
An RFID-based object localisation framework

Kirti Chawla* and
Gabriel Robins

Department of Computer Science,
University of Virginia,
Charlottesville, 22904, USA
Email: kirti@virginia.edu
Email: robins@virginia.edu
*Corresponding author

Abstract: Numerous ubiquitous computing applications depend on the ability to locate objects as a key functionality. We show that Radio Frequency Identification (RFID) technology can be leveraged to achieve object localisation in an inexpensive, reliable, flexible, and scalable manner. We outline the challenges that can adversely affect RFID-based localisation techniques, and propose practical mitigating solutions. We present several new algorithms for RFID-based object localisation that compare favourably with previous methods in terms of accuracy, speed, reliability, scalability, and cost.

Keywords: RFID; RFID-based positioning; object localisation; localisation algorithms; power-distance relationship.

Reference to this paper should be made as follows: Chawla, K. and Robins, G. (2011) 'An RFID-based object localisation framework', *Int. J. Radio Frequency Identification Technology and Applications*, Vol. 3, Nos. 1/2, pp.2–30.

Biographical notes: Kirti Chawla is currently a PhD student in the Department of Computer Science at the University of Virginia. He received an MTech in Information Technology from the Indian Institute of Technology (IIT) in 2003. His research interests include RFID, ubiquitous computing, computer security, wireless sensors, and embedded systems. He co-authored five refereed publications and three patents. From 2003 to 2007 he held software engineering and research positions at Samsung Electronics (India), Samsung Semiconductor (South Korea), and the Prabhu Goel Research Centre (India). He is a member of IEEE, ACM, and the Cryptology Research Society of India.

Gabriel Robins is Professor of Computer Science at the University of Virginia. He received a PhD in Computer Science from UCLA in 1992. His research interests include algorithms, optimisation, RFID, VLSI CAD, and bioinformatics. He co-authored a book and almost 100 refereed papers. His recognitions include a Packard Foundation Fellowship, a National Science Foundation Young Investigator Award, the SIAM Outstanding Paper Prize, and several teaching awards. He served on the US Army Science Board, and on the editorial boards and technical programme committees of several leading journals and conferences. He consults as an expert witness in major intellectual property litigations.

1 Introduction

The confluence of Radio Frequency Identification (RFID) and other wireless technologies lies at the heart of many emerging applications, such as remote medicine, robotic teams, wireless sensing, early warning systems (e.g. for tsunamis, earthquakes, chemical spills, etc.), locating points of interests (e.g. ATMs, banks, hospitals, etc.), and automated inventory management (Abowd and Mynatt, 2000; Hightower and Borriello, 2001; Mattern, 2001; Satyanarayanan, 2001; Estrin et al., 2002; Romer and Domnitcheva, 2002; Vogt, 2002; Fontelo et al., 2003; Schilit, 2003; Merrell et al., 2005; Muthukrishnan et al., 2005; Romer et al., 2005; Blewitt et al., 2006; Liu et al., 2006; Wang et al., 2007; Want, 2008). Such applications require capabilities that include real-time object identification, object tracking, and position localisation.

While typical RFID technology is sufficient for object tracking (i.e. registering the presence/absence of an object in a radio field) and identification (i.e. matching an onboard RFID tag ID with a trusted database), it does not normally provide object localisation capabilities (i.e. precisely locating the position of an object). Several RFID-based localisation techniques for stationary and mobile objects have been proposed (Ni et al., 2003; Alippi et al., 2006; Senta et al., 2007; Milella et al., 2009). However, these techniques tend to compromise key requirements such as accuracy, speed, cost, scalability, and reliability, thus severely degrading the utility of these methods. Moreover, some previous localisation methods also require cumbersome non-RFID technologies such as ultrasonic sensors, vision sensors, cameras, and lasers, which again make them unsuitable for practical use in typical environments.

We address these limitations by developing a scalable and reliable RFID-based localisation framework that accurately and rapidly determines the positions of stationary and mobile objects. Our approach consists of separate techniques to localise target tags, as well as localise readers attached to mobile objects. To localise stationary and mobile target tags, we vary the reader power levels over a set of calibrated reference tags having known sensitivities. Separately, we determine the positions of target mobile readers by measuring their proximity to known reference tags. Moreover, these two approaches can be combined to yield even higher accuracy and efficiency.

We implemented, tested, and evaluated the proposed approach to confirm its general applicability, scalability, and reliability. Our approach suits a wide range of requirements and trade-offs including accuracy, speed, and cost. We have also identified several key challenges (e.g. environmental interferences, tag sensitivity, spatial arrangement of tags, etc.) that adversely affect the performance of RFID-based object localisation, and we propose mitigating techniques.

This paper is organised as follows. Section 2 describes related research work in RFID-based object localisation. We formulate the problem of object localisation using RFID in Section 3. Section 4 presents several localisation challenges and mitigating techniques. We describe our object localisation framework in Section 5, and discuss the experimental evaluation and results in Section 6. Section 7 outlines key lessons learned, and Section 8 concludes with future research directions.

2 Related work

Recent advances in ubiquitous computing have necessitated RFID-based object localisation capabilities, with research efforts specifically targeting the positioning of either stationary or mobile objects. RFID-based localisation techniques can be broadly classified as reader and tag-based approaches. In reader-based localisation techniques, the positions of RFID readers are ascertained, while in tag-based localisation techniques, the positions of RFID tags are determined. Note that RFID tags and readers can each be either stationary or mobile. In this paper, we focus on pure-RFID object localisation techniques, utilising only the interaction between RFID readers and tags (i.e. other RF-based approaches utilising near-field propagation, surface acoustic waves, microwaves, cameras, ultrasonics, etc., are outside the scope of this work, and arguably are not as useful in many RFID scenarios). Existing RFID-based stationary object localisation techniques are described below.

Ni et al. (2003) propose placing active reference tags and determining the Euclidean distance between the reference and the target tags. K -nearest reference tags are used to determine the position estimates of a target tag, with a maximum localisation error of less than two metres. Alippi et al. (2006) model the indoor localisation problem as a non-linear stochastic inversion problem. Their experimental 2D environment has multiple readers at fixed locations and tags at unknown locations. Data is gathered using multiple antennas at different orientations. A conditional probability-based model is used, wherein tag detection probabilities vary at different power levels, yielding an average localisation error of 0.68 metres. Bekkali et al. (2007) use two mobile readers, a probabilistic RFID map, and a Kalman filter-based technique to minimise the localisation error variance. Position estimates of the target tags are determined using a Received Signal Strength Indicator (RSSI)-based metric, and a probability density function generates the probability map for each reference tag. The localisation error of this approach has a root mean square in the range of 0.5 to 1 metres.

Joho et al. (2009) develop a probabilistic sensor model based on the tag RSSI measurement, the antenna orientation, and tag location. A mobile reader moves through the environment to gather tag measurements and correlates them with the true locations. Multiple iterations are required to improve the tag position estimates, resulting in an average localisation error of 0.375 metres. Zhang et al. (2007) introduce the concept of virtual tags and a proximity map. Their key idea is to consider the presence of virtual tags with the reference tags. The RSSI values of virtual tags from each reader are calculated using a linear interpolation algorithm. Different proximity maps are generated for each reader, and the intersection of these maps is used to determine the location of the target tags. The localisation error of this approach is in the range of 0.14 to 0.29 metres.

Wang et al. (2007) propose a 3D tag positioning scheme, wherein reference tags are placed either on the floor or ceiling and at least four readers are placed on the vertices of a hexahedron. Readers gradually increase their transmission power until responses are received from the reference and target tags. Statistical averaging and the simplex method are used to reduce the localisation error to a range of 0.1 to 0.9 metres, but at the cost of high hardware expense and long positioning times. Choi and Lee (2009) study the characteristics of a passive UHF RFID system and propose an RSSI-based localisation

approach using passive tags. The K-nearest neighbours algorithm is utilised to compute the differences of the RSSI-based metric of various reference tags in order to localise a single target tag, with an average localisation error of 0.21 metres.

Hekimian-Williams et al. (2010) utilise the phase difference of the signals received at two separate antennas to localise the active tags. Additionally, they make use of software-defined radios coupled with accurately sampled clocks to implement various phase difference estimation algorithms. Thus, clock precision is an important factor in determining the localisation accuracy. While their system yields high accuracy under ideal conditions, they do not take into account key factors such as multi-path scattering and tag sensitivity. Jin et al. (2006) propose to improve the localisation accuracy of the LANDMARC system (Ni et al., 2003) by selecting only a few reference tags that have the least distance from a target tag. They utilise multiple readers to localise the target tags to within an average localisation accuracy of 0.72 metres. Zhang et al. (2007) propose using the direction of arrival of tag responses in order to localise the target tags. Simulations indicate an average localisation error of 1 metre. However, the effects of multi-path scattering, environmental interferences, and tag sensitivity variations are not considered.

