

I Sense a Disturbance in the Force: Unobtrusive Detection of Interactions with RFID-tagged Objects

K. Fishkin, B. Jiang, M. Philipose, S. Roy

IRS-TR-04-013

June 2004

DISCLAIMER: THIS DOCUMENT IS PROVIDED TO YOU "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NON-INFRINGEMENT, OR FITNESS FOR ANY PARTICULAR PURPOSE. INTEL AND THE AUTHORS OF THIS DOCUMENT DISCLAIM ALL LIABILITY, INCLUDING LIABILITY FOR INFRINGEMENT OF ANY PROPRIETARY RIGHTS, RELATING TO USE OR IMPLEMENTATION OF INFORMATION IN THIS DOCUMENT. THE PROVISION OF THIS DOCUMENT TO YOU DOES NOT PROVIDE YOU WITH ANY LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS

I Sense a Disturbance in the Force: Unobtrusive Detection of Interactions with RFID-tagged Objects

Kenneth P. Fishkin¹, Bing Jiang^{1,2}, Matthai Philipose¹, Sumit Roy²

¹Intel Research Seattle, 1100 NE 45th St,
Seattle, Washington, 98105 USA
{Kenneth.p.fishkin, Matthai.philipose}@intel.com
²Department of Electrical Engineering, University of Washington,
Seattle, Washington, 98195 USA
{bjiang, roy}@ee.washington.edu

Abstract. A novel method to infer interactions with passive RFID tagged objects is described. The method allows unobtrusive detection of human interactions with RFID tagged objects without requiring any modifications to existing communications protocols or RFID hardware. The object motion detection algorithm was integrated into a RFID monitoring system and tested in laboratory and home environments. The paper catalogs the experimental results obtained, provides plausible models and explanations and highlights the promises and outstanding future challenges for the role of RFID in ubicomp applications.

1 Introduction

Context inferencing is a cornerstone of ubiquitous computing [1,2]. A major component of context inferencing is activity inferencing – attempting, via the use of sensor networks, to infer the current activity of a person or group. Recently, several papers have suggested [3,4, 5, 6] that a fruitful method to infer these activities is by detecting *person-object* interactions: when a person picks up, touches, or otherwise uses an object in their daily domain at work or home. We, like other researchers, will focus on the home environment here, as it is particularly rich in objects of many types, but the method is not limited to that domain.

These techniques both work by affixing sensors to the objects of interest, but vary in the particular sensors employed. In one approach [3,4], “stick-on” sensors with an accelerometer, clock, and local memory are placed on the objects. Accelerometers detect touching, and record that to local memory. When the sensors are later removed, the data can be analyzed to reconstruct the set of all object touches and thereafter attempt to infer activity. This approach has the advantage of very high accuracy. False negatives are nearly impossible, and false positives only occur when the object is jostled without being truly used. However, the required stick-on sensors are custom-made, difficult to hide, and do not support *in situ* analysis.

In a second approach [5,6] the sensors employed are passive RFID (Radio Frequency Identification) tags. Passive RFID tags are an increasingly popular (cf [7,8])

sensor that consists of a batteryless transponder coupled to an IC chip. An external reader emits a low-power radio signal through its antenna to the passive tag. Upon receiving this signal via its own antenna, the IC in the RFID tag extracts the necessary power to energize an IC and then reflect back a modulated signal that carries some information, typically a globally unique tag ID (for more detail, see [9]). The reader then provides the list of sensed IDs back to the host application. This list is binary for each ID – it is either sensed, or not. These tags have the advantages of being cheap (\$0.30 each and falling rapidly), ubiquitous (forecasted deployments of billions per year), robust (they were originally developed to track livestock), hidable (under surfaces, in clothing, etc.), capable of wireless communication, and batteryless.

However, in this approach, how is the person-object interaction detected? The tag, as constructed, has no way of knowing that it is being moved; even if it did, it can communicate this information only when queried by a distant reader. Accordingly, in this approach the user wears an enhanced “glove”, a wearable device that fits over the palm and contains a small RFID reader with an antenna in the palm [5,6]. The reader is continually polling for nearby tags. The reader range is small (less than a cm), and so detection of an RFID tag can serve as a high-confidence indicator that the tagged object is about to be interacted with. This technique has the advantages of easy and cheap deployment and high accuracy, but has the significant disadvantage of requiring a user to employ a wearable.

