

Super-Distributed RFID Tag Infrastructures

Jürgen Bohn and Friedemann Mattern

Institute for Pervasive Computing
ETH Zurich, Switzerland
{bohn|mattern}@inf.ethz.ch

Abstract. With the emerging mass production of very small, cheap Radio Frequency Identification (RFID) tags, it is becoming feasible to deploy such tags on a large scale. In this paper, we advocate distribution schemes where passive RFID tags are deployed in vast quantities and in a highly redundant fashion over large areas or object surfaces. We show that such an approach opens up a whole spectrum of possibilities for creating novel RFID-based services and applications, including a new means of cooperation between mobile physical entities. We also discuss a number of challenges related to this approach, such as the density and structure of tag distributions, and tag typing and clustering. Finally, we outline two prototypical applications (a smart autonomous vacuum cleaner and a collaborative map-making system) and indicate future directions of research.

1 Introduction

In industry, the potential of radio frequency identification (RFID) technology was first recognized in the 1990s, stimulating the desire for RFID-supported applications such as product tracking, supply chain optimization, asset and tool management, and inventory and counterfeit control [22]. Besides these “conventional” application areas, passive RFID tags are also suited to augmenting physical objects with virtual representations or computational functionality, providing a versatile technology for “bridging physical and virtual worlds” in ubiquitous computing environments, as Want et al. showed [23].

Currently, the proliferation of RFID technology is advancing rapidly, while RFID reader devices, antennas and tags are becoming increasingly smaller and cheaper. As a result, the deployment of RFID technology on a larger scale is about to become both technically and economically feasible. Hitachi, for instance, is about to commence mass production of the *mu-chip* [11], which is a miniature RFID tag with a surface area of 0.3 mm². Further, the Auto-ID Center has proposed methods which could lower the cost per RFID chip to approx. five US cents [19].

In the conventional process of RFID tag deployment prevailing today, only a limited number of passive tags are placed in the environment in a deliberate and sparse fashion. Typically, RFID tags are mainly used for identifying objects [6, 24] and for detecting the containedness relationships of these objects [14]. Explicitly placed stationary tags embedded in the environment also serve as dedicated artificial *landmarks*. They can be detected by means of a mobile RFID reader and are used to support the navigation of mobile devices and robots [13, 16, 17], or to mark places and passageways [10].

In this paper, we present the concept of *super-distributed RFID tag infrastructures*, which differs from the conventional means of RFID tag deployment and utilization. We

advocate *massively-redundant* tag distributions where cheap passive RFID tags (i.e. tags without a built-in power supply) are deployed in large quantities and in a highly redundant fashion over large areas or object surfaces. We show that, in so doing, the identity of a single tag becomes insignificant in exchange for an increased efficiency, coverage, and robustness of the infrastructure thus created as a whole. We further demonstrate that such an approach opens up a whole spectrum of possibilities for creating novel RFID-based services and location-dependent applications, including a new means of cooperation between mobile entities. We also discuss some of the technological opportunities and challenges, with the intention of stimulating further research in this area.

The remainder of the paper is organized as follows: In Section 2 we introduce the concept of super-distributed RFID-tag infrastructures and describe its particular qualities in detail. In Section 3, we discuss different means of deploying RFID tags efficiently and redundantly on a large scale. Then, in Section 4, we outline two prototypical applications (a smart autonomous vacuum cleaner and a collaborative map-making system) and indicate future directions of research.

2 Super-Distributed RFID Tag Infrastructures

Passive RFID tags typically incorporate a miniature processing unit and a circuit for receiving power if the tag is brought within the field of an RFID reader. The tags are usually attached to mobile objects such as supermarket goods or other consumer products, and they send their identity to the reader over distances ranging from a few centimeters up to a few meters, depending on the type of tag.

RFID tags that are spread across a particular space in large redundant quantities can in turn be regarded as a “super-distributed” collection of tiny, immobile smart objects. The term “super-distribution” refers to the fact that a vast number of tags are involved, similar to the notion of “super-distributed objects” in [20]. Accordingly, we will refer to such a highly redundant tag distribution as a *super-distributed RFID tag infrastructure* (SDRI).

