Recognizing Daily Activities with RFID-Based Sensors

Michael Buettner*, **Richa Prasad***
*University of Washington
Seattle, USA
{buettner, richa}@cs.washington.edu

Matthai Philipose†, David Wetherall†*
†Intel Labs Seattle
Seattle, USA
{matthai.philipose,
david.wetherall}@intel.com

ABSTRACT
We explore a dense sensing approach that uses RFID sensor network technology to recognize human activities. In our setting, everyday objects are instrumented with UHF RFID tags called WISPs that are equipped with accelerometers. RFID readers detect when the objects are used by examining this sensor data, and daily activities are then inferred from the traces of object use via a Hidden Markov Model. In a study of 10 participants performing 14 activities in a model apartment, our approach yielded recognition rates with precision and recall both in the 90% range. This compares well to recognition with a more intrusive short-range RFID bracelet that detects objects in the proximity of the user; this approach saw roughly 95% precision and 60% recall in the same study. We conclude that RFID sensor networks are a promising approach for indoor activity monitoring.

Author Keywords
Activity Detection, RFID, Sensor Networks, WISP

ACM Classification Keywords
H.5.2 Information interfaces and presentation (e.g., HCI): Miscellaneous.

General Terms
Experimentation, Measurement, Performance

INTRODUCTION
The ability to recognize the indoor activities of people has long been seen as an essential capability for ubiquitous computing [26]. It has applications to elder care [12], smart hospitals [2], medication adherence [15], language learning [3], smart kindergartens [20], smart homes [5], and more [25]. However, general purpose activity recognition has proved elusive, with the development of the activity recognition component being one of the most specialized and time-consuming aspects of developing these applications. Frameworks for activity recognition that can be easily retargeted across a variety of daily activities are therefore of great interest.

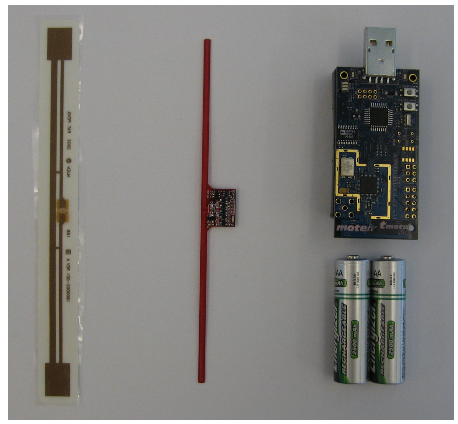

Figure 1. A standard UHF Class 1 Gen 2 RFID tag, Intel WISP, and Telos Mote (left to right).

One promising retargetable approach is the use of *dense sensing*, where wireless sensors are attached to everyday objects to monitor their use. This is valuable because the knowledge of what objects are used is sufficient to identify many daily activities [16]. Two kinds of sensors have seen some level of validation in the past few years. Most common [21, 27, 1] are battery-powered matchbox-sized wireless accelerometers costing tens of dollars each. These are attached to large, relatively permanent objects such as cupboard, microwave and refrigerator doors. They have the advantage that they do not require the person being monitored to wear or carry any new technology.

The alternate approach is to attach postage-stamp sized Radio Frequency Identification (RFID) tags to the objects and have subjects wear a bracelet [11] that has an integrated short-range RFID reader. The tags are inexpensive (roughly 40 cents each) and maintenance free given that they do not contain batteries. They allow objects that are smaller and have shorter lifetimes (e.g., toothbrushes) to be monitored. However, they require users to remember to wear a bracelet and may fail to detect the handling of objects with small grasping surfaces (e.g., books) due to the 20-30 cm range of the bracelet. Existing solutions are therefore either relatively large, expensive and require regular battery replacement or they require user involvement via a wearable and objects with grasping surfaces.

In this paper, we explore a third class of sensors that aims to provide the best of both worlds: RFID sensor networks based on Wireless Identification and Sensing Platforms (WISPs) [6]. WISPs [18] combine passive UHF RFID technology with traditional sensors. A current WISP is shown alongside a commercial UHF RFID tag and a common wireless sensor node (mote) in Figure 1. The WISP has an antenna, power-harvesting circuitry, and can be read by standard long-range RFID readers. It uses the signal from the reader both for communication and to obtain the entirety of its operating energy. In addition, WISPs contain an ultra-low power processor, memory and low power sensors such as an accelerometer that give them capabilities similar to traditional sensor network devices.

From their inception, WISPs have been targeted at uses that include indoor activity recognition via dense object-use monitoring [17]. While WISPs are currently assembled from discrete components that have a cost of roughly $25, they are intended to be mass manufactured like RFID tags at price points closer to $1. Further, integrated "reader-on-a-chip" components that cost around $40 are now on the market resulting in long-range RFID readers the size of a PCMCIA card[23]. We expect that devices based on this technology will cost on the order of a few hundred dollars. These advances suggest that RFID sensor networks will become a cost effective technology for dense sensing in the near future.

To date, however, performance limitations have made WISPs unsuitable for activity recognition. The primary limitations have been low power harvesting efficiency and high energy usage. This resulted in WISPs that operate only when very near the reader. Consequently, every WISP-based system presented in the literature to date has involved a single tag with a proximal (typically 1-2 meters away) reader, e.g., [19, 14]. Fortunately, steady hardware and software improvements (e.g., a lower power microprocessor and more efficient duty cycling) have resulted in WISPs that provide sensor reads at moderate range (up to 3 meters) with reasonable levels of reliability. This in turn holds the potential for activity recognition that has performance comparable to the existing systems based on wearables or that use battery-powered sensors, but without the drawbacks of user-involvement, expense, or maintenance.

To evaluate RFID sensor networks for activity recognition, we prototyped a system that gathers object-use data in an apartment from WISPs and then infers daily activities with a simple Hidden Markov Model (HMM). We do this through a deployment of 25 WISPs and three RFID readers in a model studio apartment. We recruited 10 users to perform 14 daily activities in our model apartment and we compared our activity recognition with a bracelet approach.