In this paper, we attempt to overcome the disadvantages of these two approaches. Our goal is to see if we can unobtrusively detect interactions between an RFID-tagged object and a person without requiring said person to wear any special device at all. By combining the advantage of easy sensor deployment with that of unobtrusive sensor detection, we enable a powerful and attractive inferencing technique.

In the following sections, we explain and report tests of the technique. We first briefly discuss techniques that could potentially detect interactions, and motivate our choice. We then describe that basic technique in more detail, and show how it can be realized on existing unmodified RFID readers. Next, we present a series of experiments characterizing the technique with today’s equipment along a number of dimensions, including performance on some representative common use scenarios. We close with a summary and discussion of open areas for future improvement.

2 Constraining the solution

Given that we wish to detect interactions with RFID-tagged objects without employing a wearable, two general solutions are available:

- 1) *Enhance a tag.* If we “blended” the two types of sensors described above, we could use the accelerometer from the first to detect interactions and the RF protocol of the second to report them wirelessly. This method has potential, but has the disadvantage of being incompatible with the billions of tags already in existence and those planned for the next few years. We would prefer techniques that can work with existing tags.
- 2) *Reader energy analysis.* If we wish to use unmodified tags without a handheld reader, we must explore techniques that operate by interacting with long-range

readers. This is the avenue we pursue here. Before we explore it in detail, we must first investigate the nature of the reader-tag communication.

Consider the communication between an RFID reader and a passive tag, as sketched earlier. The RFID reader uses its own antenna to transmit a signal – for the rest of this paper, “reader” refers to the reader coupled with its antenna(e) for simplicity. The tag upon being energized by the impinging signal reflects a modulated version back to the reader. For successful detection by the reader, the tag antenna must capture enough energy from the reader to first energize itself and then send a sufficiently strong modulated signal back to the reader. The amount of energy that makes it back to the originating reader is a function of many parameters, of which the following three are primary: the energy emitted by the reader, the distance between reader and tag (the energy dissipates with distance), and the angle between reader and tag (this affects how much energy is captured by the antennae involved). Our technique relies on the observation that when an object is interacted with, typically *these last two parameters change*: the distance and/or angle of the tags with respect to the reader change. Therefore, when a tagged object is interacted with, the signal strength received by the reader will change.

2.1 Response Rate (α)

Unfortunately, today’s RFID readers only report a binary “seen/not-seen” for tags in their range. It is quite possible that a tag may be interacted with while staying “seen” throughout the interaction. Furthermore, today’s long-range RFID readers (we examined two of the most popular, those of Alien™ and of Matrics™) do not report or allow direct knowledge of the true back-scattered signal strength at the reader. It is also very difficult to tap in to the signal received by the antenna. For example, most readers use the same antenna to transmit and receive, and the much higher-energy, much noisier transmit signal is very difficult to separate from the received signals. Some readers (like Matrics) do use distinct antennae, but in this case the received signals are so weak as to be very difficult to detect without very expensive hardware.

However, we have found an approximation to this that works well and requires no modifications to current tags or readers, and which we employ in the rest of the paper. Existing readers support a “poll” command, wherein the reader transmits N poll commands per second to tags and reports the number of received responses for each tag. We therefore define a *response rate* α as the ratio of responses to polls. α is thus a scalar on $[0..1]$. When 0, the tag cannot be seen at all. When 1, the tag is always seen. We will investigate the choice of N later. Fig. 1 (left) shows the response rate of a tag at 4 different distances from a reader: generally, the farther the tag, the lower the response rate. This relationship is analogous to that of received RF signal power with the distance [10]. In other words, the response rate can be used to approximate the RF signal strength and is the basis of our subsequent processing algorithms.

