

Reasoning about Uncertainty in Location Identification with RFID

James Brusey*

Institute for Manufacturing
Cambridge University
Cambridge, UK, jpb54@cam.ac.uk

Christian Floerkemeier*

Institute for Pervasive Computing
ETH Zurich, Zurich
Switzerland, floerkem@inf.ethz.ch

Mark Harrison*

Institute for Manufacturing
Cambridge University
Cambridge, UK, mgh12@cam.ac.uk

Martyn Fletcher

Agent Oriented Software Ltd
Mill Lane, Cambridge
martyn.fletcher@agent-software.co.uk

Abstract

Radio Frequency Identification (RFID) is set to revolutionise industrial control as it holds the potential to simplify and make more robust the tracking of parts or part carriers through manufacture, storage, distribution and ultimately the supply chain. RFID control is based on unique RFID transponder tags being attached to parts and used to identify the part as it moves through the factory or warehouse. Although RFID dramatically simplifies the process of tracking parts, there are certain situations that can lead to uncertainty about the true location of the part. This paper looks at two such situations: a robotic storage stack and a medicine cabinet. Both cases of uncertainty are successfully resolved by using a statistical filter. This work may lend itself to extensions and generalisations using Partially Observable Markov Decision Process (POMDP) models.

1 Introduction

Radio Frequency Identification (RFID) systems have become common in places where access control and tracking of physical objects is required. Examples include cattle herding, car immobilisers, and transport ticketing [Finkenzeller, 2000]. More recently, RFID systems have begun to find greater use in the consumer object identification market, in industrial automation, and in supply chain management. The use of RFID systems in these application domains has been promoted by efforts to develop low cost RFID tags as an economical replacement of bar-codes [Sarma *et al.*, 2000].

RFID systems typically consist of radio frequency (RF) tags, RF tag readers, and some software to process the tag reads. The tags typically respond to an RF broadcast by the tag reader by sending their serial number or other data stored

in their memory to the reader. Compared to optical bar code systems, RFID tags have the advantage that they can be read without line of sight through non-conducting materials and that multiple tags can be detected at once. As such, RFID systems are a useful tool in tracking the location of physical objects.

However, due to the low-cost and low-power constraints of RFID tags, reliability concerns arise under certain circumstances. In particular, we have noticed two types of undesirable effects:

- False negative reads, where RFID tags might not be read at all, leading to the mistaken belief that the object is not present, and,
- False positive reads, where RFID tags might be read when they are outside the region normally associated with the location of the RFID reader, leading to a mistaken belief that the object is present.

One approach to tackle the above problems would be to design more sophisticated tags that are not as vulnerable to shielding or interference, or to shield parts of the environment that are not intended to be seen. However, for RFID tags to become ubiquitous, they must be cheap and able to be used in hostile environments [Weis *et al.*, 2003]. For example, Sarma [Sarma, 2001; Sarma *et al.*, 2000] has proposed the design of a passive tag limited to one hundred bits of storage and between 500 to 5000 gates, to allow the per tag cost to be reduced to around 5 cents. Given these restrictions on the tag, we propose the use of time-based filters to avoid the false positive and negative reads mentioned above.

The rest of this paper is arranged as follows. The following section gives an introduction to RFID systems components and describes the interface between tags and readers in more detail. Section 3 looks more specifically at the relationship between RFID and uncertainty reasoning in robotics. Section 4 presents a case study of the uncertainty faced in low-cost RFID systems and how this uncertainty is currently resolved. Section 5 offers a robotics case study which details the problem of false positive reads rather than false negative

*Supported by the Auto-ID Center
(<http://www.autoidcenter.org>)

reads as described in the previous case study. Section 6 proposes a more generalised approach to reasoning about uncertainty in RFID-based robotics applications.