We make two main contributions. First, we characterize the behavior and performance of an ensemble of WISPs when deployed in a realistic scenario. We find that WISP behavior is well within the regime where practical dense sensor deployments are feasible. Most importantly, the measured, consistent 3-4 meter read range implies that spaces to be monitored can be covered with a small number of readers.

Our second contribution is the experimental study using our deployment for activity detection. Over the ten users and 14 tasks, recognition rates are in the 90% range for precision and recall. This indicates that RFID sensor networks are promising for activity recognition. In comparison, a parallel bracelet-based system using the same user activity achieves higher precision but lower recall; we provide a more detailed performance breakdown and comparison later in the paper.

The rest of this paper is organized as follows. We describe our activity recognition system in the next section, followed by the details of our deployment in Section 3. Then we evaluate the read range and reliability of the WISP deployment in Section 4, followed by the overall activity recognition performance in Section 5. We discuss future directions and related work, then conclude.

ACTIVITY DETECTION SYSTEM DESIGN

Our behavioral monitoring system comprises two parts, an RFID sensor network (RSN) that gathers sensor data from the environment and an inference engine that classifies activities based on this sensor data. The RSN consists of WISPs and RFID readers. Below, we provide more specifics on the components of our system.

Wireless Identification and Sensing Platform (WISP)

Prototype WISPs have been developed by Intel Research Seattle, and we expect to see sensing and general purpose computation in commercial tags in the future. The most recent Intel WISP, shown in Figure 1, is a battery-free, wirelessly powered platform for sensing and computation. WISPs are powered by and communicate with EPC "Gen 2" RFID readers, and so leverage a widely adopted infrastructure.

Computation is provided by a fully programmable ultra-low-power 16-bit flash micro-controller with an analog to digital converter. This WISP includes 32K of flash program space, a 3D accelerometer, temperature sensor, and 8K serial flash. Small header pins expose micro-controller ports for expansion daughter boards, external sensors and peripherals. Application software is written in C. For our behavioral monitoring system, WISPs transmit unique identifiers along with their most recent accelerometer reading.

RFID Infrastructure

RFID readers provide power to the WISPs and query them for sensor data. "Gen 2" RFID readers operate in the 900 MHz ISM band and are designed to power and communicate with RFID tags at up to 10 meters. Because of this long range, they can be located out of the way along walls or in the ceiling.

For gathering sensor data, we developed a Low-Level Reader Protocol (LLRP)[10] based application that harnesses multiple readers and integrates the results of their queries. LLRP is a vendor agnostic standard for communicating with RFID readers, so our application can be used with any LLRP compliant reader.

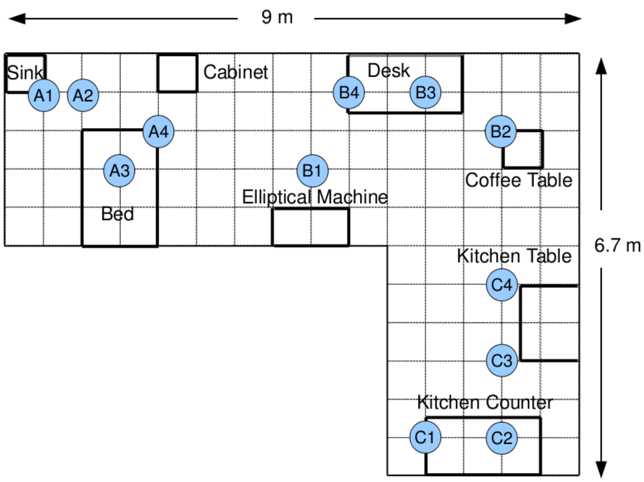

Figure 2. Apartment layout with antennas

Figure 3. Instrumented objects on kitchen table

The application continuously gathers accelerometer data from the WISPs and filters the data to detect objects that are moving. This is done by detecting significant changes between subsequent accelerometer readings. The output is a sequence of IDs indicating the objects that moved.

Inference Engine

Given the sequence of objects o_t that have been moved, we infer the sequence of activities A_t being performed at the corresponding timesteps using a Hidden Markov Model (HMM). We use the GMTK package [4] to estimate parameters of the model using fully labeled data and a Dirichlet prior (with $\alpha_i = 5$) via the EM algorithm. Inference is exact and performed offline. The reasoning infrastructure is thus quite simple. We consider this simplicity a strength of the object-use based activity recognition approach enabled by dense sensing infrastructure such as RSNs.

SYSTEM DEPLOYMENT

Throughout this study we make use of a model studio apartment. We have taken care to make the apartment as realistic as possible, and it includes a kitchenette area, a living area, a bathroom including a shower and sink and a bed.

We use three Impinj Speedway RFID readers with 4 antennas per reader to cover the apartment with RFID signal. We place the readers and antennas in the ceiling of the apartment with antennas facing directly downward. The reasons for this configuration will be presented in a later section.

The apartment layout and antenna placements are shown in Figure 2. Although requiring roughly one RFID reader per living area may seem expensive (given that our readers cost roughly \$1500 each; antennas cost \$100 each), two factors mitigate the cost. First, the cost of RFID readers is rapidly decreasing, with readers now available for less than \$1000 [22]. Second, once the readers are in place, they can be used to monitor large numbers (e.g., hundreds) of WISPs, so that reader cost may be amortized across many applications.

The apartment contains objects commonly found in a household. We attached WISPs to 25 of these objects, listed in Table 1, that are used for everyday activities such as preparing food, taking part in leisure activities, and performing self-care. Figure 3 shows the kitchen table in the apartment along with a few instrumented objects. The WISPs can be seen clearly on the butter dish and the bowl. The white squares on the cereal and milk jug are short-range RFID tags, which we use for our baseline experiments with the iBracelet. All objects were instrumented with both WISPs and short range tags.

