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):

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

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,

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.