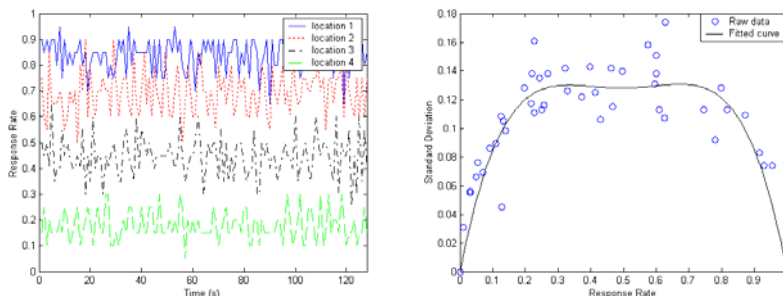


Fig. 1. Left: The response rate at 4 different distances, with $N=20$. Right: The relationship between mean and standard deviation of response rate ($N=10$)

Fig. 1 (left) also demonstrates that α is a noisy signal; some smoothing of the raw response rate is desirable. Thus, in subsequent figures, α denotes a suitably smoothed version of the raw response rates. To disambiguate between signal (a true object interaction) and noise (ambient jitter), the sample set N_s should be large enough to provide a good estimate of the mean signal, while small enough to detect changes quickly. We will amplify on these issues later.

To disambiguate between signal and noise, we need to know the standard deviation of an α set. The data shown in Fig. 1 (right) is derived based on $N_s = 256$ consecutive polls at a given location by varying the distance between reader and tag. Three different tag types were used to produce the plot that shows that the standard deviation is lowest when α is at the extreme values (0 or 1).

This affects our noise-signal disambiguation algorithms – a small change in α is more likely to be significant if α was at an extreme. The distances and orientations at which α is 1, 0, and in-between vary depending on the tag, the reader, the environmental conditions, and a host of other lesser factors. As a rough rule of thumb, we have found that with today's Alien™ readers α can detectably change with a motion as small as 3 cm and a rotation as small as 5 degrees.

Fig. 2 (left) illustrates how α changes as a function of the distance between tags and readers, for three different tag types. If the tag is too close (distance less than 150 cm), α is “saturated” at 1 – motions within that range can't be inferred. If the tag is too far (distance greater than about 275 cm in this case), α is saturated at 0. Again, motions within that region can't be inferred. (Fortunately, as we will see below, rotation works much more reliably). We stress that the exact values for the saturation regions are highly dependent on the particular reader, tags, and environmental conditions employed. This makes it vital that an interaction algorithm pay more attention to *changes* in α , rather than its absolute value.




Fig. 2. Response rate as a function of distance from the reader antenna (left), and as a function of angle between the reader and tag antennae (right).

Fig. 2 (right) illustrates how α changes as a function of the angle between the tag and reader antennae. This curve tends to be much smoother than the translation curve, as in this case we are basically reflecting the cosine-wave falloff pattern in how much energy reaches the tag. Inferring interactions from α will accordingly infer interactions which involve a rotation better than those which only involve a translation. Our experiments later will show that in our experience most interactions with an object have a significant rotational component.

We now discuss two other issues that impact on how we disambiguate true from false signals, namely flooring conditions and the presence of nearby metal.

Fig. 3 has the same axes as Fig. 2(left) (which showed α varying monotonically with distance), but here the relationship is clearly non-monotonic, and much more high-frequency. We have found that this behavior is often obtained in environments with metal floors (such as newly built laboratories). It appears that the metal “slats” in the floor serve as a wave guide, serving to increase the range in which α is non-zero, while making it much more volatile. This volatility actually aids our algorithm, as changes in α are more sensitive to small changes in distance. We have found that a similar, but coarser effect can be obtained by laying aluminum foil strips down on a wooden floor: one could therefore cheaply modify an existing room for improved reader distance.




Fig. 3. Response rate as a function of distance from the reader, with a metal floor

A subtler, rarer effect is caused by the proximity of tags to each other or to nearby metal, which we term the “coupling effect”. When two tags are placed on top or in front of each other, the top/front tag occludes the return signal from the other. Thus α for a tag could decrease without it being moved (negative coupling). A rarer occurrence is that if two tags are placed at a particular distance and relationship to each other, α for the non-moving tag can actually *increase* if the moving tag helps reflect extra energy onto the second tag (positive coupling). Fig. 4 shows both these cases – note the small increase in response rate at about 8 cm distance in the left example, and at about 32 sec elapsed time in the right example. This can also happen with metal objects other than tags – for example, if a person with a metal belt buckle is at just the right distance, this can cause the α for a nearby tag to increase. In our experience, positive coupling is extremely rare and will henceforth be ignored; however negative coupling is accounted for in our algorithm.