Although the WISPs used in our study have rigid PCB antennas, WISPs are available with flexible wire antennas. This would make instrumenting objects with curved surfaces (e.g., bowls) less awkward. While flexible antennas may perform more poorly than rigid antennas, particularly when bent, commercially available RFID tags generally use flexible antennas. We expect commercially manufactured WISPs to have the same sticker-like form factor as standard “Gen 2” tags.

We use this deployment both to evaluate the effectiveness of our activity detection system and to characterize the performance of WISPs and RSNs in a realistic home environment.

RFID SENSOR NETWORK CHARACTERIZATION

The use of WISPs in dense deployments (such as for indoor activity recognition) raises a number of questions about their performance characteristics. The questions derive from two sources. First, because WISPs consume substantially more

Area	Objects
Counter	Coffee, Mug, Cream, Sugar, Jug, Glass, Koolaid
Kitchen Table	Breadbox, Plate, Butter, Cereal, Milk, Bowl
Coffee Table	Book, TV Remote
Desk	Windex, Towels, Plant Food, Water can, Phone
Cabinet	Vitamins, Antacids, Toothpaste
Misc.	Elliptical, Bed spread

Table 1. Tagged Objects

power than stock RFID tags, some of their key performance characteristics such as range, spatial density of deployment and orientation sensitivity have traditionally lagged behind commercial RFID tags. Second, sensing applications in the home impose different (and sometimes more challenging) requirements than traditional RFID applications such as inventory management. In particular, since we seek to track the use of objects handled by people, occlusion by the human body is a persistent issue. Further, in cases where we are interested in kinematics (where we would like to understand *how* the object moves in addition to *whether* it does), it is essential to read sensor data at high rates off the tag (e.g., accelerometers may yield 16 bits at 10-100Hz as opposed to the single bit once a second that is required for pure object use detection).

Given the above concerns about the suitability of WISPs for activity recognition, we performed a series of experiments on RSNs to answer the following questions:

- **Range** Can WISPs be read from sufficiently far away that a "reasonable" number of readers will provide good coverage of the whole home?
- **Spatial Density** Can a reader maintain reasonable WISP motion detection rates even when a relatively large (one to two dozen) number of WISPs is active in front of it?
- **Orientation Sensitivity** Given that day-to-day objects will be positioned haphazardly with respect to the WISP reader while at rest and while in use, and RFID technology is famously sensitive to tag orientation, can we expect readers to provide adequate read rates when monitoring activity?
- **Occlusion** To what extent is the ability to read tags compromised in practice due to proximity of humans handling tagged objects? Shielding by, and close proximity to, liquids are both known to dramatically reduce RFID tag visibility to readers.
- **Noise** How noisy is the accelerometer signal from a WISP sitting "at rest"?

Read Range and Reader Density

We first examine how the read rates from WISPs vary with range from reader. While less noisy metrics for characterizing RFID deployments have been proposed[13], we focus on read rate as it directly captures the sensing rates of the WISPs. After characterizing the performance of WISPs in isolation, we then examine the performance of a particular RSN deployment with one four-antenna reader per living area of our model apartment, which we believe is a "reasonable" reader density.

How WISP Read Rate Varies With Range

We placed three WISPs 15 cm apart from each other at an increasing distance from the reader. All WISPs had their antennas oriented parallel with the reader antenna (this is the ideal orientation). Figure 4 shows the read rate for each of the three tags. The first thing we notice is that the rates vary widely for the different WISPs. At 60 cm WISP 1 has a rate of nearly 20 reads per second while WISP 2 has a rate of only four reads per second. However, at 1.2 meters this trend is reversed and WISP 2 is read nearly 16 times per second with WISP 1 being read only 3 times per second. This level of variation is fundamental to RFID technology, and must be considered when deploying any RFID system.

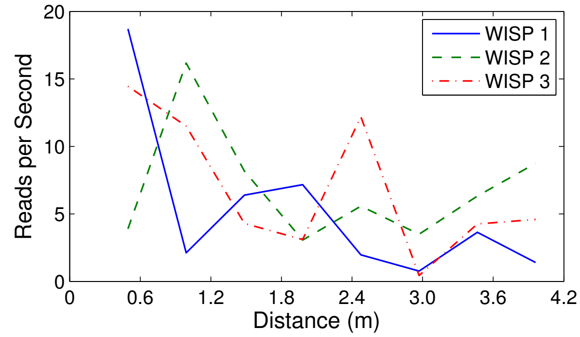

Figure 4. Read rates for WISPs 15 cm apart

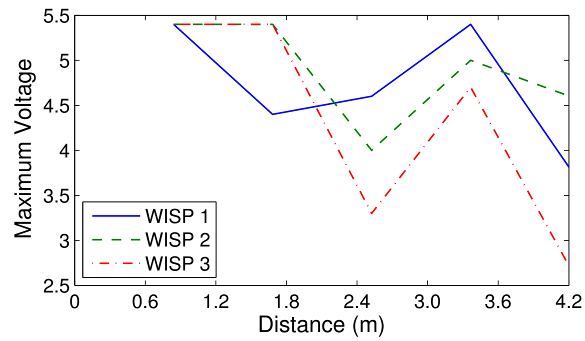

Figure 5. Maximum stored voltage for WISPs 15 cm apart

The position dependent variation in read rates is due to small scale fading in the RF environment. To clearly see this effect we placed three WISPs 15 cm apart and measured the maximum voltage rectified by each device at increasing distances. The maximum voltage indicates the amount of energy available to the WISP for computation and communication. Figure 5 shows that WISPs harvest different amounts of energy depending on their position. The amount of energy a WISP can harvest greatly affects its read rate.

Although such dramatic performance variation across small distances may trigger concerns of many tags being left unobserved, the small-scale variation works in our favor for this application. The coherence distance, the distance between areas of high amplitude and low amplitude "nulls" in the RF environment [24], is on the order of a quarter wavelength; approximately 8 cm for 915 MHz. Changes in position greater than the coherence distance will often trigger drastic changes in performance, even for standard tags. Given that we expect objects used during daily activities to move more than 8 cm, we expect WISPs to move through a high amplitude area while in use, or at least out of a "null".