2 RFID Primer

RFID systems are composed of two main elements:

- the RFID tag or transponder, which usually carries a serial number and potentially other data and
- the RFID reader or interrogator, which detects tags and reads from and writes to the tags

RFID tags, which are attached to objects that need to be identified, consist of an antenna that is used to communicate with the reader, and a microchip, which stores, among other things, the identifying bit sequence. The RFID reader interrogates tags for their data using wireless communication. The reader contains the RF interface to the tags, internal storage, processing power and an interface to a host computer system to transfer the data sensed. Different RFID systems can be distinguished by:

- the amount of data that can be stored on the tags. This can be anything from a single bit, such as with Electronic Article Surveillance Tags, to thousands of bytes. Tags with larger capacities are typically battery powered.
- the type of power supply. Low-cost RFID systems are usually passive, which means that they use the field emitted by the RFID reader to power the microchip. Active tags, on the other hand, have their own power supply and can transmit over larger distances.
- the operating frequency. Common frequency bands for RFID systems are 135 kHz, 13.56 MHz, 915 MHz and 2.45 GHz. These frequencies are set by governmental bodies that control the electromagnetic spectrum in a region. For each band, government regulations specify the maximum radiation power and bandwidth.
- the anti-collision algorithm used. RF collisions occur when multiple tags respond simultaneously to a request from the reader. Their signals can interfere with each other, preventing the reader from identifying any of the tags.

Finkenzeller [Finkenzeller, 2000] gives a more in-depth classification of RFID systems. The case studies described in this paper use low cost, passive RFID tags operating at 13.56 MHz that only store a unique serial number.

3 RFID and Uncertainty in Robotic Control

RFID is particularly useful in industrial robotic control applications since the data that it provides is precise and, on the surface, easy to interpret. It is more precise, in terms of identifying objects than, say, computer vision, as it can identify which physical instance of object is involved, not just which class of object. When a tagged part arrives at a processing robot, it is possible to determine the history of the part in terms of where the unique part has been sensed and at what time. This is important from the point of view of being able to produce highly customised products from an

automated production line. For machine B to know how to treat a part, it may need to know what actions machine A has performed previously, and specifically, what actions were performed on that part. This is made possible with RFID by using the unique tag number stored on the tag as a database key to look-up information about the part.

RFID may also be useful in mobile robotics as a substitute for other forms of landmark detection as a basis for navigation [Lazanas and Latombe, 1992]. Specifically, a mobile robot might use an on-board RFID tag reader to read tags embedded in the floors or walls of corridors and rooms, and thus be able to deduce its rough location. An immediate advantage of using RFID is that a clear line-of-sight to the landmark is not required.

Although RFID provides highly reliable information about tags being near readers at particular times, it is still important to interpret this information correctly. In this respect, there are problems associated with reasoning about uncertain RFID data that are similar to those difficulties when using other types of robotic sensors. On the other hand, although existing work on reasoning about the uncertainty in other types of sensor data might be generalised to include RFID sensors, we suspect that this might not take full advantage of the fact that RFID is accurate in terms of the *identity* of the object being sensed.

As stated in the introduction, we have noticed two classes of problems with the way RFID tag reads are interpreted: false negative, and false positive reads. These problems have implications for robotic control that is based on this data. To give an example of the problem with false negative reads, consider a conveyor system with tagged shuttles that has tag readers prior to gates in the conveyor that switch the shuttles between different tracks. (This is similar to a train being switched between tracks using railway points.) If a tagged shuttle on a conveyor passes a tag reader but is not sensed (an example of a false negative), this may mean that the subsequent conveyor gate is not switched correctly. In the worst case, this may cause the conveyor to jam as the control system does not know of the presence of the shuttle.

False positive reads have similar consequences for robotic control. Where two lines of conveyor track run close together, a shuttle may appear to “jump tracks” if a tag reader intended to sense shuttles on one track reads tags from a shuttle that is on the other track. Again, this may cause the behaviour of the control system to be unpredictable, however the root cause is in the RFID sensor data and the interpretation of that data.

In the following case studies, we develop filtering mechanisms to cope with these two problems. Although these mechanisms are somewhat *ad hoc*, we provide the following rules of thumb for solving the problems more generally. The first rule is that, where it can be expected that tagged objects mostly sit statically in range of the reader, then they should only be considered to be not present when no read has occurred for some period of time. This period of time must be adjusted to compromise between responsiveness and the reliability of the resulting control system. Note that it may be necessary to adjust the environment so that objects do sit in read range for some period to ensure that they have a good probability of being read. For example, in the example given



Figure 1: The Smart Medicine Cabinet Application.

above of shuttles on a conveyor, it may be necessary to slow down or briefly stop the shuttles as they move past the tag reader.

The second rule of thumb that we apply here is that, where a tag reader should only read one tagged object, but sometimes registers another, more distant object, then the decision of which is the object that is actually present can be made by integrating over time, weighting recent tag reads more highly than old ones.

These rules of thumb are applied in the following case studies, with a more formal model being alluded to in section 6.

