion. Our estimations are that on a light-weight battery of 600 mAh, the prototype would be able to log at least for two days continuously, assuming a highly active daily schedule.

The main obstacle for several days' worth of *continuous* logging is currently the formfactor of the sensor. Before and during the development of the bracelet we held short-term user trials which favored the oval design over the round one conform to [16], mostly due to it being easier to put it on for subjects with bigger hands. Although the weight and size of the bracelet sensor were found to be acceptable by these user trials during initial development, the utilized straps were found to be hard to adjust and especially uncomfortable to wear during the night in the long-term study. For this reason, several designs are planned that will be evaluated over day-long studies that wrap the current prototype in resin casings with more suitable wrist-bands built in.

The data from this experiment is currently analyzed for appropriate fusion methods that take the detected tags and accelerometer data and produce a log of object interactions that can occur in daily life. Other data sets such as the previously discussed gardening data set will be evaluated as well to test the capabilities of higher-level classifiers to replace our current nearest-neighbor-based method. After this, a multi-user study is planned in a scaled-up setting in terms of number of tagged objects, number of users, and logged time period per user.

Figure 12. Some of the objects from the longer-term domestic setting study. Tagged objects and tools were chosen for their interactions, whereby 43 tags have been deployed, and scattered around the living environment of the test subject.

CONCLUSIONS

This paper presents a lightweight bracelet-like sensor that continuously detects *which* objects are handled, and *how*. It does this by combining a small-scale RFID reader and a 3D accelerometer: tags on the objects reveal what they are, while motion patterns performed while holding the objects characterize the type of interaction. Several prototypes of these bracelet-like sensor devices have been proposed previously, but we contribute in this paper especially on the aspects in this research to make it deployable in real-world environments for longer periods, balancing between sensor data richness and power consumption of the wearable sensor.

Apart from detailing the crucial hardware design choices to achieve both good detections and a long battery lifespan, we also suggest the use of a small practical benchmark study, the "box test", to test various parameters used in the wrist-worn RFID antenna circuit, and an efficient way of approximating the accelerometer data. It allows to test various antenna designs in a limited location, with a wide range of objects and test subjects. We used this benchmark to evaluate different antenna shapes, various matching circuit parameters (via the Q value), and a range of RFID sampling frequencies. A technically optimal configuration was found to be an oval-shaped antenna, with a Q-value of 30, searching for nearby tags once every second.

Two data sets were recorded to test this prototype 'in the field'. A short study following a gardening scenario focused on the use of lots of objects and various interactions per object, to evaluate good approximations for the accelerometer data. Using a segmenting algorithm from the time series data mining community, we presented a simple yet efficient technique which can be employed to match patterns in inertial data, and showed that the produced segments result indeed in less data volume. A preliminary data set was recorded over several days to validate the long-term functioning of our light-weight bracelet in practice, identifying remaining challenges in the bracelet's strap design but showing that continuous operation for longer periods is feasible.

The source code, data sets, and design files discussed in this paper are reachable at `http://porcupine2.sf.net` or by contacting the first author, to encourage reproducing these results.

ACKNOWLEDGEMENTS

We would like to thank Dr.-Ing. Oliver Soffke and Dr.-Ing. Thomas Hollstein for their expertise in RFID antenna circuits and the use of their network analyzer.

This work was sponsored by the Project "Methods for Activity Spotting With On-Body Sensors" from the German Research Foundation (DFG).