Some RFID-based positioning techniques are specifically designed to localise *mobile* objects (as opposed to stationary ones). For example, Chae and Han (2005) propose a two-step approach to localise mobile robots in an indoor environment. In their first step, an onboard RFID reader is coarsely localised with respect to neighbourhood active reference tags. In the second step, a vision sensor combined with a feature detection algorithm identifies key environmental features to minimise the average localisation error to 0.23 metres. Their approach is less applicable in different scenarios since the onboard vision sensor requires a sufficiently illuminated environment and objects must be within line-of-sight (a fundamental drawback that RFID technology was intended to eliminate in the first place).

Choi and Lee (2009) propose to localise mobile robots in an indoor environment by utilising ultrasonic sensors in combination with an onboard reader. Their localisation approach has two stages. In the first stage, the global position of the mobile robot is estimated through onboard reader localisation with respect to the neighbourhood passive reference tags. The second stage uses ultrasonic sensors for local position estimates. While their approach can yield higher accuracy, it is inherently not a pure RFID-based method, but rather a sound-based approach and is thus highly limited by issues such as environmental noise, line-of-sight, echoes, etc.

Hähnel et al. (2004) propose a laser range scanner combined with an RFID reader onboard a mobile robot. The laser range scanner is used to learn a map comprised of reference tags, which in turn is used to estimate the position and orientation of mobile robots. However, this approach imposes line-of-sight constraints, and moreover tag orientation issues degrade the detection probability of the reference tags, resulting in high localisation errors in the 1 to 10 metres range. Han et al. (2007) propose a mobile object localisation technique by using reference tags and onboard mobile readers. They show that the particular spatial arrangement of tags affects the localisation error and propose a triangular tag arrangement scheme to minimise it. Their approach yields an average localisation error of 0.09 metres in a small test region of one metre square.

Milella et al. (2009) utilise an onboard monocular camera, a reader and a tag bearing estimation technique based on a ‘fuzzy inference system’ to localise mobile robots to within an average error of 0.64 metres. Senta et al. (2007) present a mobile robot localisation technique based on reference tags, onboard readers, and a support vector machine (SVM)-based machine learning approach. This method yields localisation errors of over 0.2 metres, and is limited by the tag spatial arrangement, measurement noise, and tag-reader proximity. Seo and Lee (2008) describe a mobile object localisation system that transmits an RFID signal from an onboard reader to the neighbourhood beacon, which in return responds with an ultrasonic signal. The estimated distance is computed based on the time difference between transmitted and received signals, with an average localisation error in the range of 0.2 to 1.6 metres. Vorst et al. (2008) present a mobile object localisation approach using reference tags, onboard readers, and a particle filter-based technique. They compare prior-obtained training data with real-time RFID measurements to yield an average localisation error in the range of 0.2 to 0.6 metres.

Currently, the effectiveness of several of the previous approaches is hindered by reliance on line-of-sight techniques, combining multiple non-RFID (e.g. ultrasonic sensors, cameras, lasers, etc.) and RFID components in an ad-hoc manner, large number of onboard components, and high localisation delays (Chae and Han, 2005; Hähnel et al., 2004; Choi and Lee, 2009; Milella et al., 2009). Moreover, some of the above methods are too expensive or unwieldy due to the cost, size, and weight of the required infrastructure. Finally, the above approaches ignore the key issue that the RFID equipment itself can introduce significant amount of experimental errors. For example, previous works ignore the fact that ‘identical’ tags can have widely varying detection sensitivities, which can greatly affect the experimental outcomes (Chawla et al., 2010a; Chawla et al., 2010b). Thus, instead of addressing and mitigating these basic principles (as we do in our approach), previous research works resort to Herculean efforts in order to reduce the errors on other fronts, while ignoring bigger error sources, resulting in a hodgepodge of ad-hoc and sometimes ineffectual techniques.

3 Problem statement: object localisation using RFID

We address the problem of localising stationary and mobile objects by utilising ‘only’ RFID-based technology (as opposed to relying on non-RFID technology such as lasers, ultrasonic sensors, cameras, etc.). In this section, we describe the underlying principles of the proposed approach and the key performance parameters for optimisation. RFID-based object localisation requires determining the positions of stationary and mobile objects affixed with tags and/or readers. Radio signal properties such as power-distance relationships can ascertain these locations. Theoretically, the radio wave’s power-distance relationship can be characterised based on the Friis transmission equation as follows (Finkenzeller, 2003):

$$\frac{P_R}{P_T} = G_R G_T \left(\frac{\lambda}{4\pi D} \right)^2 \quad (1)$$


Here, P_R is the power transmitted by the reader, P_T is the power received at the tag, G_R and G_T are the respective antenna gains of the reader and the tag, λ is the radio wave wavelength, and D is the distance between the tag and reader. For a typical RFID system,

variables such as λ , G_R and G_T are some of the main design parameters. Thus, by knowing the power levels at the reader and the tag, the distance between them can be estimated. Alternatively, if the distance between the readers and tags are known, then the received power level at the tags can be determined.

Thus, our overall RFID localisation strategy is as follows. We slowly increase the reader's power level from low to high in small increments. When a given tag becomes detectable to a reader for the first time, the power level at which this first detection event happens indicates the tag's distance from that reader. As different readers perform such readings (from different directions), the tag's position can be estimated with increasing accuracy by considering the intersections of these detection regions.

Figure 1 illustrates a shared region induced by the geometric intersections of the radio wave lobes used to detect a tag by several readers. Such intersection regions, if small enough in size, can help minimise the error in position when estimating the locations of target objects using reference tags (i.e. regions overlapped by more radio wave lobes have a smaller area than other more peripheral regions covered by fewer lobes, resulting in increased localisation accuracy). Note, however, that this intuitive intersection-of-regions analogy is only a conceptual explanatory tool. Our system does not explicitly compute geometrical regions, nor is it even particularly aware of geometry in general. Rather, our system compares the detection power levels of target tags with those of known reference tags, in order to infer the target tag's position. In other words, our approach is 'relativised' in that it tries to match the behaviours of known and unknown tags, under the key assumption that *if the behaviours and responses of two tags are very similar, then their positions must be very close as well*.

Figure 1 A shared region induced by the intersection of radio wave lobes



At first glance it may seem contradictory that a *positioning* system can be mostly oblivious to geometrical considerations. However, because of all the real-world factors that interfere with accurate RF transmission and reception, correlating a complex geometry with precise levels of RF receptivity is difficult. Our system sidesteps these complicated issues by ignoring the geometry, and instead takes a pragmatic *relative* approach by *observing and comparing* behaviours rather than trying to accurately *predict* them. Note that such an *empirical* approach naturally adapts and automatically calibrates to unknown conditions and unexpected effects, since these would presumably affect (identical) target and reference tags in a very similar way. Thus, the geometry-obliviousness feature of our system is not a weakness but rather a deliberate capability that yields performance advantages.

In real-world scenarios, various ad-hoc interfering factors (e.g. environmental conditions, multi-path scattering, and RF occlusions due to liquids and metals, etc.) affect signal strengths and received power levels. Moreover, variability in detection sensitivities across ‘identical’ tags poses a unique challenge in reliably establishing and leveraging the empirical power-distance relationship. To understand the implications of this variability, consider two tags of the same type (e.g. ‘Impinj Dogbone Monza 3’ UHF passive tag) having different sensitivities (due to manufacturing variations).

These tags, when kept at the same fixed location from the reader, will be initially detected at different reader power levels, thereby skewing the observed empirical power-distance relationship. Our proposed object localisation framework considers these challenges and takes the pragmatic approach of only using *uniformly sensitive* reference tags to establish the empirical power-distance relationship. Section 5.1 below will discuss this sensitivity analysis in greater detail.

4 Localisation challenges

As discussed above, all RFID-based object localisation techniques have inherent position estimate errors due to various external (e.g. environmental interferences) and internal (e.g. RFID tags and reader related) factors. This section describes several key challenges that could induce localisation errors and our proposed techniques to mitigate them.

4.1 Interference and RF occlusion

Environmental factors such as radio noise and occlusions by liquids or metals (which tend to be opaque to RF signals) can cause radio wave scattering and attenuation, which in turn can result in localisation errors. Mitigating techniques such as electrostatic shielding, full Faraday cycle analysis, and path-loss contour mapping can help reduce the impact of such factors on localisation accuracy (Sweeney, 2005). Deploying more tags and readers in the region of interest can also reduce adverse effects due to interferences and occlusions.

4.2 Tag sensitivity

Tag detection sensitivity is characterised by the minimum power needed to read the tag at a particular distance. It is a function of chip threshold power sensitivity, tag antenna gain, and chip’s high impedance state (Nikitin and Rao, 2008). Moreover, tag manufacturing variability can dramatically affect the detection sensitivities of tags. Thus, tags with low sensitivities become invisible at shorter distances than their higher sensitivity counterparts, leading to localisation errors. To address this issue, we propose a pre-processing step of sorting (i.e. ‘binning’) the tags based on their detection sensitivities. We thus classify tags as ‘highly sensitive’, ‘average sensitive’, and ‘low sensitive’ using read measurements over different power and distance combinations (Chawla et al., 2010a; Chawla et al., 2010b), as detailed in Section 5 below. This enables only uniformly sensitive tags to be deployed in the same experiment, resulting in more consistent and meaningful experimental results. Curiously, previous works all seem to ignore this critical issue.