4 Case Study: Smart Medicine Cabinet

The Smart Medicine Cabinet is an application where an RFID reader is integrated into a medicine cabinet and each medicine packet is equipped with an RFID tag [Floerkemeier *et al.*, 2003]. The medicine cabinet is “smart” as it can sense its own content, as well as when contents are added or removed. It can also report this information to the user via a graphical user interface (GUI) or via audio output. Figures 1 and 2 show the cabinet and GUI in more detail. The medicine cabinet is equipped with an RFID system that provides a maximum read range suitable for the application and that employs an Aloha-like anti-collision protocol [Vogt, 2002]. During the operation of this medical cabinet, false negative reads occur for various reasons:

- RF collisions occur leading to one or more tags not being detected even though they are in range,
- a tag is not detected due to radio frequency interference or due to metal shielding,

- the RFID reader fails to detect any tags on a single read because of unknown operational problems of the RFID reader,

Figure 3 shows the output from the tag reader, when six tags are continuously present in its read range. The missing reads demonstrate the false negative reads mentioned above. Compared to the stack reader case study (see section 5 for details) the above situation is different because we are dealing with false negative reads rather than false positive reads. In this case study, there are by default multiple tags in the read range, but we tend to get false negative reads. That is, the tags are within the read range of the RFID reader, but are not detected.

To the user these false negative reads appear as “flickering” on the graphical user interface. When the addition or removal of items triggers speech output, such as “You just removed a pack of Aspirin”, the false negative reads tend to corrupt the entire operation of the system.

To address these unwanted effects we are currently using a top-hat function that excludes all reads that are Δt_{hat} older than the current time, t_{now} with Δt_{hat} being a constant,

$$f_{\text{hat}}(t) = \begin{cases} 1 & |t_{\text{now}} - t| < \Delta t_{\text{hat}} \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

If there is a tag read for an item at time t such that f_{hat} is non-zero, then that item is considered to be present.

This approach represents a compromise between the amount of flickering that occurs and the responsiveness of the user interface since the time it takes to detect the removal of an item is increased to Δt_{hat} . It falls short of an optimal solution, since it makes no use of *a priori* knowledge, such as:

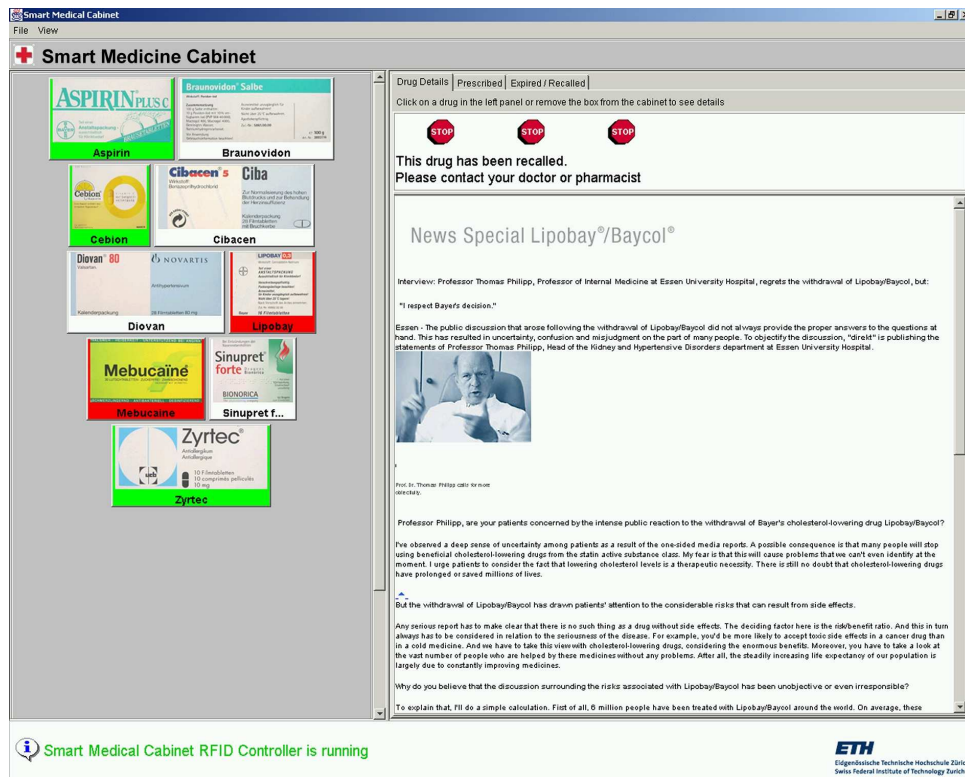


Figure 2: Screen-shot of the Smart Medicine Cabinet Application.