A highly redundant and dense distribution of tiny objects is also a common characteristic of wireless sensor networks which consist of a large number of very compact, autonomous sensor nodes. However, the two concepts differ fundamentally: in contrast to a fixed structure of independent and passive tags as part of an SDRI, wireless sensor networks are based on the “collaborative effort of a large number of nodes” [3]. Further, the topology of wireless sensor networks may change due to mobility on the part of its nodes. In addition, wireless sensor nodes carry their own power supply used to enable active sensing, data processing, and communication with other sensor nodes, whereas passive RFID tags only have very limited functionality, generally restricted to reading and writing a small amount of data. Also, compared to typical wireless sensor networks with nodes communicating over distances of tens of meters or more, mobile RFID antennas generally operate at a much shorter range.

By deploying an SDRI in an area, the overall physical space is divided into tagged and thus uniquely identifiable physical locations. This means that each tag can be used as an identifier for the precise location it covers, where coverage is pragmatically defined as the reading range of the tag. What we thus obtain can be described as an approx-

imate “discrete partitioning” of the physical space, of which different implementations are possible. If the tags of an SDRI are distributed according to a regular grid pattern, for example, then the partitioning of the physical space itself can be considered a physical grid of uniquely addressable cells approximating to the concept of a regular occupancy grid, as applied in the field of mobile robot navigation [7], for instance. If, on the other hand, the tags of the SDRI are distributed in a random fashion, we obtain an irregular pattern of uniquely addressable cells. Ideally, these cells are non-overlapping and cover the whole area. In practice, one can only approximate these properties.

In addition to the massive potential redundancy of RFID tags, two particularly interesting qualities of SDRIs are that they enable mobile devices to interact with their local physical environment, and that such an interaction can be performed in a highly distributed and concurrent manner. In the following, we explain these qualities in more detail.

2.1 Local Interaction with Physical Places

An SDRI allows mobile objects to store and retrieve data in the precise geographic location in which they are situated by writing to or reading from nearby RFID tags. Independent, anonymous entities are thus in a position to share knowledge and context information *in situ*.

One potential application of this quality is self-describing and self-announcing locations, where mobile devices can gain contextual or topological information on the spot simply by querying the local part of the RFID tag infrastructure. For instance, mobile GPS-enhanced vehicles could locally store positioning information while moving within an SDRI. Once a sufficiently large proportion of an affected area has thus been initialized, other mobile devices can be helped to recalibrate their GPS receivers, and GPS-less devices can be enabled to determine their position without using a dedicated positioning system themselves. Positional information stored in the SDRI can also be used to establish a fall-back service in case the primary positioning service is temporarily unavailable, thus increasing the overall availability of positional information in the area. Further, an SDRI facilitates the definition of arbitrary regions within physical spaces: virtual zones, barriers and markers can easily be defined by marking particular tags (or the tags along a border line) in the SDRI accordingly.

Furthermore, SDRIs in general offer *physical anchor points* which can serve as *entry points into virtual spaces* by allowing mobile devices to leave data traces, messages or links to virtual information (residing in a background infrastructure) wherever they roam. It is therefore possible to use the RFID tags of an SDRI as an alternative medium for implementing physical hyperlinks [12, 18] in virtual spaces, or as a means of attaching virtual annotations [21] to physical places.

By providing a means for roaming mobile objects to anchor and thus persistently store location-dependent sensor information on the spot, SDRIs also constitute a self-sufficient alternative to services such as GeoWiki [8], where virtual information is linked to a geographical address, but which require explicit knowledge of the current geographic location or the continuous availability of a location service of a sufficiently high resolution and accuracy.




Fig. 1. A mobile object has left a virtual trace on an RFID tagged floor space. Here, each tag that is part of the trace not only contains information about the identity of the tracked physical entity, but also indicates changes in its direction.

Mobile objects may also *leave virtual traces* in the physical space they traverse. By writing an anonymous ID to writable tags¹ in the floor space, a vehicle, for example, can leave a trace which can subsequently be retraced and followed by other mobile objects (see Fig. 1). These traces could later be overwritten by other vehicles or persons, so that they would fade away over time. By exploiting tag redundancy and enforcing suitable tag writing strategies, it should be possible to prevent the immediate deletion of a newly laid trace by following objects. For example, a moving vehicle could randomly choose one tag out of k available tags per location. Furthermore, a single tag could store the IDs of several different traces.