4.3 Tag spatiality


RFID-based object localisation techniques typically utilise reference tags placed at known locations. The positions and arrangements of these reference tags can significantly affect the localisation accuracy. Regular placements of the reference tags (as opposed to random arrangements) tend to yield lower positioning errors (Han et al., 2007).

4.4 Tag orientation

Tag orientation significantly affects tag and reader interaction. For example, Bolotnyy and Robins (2007a, 2007b, 2009) analysed how tag orientation impacts the tag detection probability. In particular, they discovered that when multiple tags are placed on the same object, orthogonal orientations yield much higher detection probabilities than parallel orientations.

Figure 2a, shows a 3D object with multiple orthogonally oriented RFID tags. Figure 2b shows orthogonal planar (i.e. horizontal and vertical) orientations of two tags. In Section 5, our experiments indicate that horizontal planar orientations increase the tag's sensitivity. Thus, utilising multiple tags in orthogonal spatial and horizontal planar orientations tends to improve the overall localisation accuracy.

Figure 2 Tag orientations



4.5 Reader locality

Theoretically, the usable power in the radio waves emitted by the reader attenuates inversely proportional to the cube (for near-field) and square (for far-field) of the distance (as given by the Friis transmission equation). This determines the operating/detection region for the tags with respect to the readers. Thus, the reader's location and proximity to a tag impacts the tag's localisation accuracy. We propose that more tags should be placed in regions likely to be nearer to the objects being localised in order to improve the overall localisation accuracy.

The main guiding principle behind all the above mitigating techniques is to *identify and minimise possible errors at the sources where they arise*. This leads to efficient localisation techniques, fewer onboard components, higher localisation accuracy and speed. In the following section, we use this principle with the proposed object localisation framework to improve the localisation accuracy and speed.

5 Object localisation framework

The proposed localisation approach utilises two different techniques. In the first technique, an onboard reader and reference tags embedded in the environment are used to coarsely localise the mobile object. The second technique varies the power levels of environment-embedded readers to localise the onboard tag via the empirical power-distance relationship (calibrated using reference tags at known positions). To ensure uniform behaviour from the tags, we test, sort, and select them based on their (similar) detection sensitivity. Also, by employing multi-tags (Bolotnyy and Robins, 2007a; Bolotnyy and Robins, 2007b; Bolotnyy and Robins, 2009), we reduce the uncertainties when inferring tag positions. Finally, we combine these localisation techniques and propose several heuristics for significantly improving the localisation accuracy.

While tags are sorted, placed, and calibrated as part of an offline pre-processing phase, the actual localisation and error minimisation heuristics are performed in real time. The calibration process may be repeated occasionally, in order to adjust the system to varying environmental conditions. Re-calibration may also be performed in parallel with actual localisation operations to accommodate ‘drifts’ in the empirical power-distance relationship. Below we describe key aspects of the proposed localisation approach.

5.1 Tag sensitivity analysis

Tag manufacturing variability can dramatically affect the detection sensitivity of tags (i.e. the minimum reader power level needed to successfully read a tag at a given location). In fact, a small fraction of any commercially obtained batch of tags are typically even ‘dead’ (i.e. non-functional) altogether. While the localisation speed will increase with higher tag sensitivities, the accuracy of the proposed localisation framework depends on the uniform detection sensitivities of the tags. Thus, an offline pre-processing quality-control check provides a characterisation of the sensitivities to ensure that only tags with uniform (and reasonably high) sensitivities contribute to our subsequent localisation experiments.

Our experimental evaluation showed that tag sensitivity varied considerably across a group of 243 tags of the same type. We have characterised the tag sensitivities based on the read counts under different reader power levels and distance combinations. Thus, given a fixed reader power level, if a tag has low read counts among its peers, we call it ‘low sensitive’. Similarly, tags with high read counts are labelled as ‘highly sensitive’, and tags having equal read count are called ‘average sensitive’. We performed two experiments to quantify single tag sensitivities by varying the power levels and distances between the readers and the tags. While these experiments use EPC Gen2 passive tags, this *tag binning* approach is equally applicable to other types of RFID tags.

5.1.1 Single tag calibration

In this experiment, a batch of four tags was placed at a distance of 2.54 metres from the reader. We varied the reader power level from 25.6 dBm to 31.6 dBm, in steps of 3 dBm. We recorded the cumulative read counts of each tag for 60 seconds (i.e. 3 read iterations lasting 20 seconds each). We found that 114 out of 243 tags had cumulative read counts of zero at 25.6 dBm, while remaining tags had read counts in the range of 3 to 9 (some

tags had read counts as high as 10). Moreover, at a reader power level of 28.6 dBm, most of the tags had cumulative read counts in the range 6 to 11, and the cumulative read counts ranged between 5 and 11 at 31.6 dBm.

We labelled tags as ‘low sensitive’ only if they had zero cumulative read counts at a power level of 25.6 dBm. Also, tags were labelled as ‘low sensitive’ at 28.6 dBm only if they were also labelled as ‘low sensitive’ at 25.6 dBm. Similarly, we labelled tags as ‘highly sensitive’ at 25.6 dBm only if they were also labelled as ‘highly sensitive’ at 31.6 dBm. While the combination of power levels and distance ranges was comparatively small, variations in tag sensitivities were evident even at this scale. Using this process, 89 out of 243 tags were classified as highly sensitive, 133 tags ranked as average sensitive, and the remaining tags were considered to be low sensitive (and some tags were dead altogether). Thus, this experiment classified 243 tags into three sensitivity categories based on the reader power levels required to detect them.

Similarly, in another set of experiments, we kept the reader power level constant and varied the distance between the tags and the reader with the same increments as above. This process classified 61 out of the 243 tags as low sensitive, 161 tags as average sensitive, and 21 tags as highly sensitive. We then combined the outcomes of these two sensitivity experiments by taking the intersection of the ‘average sensitive’ tag sets from each experiment, thus classifying 133 tags as overall ‘average sensitive’.

In the ensuing localisation experiments, we selected all the reference and target tags from this overall average sensitive tag set. We also constructed four-way multi-tag platforms, made from the average sensitive tags that provide higher operational reliability. See the work (Bolotnyy and Robins, 2007a; Bolotnyy and Robins, 2007b; Bolotnyy and Robins, 2009) for a more detailed analysis of multi-tags and their advantages over ordinary tags.

Figure 3 illustrates the design of our four-way multi-tag platform consisting of four ‘Impinj Dogbone Monza 3’ UHF passive tags having overall average sensitivity, mounted on a vertical stand made of Lego bricks (our choice of Lego components is based on the versatility of Lego bricks as well as the transparency of their plastic material to radio waves). We built 33 such multi-tag platforms, and each tag on these platforms was calibrated separately according to the power-distance relationship discussed above (Chawla et al., 2010a; Chawla et al., 2010b). Following our usual ‘abundance of caution’ philosophy, we performed two extensive platform calibration experiments to ensure uniform detection sensitivity for the entire multi-tag platform across such variables as tag rotation and reader proximity. These experiments are described in more detail below.

Figure 3 A four-way multi-tag platform constructed using Lego Bricks



5.1.2 *Multi-tag calibration*

In this calibration experiment, we ensured that the four-way multi-tag platforms consisting of four proximate equally sensitive tags all have similar sensitivities. This goal was achieved by determining the average read count of constituent tags having matching orientations with respect to the reader's antennas. Thus, tags at position one, two, three, and four were oriented towards antenna one, two, three, and four, respectively. We kept the reader power level constant at 31.6 dBm and varied the distance between the reader and the multi-tag platform within the range of 1.27 to 3.81 metres. We also varied the reader power level within the range of 25.6 to 31.6 dBm, keeping the distance between them constant at 2.54 metres. We repeated this calibration experiment three times and computed the average.

We also analysed the sensitivity of multi-tags with respect to rotation. Using similar power and distance combinations as above, we measured the average read counts of each multi-tag in each of the four possible 90-degree rotations. These experiments showed that aside from minor variations in average read counts (possibly due to the long oval shape of the radio wave lobe emitted by the reader's antenna), all multi-tags performed uniformly with respect to distance, power level, and rotation.

In conclusion, these experiments show that combining individually calibrated average-sensitive tags results in equally sensitive multi-tags. On a philosophical note, while this outcome is happily the expected one, we still followed sound 'scientific method' principles and carefully verified this hypothesis experimentally. This kind of careful, methodical controlled experimentation is not wasted effort, since it sometimes can uncover startling and unexpected facts, such as the huge variance in sensitivity of ostensibly 'identical' commercial tags discussed above.

5.2 *Localisation algorithms*

We now describe our proposed localisation algorithms for estimating positions of stationary and mobile objects.