- the typical read rate of a specific RFID reader in the application environment (for example, in a metal cabinet).
- the typical read rate associated with the tag when attached to a particular object. For example, a tag on a medication pack containing a plastic bottle is likely to have a higher read rate than that on a medication pack containing a large number of metal-coated blister packs.
- the typical read rate of the object in its “logistical unit”. For example, blister packs equipped with RFID tags and packed densely into a folding box are certainly less likely to be detected than a single RFID tag on the folding box itself.
- the average duration the tag stays in the read range.
- the typical state transitions occurring. For example, it is unlikely that all medication items are removed at the same point of time.

The above likelihoods could be integrated into the simple scheme shown above by adapting Δt_{hat} appropriately. However, a more sophisticated approach as outlined in section 6 seems to be the more suitable approach. In the following section, we look at the converse problem, that of false positive reads.

5 Case Study: Stack Reader

A demonstration automated packing cell is being developed at the Cambridge University Institute for Manufacturing [McFarlane, 2002b; Brusey *et al.*, 2003]. The system currently

consists of a single anthropomorphic Fanuc¹ robot combined with three Montech² conveyor loops. A central aim behind the development of this system is to demonstrate how RFID integrates with control in an automated manufacturing environment. In this aspect, it is similar to recent work by Kärkkäinen *et al.* [Kärkkäinen *et al.*, 2003] that uses RFID to store all object data for the part being manipulated. Our work differs in that only an identifying bit sequence is stored on the tag and this tag is used as a database key to access data about the object.

The packing cell takes men’s shaving items, such as razors, bottles of foam, gel, and deodorant, and packs them into gift boxes. To simplify the task, individual items are first put into plastic carriers. Items come into the system on shuttles and are sorted into one of four stacks based on the product class. Items are later pulled from the bottom of the stack and packed into empty gift boxes.

This system is quite flexible, allowing each gift box to contain three items from any of the four product classes in any order. Furthermore it is quite robust to disturbances, such as manual addition or removal from the stacks. The ability to sense the identity, or at least the product class, of the item at the bottom of the stack is key to ensuring that this system is robust.

The detailed operation of the stack, which is shown in fig-

¹www.fanuc.com

²www.montech.ch




Figure 3: Example of a series of tag reads showing regular reads by an RFID reader of six tags which are present in its read range. The “missing” reads indicate that not all tags are detected on each scan (false negative reads). It also shows the different causes of false negative reads such as unknown operational and interference problems (1) and RF collisions (2).

ure 4, is as follows. When an item is put into the stack it falls down until it is either at the bottom of the stack or sitting on top of another item. Items are only fed in from the top and only taken out from the bottom. A Checkpoint Performa³ RFID tag reader is used to sense the item at the bottom of the stack. A key problem with the operation of the stack is that sometimes tags on items above the bottom of the stack are read. This tends to happen intermittently and may be affected by interference from other readers nearby. In addition, the physical characteristics of the item can change the shape of the RF field, thus leading to changes in the effective range of the tag reader.

Intermittent tag reads for items above the bottom of the stack may lead to incorrectly assuming that the bottom item has been removed. We consider this type of tag read to be a *false positive*, since the item is not in the position that we associate with the tag reader. These false positive reads can be particularly problematic if they are used as a trigger for other actions. To counter false positive reads, it seems necessary to take into account previous tag reads and to filter the tag reads in some way. Obviously, not all previous tag reads will be relevant. For example, if there are 100 tag reads and the first 60 are for tag *a* and the rest for tag *b*, then it is probably the case that tag *a* has been removed and been replaced by *b*. One simple way to filter this information is to decay the weighting associated with each read based on how long ago it occurred.

Figure 5 shows an extract from the log of a particular shelf reader. The time shown on the graph is based on the time that the RFID software receives the read from the reader rather than the true time of the read. Note that although tag *b* is sometimes read twice in succession, tag reads for *a* are more common throughout the period shown in the graph.