Fig. 4. The coupling effect: the response rate of a fixed tag changes as another tag moves just in front of it. In the left, in case “|” the tag is moved along the normal direction to the plane of the fixed tag; in case “—” the direction is parallel to that plane. On the right, the direction is again parallel to the plane, and the X axis now shows time as the tag moves

All figures in this section report results obtained by using the commonly available Alien Technology 915 MHz RFID system with the “I”, “C”, and “S” shaped passive tags, and their 2.45 GHz RFID system with “I” shaped tags. Circularly polarized reader antennae were used.

2.2 Using multiple tags/readers to increase accuracy

In the previous section we outlined how the response rate α changes as tags move or are rotated. Unfortunately, this is not that only thing that can cause α to change. RFID signals in this spectrum band are also reflected by metal, or more importantly blocked by water. If large bags of water such as humans move between a reader and a tag, α will plummet, just as it will when the tag is moved away from the reader – how can we disambiguate between these two cases? This is an inherently unsolvable problem in the base case – there simply is not enough information. We propose a novel

method for adding information to help with this problem, namely, by using *multiple* tags on an object and/or *multiple* readers placed with proper topology. By placing multiple tags on the same object at right angles to each other, we can cross-correlate their respective return rates. If α goes up for one tag, but down for another, then we infer that the tag is being rotated but not occluded (occlusion causes *all* alphas to decrease). Similarly, if we use readers located perpendicular to each other, most occlusions will cause only one set of alphas to drop (those obtained at the blocked reader), but those at the unblocked readers will stay approximately constant. Hence multiply tagged objects along with multiple readers can improve algorithm performance considerably, as shown in section 4.4 and 4.5.

3 Algorithms

Given the broad outlines of our approach, and initial signs that such an approach has the potential to robustly detect interactions with tagged object, we now describe how to turn the approach into a working algorithm. The algorithm must be able to detect interactions *reliably*, and as *quickly* as possible after the event has occurred.

3.1 Selection of N and N_s

The response rate measure α is itself derived from two parameters. From taking N polls per second we derive a sampling of response rate per second. By then taking the mean over N_s seconds we filter out noise and obtain the final value α . A small N means quick sampling, so sudden changes can be detected more quickly. However, the smaller N is, the less the granularity and resolution of α . Accordingly, a time vs. accuracy tradeoff occurs.

Similarly, N_s should be small enough for a quick decision, yet large enough to allow the final α to be a stable and accurate determination. The mean value variation of α decreases exponentially as N_s increases. Generally, the error is less than 5% when N_s is > 10 , the value we use hereout.

3.2 Excluding false positives due to mutual coupling

The negative coupling effect (the change in tag T2's response rate when tag T1 moves near it) can cause a false positive for tag T2. Although there is no fool-proof way to differentiate this effect from a real interaction, it is possible to look at the temporal relationship between the response rates, and use correlation analysis to exclude a false positive when there is a high correlation, as it is extremely unlikely that two tags on two different objects will exhibit the same alpha signatures over time. The correlation coefficient is defined as

$$r_{ce} = \frac{\sum_{i=1}^{N_s} (\alpha_{1i} - \alpha_{1mean})(\alpha_{2i} - \alpha_{2mean})}{(N_s - 1)\sigma_1\sigma_2} \quad (1)$$

where r_{ce} is the correlation coefficient; α_{1i} and α_{2i} are the response rates for tags 1 and 2, and σ_1 and σ_2 are the standard deviations of those rates. Fisher's Z-transformation [11] is used to prove the significance of r_{ce} :

$$\frac{1}{2} \ln \left(\frac{1+\rho}{1-\rho} \right) = \frac{1}{2} \ln \left(\frac{1+r_{ce}}{1-r_{ce}} \right) \pm 1.96 \sqrt{\frac{1}{N_s - 3}} \quad (2)$$

If the two solutions of (2) have the same sign, there is a statistically significant relationship between the two series at the 95% confidence level, giving us a 95% accurate way to detect and exclude negative coupling. Fig. 4 (right) shows an example of a false movement caused by the coupling effect. The statistical test successfully detects this case ($r_{ce} = 0.88$, $\rho_1 = 0.78$, and $\rho_2 = 0.94$).