5.2.1 *Tag localisation algorithm I: Linear Search*

The first localisation algorithm linearly increments the reader power level from lowest to highest in order to determine the minimum power level required to detect a tag. While this approach finds the minimum tag detection power level, it may take some time to converge due to the linear step-wise power incrementing approach.

Alternatively, we can instead vary the power level from highest to lowest in order to detect tags, since tags are typically not located very near the readers, but rather farther away from them. Thus, stepping down the power level (i.e. from highest to lowest) will tend to lower the average number of iterations required to determine the minimum tag detection power level. We call this algorithm 'Linear Search' and outline its pseudo-code in Figure 4.

Figure 4 Tag localisation algorithm I: Linear Search

```

Input:
Unique tag ID TagID
Power varying step PowerStep
Power varying direction PowerDirection

Output:
Minimum power level Power, MinPower required to detect tag TagID

```

```

if (PowerDirection = LowToHigh) then
  Power = MinPowerLevel
  repeat
    if (Power > MaxPowerLevel) then
      return NotFound
    end
    set reader power level to Power
    search for tags until successful or time-out
    if (TagID is found) then
      return Power
    end
    Power = Power + PowerStep
  end
else
  Power = MaxPowerLevel
  MinPower = MaxPowerLevel
  repeat
    if (Power < 0) then
      return NotFound
    end
    set reader power level to Power
    search for tags until successful or time-out
    if (TagID is found) then
      MinPower = Power
    else
      return MinPower
    end
    Power = Power - PowerStep
  end
end

```

We control the power level step using the input parameter ‘PowerStep’. The direction of reader power level increment (e.g. from lowest to highest) is controlled using the input parameter ‘PowerDirection’. At each power level, a single tag having the unique tag ID (specified by the input parameter ‘TagID’) is searched until either it is found, or else a search timeout occurs. Thus, given a set of tags to be found this Linear Search algorithm operates in a serial manner, to determine the minimum tag detection power levels for each given tag. This algorithm provides highly accurate tag detection power level, at the possible expense of longer overall running times (due to the additional number of RFID read operations that may be required).

5.2.2 Tag localisation algorithm II: Binary Search

In the second localisation algorithm, we start at a mid-value power level, and then step-up or step-down the power, based on the reader's ability to detect any tags at that stage. This binary search-based approach will tend to converge faster on the minimum power levels required to detect tags. We control the direction of the power level increment by using the variable 'TagFound'. Figure 5 depicts this proposed Binary Search algorithm.

Figure 5 Tag localisation algorithm II: Binary Search

```

Input:
Unique tag ID TagID
Power varying step PowerStep

Output:
Minimum tag detection power level TagPower

```

```

TagPower = MaxPowerLevel + 1
TagFound = False
repeat
  Power = (MinPowerLevel + MaxPowerLevel) / 2
  set reader power level to Power
  search for tags until successful or time-out
  if (TagID is found) then
    TagFound = True
    if (TagPower > Power) then
      TagPower = Power
    end
  end
  if (TagFound) then
    MaxPowerLevel = Power
  else
    MinPowerLevel = Power
  end
  if ((MaxPowerLevel - MinPowerLevel) ≤ PowerStep) then
    return TagPower
  end
end

```

When the minimum and maximum power levels differ by only the power step size, we return the tag detection power levels (using the variable 'TagPower'). While the Binary Search algorithm searches for tags in an exponentially converging manner (as opposed to linearly), it requires less time to terminate than the Linear Search algorithm. However, the power level resolution for Binary Search algorithm is sometimes less accurate than that of the Linear Search approach, which uses a finer-granularity power step size increment. Thus, the Binary Search algorithm trades off run time against accuracy in the minimum tag detection power level.

5.2.3 Tag localisation algorithm III: Parallel Search

The Linear Search and the Binary Search algorithms search for only a single tag per execution cycle (i.e. they operate in a serial manner). Our third algorithm addresses this limitation by determining the minimum tag detection power levels in parallel, for all the

tags in the reader's vicinity. This 'Parallel Search' strategy is analogous to running the Linear Search algorithm in parallel for all the tags. Figure 6 illustrates this Parallel Search algorithm, which takes as input (through the parameter 'TagIDSet') a set of tags to be found and returns the minimum tag detection power levels (via the variable 'PowerSet'). A key step in this algorithm is determining which tags have reached their optimal minimum detection power level by checking the power levels of tags whose detectabilities have not changed for a given number of iterations.

Figure 6 Tag localisation algorithm III: Parallel Search

Input:
Set of unique tags *TagIDSet*
Power varying step *PowerStep*

Output:
Set of minimum tag detection power levels *PowerSet*

```

set all tags in TagIDSet to MaxPowerLevel
Power = MaxPowerLevel
repeat till Power  $\geq$  0
| set reader power level to Power
| search for tags until successful or time-out
| update power levels in PowerSet for all detected tags in TagIDSet
| check for tags in TagIDSet with no power level update in PowerSet
| Power = Power – PowerStep
end
return PowerSet

```

Note that the underlying RFID equipment and protocols are already designed to detect/read multiple tags in a single read phase (including collision resolution). We thus rely on the embedded RFID hardware functionality to achieve the required parallelism (i.e. the simultaneous detection of many tags in a single reader operation). The proposed Parallel Search algorithm sweeps from the highest to lowest power levels, since more tags tend to be farther away from a given reader than closer to it. Since Parallel Search can determine the minimum power levels of many tags in parallel, it enables the simultaneous localisation of multiple stationary and mobile objects, resulting in faster overall run times.

5.2.4 Reader localisation algorithm: Measure and Report

While the previous algorithms localise tags attached to the objects, our fourth algorithm uses readers located onboard the objects themselves, in order to estimate their positions. In particular, we detect and record the unique tag IDs encountered by an object along a motion path. We associate a timestamp with each such measurement, resulting in a list of tuples of the form $\langle \text{TagID}, \text{Timestamp} \rangle$. Thus, we can determine the trajectory of the objects by knowing their proximity to the locations of reference tags, along with the measurement times. We call this algorithm 'Measure and Report', as described in Figure 7.

Figure 7 Reader localisation algorithm: Measure and Report

```

Input:
Unique tag ID TagID

Output:
Two-tuple having unique tag ID TagID and current date-time TimeStamp

```

```

if (TagID is found) then
  | return (TagID, TimeStamp)
else
  | return NotFound
end

```

Table 1 describes the runtime complexity of each algorithm. The Measure and Report algorithm is the fastest algorithm, while the Linear Search and Binary Search algorithms take considerably more time due to their serial manner of operation. The Parallel Search algorithm takes less time than either the Linear Search or Binary Search algorithm, as its run time is independent of the number of tags and only depends on the number of power levels used during the searching.

Table 1 Time complexity of localisation algorithms

<i>Localisation technique</i>	<i>Localisation algorithm</i>	<i>Time complexity</i>
Reader localisation	Measure and Report	O(1)
Tag localisation	Linear Search	O(N·P)
	Binary Search	O(N·LogP)
	Parallel Search	O(P)

Notes: N = Number of tags, P = Number of reader power levels used.

Moreover, the algorithms that require more time tend to generate higher resolution minimum tag detection power levels, whereas the faster algorithms trade off localisation accuracy for speed. We experiment with different combinations of these algorithms to observe the tradeoffs between localisation accuracy and running speed. Furthermore, localisation errors can occur due to (a) the onboard reader operating range, (b) matching the onboard target tags with the nearest reference tags, and (c) the inherent power level estimation errors of the algorithms (e.g. due to varying environmental conditions). We next discuss these errors along with mitigating techniques.

5.3 Error heuristics

Localisation errors occur in the first technique (i.e. onboard reader localisation) due to limitations in the read-range of the onboard reader. Since mobile objects can move arbitrarily, an inexpensive and reliable way to reduce this type of error is by placing more densely/regularly arranged reference tags throughout the expected regions of motion. In the second technique (i.e. onboard tag localisation), errors in the position estimates can occur by identifying the onboard multi-tag with the nearest neighbourhood reference tags.


Figure 8 A source of localisation errors

Figure 8 depicts four antennas emitting radio waves, conceptually forming an intersection region where the onboard target tags may be found. In the process of matching the behaviour of an unknown tag to those of known reference tags, potential ambiguities and ‘round off’ errors may occur. In order to minimise possible localisation errors, we have developed 11 error-reducing heuristics, as described in following sub-sections.