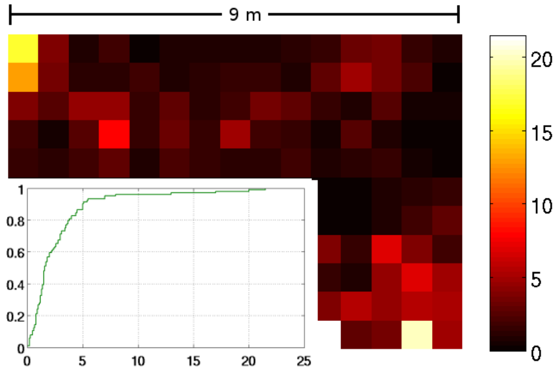

Figure 6. Apartment coverage in reads per second at each location. Apartment layout is shown in Figure 2.

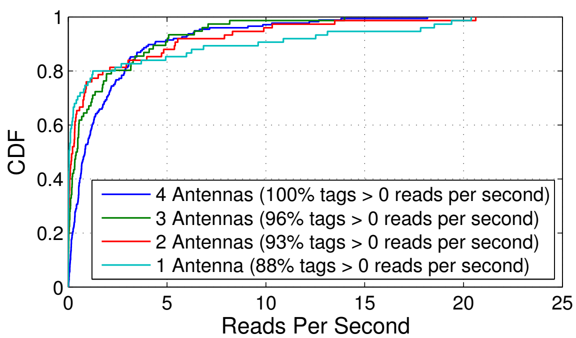

Figure 7. Average WISP read rates when disabling reader antennas

Read Performance for a Reasonable Deployment

Given robust read performance at over 3 meters, we instrumented the apartment shown in Figure 2. Reader antennas were placed in the 3 meter "false" ceiling facing downwards.

To determine how well our RFID readers cover the model apartment we performed an experiment moving a single WISP across the room in 60 cm increments. The WISP was held 1 meter above the ground on a wooden dowel to best avoid occlusion, and the read rate was measured over 30 seconds. Figure 6 shows a heat-map of the WISP read rate at each location, with the inset showing a cumulative distribution function (CDF) of the read rates. The WISP could be read at nearly every point in the room excepting one point directly behind a metal door. The median and mean read rates are 1.6 and 2.8 reads per second, and the maximum was greater than 20 reads per second.

We also looked at how reducing the number of antennas at each reader affects performance. We performed three sets of experiments, with each set having one fewer antenna on each reader. Each set had three trials with the objects being moved slightly between trials to smooth out effects due to small scale fading. We did not attempt to keep the "best" antennas enabled, but instead disabled antennas in descending order as pictured in Figure 2. Figure 7 shows the CDF for read rates for a decreasing number of antennas.

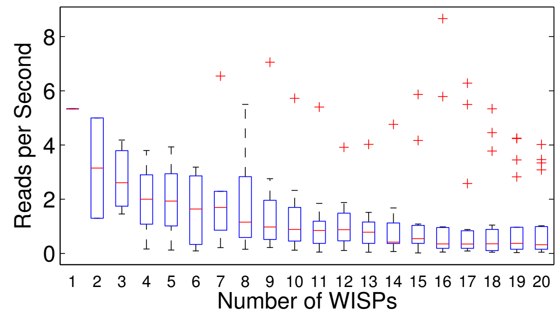

Figure 8. Distribution of read rate for an increasing number of WISPs

As the number of antennas is reduced the number of WISPs with high read rates goes down as one would expect; the median read rate is monotonically decreasing. However, the maximum read rate increases as the number of antennas is decreased. This is because the reader does not transmit on all antennas simultaneously. Instead, it transmits on one antenna at a time in a round robin schedule. If a WISP is only in range of one antenna it is read less frequently as more time is spent on other antennas. Surprisingly, even with only one antenna per reader more than 88% of the WISPs were read at least once.

These results suggest that a more intelligent use of multiple antennas could be used to "zero in" on WISPs of interest, while still providing good coverage. Overall, we find that our deployment provides useful functionality (read rates at least every few seconds) with a reasonable infrastructure deployment (3 readers with 4 antennas each).

Spatial Density

We first look at WISP performance as an increasing number of devices are added to the environment. We deployed a single reader antenna in the ceiling and WISPs on a table directly below. We added WISPs to the table one by one, spacing them approximately 15 cm apart in a 60 cm by 90 cm grid. Figure 8 shows the distribution in read rates as we add up to 20 tags to the table.

While per-WISP read rates tend to decrease as more WISPs are added, we found that the dominant factor in the read rate was physical position. For instance, there are a number of outliers added, namely the 7th, 15th, 17th, and 20th WISPs. These WISPs saw very high read rates but their introduction had a minimal effect on the read rates of other tags. In fact, the net read rate was seen to increase up to 20 tags and we were unable to saturate the channel.

It seems that RSNs should comfortably be able to support the roughly one hundred WISPs per home (or tens per living area) that we believe will make for useful deployments.

Orientation Sensitivity

Figure 9 shows the effect of WISP orientation on read rate. The WISPs we are using have rigid PCB dipole antennas.