3.3 The Final Algorithm

Given the considerations outlined earlier, we arrive at the final algorithm which consists of a series of initial checks applied to a new α for rapid detection of obvious motion. If no check is true, more response data is collected for reliable motion detection. If a check is true, then processing continues with some final screens. The cases checked in order are:

- 1) *Jump away/to $\alpha = 1$.* $\alpha = 1$ when the tag is close to the reader. If α jumps away from, or to, a value of 1, this check is set. We presently define a "jump" as being a delta of ≥ 0.1 , a threshold that was determined experimentally.
- 2) *Jump away/to $\alpha = 0$.* $\alpha = 0$ when the tag is far from the reader and/or at an unreadable orientation to it: the signal is at minimum strength. If α jumps away from, or to, a value of 0, this check is set. The "jump" delta is set at 0.1, a threshold that was determined experimentally.
- 3) *A large change in α .* This check is set if α is greater than 3 standard deviations away from the mean of the preceding set of N_s samples.

If none of these checks has been set, it indicates that a single α value is insufficient to reach a conclusion. We therefore collect a new *set* of N_s samples, with mean of α_2 and apply the following:

- 4) *Jump away/to $\alpha_2 = 1$.* Analogous to check #1.
- 5) *Jump away/to $\alpha_2 = 0$.* Analogous to check #2.
- 6) *Edge detected.* α_2 is essentially a low-pass filter on the N_s most recent samples. This is compared to a low-pass filter on the N_s samples before that: if they differ significantly, then we conclude that a significant change (an edge crossing) has occurred, and the check is set.

If any check is set, we have a reading of interest. Two final screens are then made:

Intel Research Seattle tech memo IRS-TR-04-013. To appear in Ubi-comp 2004

- *Coupling check.* As discussed in section 3.2, we check to see if the data is probably representing a coupling effect, rather than a true interaction. If it does, then no positive is signaled.
- *Occlusion check.* As discussed in section 2.2, if possible we check across multiple tags on the same object, and/or the readings for the same tag across multiple antennae, to see if we can rule out occlusion. If we can, then a positive is signaled.

If neither check has been passed, the algorithm concludes that an interaction has probably occurred, though occlusion is possible. In this case, a positive detect is signaled, with an additional bit raised indicating that occlusion is possible.

4 Experiments and Scenarios

In the previous sections, we have outlined the technique and discussed its many parameters. In this section we characterize the technique and show the effect of these parameters in practice. One difficulty of measuring any RFID technique is that RFID signal strength is impacted by many variables, e.g.:

- 1) *Flooring.* As mentioned in section 2.1, this has a significant impact.
- 2) *Distance between tag and reader.* As shown in Fig. 2.
- 3) *Number of tags on the object and their placement on object*
- 4) *Number of readers and their deployment topology*
- 5) *Number of nearby tags*
- 6) *Number of objects moved simultaneously*
- 7) *Tag orientation.* As shown in Fig. 2.
- 8) *Amount and direction of tag rotation*
- 9) *Tag Type.* Many different tags exist, optimized for different circumstances
- 10) *Type of Reader.* Different readers, especially those from different manufacturers, have vastly different performance

There are far too many possible combinations of conditions to exhaustively test each. Instead, in this section we rigorously show the performance of the algorithm on a “base condition”, and then show the effect of varying each of these 10 parameters independently. To make this more concrete, we then conclude with two parameter “bundles” representing common deployment scenarios.

4.1 The Base Condition

For the base condition, we tagged a single object - a cardboard cube 15 cm on a side that was read by an Alien 915 MHz reader with a circularly polarized antenna. It had a single Alien “I” tag placed on it at the same height as the reader center (100 cm off the ground) at a distance of 50 cm. The object and readers were deployed in a room in one of the authors’ houses, with linoleum flooring over concrete.

We then performed 8 experiments, each repeated 10 times. Assume a left-handed (X,Y,Z) coordinate system with the origin in the center of the object, Z pointing to-

wards the ceiling, and Y pointing towards the reader. The “I” tag was placed in two different orientations: once parallel to the Z axis and one parallel to the X axis: 100 events were done for each orientation. The activities can then be described as follows:

- 1) Rotate 90 degrees about Z
- 2) Rotate 90 degrees about X
- 3) Lift up (Z + 20 cm)
- 4) Pull away (Y – 20 cm)
- 5) Slide right (X + 20 cm)
- 6) Wave hand in front of tag
- 7) Walk in front of tag
- 8) Do nothing.