In our environment, we wish to monitor the state of the stack. However our initial implementation spuriously registered changes of state when no physical movement had actually occurred. Considering the occasional readings of the item above the base position as spurious, our second implementation filtered the tag reads by finding the *mode* or most popular tag read to determine which item was truly at the bottom of the stack. This approach has the problem that it does not respond correctly when an item is removed. Specifically, the *mode* yields the previous item until the number of new tag reads exceed the number accumulated for the previous tag.

To solve this problem, we modified our approach to attach greater weight to more recent tag reads. A weighted average was used based on the age of the tag reading. Firstly a top-hat function f_{hat} , as defined in equation (1), excludes all events that are Δt_{hat} seconds older than the current time, t_{now} . Secondly, the Gaussian weighting function,

$$f_{\text{Gaussian}}(t) = e^{-\frac{(t_{\text{max}}-t)^2}{2\sigma^2}}, \quad (2)$$

was applied, with a Gaussian half-width of σ seconds, centred on the time of the most recent of the timestamped readings

³<http://www.checkpointsystems.com>



Figure 4: A close-up view of the stack and associated tag readers.




Figure 5: Example of a series of tag reads showing regular reads from the base item nearby (corresponding to tag *a*), with intermittent reads for an item further away (corresponding to tag *b*).

t_{\max} . Given sets of timed tag reads R_i for each identity i , the weight for tag identity i is defined as,

$$w(i) = \sum_{t \in R_i} f_{\text{hat}}(t) f_{\text{Gaussian}}(t) \quad (3)$$

Normally, the identity i with the largest weight $w(i)$ is considered to belong to the tag at the base of the stack. A special case is where the weight is zero for all identities, and this implies that no items are left in the stack. This is required so that it is possible to detect when all items have been removed from the stack. Note that the top-hat function ensures that the weights will be zero when no reads have occurred in the last Δt_{hat} seconds.

By trial and error, we found that setting both Δt_{hat} and σ to 2 seconds produced a good compromise between filtering out unwanted tag reads and being responsive to changes in the state of the stack. Note that with these tag readers, we typically average around 13 tag reads per second. The resulting filter was tested by packing boxes and checking that the identity of the items seen during packing matched those of the items actually packed. Before implementing the filter, mismatches were commonplace, while afterwards, all trials produced a match between the items seen by the stack reader during packing and the resulting packed box.

6 Future Work

There are several problems with filtering the data in the way described so far. Perhaps the most important is that there is no strong theoretical basis for the design of the filters. A second issue is that the filters may not be suitable in other situations. Thirdly, the filters may need to be tuned and as yet we have no method other than trial and error to do so.

We believe all of these problems might be addressed by applying a Bayesian approach, such as that used in work on Markov Localisation [Fox, 1998; Thrun *et al.*, 2001]. Taking the example of the stack used for the second case study, the system can be modelled as a Partially Observable Markov Decision Process (POMDP) [Cassandra *et al.*, 1996]. It is partially observable since items above the base cannot be reliably detected by the tag reader. However, it might be possible to infer the state of the stack from information about what actions have been taken. Some progress on dealing with partial observability and RFID has already been made by McFarlane [McFarlane, 2002a]. As with Markov Localisation, rather than extracting the complete state of the system from the tag reader, tag reads and actions taken are integrated over time to build a probability distribution over possible states. A significant challenge with applying this approach to RFID is that the state space is potentially very large, and POMDP solving algorithms tend to be infeasible except for quite small state spaces. Note that the state space is large because there may be many unique tags in the system. Therefore, as with the Monte Carlo extensions to Markov Localisation [Thrun *et al.*, 2001] it may be necessary to work with a random sample of all possible states.

There also exist other alternatives to using POMDP based methods, such as using Monte Carlo methods, or Hidden Markov Models [Liu and Chen, 1998], that might be more

computationally feasible. In addition, it may be possible to perform high-level reasoning about tag reads based on geometric and logical constraints.

To address the fact that the results presented here are somewhat preliminary, we intend to enhance our experimentation to analyse the effectiveness of the algorithms in dealing with the false positive and false negative problems. The outline of this experiment is as follows. Given two stacks (as shown in figure 4), initially one filled, the other unfilled, a robot repeatedly removes an item from the base of one stack and deposits it on the top of the other. Therefore, the stacks at any given instant, may contain zero, one, two or three items. A goal of our reasoning system is to accurately identify the complete state of the stacks at any point in time. We will use two approaches to identify the state of the two stacks. The first will use unfiltered RFID data to determine which items are in which positions in the stack. The second will use filtered data to do this. From this we will be able to analyse the relative effectiveness of these two approaches in reducing false positive and negative reads.