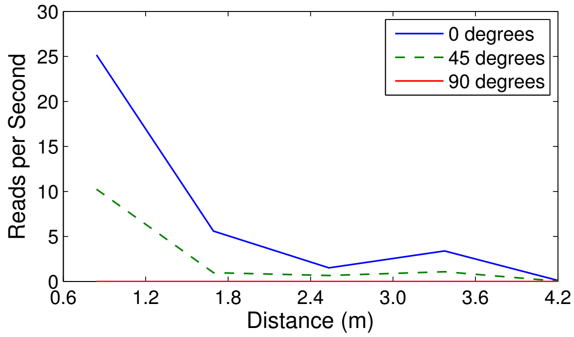

Figure 9. Effect of WISP orientation (WISP antenna axis relative to the plane of the reader antenna)

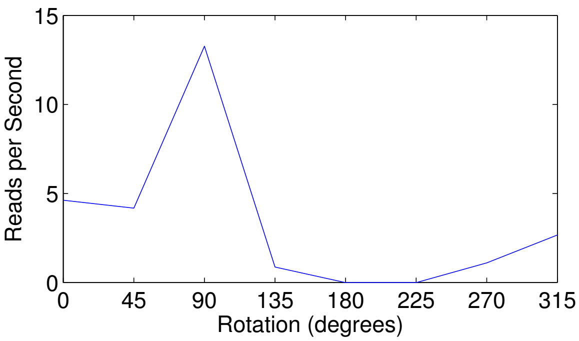

Figure 10. Occlusion due to the body at 2 meters

Such antennas are inefficient when oriented perpendicular to the plane of the reader antenna. The figure shows read rate over distance of a WISP oriented with different angles relative to the reader antenna.

The read rate decreases as the angle moves from 0° (parallel) towards 90° (perpendicular); even at one meter the WISP was never read when oriented at 90°. This has implications for how WISPs should be affixed to objects. For example, if the reader antennas are in the ceiling, and you have a tall thin object, attaching a WISP will be difficult as the orientation that matches the shape of the object will result in the WISP being badly aligned. On the other hand, we expect objects that are *in use* to experience substantial changes in orientation (45° seems quite plausible for many objects). Thus as long as there isn't a very large dead spot around the perpendicular orientation, WISPs should be adequate.

Omni-directional antennas for RFID tags exist and have been prototyped for future WISPS. These antennas make tags largely orientation insensitive, but come at the cost of reduced gain in any one direction.

Occlusion During Normal Use

A common occurrence when deploying WISPs is occlusions due to people and objects. The most likely source of occlusion is the subject themselves. To show this effect we had a subject hold a cup with a WISP attached and we measured

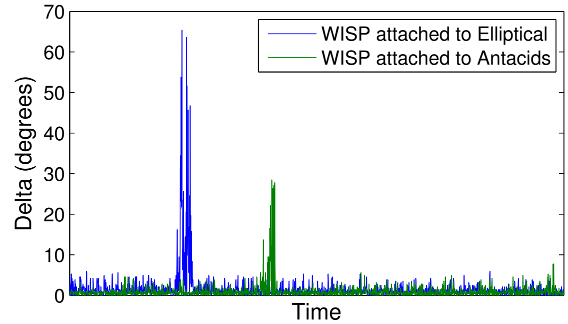

Figure 11. Maximum change along any axis for two WISPs

the read rate as the subject rotated their body. The reader antenna was on a vertical stand 1.2 meters off the ground and located 2 meters away. The WISP was oriented vertically so it was always ideally oriented towards the reader antenna.

Figure 10 shows that as the body begins to occlude the WISP the read rate drops to zero; as the subject completes the rotation the WISP is again read. Another point of interest is that at 0° the hand is partially occluding the WISP, which only has a clear line of sight at 90°.

Based on these results, we opted to place our readers in the ceiling. This deployment provides good RF coverage that reduces occlusions and, because the reader antennas are hidden in the ceiling, is completely inconspicuous. For most object manipulation this will provide a clear line of sight unless the object is below the hand. As the read rates of WISPs tend to drop off much beyond 3 meters, and ceilings are generally about that height, we do not lose significant range by placing antennas in the ceiling.

Sensor Noise

To detect moving objects based on accelerometer data we need to determine the amount of noise in the sensor. The WISP accelerometer samples the orientation along all three axes, and movement detection in our system is done by comparing subsequent sensor readings, and looking for significant changes along any one axis. The key parameter for this is the threshold for considering a change as "significant". The parameter must be well above the noise of the sensor but low enough to avoid false negatives.

Figure 11 shows the maximum change along any axis for two WISPs over time. The change is measured in degrees around the axis. We found persistent noise from all WISPs with a magnitude generally no greater than 8°. The two groups of spikes are when the instrumented objects are interacted with. The group to the left is from a WISP attached to the arm of an elliptical machine, and the group on the right is from a WISP attached to a container of antacid tablets. While the antacids are moved more subtly, the movement is still clearly detectable.

While we currently use the accelerometer data only to detect movement, the data is surely rich enough to be used more

Figure 12. The iBracelet wrist-worn RFID reader. To facilitate comparison with WISPs (which do not require any wearables), we tagged our objects with both stock RFID tags and WISPs. We required all subjects to wear an iBracelet on their dominant hand. Subject activity resulted in synchronized streams of object-use data from iBracelets and WISPs.

effectively. However, for our purposes, we simply consider changes of greater than 8 ° as movement.

ACTIVITY DETECTION EVALUATION

The goal of our activity detection system is to accurately detect what activity a subject performed using our RFID sensor network as a data source for our inference engine. Beyond simply stating the performance of our system, we do a direct comparison between our system and an existing solution for object based activity detection. Specifically, we concurrently deploy an RSN and an iBracelet (Figure 12) short range RFID system. As the technologies operate on different frequencies (900 MHz for the RSN and 13.56 MHz and 2.4 GHz for the iBracelet) we are able to gather data from both technologies during each user trial. The same inference engine is then used to detect activity for both data streams.

As will be shown in the following sections, activity detection using the RSN compares favorably to using the iBracelet. For nearly all activities the RSN saw better *recall* and the *precision* is on par with the iBracelet. This is a clear win as the RSN does not require the subject to wear any equipment, and precision may be enhanced at a higher layer, i.e. using other knowledge such as context.

Experimental Setup

Our user study consists of 10 subjects with each subject performing 14 specified activities. Twenty five objects were instrumented with WISPs and short range RFID tags.

The iBracelet queried for short range tags every second and transmitted tag IDs to the host computer, and the ceiling mounted RFID readers continuously read the WISPs and logged the accelerometer readings. The accelerometer readings were then filtered and a sequence of moved objects was generated. All object events were time stamped and logged at the host computer.