5.3.1 Error heuristic I: absolute difference

This localisation error-reducing heuristic computes the absolute difference of the reader power levels between the neighbouring tags and the onboard tag (all of which are multi-tags in our experiments). We propose four such possible heuristic variants, as follows:

$$H_1 : \text{Min}_{\forall J} \left(\sum_{I=1}^M \Delta_I (R_J) \right) \quad (2)$$

$$H_2 : \text{Min}_{\forall J, K; J \neq K} \left(\sum_{I=1}^M \Delta_I (R_J) + \sum_{I=1}^M \Delta_I (R_K) \right) \quad (3)$$

$$H_3 : \text{Min}_{\forall J, K \in B; J \neq K} \left(\sum_{I=1}^M \Delta_I (R_J) + \sum_{I=1}^M \Delta_I (R_K) \right) \quad (4)$$

$$H_4 : \text{Min}_{\forall J, K \in B; J \neq K} \left(\sum_{I=1}^M \Delta_I (R_J) + \sum_{I=1}^M \Delta_I (R_K) \right) \text{ s.t. } \left(\sum_{I=1}^M \Delta_I (R_J) < \sum_{I=1}^M \Delta_I (R_K) \right) \quad (5)$$

5.3.2 Error heuristic II: minimum power reader selection

This error-reducing heuristic computes the absolute difference of the power levels between the neighbouring reference tags and the onboard tag using the minimum power levels of two orthogonally positioned readers. We propose two such possible heuristic variants, as follows:

$$H_5 : \text{Min}_{\forall J, K; J \neq K} (\Delta_J (T) + \Delta_K (T)) \quad (6)$$

$$H_6 : \text{Min}_{\forall J, K \in B; J \neq K} (\Delta_J (T) + \Delta_K (T)) \quad (7)$$

5.3.3 Error heuristic III: root sum square absolute difference

This error-reducing heuristic computes the square root of the sum of the squares of the absolute difference (in reader power levels) between the neighbouring reference tags and the onboard tag. We propose four such possible heuristic variants, as follows:

$$H_7 : \text{Min}_{\forall J} \left(\sqrt{\sum_{I=1}^M \Delta_I (R_J)^2} \right) \quad (8)$$

$$H_8 : \text{Min}_{\forall J, K; J \neq K} \left(\sqrt{\sum_{I=1}^M \Delta_I (R_J)^2} + \sqrt{\sum_{I=1}^M \Delta_I (R_K)^2} \right) \quad (9)$$

$$H_9 : \text{Min}_{\forall J, K \in B; J \neq K} \left(\sqrt{\sum_{I=1}^M \Delta_I (R_J)^2} + \sqrt{\sum_{I=1}^M \Delta_I (R_K)^2} \right) \quad (10)$$

$$H_{10} : \text{Min}_{\forall J, K \in B; J \neq K} \left(\sqrt{\sum_{I=1}^M \Delta_I (R_J)^2} + \sqrt{\sum_{I=1}^M \Delta_I (R_K)^2} \right) \text{s.t.} \sqrt{\sum_{I=1}^M \Delta_I (R_J)^2} < \sqrt{\sum_{I=1}^M \Delta_I (R_K)^2} \quad (11)$$

5.3.4 Error heuristic IV: all heuristics minimum

This error-reducing meta-heuristic computes the minimum over all the power levels obtained using the above heuristics. Thus, this ‘meta-heuristic’ yield the least tag detection power levels, denoted as follows:

$$H_{11} : \text{Min}_{\forall L} (H_L) \quad (12)$$

The variables in all the expressions above have the following denotations:

T = Target tag

R_J = Reference tag J

H = Heuristic

Power = Minimum detection power level for a tag

M = Number of readers

B = Set of neighbours

$\Delta_I(R) = |\text{Power}(T) - \text{Power}(R)|$ using reader I

J, K = Iteration variables for neighbourhood tags

L = Iteration variable for heuristics

The above error-reducing heuristics contribute to an online post-processing step in our object localisation framework, once the minimum detection power levels of the reference and onboard tags have been determined. Employing different combinations of localisation algorithms and error-reducing heuristics can achieve a desired level of accuracy. Another key feature of the proposed framework is the flexibility to incorporate new localisation algorithms and error-reducing heuristics that may be developed in the future, enabling this framework to localise objects with even higher accuracy and speed.


6 Experimental evaluation

This section presents the implementation details, experimental evaluation methodology, localisation accuracy data, and run times. We also compare the proposed localisation approach with existing localisation techniques.

6.1 Experimental set-up

We evaluated the proposed object localisation framework in an indoor environment using one reader connected to four antennas, 33 reference multi-tags, one onboard reader, and one four way multi-tag platform per mobile robot that acted as either a stationary or a mobile object. We constructed two mobile robots using Lego Mindstorms kits, with the onboard controller being an HP iPAQ hx2490 Pocket PC, connected to an iDtronic Voltaire portable RFID reader, as shown in Figure 9.

Figure 9 Experiment components: mobile robot platform and track



The Lego railroad track is used to conveniently move the mobile robots along predetermined paths. Also shown are the four-way multi-tag platforms used as reference tags. To localise any object/tag, we vary the stationary reader power levels to localise the onboard multi-tag, using in turn all of the localisation algorithms described above. Alternatively, we also localised objects with the onboard reader to read the reference tags encountered during motion and transmit the tag IDs to a backend workstation with the onboard bluetooth link. A key aspect of the proposed localisation approach is to limit the use of onboard non-RFID components, while still obtaining good localisation accuracy and speed.

Table 2 describes our experimental set-up and implementation details. Note that the mobile robots utilising only RFID technology for localisation have somewhat bounded speed, due to the computational delays inherent in determining minimum tag detection power levels, as well as the reader's operational speed (Han et al., 2007). Thus, while mobile robots may move several metres per second while using onboard RFID readers

for localisation, more precise localisation requires slower speeds. Such locomotion speed limitations are also present in other existing mobile object localisation techniques. We believe that as RFID technology keep advancing, faster readers, identification protocols, combined with more sophisticated tags can help achieve localisation for mobile objects moving at higher speeds.

Table 2 Experimental set-up

<i>Type</i>	<i>Component</i>	<i>Details</i>
Backend Workstation	CPU	AMD Athlon 64 @ 2GHz
	RAM	1 GBytes
	Hard Disc	100 GBytes
	OS	Windows XP
	PL	C++/C#
	API	M4 LIB
	RFID Equipment	Reader
	Antenna	Linear with 6 dBi gain
	Tag	Impinj Dogbone Monza 3 93 × 23 mm
	Protocol	EPC Gen2
Environment	Map Area	8 metres
	Room Volume	41 cubic metres
	Reference Tags	132 single tags (33 Multi-tags)
Robots	Kit	Lego Mindstorms
Onboard Control	PDA	HP iPAQ hx2490
Onboard Link	Wireless	Bluetooth dongle

6.2 *Experimental results*


We divide the proposed localisation approach into the set-up phase and the localisation phase. In the set-up phase, we instrument a 2D region with 33 reference multi-tags and determine their empirical power-distance relationships. In the localisation phase, we estimate the positions of target objects by comparing the prior-stored empirical power-distance relationship of the reference tags with real-time measurements obtained from the onboard target tags.

6.2.1 *The empirical power-distance relationship*

In Section 3, we discussed the ideal power-distance relationship based on the Friis equation. We now describe our methodology for empirically determining the power-distance relationship, while accommodating various error sources, such as multi-path scattering, environmental interferences, RF occlusions, etc.

Figure 10 illustrates the power-distance relationship between the uniformly sensitive reference tags and the readers. We utilised the Linear Search algorithm to calibrate the empirical power-distance relationship, since it provides highly accurate minimum tag detection power levels. Note that as the distances between the tags and readers increase, more reader power is required to detect the tags. Such a relationship also implicitly determines the read-range of the reader (e.g. in illustration given in Figure 10, the operating region is 4 metres long and 2 metres wide).

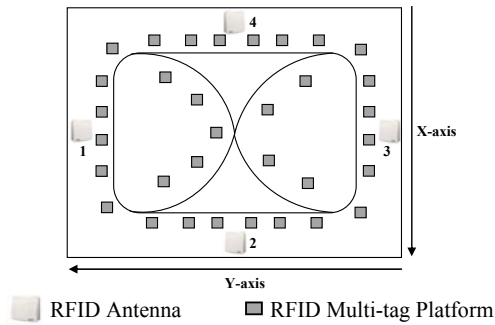
Figure 10 The empirically observed power-distance relationship