7 Summary

This paper examined two key problems with using RFID based sensors to identify the location of physical objects, which we term false positive and false negative reads. These problems were examined in the context of two real applications of RFID, one a smart medicine cabinet and the other a robotic packing cell. Filtering tag reads over time was found to be effective in reducing both false positive and false negative reads. There were, however some limitations to the filtering approach used and we plan to address these limitations in future work.

Acknowledgements

The authors would like to thank Andy Garcia, Alan Thorne, and Steve Hodges for designing the smart stack used in section 5.

References

- [Brusey *et al.*, 2003] James Brusey, Martyn Fletcher, Mark Harrison, Alan Thorne, Steve Hodges, and Duncan McFarlane. Auto-id based control demonstration phase 2: Pick and place packing with holonic control. Technical Report CAM-AUTOID-WH-011, Auto-ID Center, 2003. <http://www.autoidcenter.org/research/CAM-AUTOID-WH-011.pdf>.
- [Cassandra *et al.*, 1996] Anthony R. Cassandra, Leslie Pack Kaelbling, and James A. Kurien. Acting under uncertainty: Discrete bayesian models for mobile-robot navigation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1996.
- [Finkenzeller, 2000] Klaus Finkenzeller. *RFID Handbook: Radio-Frequency Identification Fundamentals and Applications*. John Wiley & Sons, 2000.
- [Floerkemeier *et al.*, 2003] Christian Floerkemeier, Matthias Lampe, and Thomas Schoch. The smart box concept for

- ubiquitous computing environments. In *Proceedings of SOC'2003 (Smart Objects Conference)*, Grenoble, May 2003.
- [Fox, 1998] Dieter Fox. *Markov Localization: A Probabilistic Framework for Mobile Robot Localization and Navigation*. PhD thesis, Institute of Computer Science, University of Bonn, Germany, 1998.
- [Kärkkäinen *et al.*, 2003] Mikko Kärkkäinen, Jan Holmström, Kary Främling, and Karlos Artto. Intelligent products—a step towards a more effective project delivery chain. *Computers in Industry*, 50(2):141–151, February 2003.
- [Lazanas and Latombe, 1992] Anthony Lazanas and Jean-Claude Latombe. Landmark-based robot navigation. In *Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI-92)*, pages 816–822, San Jose, California, 1992. AAAI Press.
- [Liu and Chen, 1998] Jun S. Liu and Rong Chen. Sequential Monte Carlo methods for dynamic systems. *Journal of the American Statistical Association*, 93(443):1032–1044, 1998.
- [McFarlane, 2002a] Duncan McFarlane. Product identity and its impact on discrete event observability. In *Proceedings of the European Control Conference (ECC'02)*, Cambridge, UK, 2002.
- [McFarlane, 2002b] Duncan McFarlane. Towards Auto-ID based control systems. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics (IEEE SMC'02)*, Hammamet, Tunisia, 2002.
- [Sarma *et al.*, 2000] Sanjay Sarma, David L. Brock, and Kevin Ashton. The networked physical world – proposals for engineering the next generation of computing, commerce & automatic identification. Technical Report MIT-AUTOID-WH-001, MIT Auto-ID Center, 2000. <http://www.autoidcenter.org/research/MIT-AUTOID-WH-001.pdf>.
- [Sarma, 2001] Sanjay E. Sarma. Towards the five-cent tag. Technical Report MIT-AUTOID-WH-006, MIT Auto-ID Center, 2001. <http://www.autoidcenter.org/research/MIT-AUTOID-WH-006.pdf>.
- [Thrun *et al.*, 2001] Sebastian Thrun, Dieter Fox, Wolfram Burgard, and Frank Dellaert. Robust monte carlo localization for mobile robots. *Artificial Intelligence*, 128:99–141, 2001.
- [Vogt, 2002] Harald Vogt. Efficient object identification with passive RFID tags. In F. Mattern and M. Naghshineh, editors, *International Conference on Pervasive Computing*, volume 2414 of *Lecture Notes in Computer Science*, pages 98–113, Zurich, August 2002. Springer-Verlag.
- [Weis *et al.*, 2003] Stephen A. Weis, Sanjay E. Sarma, Ronald L. Rivest, and Daniel W. Engels. Security and privacy aspects of low-cost radio frequency identification systems. In *First Annual Conference on Security in Pervasive Computing*, 2003.