The host computer also presented a randomly ordered list of the 14 activities to the subjects. After completing each activity the subject indicated which activity they had just completed. This gave us a time stamped "ground truth" account of which activities were performed and at what times. This allowed us to automatically label our data indicating which objects were interacted with during each activity.

Data Analysis

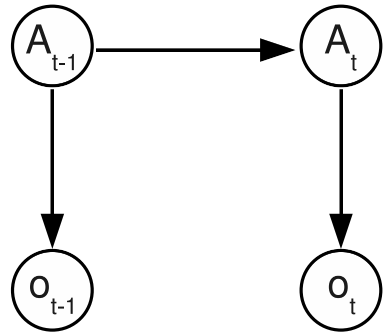

Figure 13. HMM mapping objects o_t to activities A_t

The Graphical Models Toolkit HMM, shown in Figure 13, is trained on a set of subject traces, each of which is a sequence of recorded object events. An iBracelet object event is an instance of an object being read, and for the RSN an event is an instance of detected movement. Each object event is labeled with the activity that was being performed when the event was recorded. The model is then tested on an unlabeled sequence, with the output being that same sequence annotated with the inferred activity for each object.

As the sampling rate is different for the two technologies, to compare between them we preprocessed the data traces by compacting repetition subsequences into "episodes". For example, if the subject handles an object for 45 seconds the iBracelet may report 45 events while the RSN may record only ten. After testing, each event will be labeled with an inferred activity, and we generally see long sequences with the same annotation. If we calculate per event false positives and negatives, and both technologies infer incorrectly, the technology with the higher sampling rate will be unduly penalized.

To avoid this, we consider these sequences as a single "episode", with each episode having a starting and ending time. In other words, an episode is the time during which our inference engine thinks a particular activity was occurring. If, based on the labeled data, the episode spanned two activities, we consider the episode to have taken place in both, i.e., it would be counted as both a true and a false positive.

To calculate *Precision* we compare the number of episodes that were correctly inferred with the total number of episodes (True Positives / (True Positives + False Positives)). *Recall* is calculated by determining how many of the 14 activities saw at least *one* correctly inferred episode (True Positives / Total number of Activities); multiple correct episodes during

Activity	RSN		iBracelet	
	Precision(%)	Recall(%)	Precision(%)	Recall(%)
Make Cereal	92	100	100	50
Make Sandwich	72	80	90	90
Make Coffee	92	100	100	90
Make Kool-aid	100	100	100	90
Read Book	85	100	88	90
Watch TV	83	40	100	90
Clean Windows	100	100	60	30
Tend to Plants	100	100	100	70
Use Telephone	100	100	100	20
Use Elliptical	100	100	100	30
Take Vitamins	73	80	100	30
Take Antacids	91	100	100	40
Brush Teeth	91	100	100	100
Go to Sleep	88	70	100	20
Totals	90	91	95	60

Table 2. Activity detection success

the same activity were counted only once. We performed a 10 fold cross validation on our 10 subject traces.

Results

Aggregate results for all subjects are shown in Table 2 for each of the 14 activities. As an example, consider the RSN results for making cereal. During the times when the subjects were making cereal, 92% of the episodes were labelled correctly and at least one episode was correctly inferred for each subject.

The RSN achieved 90% *precision* and 91% *recall*. The iBracelet, on the other hand, saw very high *precision* but low *recall*. The poor recall is largely because many objects do not have obvious grasping surfaces so tag placement is difficult. The iBracelet needs to be within 10 cm of a tag, and with a clear line of sight, in order to reliably read the tag.

For example, the only surface on the telephone handset where a tag could be placed was hidden from the iBracelet by the hand. In other cases, such as the cereal box, the subjects had many options for grasping and were likely to hold the box on an edge away from the tag. However, when a tag was in range of the iBracelet there was little chance of it being the wrong object. This results in high *precision*.

Table 3 shows the confusion matrix for the RSN results, calculated based on episodes. For example, while making cereal 11 episodes were correctly labeled while one was inferred as cleaning windows. Eleven correctly identified episodes with only ten subjects is caused by one subject seeing three episodes during the activity, with the first being correctly labeled, the second being "Windows", and the third again being correctly labeled.

Activities that were performed in the same area are often confused with each other, e.g. during the "Vitamins" activity two episodes were labeled "Antacids" and one was inferred as brushing teeth. In this case, the objects for the three activities are all located in the same cabinet and were likely shuffled unintentionally, or moved by the subject. In other cases, the root cause of mislabeling is unclear, such as labeling an episode as brushing teeth when the subject was make a sandwich.

To better understand the results using the RSN we looked at what objects were seen during each activity; this is shown in Table 4. The percentage of subjects that used the object is listed in parentheses and the average number of events during the activity is also shown.

The table shows what are likely spurious events, e.g. the plant food being used when making a bowl of cereal or using sugar when brushing teeth. These events are seen for only a single subject (10% of the subjects) and only one event is recorded in each case. Spurious reads of the plant food are seen quite frequently, suggesting that the noise in the sensor may be greater than 8°, and more sophisticated techniques should be used for detecting movement. However, it is interesting to note that there were no false positives for "Tending Plants" suggesting that our inference engine considered occurrences of the plant food as noise.

The most difficult activity to classify, watching TV, was due to the remote control being read during only six of the ten subject trials, and even then with a relatively low read rate. This could only be remedied by having better reader coverage.

Overall, we are pleased with the performance of our RSN based activity detection system and find the results promising. A system with better than 90% *precision* and *recall* is well above the bar for usefulness, particularly as it does not require instrumenting the subject. With that said, a far more extensive user study must be conducted, particularly in an actual home environment, before the real world viability of our system can be established.