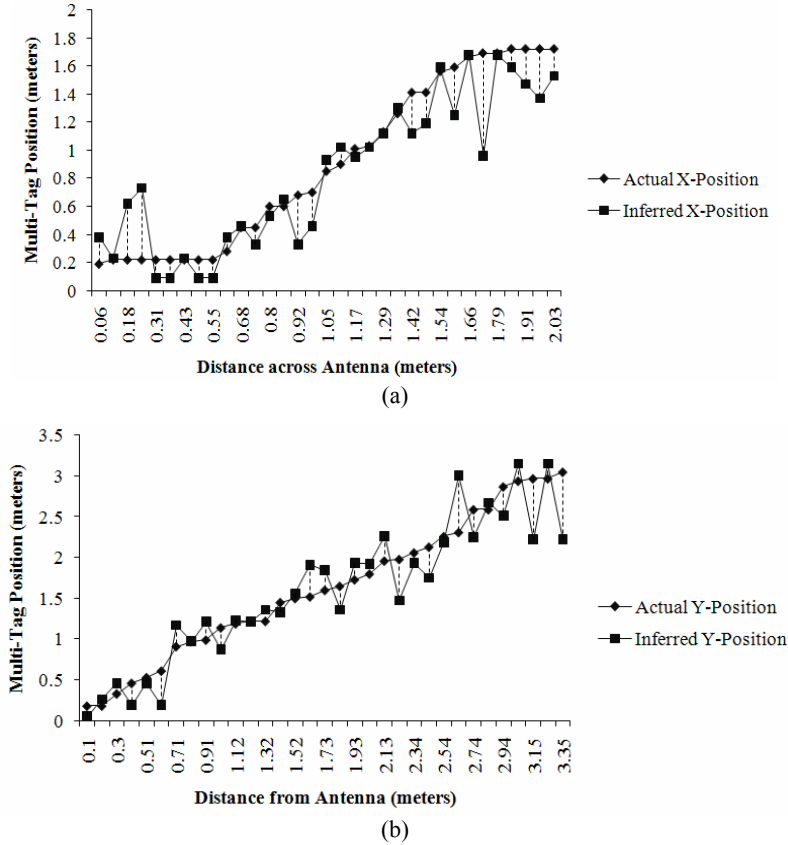
6.2.2 Localisation accuracy

We now describe our experimental results pertaining to localisation. We utilised the proposed algorithms in different combinations during the set-up phase and the localisation phase, in order to determine which algorithm combinations yield the smallest localisation errors. In our experiments, reference tags are placed throughout the 2D experimental region at positions specified by their distances from antenna four (X-axis) and from antenna three (Y-axis), as illustrated in Figure 11.

Figure 11 A top-view of experimental set-up



The localisation accuracy along X and Y axis is depicted in Figures 12a and 12b. The proposed localisation approach closely approximates (along both axes) the actual positions of target objects. We used the Linear Search algorithm in the set-up phase and the Binary Search algorithm in the localisation phase in order to yield the desired localisation accuracy. Note that the Parallel Search algorithm and the Measure and Report algorithm achieve comparable accuracy with higher localisation speed. Importantly, the Linear Search algorithm in the low-to-high mode (i.e. where the power levels are incremented from lowest to highest) provides more accurate minimum tag detection power levels, which improves the localisation accuracy (see Figure 13a).

Figure 12 Localisation accuracy: (a) along X-axis, and (b) along Y-axis

6.2.3 Localisation speed


Localisation speed is heavily dependent on the time required to estimate minimum tag detection power levels. While the Linear Search algorithm (in high-to-low mode), the Binary Search algorithm, and Parallel Search algorithm all provide less accurate estimates of the minimum tag detection power levels, they also run reasonably quickly, as illustrated in Figure 13b. On the other hand, the Linear Search algorithm in low-to-high mode requires more time to run than the other algorithms, but is more accurate. This data confirms our hypothesis that varying the power levels from high to low is typically more efficient for localising tags farther away from the reader (i.e. where one would typically expect to find them). Thus, by combining the different algorithms and their variants, appropriate application-driven trade-offs can be chosen between localisation accuracy and speed.

6.2.4 Power step vs. localisation accuracy

In Section 5, we discussed the parameter ‘PowerStep’ which specifies the power level increment. We measured the impact of this parameter on the minimum tag detection power levels. This is accomplished by determining for a given target tag, the minimum

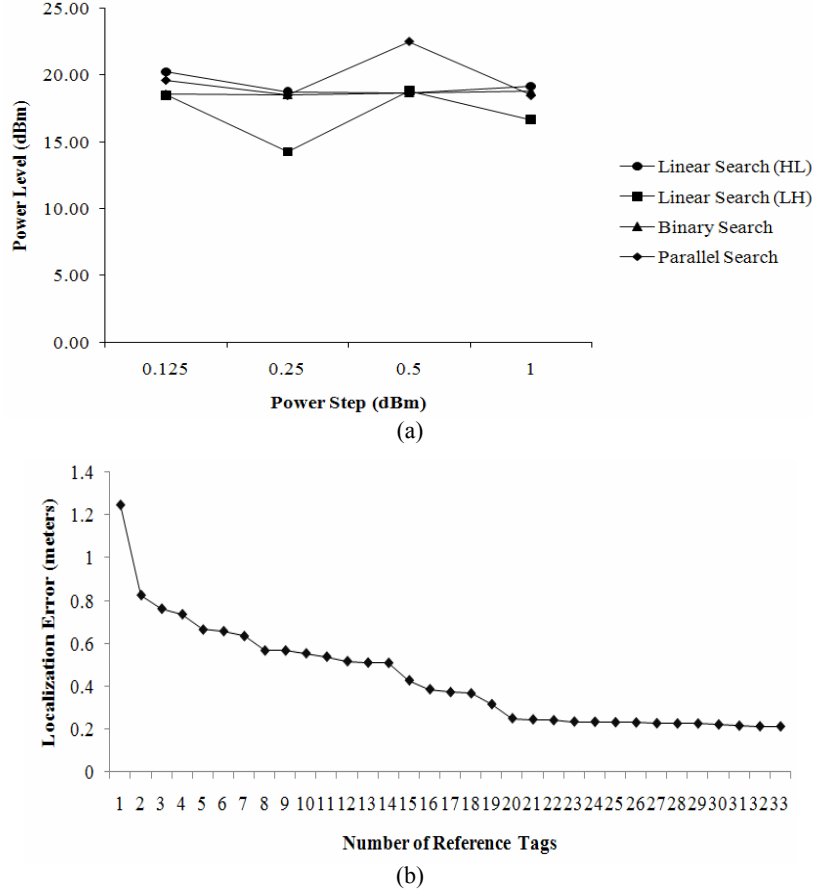
tag detection power levels using different power steps. Figure 14a illustrates the minimum detection power levels of a tag for four different power steps, measured using the three localisation algorithms and two orthogonally placed antennas. Linear Search algorithm in low-to-high mode reports the lowest minimum tag detection power level, while Parallel Search yields the highest minimum tag detection power level for the same tag. Since localisation accuracy is based on determining the minimum tag detection power levels (i.e. if an algorithm reports a lower minimum tag detection power level than other comparable algorithms then that algorithm yields higher accuracy), these algorithms are able to trade off accuracy and speed.

Figure 13 Algorithmic variability: (a) Minimum tag detection power levels, and (b) Localisation time



6.2.5 Reference tag density vs. localisation accuracy

We hypothesised that increasing the number of reference tags increases the localisation accuracy only up to a certain point. To test this hypothesis, we varied the number of reference tags from 1 to 33 and determined that the range of localisation error varied from 1.2 to 0.2 metres, as depicted in Figure 14b. Thus, adding inexpensive passive reference tags increases the localisation accuracy only up to a point. For example, with only 22 reference tags, 88% of the maximum possible localisation accuracy was achieved (Chawla et al., 2010a; Chawla et al., 2010b).

Figure 14 The impact on localisation accuracy: (a) Power step size, and (b) Number of reference tags

6.3 Comparative analysis and visualisation

We now compare the proposed localisation approach to existing stationary and mobile object localisation techniques. For each proposed technique we measured separately the time required by the set-up phase and the localisation phase (Table 3). Curiously, most previous localisation techniques do not report either of these times (i.e. the set-up time nor the localisation time). These omissions bring into question the practicality of previous approaches, since excessive run times can severely limit the utility of a localisation method. Some previous localisation techniques also heavily utilise non-RFID based components (e.g. cameras, ultrasonic sensors, lasers, etc.). Thus, the accuracy of such methods comes at an onerous price of resorting to cumbersome, ad-hoc combinations, and expensive technology/equipment, tending to further diminish the practicality of such methods.

Table 3 also reports the size of the test area, and the average localisation error. Note that the maximum time required by any of the proposed localisation approaches during the set-up phase is due to the Linear Search algorithm operating in the low-to-high mode. However, it should be noted that the set-up phase is an offline pre-processing step that is

executed only infrequently. The proposed Parallel Search localisation approach can simultaneously localise multiple stationary and mobile objects in 1.67 minutes with an overall average accuracy of 0.35 metres. Moreover, combining the different localisation algorithms in the set-up and the localisation phases can achieve a combined overall average accuracy of 0.18 metres. We believe that additional fine-tuning of our methodology can further improve these results.