REFERENCES

1. M. Boronowsky, T. Nicolai, C. Schlieder, and A. Schmidt. Winspect - a case study for wearable computing supported inspection tasks. In *ISWC'01: Proceedings of the 5th IEEE International Symposium on Wearable Computers*, page 163, 2001.
2. A. Feldman. Reachmedia - on the move interaction with everyday objects. Master's thesis, Master of Science in Media Arts and Sciences at the Massachusetts Institute of Technology, 2004.
3. A. Feldman, E. M. Tapia, S. Sadi, P. Maes, and C. Schmandt. ReachMedia: On-the-move interaction with everyday objects. In *ISWC '05: Proceedings of the 9th IEEE International Symposium on Wearable Computers*, pages 52–59, 2005.
4. K. P. Fishkin, M. Philipose, and A. Rea. Hands-On RFID: Wireless Wearables for Detecting Use of Objects. In *ISWC '05: Proceedings of the 9th IEEE International Symposium on Wearable Computers*, pages 38–43, 2005.
5. M. R. Hodges and M. E. Pollack. An 'Object-Use Fingerprint': The Use of Electronic Sensors for Human Identification. In *UbiComp '08: 9th International Conference on Ubiquitous Computing*, pages 289–303. Springer, 2007.
6. E. J. Keogh, Chu, Hart, and Pazzani. An online algorithm for segmenting time series. In *ICDM '01: Proceedings of the 2001 IEEE International Conference on Data Mining*, pages 289–296, 2001.
7. K. V. Laerhoven and A. K. Aronsen. Memorizing what you did last week: Towards detailed actigraphy with a wearable sensor. In *ICDCSW'07: Proceedings of the 27th International Conference on Distributed Computing Systems Workshops*, page 47, 2007.
8. K. V. Laerhoven and E. Berlin. When did this happen? Efficient subsequence representation and matching for wearable activity data. In *ISWC'09: Proceedings of the 13th International Symposium on Wearable Computers*, pages 101–104. IEEE Press, 2009.
9. B. Logan, J. Healey, M. Philipose, E. M. Tapia, and S. S. Intille. A long-term evaluation of sensing modalities for activity recognition. In *UbiComp '07: 9th International Conference on Ubiquitous Computing*, pages 483–500, 2007.
10. Y. Medynskiy, S. Gov, A. Mazalek, and D. Minnen. Wearable RFID for Play. In *Tangible Play Workshop, 2007 Conference on Intelligent User Interfaces*, pages 45–48, January 2007.
11. D. J. Patterson, D. Fox, H. A. Kautz, and M. Philipose. Fine-grained activity recognition by aggregating abstract object usage. In *ISWC'05: Proceedings of the 9th IEEE International Symposium on Wearable Computers*, pages 44–51, 2005.
12. M. Philipose, K. P. Fishkin, M. Perkowitz, D. J. Patterson, D. Fox, H. Kautz, and D. Hahnel. Inferring activities from interactions with objects. *IEEE Pervasive Computing*, 3(4):50–57, 2004.
13. A. Schmidt, H.-W. Gellersen, and C. Merz. Enabling Implicit Human Computer Interaction: A Wearable RFID-Tag Reader. In *ISWC '00: Proceedings of the 4th IEEE International Symposium on Wearable Computers*, pages 193–194, 2000.
14. K. O. Soffke. *Modellierung, Simulation und Entwurf induktiv gekoppelter Transpondersysteme*. PhD thesis, TU-Darmstadt, Fachbereich Elektrotechnik und Informationstechnik, 2007.
15. T. Tanaka, S. Yamashita, K. Aiki, H. Kuriyama, and K. Yano. Life microscope: Continuous daily-activity recording system with tiny wireless sensor. In *INSS 2008: 5th International Conference on Networked Sensing Systems*, pages 162–165, 2008.
16. S. Wang, W. Pentney, A. Popescu, T. Choudhury, and M. Philipose. Common sense based joint training of human activity recognizers. In *IJCAI 2007: Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pages 2237–2242, 2007.
17. R. Want. An Introduction to RFID Technology. *IEEE Pervasive Computing*, 5(1):25–33, Jan.-Mar. 2006.
18. A. Zinnen, K. V. Laerhoven, and B. Schiele. Toward recognition of short and non-repetitive activities from wearable sensors. In *European Conference on Ambient Intelligence (AmI'07)*, pages 142–158. Springer, 2007.