Ground Truth ->	Cereal	Sandwich	Coffee	Koolaid	Book	TV	Windows	Plants	Telephone	Elliptical	Vitamins	Antacids	Teeth	Sleep
Cereal	11	1	0	0	1	0	0	0	0	0	0	0	0	0
Sandwich	0	8	0	0	0	0	0	0	0	0	0	0	0	0
Coffee	0	0	11	0	0	0	0	0	0	0	0	0	0	0
Koolaid	0	0	1	11	0	0	0	0	0	0	0	0	0	0
Book	0	0	0	0	11	1	0	0	0	0	0	0	0	0
TV	0	1	0	0	1	5	0	0	0	0	0	0	0	0
Windows	1	0	0	0	0	0	10	0	0	0	0	0	0	0
Plants	0	0	0	0	0	0	0	10	0	0	0	0	0	0
Telephone	0	0	0	0	0	0	0	0	10	0	0	0	0	0
Elliptical	0	0	0	0	0	0	0	0	0	10	0	0	0	0
Vitamins	0	0	0	0	0	0	0	0	0	0	8	0	0	0
Antacids	0	0	0	0	0	0	0	0	0	0	2	10	0	0
Teeth	0	1	0	0	0	0	0	0	0	0	1	1	10	1
Sleep	0	0	0	0	0	0	0	0	0	0	0	0	1	7

Table 3. Confusion Matrix for RSN

Activity	Objects (% Subjects): Average Number of Events					
Make Cereal	Milk(100): 5	Bowl(50): 2	Cereal(30): 2	BreadBox(10): 1	Towels(10): 2	Glass(10): 1
Make Sandwich	BreadBox(90): 2	Butterdish(40): 2	Milk(30): 2	Bowl(10): 1	Plate(10): 1	
Make Coffee	Coffee(100): 8	Cream(90): 6	Mug(30): 1	Sugar(30): 2	Jug(20): 5	Glass(10): 1
Make Koolaid	Jug(100): 18	Koolaid(100): 10	Glass(80): 3	Plant food(10): 1		
Read Book	Book(100): 12	Plant food(20): 1	BreadBox(10): 2	Remote(10): 2		
Watch TV	Remote(60): 4	BreadBox(10): 1	Plant food(10): 1	Towels(10): 1	Jug(10): 1	Koolaid(10): 1
Clean Windows	Windex(100): 10	Towels(80): 7	Plant food(30):1	Water Can(10): 2		
Tend Plants	Water Can(100): 7	Plant food(100): 10				
Use Telephone	Phone(100): 4					
Use Elliptical	Elliptical(100): 64	Plant food(10): 1				
Take Vitamins	Vitamins(90): 4	Antacids(20): 5	Toothpaste(10): 3	Plant food(10): 1		
Take Antacids	Antacids(100): 13	Toothpaste(10): 2	Vitamins(10): 1			
Brush Teeth	Toothpaste(100):9	Plant food(20):1	BreadBox(10):1	Antacids(10): 2	Bed spread(10): 4	Sugar(10): 1
Go to Sleep	Bed spread(80): 6	Plant food(10): 1	Towels(10): 1	Antacids(10): 1		

Table 4. Object Usage

DISCUSSION

Our behavioral monitoring system demonstrates that simple RFID-based sensors are useful for indoor activity recognition. Improvements in WISP technology will improve these systems through better coverage and more reliable sensor reads. Beyond this, however, we believe that the capabilities of RFID sensor networks can be leveraged for more effective activity recognition in at least two ways.

We currently threshold accelerometer data to detect movement. This throws away most of the data in the process. Instead, we could use the raw accelerometer data to extract details about object motion and classify gestures, e.g., the use of tools for their primary versus other purposes [9]. The 3D accelerometer data is particularly useful because it provides the orientation of the object with respect to gravity. The WISP also has other sensors, e.g., temperature and light, that could be leveraged in some settings, as well as the ability to be fitted with custom sensors, e.g., a capacitive sensor to measure the fill level of a container[28] or a neural sensor for recording the firing patterns of neurons[14].

We can use computation at the WISP to our advantage as well. Our system filters accelerometer data at the host to detect movement, but this functionality could instead be done at the WISP; only movement events (or even larger recognized gestures per the above) would then need to be communicated to the reader. This is advantageous because communication uses more energy on the WISP than computation, and it is likely that more sophisticated filtering can be supported with WISP rather than host computation. This would additionally reduce wireless traffic and enable higher rate sensing on the shared communication channel for selected objects.

Lastly, the read rates of the WISP tags in the system are coupled and there is the potential to improve the system as a whole. For example, in our study we found that reducing the number of antennas caused more energy to be directed towards a particular area. This caused some WISPs to experience increased read rates. Thus, the ability to control reader behavior, e.g., via LLRP applications, can be used to direct more energy to objects of interest. Ultimately, the emergence of flexible ”Gen 2” platforms, e.g., [7], will enable new protocols that are tailored to RFID sensor networks.

RELATED WORK

While prior work related to activity recognition and dense sensing technologies was covered in the introduction, there has also been work using the WISP platform. The advantage of sensor extensibility was shown in [14], where a WISP was fitted with a sensor capable of detecting the firing patterns of neurons, and [28] where a capacitive sensor was used to measure the fill level of liquid filled containers. Additionally, [8] demonstrated the computational capabilities of the WISP by implementing the first strong encryption (RC5) on a passive RFID device, and [9] showed that with high enough read rates gesture recognition could be performed on the WISP. The concept of RFID sensor networks, where large collections of WISPs are applied to traditional "smart-dust" applications, was introduced in [6]. This paper presents the first realization of this vision.

CONCLUSION

The ability to recognize indoor activities in a retargetable, inexpensive, and unobtrusive manner is a key capability that will support many ubiquitous computing systems. In this paper, we have presented the first study of dense sensing based on RFID sensor networks. Here, everyday objects are tagged with WISPs to detect when they are in use, and a simple HMM is used to convert object traces into high-level daily activities. We deployed 25 WISPs and three readers in a model apartment, then recruited ten users to perform 14 activities. Our recognition system based on an RFID sensor network was able to recognize tasks with roughly 90% precision and 91% recall. To put these numbers in context, we also ran a recognition system with a short-range RFID bracelet based on previous work. This system delivered 95% precision but only roughly 60% recall due to missing instances of object-use. We conclude that RFID sensor networks are already promising for indoor activity recognition. Moreover, we only expect their abilities to improve with time. This will come not only from better coverage and reliability, but also from making better use of the capabilities of WISPs.