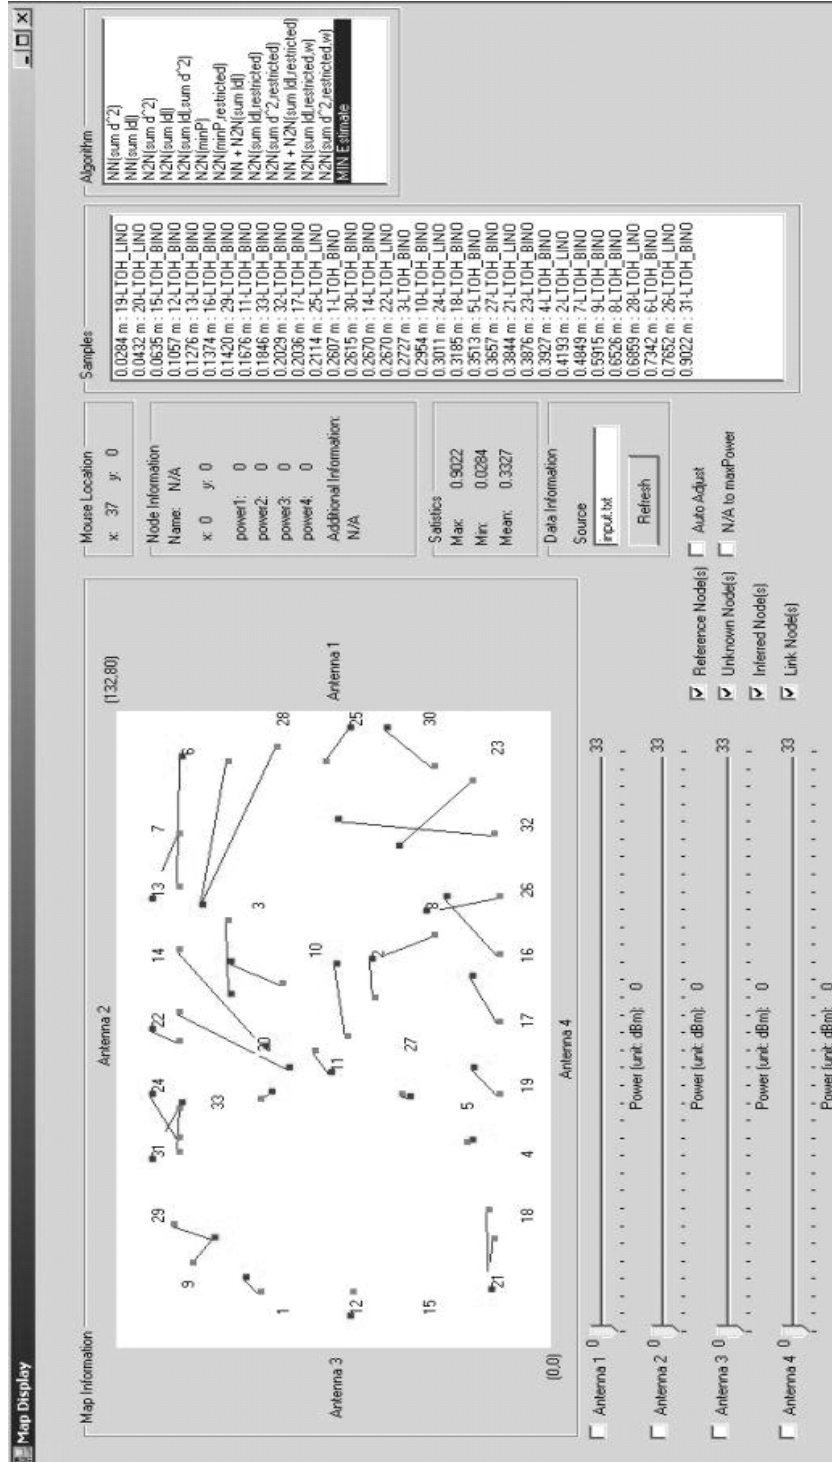
Table 3 Comparison of proposed RFID-based object localisation techniques with existing methods

<i>Technique</i>	<i>Average Time (min)</i>		<i>Test area (m²)</i>	<i>Localisation error (m)</i>
	<i>Set-up phase</i>	<i>Localisation phase</i>		
Ni et al., 2003 ^a	NR	NR	NR	2
Alippi et al., 2006 ^a	NR	NR	20	0.68
Bekkali et al., 2007 ^a	NR	NR	9	0.5–1.0
Joho et al., 2009 ^a	27	NR	NR	0.375
Zhao et al., 2007 ^a	NR	NR	20	0.14–0.29
Wang et al., 2007 ^a	NR	NR	NR	0.1–0.9
Choi et al., 2009 ^a	NR	NR	NR	0.21
Jin et al., 2006 ^a	NR	NR	40	0.72
Zhang et al., 2007 ^a	NR	NR	NR	1
Chae and Han, 2005 ^b	NR	NR	48.4	0.23
Choi and Lee, 2009 ^b	NR	NR	14.4	0.016–0.024
Hähnel et al., 2004 ^b	NR	NR	784	1–10
Han et al., 2007 ^b	NR	NR	1	0.09
Milella et al., 2009 ^b	NR	NR	70	0.64
Senta et al., 2007 ^b	NR	NR	2	0.2
Seo and Lee, 2008 ^b	NR	NR	5	0.2–1.6
Vorst et al., 2008 ^b	NR	NR	125	0.2–0.6
<i>Linear Search (HL)</i>	<i>29.78</i>	<i>1.42</i>	<i>8</i>	<i>0.29</i>
<i>Linear Search (LH)</i>	<i>161.23</i>	<i>5.28</i>	<i>8</i>	<i>0.27</i>
<i>Binary Search</i>	<i>47.24</i>	<i>1.95</i>	<i>8</i>	<i>0.31</i>
<i>Parallel Search</i>	<i>1.67</i>	<i>1.67</i>	<i>8</i>	<i>0.35</i>
<i>Measure and Report</i>	<i>0</i>	<i>0</i>	<i>8</i>	<i>0.25</i>
<i>Combined Approach</i>	<i>161.23</i>	<i>10.32</i>	<i>8</i>	<i>0.18</i>

Notes: NR = Not Reported, ^aStationary Localisation Technique, ^bMobile Localisation Technique, LH = Low to High, HL = High to Low.

We also developed a visualisation tool that provides a real-time view of the object localisation environment. The visualisation user interface contains five key regions, which enable the end-users to select different error heuristics, view localisation error per location (both graphically and analytically), and control experimental variables (e.g. antenna power, reference nodes, etc.). Using this tool, practitioners can determine the localisation accuracy at various points in the environment. Regions with low localisation accuracy can be enhanced with additional low-cost passive tags. Thus, the visualisation tool (see Figure 15) also serves as a feedback mechanism that can help enhance the localisation accuracy, as well as controlling the experimental testbed.

Figure 15 The visualisation user interface



7 Key lessons learned

This section outlines some key insights and lessons learned during the conceptualisation, design, and implementation of the proposed object localisation framework.

7.1 Calibration of tags

Typically, RFID-based localisation techniques use off-the-shelf tags as-is in their experiments. We have discovered that even in a brand-new batch of tags, tag sensitivities varied considerably, although previous works are surprisingly oblivious to this important issue. Identifying and systematically mitigating this source of localisation errors enabled us to minimise its impact on localisation accuracy.

7.2 Empirical power-distance relationship

While the ideal power-distance relationship can be described using, e.g., the Friis transmission equation, factors such as multi-path scattering, environmental interferences, RF occlusions and variability in tag sensitivities can substantially skew this relationship. Nevertheless, the power-distance relationship can be reliably characterised by direct empirical calibration measurements, rendering moot many of the unpredictable and hard-to-characterise variabilities.

7.3 Simplicity in algorithmic design

We emphasised simplicity (i.e. ‘Occam’s razor’) in designing the localisation algorithms and error heuristics yielding effective and efficient techniques. Our streamlined methods can be implemented with relatively little overhead, using widely available commercial RFID equipment.

7.4 Tag orientation and placement strategies

Tag orientation and placement are key factors that impact effective detection. Moreover, multi-tags have superior detection sensitivities over ordinary single tags. We experimentally determined the preferable (multi-)tag orientations and placement strategies (e.g. orthogonal orientations yield higher detection probabilities).

7.5 Minimise onboard components

While designing the RFID testbed mobile platforms, we focused on simplifying the onboard components and avoiding non-RFID components, in order to reduce the overall complexity. The resulting object localisation framework is efficient, robust, scalable, and economically viable.

7.6 Visualisation and control tools

We designed a graphical user interface to visualise the operation, performance, and localisation accuracy of the experimental environment. This GUI enabled the quick identification of various issues, including the identification (and enhancement) of low-accuracy regions, the discovery of improved tag placement strategies, etc.

7.7 Application scenarios

The proposed object localisation framework was designed to operate in variety of scenarios, such as the 3D positioning of objects, the simultaneous localisation of multiple stationary and mobile objects, localisation in outdoor settings, etc. The design principles in our framework aim towards practicality, relevance, and generality.

8 Conclusions and future work

We proposed an accurate, scalable, reliable, streamlined RFID-based approach for localising objects. We have outlined several localisation challenges and developed practical techniques to mitigate them. We deployed uniformly sensitive four-way multi-tags in order to improve the object localisation accuracy and speed. Furthermore, we analysed the impact of the reference tag density on the localisation accuracy, and compared the proposed localisation approach to existing methods. While these experiments were performed in an indoor setting, the proposed localisation approach is general enough to be widely applicable to other scenarios (e.g. 3D regions, outdoor environments, multiple stationary, mobile object localisation, etc.). Future research can also strive for further improvements in the overall localisation speed, especially for moving objects.