The results are graded as either “hits” or “misses”. For activities 1-5, a “hit” represents a correctly signaled interaction. For activities 6-8, a “hit” represents a correctly *non*-signaled interaction. The “occlusion is possible” bit from Section 3.3 was not included in the analysis. The results are as follows:

Table 1. Base condition: 1 tag, 1 reader, 50 cm distance

Activity	Hits:Misses	Accuracy
Z rotation	0:10 (Z tag)	0%
	10:0 (X tag)	100%
X rotation	10:0 (Z tag)	100%
	0:10 (X tag)	0%
Lift Up	0:20	0%
Pull away	0:20	0%
Slide Right	0:20	0%
Wave Hand	20:0	100%
Walk	4:6 (Z tag)	40%
	9:1 (X tag)	90%
Nothing	20:0	100%

In the base condition, the algorithm can detect most rotations, and is robust against most false positives, but surprisingly is completely unable to detect translations. In our experience, as the later tables will show, this is largely a function of the particular reader (Matrices readers, which we have not yet exhaustively tested, seem to perform much better than Alien readers) and the environmental conditions, particularly the flooring effect mentioned in Section 2.1. Given this behavior in the base condition, we now show the effects of altering individual parameters of the scenario, to aid in characterization, starting with the just-mentioned floor effect.

4.2 Varying Floor

We repeated the Lift/Pull/Slide activities within a building with a raised metal floor (as is common in many computing environments). The prominent nearby metal acts as a waveguide as discussed above at both 50 and 100 cm distances and significantly improves the accuracy of correct detection.

Table 2. The effect of floor type on accuracy

Activity	Hits:Misses 50 cm	Accu- racy	Hits:Misses 100 cm	Accuracy	Hits:Misses 200 cm	Accu- racy
Lift Up	10:0	100%	10:0	100%	0:10	0%
Pull away	10:0	100%	10:0	100%	0:10	0%
Slide Right	10:0	100%	10:0	100%	0:10	0%

The metal floor had a huge impact on translations, which now are detected with complete accuracy. We have found this can also be replicated by running a strip of tin foil along a floor.

4.3 Varying Distance

The base condition was performed, but now with 100 cm and 200 cm tag distances from the reader. The 100 cm results were identical to the base condition. At 200 cm, the results were identical except as follows:

Table 3. The effect of 200 cm distance

Activity	Hits:Misses 200 cm Z (X)	Accuracy Z (X)
Wave Hand	6:4 (10:0)	60% (100%)
Walk	4:6 (4:6)	60% (60%)

As the distance increases, the alpha values can get so low that it becomes more difficult to exclude a true positive from a false one when waving a hand.

4.4 Varying Tags Per Object

In the first variant, two tags were placed on the object, one parallel to the Z axis, one parallel to the X axis. In the second variant, a third tag was added, parallel to the Y axis. The results were equal to the base condition, except for the rotation activities, which were now detected with 100% accuracy, and the “Walk” activity, which was detected with 50% accuracy with 2 tags and 60% accuracy with 3 tags.

4.5 Two Readers

The base condition was repeated, but now employing two readers, located perpendicular to each other. The results were equal to the base condition, except for the

Intel Research Seattle tech memo IRS-TR-04-013. To appear in Ubi-comp 2004

rotation activities, which were now detected with 100% accuracy, and the “Walk” activity, which was detected with 90% accuracy.

4.6 Multiple Objects in the Field

The base condition was repeated with 3 tagged objects in the field of the antenna, and with 6 tagged objects in the field. The results were equal to that of the base condition, except as shown below:

Table 4. The effect of multiple objects in the field

Activity	Hits:Misses 3 objects	Accuracy	Hits:Misses 6 objects	Accuracy
Walk	21:9	70%	52:8	65%

We see that adding more tagged objects to the field had no significant effect: the algorithm should scale well as the number of tagged objects in the environment increases.

4.7 Multiple Objects Moved Simultaneously

The base condition was repeated, where two objects were moved simultaneously. No difference was detected from the base condition.