REFERENCES

1. R. Aipperspach, E. Cohen, and J. F. Canny. Modeling human behavior from simple sensors in the home. In *Pervasive*, pages 337–348, 2006.
2. J. E. Bardram and H. B. Christensen. Pervasive computing support for hospitals: An overview of the activity-based computing project. *IEEE Pervasive Computing*, 6(1):44–51, 2007.
3. J. Beaudin, S. S. Intille, E. M. Tapia, R. Rockinson, and M. E. Morris. Context-sensitive microlearning of foreign language vocabulary on a mobile device. In *Ambient Intelligence*, volume 4794 of *Lecture Notes in Computer Science*, pages 55–72. Springer, 2007.
4. J. A. Bilmes. The GMTK documentation. http://ssli.ee.washington.edu/ bilmes/gmtk.
5. A. F. Bobick, S. S. Intille, J. W. Davis, F. Baird, C. S. Pinhanez, L. W. Campbell, Y. A. Ivanov, A. Schütte, and A. D. Wilson. The kidsroom. *Communications of the ACM*, 43(3):60–61, 2000.
6. M. Buettner, B. Greenstein, A. Sample, J. R. Smith, and D. Wetherall. Revisiting smart dust with rfid sensor networks. In *Proc. Hot Topics in Networks (HotNets)*, 2008.
7. M. Buettner and D. Wetherall. An empirical study of UHF RFID performance. In *Proc. Mobicom*, 2008.
8. H. J. Chae, D. J. Yeager, J. R. Smith, and K. Fu. Maximalist cryptography and computation on the wisp uhf rfid tag. In *Proc. Conference on RFID Security*, 2007.
9. A. Czeskis, K. Koscher, J. R. Smith, and T. Kohno. Rfids and secret handshakes: Defending against ghost-and-leech attacks and unauthorized reads with context-aware communications. In *15th ACM Conference on Computer and Communications Security*, 2008.
10. EPCglobal. Low level reader protocol (llrp), version 1.0.1. 2007.
11. K. P. Fishkin, M. Philipose, and A. D. Rea. Hands-on rfid: Wireless wearables for detecting use of objects. In *ISWC*, pages 38–43, 2005.
12. K. Z. Haigh, L. M. Kiff, and G. Ho. The Independent LifeStyle Assistant$^{\text{TM}}$ (I.L.S.A.): Lessons Learned. *Assistive Technology*, 2006.
13. S. Hodges, A. Thorne, H. Mallinson, and C. Floerkemeier. Assessing and optimizing the range of uhf rfid to enable real-world pervasive computing applications. In *Lecture Notes in Computer Science - Pervasive Computing*, volume 4480, pages 280–297, 2007.
14. J. Holleman, D. Yeager, R. Prasad, J. Smith, and B. Otis. Neuralwisp: An energy-harvesting wireless neural interface with 1-m range. In *BioCAS 2008*, pages 37–40, Nov. 2008.
15. J. Lundell, T. Hayes, S. Vurgun, U. Ozertem, J. Kimel, J. Kaye, F. Guilak, and M. Pavel. Continuous activity monitoring and intelligent contextual prompting to improve medication adherence. In *EMBS*, pages 6286–6289, Aug. 2007.
16. M. Philipose, K. Fishkin, M. Perkowitz, D. Patterson, D. Fox, H. Kautz, and D. Hahnel. Inferring activities from interactions with objects. *Pervasive Computing, IEEE*, 3(4):50–57, Oct.-Dec. 2004.
17. M. Philipose, J. Smith, B. Jiang, A. Mamishev, S. Roy, and K. Sundara-Rajan. Battery-free wireless identification and sensing. *Pervasive Computing, IEEE*, 4(1):37–45, Jan.-March 2005.
18. A. P. Sample, D. J. Yeager, P. S. Powledge, and J. R. Smith. Design of an rfid-based battery-free programmable sensing platform. In *IEEE Transactions on Instrumentation and Measurement*, 2008.
19. J. R. Smith, A. P. Sample, P. Powledge, A. Mamishev, and S. Roy. A wirelessly powered platform for sensing and computation. In *Proc. Ubicomp*, 2006.
20. M. B. Srivastava, R. R. Muntz, and M. Potkonjak. Smart kindergarten: sensor-based wireless networks for smart developmental problem-solving enviroments. In *MOBICOM*, pages 132–138, 2001.
21. E. M. Tapia, S. S. Intille, and K. Larson. Activity recognition in the home using simple and ubiquitous sensors. In *Pervasive*, pages 158–175, 2004.
22. ThingMagic. Mercury4 rfid reader. http://www.thingmagic.com/fixed-rfid-readers/mercury4/2-fixed-rfid-readers/22-mercury4.
23. ThingMagic. Mercury5e rfid reader. http://www.thingmagic.com/embedded-rfid-readers/mercury5e.
24. D. Tse and P. Viswanath. *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.
25. J. Ward, P. Lukowicz, G. Tröster, and T. Starner. Activity recognition of assembly tasks using body-worn microphones and accelerometers. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 28(10):1553–1567, October 2006.
26. M. Weiser. The computer for the twenty-first century. *Scientific American*, 265(3):95–104, September 1991.
27. D. H. Wilson and C. G. Atkeson. Simultaneous tracking and activity recognition (STAR) using many anonymous, binary sensors. In *Pervasive*, pages 62–79, 2005.
28. D. Yeager, R. Prasad, D. Wetherall, P. Powledge, and J. Smith. Wirelessly-charged uhf tags for sensor data collection. In *Proc. IEEE RFID*, 2008.