Acknowledgements

We thank the anonymous reviewers for their helpful feedback. This research was supported by the U.S. National Science Foundation through grant number CNS-0716635 (PI: Professor Gabriel Robins). We thank Liuyi Zhang for his help with the experiments, and Michael Skalak for his help with proofreading.

References

- Abowd, D. and Mynatt, E.D. (2000) ‘Charting past, present, and future in ubiquitous computing’, *ACM Transactions on Computer-Human Interaction*, Vol. 7, No. 1, pp.29–58.
- Alippi, C., Cogliati, D. and Vanini, G. (2006) ‘A statistical approach to localize passive RFIDs’, *IEEE International Symposium on Circuits and Systems (ISCAS 2006)*, Island of Kos, Greece, pp.843–846.
- Bekkali, A., Sanson, H. and Matsumoto, M. (2007) ‘RFID indoor positioning based on probabilistic RFID map and Kalman filtering’, *IEEE International Conference on Wireless and Mobile Computing, Networking and Communications (WIMOB 2007)*, New York, pp.21–21.
- Blewitt, G., Kreemer, C., Hammond, W.C., Plag, H., Stein, S. and Okal, E. (2006) ‘Rapid determination of earthquake magnitude using GPS for tsunami warning systems’, *Geophysical Research Letters*, Vol. 33, No. L11309, 4pp.
- Bolotnyy, L. and Robins, G. (2007a) ‘The case for multi-tag RFID systems’, *IEEE International Conference on Wireless Algorithms, Systems and Applications (WASA 2007)*, Chicago, pp.174–186.
- Bolotnyy, L. and Robins, G. (2007b) ‘Multi-tag RFID systems’, *International Journal of Internet Protocol Technology*, Vol. 2, Nos. 3/4, pp.218–231.

- Bolotnyy, L. and Robins, G. (2009) 'Multi-tag RFID systems', in Zhang, Y. and Kitsos, P. (Eds): *Security in RFID and Sensor Networks*, Taylor and Francis Group, Chapter 1, pp.3–28.
- Chae, H. and Han, K. (2005) 'Combination of RFID and vision for mobile robot localization', *IEEE International Conference Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP 2005)*, Melbourne, Australia, pp.75–80.
- Chawla, K., Robins, G. and Zhang, L. (2010a) 'Object localization using RFID', *IEEE International Symposium on Wireless and Pervasive Computing (ISWPC 2010)*, Modena, Italy, pp.301–306.
- Chawla, K., Robins, G. and Zhang, L. (2010b) 'Efficient RFID-based mobile object localization', *IEEE International Workshop on Selected Topics in Mobile and Wireless Computing (STMWC 2010)*, Niagara Falls, Canada, pp.683–690.
- Choi, B. and Lee, J. (2009) 'Mobile robot localization scheme based on RFID and sonar fusion system', *IEEE International Symposium on Industrial Electronics (ISIE 2009)*, Seoul, South Korea, pp.1035–1040.
- Choi, J.S., Lee, H., Elmasri, R. and Engels, D.W. (2009) 'Localization systems using passive UHF RFID', *International Joint Conference on INC, IMS, and IDC*, Seoul, South Korea, pp.1727–1732.
- Estrin, D., Culler, D., Pister, K. and Sukhatme, G. (2002) 'Connecting the physical world with pervasive networks', *IEEE Pervasive Computing*, Vol. 1, No. 1, pp.56–69.
- Finkensteller, K. (2003) *RFID-Handbook: Fundamentals and Applications in Contactless Smart Cards and Identification*, 2nd ed., Wiley and Sons Inc.
- Fontelo, P., Ackerman, M., Kim, G. and Locatis, C. (2003) 'The PDA as a portal to knowledge sources in a wireless setting', *Telemedicine Journal and e-Health*, Vol. 9, No. 2, pp.141–147.
- Hähnel, D., Burgard, W., Fox, D., Fishkin, K. and Philipose, M. (2004) 'Mapping and localization with RFID technology', *IEEE International Conference on Robotics and Automation (ICRA 2004)*, Los Angeles, pp.1015–1020.
- Han, S., Lim, H. and Lee, J. (2007) 'An efficient localization scheme for a differential-driving mobile robot based on RFID system', *IEEE Transactions on Industrial Electronics*, Vol. 54, No. 6, pp.3362–3369.
- Hekimian-Williams, C., Grant, B., Liu, X., Zhang, Z. and Kumar, P. (2010) 'Accurate localization of RFID tags using phase difference', *IEEE International Conference on RFID (RFID 2010)*, Florida, USA.
- Hightower, J. and Borriello, G. (2001) 'Location systems for ubiquitous computing', *IEEE Computer*, Vol. 34, No. 8, pp.57–66.
- Jin, G., Lu, X. and Park, M. (2006) 'An indoor localization mechanism using active RFID tag', *IEEE International Conference on Sensor Networks, Ubiquitous, and Trustworthy Computing (SUTC 2006)*, Taichung, Taiwan, 4pp.
- Joho, D., Plagemann, C. and Burgard, W. (2009) 'Modeling RFID signal strength and tag detection for localization and mapping', *IEEE International Conference on Robotics and Automation (ICRA 2009)*, Kobe, Japan, pp.3160–3165.
- Liu, X., Corner, M. and Shenoy, P. (2006) 'Ferret: RFID localization for pervasive multimedia', *Lecture Notes in Computer Science*, Vol. 4206/2006, pp.422–440.
- Mattern, F. (2001) 'The vision and technical foundations of ubiquitous computing', *UPGRADE-The European Online Magazine for the IT Professional*, Vol. 2, No. 5, 6pp.
- Merrell, R.C., Merriam, N. and Doarn, C. (2005) 'Information support for the ambulant health worker', *Telemedicine Journal and e-Health*, Vol. 10, No. 4, pp.432–436.
- Milella, A., Di Paola, D., Cicirelli, G. and D'orazio, T. (2009) 'RFID tag bearing estimation for mobile robot localization', *IEEE International Conference on Advanced Robotics (ICAR 2009)*, Munich, Germany, pp.1–6.
- Muthukrishnan, K., Lijding, M. and Havinga, P. (2005) 'Towards smart surroundings-enabling techniques and technologies for localization', *Lecture Notes in Computer Science*, Vol. 3479/2005, pp.350–362.

- Ni, L., Liu, Y., Lau, Y. and Patil, A. (2003) 'LANDMARC: indoor location sensing using active RFID', *IEEE International Conference on Pervasive Computing (PerCom 2003)*, Texas, pp.407–415.
- Nikitin, P.V. and Rao, K.V.S. (2008) 'Antennas and propagation in UHF RFID systems', *IEEE International Conference on RFID (RFID 2008)*, Nevada, pp.277–288.
- Romer, K. and Domnitcheva, S. (2002) 'Smart playing cards – a ubiquitous computing game', *Journal on Personal and Ubiquitous Computing*, Vol. 6, Nos. 5/6, pp.371–377.
- Romer, K., Schoch, T., Mattern, F. and Dubendorfer, T. (2005) 'Smart identification frameworks for ubiquitous computing applications', *Journal on Wireless Networks*, Vol. 10, No. 6, pp.689–700.
- Satyanarayanan, M. (2001) 'Pervasive computing: vision and challenges', *IEEE Personal Communication*, Vol. 8, No. 4, pp.10–17.
- Schilit, B.N. (2003) 'Mega-utilities drive invisible technologies', *IEEE Computer*, Vol. 36, No. 2, pp.97–99.
- Senta, Y., Kimuro, Y., Takarabe, S. and Hasegawa, T. (2007) 'Machine learning approach to self-localization of mobile robots using RFID tag', *IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, Zurich, Switzerland, pp.1–6.
- Seo, D. and Lee, J. (2008) 'Localization algorithm for a mobile robot using iGS', *International Federation of Automatic Control – World Congress*, Vol. 17, No. 1, pp.742–747.
- Sweeney, P.J. (2005) *RFID for Dummies*, Wiley Publishing Inc.
- Vogt, H. (2002) 'Efficient object identification with passive RFID tags', *Lecture Notes in Computer Science*, Vol. 2414/2002, pp.98–113.
- Vorst, P., Schneegans, S., Yang, B. and Zell, A. (2008) 'Self-localization with RFID snapshots in densely tagged environments', *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2008)*, Nice, France, pp.1353–1358.
- Wang, C., Wu, H. and Tzeng, N. (2007) 'RFID-based 3-D positioning schemes', *IEEE International Conference on Computer Communications (INFOCOM 2007)*, Alaska, pp.1235–1243.
- Want, R. (2008) 'RFID – a key to automating everything', *Scientific American*, pp.56–65.
- Zhang, Y., Amin, M. and Kaushik, S. (2007) 'Localization and tracking of passive RFID tags based on direction estimation', *International Journal of Antennas and Propagation*, Vol. 2007, No. 17426, 9pp.
- Zhao, Y., Liu, Y. and Ni, L. (2007) 'VIRE – active RFID-based localization using virtual reference elimination', *IACC/US-NSF International Conference on Parallel Processing (ICPP 2007)*, Xi'an, China, p.56.