4.8 Orientation

The base condition was repeated with the tag original orientation varied to 30, 60, and 90 degrees off the X axis. The results were equal to that of the base condition, for the “Walk” activity, which had 50%, 50%, and 50% accuracy at orientations of 30, 60, and 90 degrees, respectively.

The 90 degree case is the most interesting. In this case, the reader normally does not see the tag, as its antenna is oriented such that it cannot catch and reflect sufficient energy. However, even in this configuration, the algorithm still detected rotation, as this rotation brings the antenna into the view of the reader.

4.9 Magnitude of Rotation

We varied the amount of rotation for the “Z-rotation” and “X-rotation” activities. Instead of being fixed at 90 degrees, 30 degree and 60 degree rotations were also used. These were tested with tags located parallel to the X, Y, and Z axes:

Table 5. Effect of varying magnitude of rotation

Activity	Hits:Misses X-parallel tag	Accu- racy	Hits:Misses Y-parallel tag	Accu- racy	Hits:Misses Z-parallel tag	Accu- racy
Z rotation, 30 degrees	0:10	0%	10:0	100%	0:10	0%
Z rotation, 60 degrees	10:0	100%	10:0	100%	10:0	100%
X rotation, 30 degrees	0:10	0%	0:10	0%	10:0	100%
X rotation, 60 degrees	10:0	100%	10:0	100%	10:0	100%

We see that 60 degrees is sufficient to detect rotation. At 30 degrees, it appeared to only work for one of the three possible tag orientations.

4.10 Tag Type

In this condition, we tested two other types of tags, the “S” and “C” tags from Alien. The “S” tag is an older tag which has been largely supplanted by the “I” tag, the “C” tag is a smaller tag optimized for placement near liquid. The results were equal to that of the base condition, except for the “Walk” activity, which was detected with 100% accuracy by the “S” tag and 50% accuracy by the “C” tag. Our conjecture is that the older “S” tags perform better as the newer tags emphasize cost savings over range.

4.11 Living Room Scenario

There are so many variables and parameters that it can be difficult to get a sense for typical real-world performance from the preceding tables. Accordingly, we tested the algorithm on two specific scenarios we felt representative of real-world settings for activity inferencing via tagged objects, namely living rooms and bathrooms in the home [12,13].

The first scenario represents a typical living-room interaction. Four items were tagged: a hardbound book (2 tags: back cover and spine), a magazine (1 tag: back cover), a deck of cards (1 tag on the box), and a TV remote control (1 tag on the back). The objects were placed in their normal positions on a dresser and magazine rack next to a living room chair. We then performed a sequence of typical interactions with these objects, 30 interactions in total: objects were picked up and/or put down 23 times, an object was motioned with while in the hand 3 times, and a hand was waved in front of each object, for a total of 4 interactions. Two readers were used, both

wall-mounted, on perpendicular walls. The algorithm was then left running overnight with no human present in the chair, to guard against false positives.

The results were as follows: all 23 pick up / put down events were detected, all 3 motions with an object were detected, and all 4 hand-waves were correctly labeled as occlusions: 100% accuracy by the algorithm. No false positives occurred.

A second experiment was then performed with the same tagged objects, but this time using only a single wall-mounted reader. This time 9 pick up / put down events were performed and 3 interactions where one object was placed atop another (the book on top of the magazine).

The results were as follows: all 9 pick up / put down events were detected, and one of the 3 placements were detected: the other two were not. No false positives occurred. Overall, 10 of 12 events were detected, for an accuracy of 83%.

4.12 Bathroom Scenario

In this scenario, four tagged items were tagged and placed on a bathroom counter: a canister of hair spray (2 tags: bottom and side), 1 drinking cup (2 tags: bottom and side), a towel (2 tags, placed at right angles to each other on the plane of the towel), and a soap dispenser (1 tag).

We then performed 6 events where a single object was picked up or put down, and 5 events where a pair of objects were picked up or put down in unison.

Results: all 11 events were correctly detected. However, one false positive occurred. Overall accuracy is therefore 11/12, or 92%

4.13 Experimental Results

The preceding experiments represent a total of 1353 tested events. We can roughly summarize the results with this particular make of RFID reader as follows:

- *Test scenarios.* The system did very well on both the “real-world” simple tests, with an overall accuracy of 94%. While the reader placement and the lack of multiple-person movement simplified the situation over a true real-world deployment, we believe that the results are encouraging.
- *Rotation.* The system could nearly always detect rotations, particularly when additional tags and/or readers were employed. As most real-world manipulations involve some degree of rotation, it appears this will be the most common detection mechanism. Another advantage of this emphasis is that the pattern obtained by an occlusion (all response rates dropping) is virtually never seen in a rotation so long as multiple tags are employed – this could greatly reduce false positives.
- *Translation.* The system was nearly unable to detect translation-only movement. This is largely due to the fact that we had to approximate received signal strength by response rate, which is often too insensitive to motion. We believe that this problem is a transient one, as future readers become more sensitive and more open to queries of signal strength.

Intel Research Seattle tech memo IRS-TR-04-013. To appear in Ubi-comp 2004

- *False positives.* This varied the most depending on configuration and setting. In general, with only one tag and only one reader, disambiguation is quite difficult. However, by adding additional tags or especially readers much better performance was obtained.

For a feasible real-world Ubicomp deployment, we would like a system where one or two readers could “strobe” an average-sized room. Present readers and tags don’t quite reach this goal due to the energy required to energize the tag, but as RFID tags continue to use Moore’s law to reduce their energy requirements, and hence increase their range (for example, in the last 5 years reader range has increased by nearly a factor of 10), we believe that in a year or two room-sized strobing will be feasible.

5 Conclusions

As ubiquitous computing matures, we will need increasingly powerful context inferring from increasingly unobtrusive sensor networks. In this paper we have described one potentially powerful aid to this goal: using long-range unobtrusive RFID detectors to detect people’s interactions with RFID-tagged objects. The algorithm works, albeit in limited circumstances, today. As RFID tags and especially RFID readers continue their exponential rates of improvement in range, size, and cost, we believe the algorithm will become more and more attractive.

There are many areas for future work. On the hardware front, we plan to explore enhanced RFID tags that can detect and report acceleration directly – this would remove the guesswork from the system, albeit at the cost of introducing a new backwards-incompatible sensor. We also plan to explore upcoming RFID readers to see if they can provide more direct measurement of received signal strength than the response rate approximation. On the software front, the algorithm can be improved through improved statistical techniques, for example by analyzing data streams to learn the best values for the “jump” thresholds used in several of the screening tests.

References

- 1 Dey, A.K., Salber, D., and Abowd, G.D.: “A Conceptual Framework and a Toolkit for Supporting the Rapid Prototyping of Context-Aware Applications”, HCI Journal 16(2-4), 2001, pp. 97-166.
- 2 Schilit, B.N., Adams, N.I., and Want, R.: Context-Aware Computing Applications, MCSA ‘94
- 3 Munguia, E.T.: “Activity Recognition in the Home Setting Using Simple and Ubiquitous Sensors”. MIT M.S. Thesis, 2003
- 4 Munguia, E.T., Intille, S.S., and Larson, K. “Activity Recognition in the Home Using Simple and Ubiquitous Sensors”. Pervasive 2004, to appear.
- 5 Reference removed for anonymity
- 6 Perkowitz, M., Philipose, M., Patterson, D.J., and Fishkin, K.P.: “Mining Models of Human Activities from the Web”, WWW 2004 (to appear).

Intel Research Seattle tech memo IRS-TR-04-013. To appear in Ubi-comp 2004

- 7 Stanford, V.: "Pervasive computing goes the last hundred feet with RFID systems," *IEEE Pervasive Computing*, vol. 2, no. 2, pp. 9-14, 2003.
- 8 Want, R.: "RFID: A Key to Automating Everything", *Scientific American*, January 2004
- 9 Finkenzeller, K.: *Rfid handbook: fundamentals and applications in ccontactless smart cards and identification*, Wiley, 2nd ed., 2003
- 10 Kraus, J.D. : "Antennas", McGraw-Hill, 1988
- 11 Mode, E. B.: *Elements of probability and statistics*, Prentice-Hall, 1966.
- 12 Morris, M., Lundell., J, Dishman, E., Needham, B: "New Perspectives on Ubiquitous Computing from Ethnographic Study of Elders with Cognitive Decline". *UbiComp 2003: 227-242*
- 13 Mynatt, E., Essa, I., and Rogers, W.: "Increasing the Opportunities for Aging in Place", *Universal Usability*, 2000 November 2000