2.2 Global Collaboration Between Mobile Objects

An SDRI can be regarded as a scalable shared medium with (almost) unlimited, independent, and highly distributed physical “access points”. In this respect, an SDRI is particularly conducive to supporting *global collaboration*, where several independent physical entities work together on a single task in a highly decentralized and concurrent fashion. This is possible since SDRIs enable mobile devices to store and access data locally at the precise location they occupy at a given moment, so that devices which are situated at different locations within an SDRI can read from or write to tags simultaneously and independently.

Consequently, SDRIs are well suited for the implementation of a number of collaborative applications. Mobile objects that gather location-dependent information can store that information directly at the respective location by means of the SDRI, resulting in teamwork between anonymous entities that can be exploited for initializing or “bootstrapping” a particular SDRI with positional or topological information, for instance. Further, mobile objects can, depending on their capabilities, participate *on the fly* to achieve such a common goal while actually pursuing a different primary objective.

One concrete example of this is the collaborative exploration of an area. Since observations are based on globally unique tag IDs, the different map-making observations

¹ Tags are physically writable if an RFID scanner can be used to write data directly onto the tags. However, it is also possible to virtually write data onto read-only tags by means of a suitable background infrastructure, as described in Section 2.4.

of independent mobile entities can be unambiguously combined to form a global map (see also the prototype description below in Section 4). Such an SDRI-aided collaborative map-making process scales well as it can be performed by an arbitrary number of concurrent entities in a highly decentralized fashion.

Of course, for some of the scenarios described to become possible, both technical challenges (e.g., reading from or writing to RFID tags at high velocities) and conceptual issues (such as the question of commonly understandable topological models and ontologies, or the problem of updating information previously stored in the SDRI in order to reflect changes in the environment) have to be addressed.

2.3 Scope of Deployment

So far we have discussed a number of scenarios where SDRIs are deployed over large floor spaces in order to provide a novel means of local interaction and global collaboration. However, the scope of SDRI deployment is not just limited to such large-scale scenarios. There are also various situations where smaller-scale SDRIs have their distinct benefits.

For a table equipped with RFID tags embedded in its surface, for example, the intrinsic qualities of SDRIs still apply: such a “small-scale” SDRI also yields uniquely addressable cells, allowing multiple devices (or users) to interact with different sections of the tabletop simultaneously. If the tag distribution of the tabletop is known or if each tag knows its position with respect to the local coordinate system of the table, a smart object on the table can also easily determine its position with regard to the tabletop, or derive the relative distance from other smart objects on top of the table by communicating and exchanging particular tag IDs or tag coordinates. Similarly, a wall whose surface is coated with a layer of RFID tags can be turned into a smart “noticeboard” featuring support for the positioning of objects that are attached to it.

2.4 Physical vs. Virtual Tags

If RFID tags support read-write operations, then they obviously enable mobile devices to store a certain amount of data directly on the physical tags themselves. As a consequence, a mobile device can read from and write to the physical matter at its respective location, literally speaking.

Accordingly, if the available physical RFID tags are of a read-only type, a mobile device cannot directly write data to the tags. However, in this case we can still use the unique ID of the tags to unambiguously map each *physical tag* of the SDRI to a corresponding *virtual tag* residing in the background infrastructure. Rather than writing to a physical tag within the direct range of the mobile device, the device instead wirelessly connects to the virtual representation of the tag. The virtual tag may either simply provide the basic data read/write operations of a physical read-write tag, or even augment its capabilities by offering additional services which cannot be implemented on the small physical tags themselves due to resource limitations. So one distinctive advantage of virtual tags over mere physical tags is that they don't suffer from physical resource limitations, enabling us to write an almost unlimited quantity of data to a virtual representation of the physical tag.

However, to ensure instantaneous read/write access to virtual tags, continuous wireless access to a suitable background infrastructure that manages the virtually written tag-data is needed. In this case, the SDRI is no longer self-sufficient in so far as it needs to be continuously connected to the background infrastructure.

3 Efficient and Redundant Large-Scale Deployment of RFID Tags

The deployment of large quantities of tiny RFID tags over large areas necessitates an efficient means of tag distribution. Instead of minutely distributing tags across an area according to rigid and well-defined patterns, which typically goes hand in hand with a time-consuming calibration of these tags, we think that a *highly redundant random distribution of tags* is often more favorable. Such a random distribution can be achieved in various ways. For instance, RFID tags could be randomly mixed with various building materials such as paint, floor screed, concrete, etc., which is similar to the idea of mixing computer particles with “bulk materials” as described by Abelson et al. in [1]. Of course, in some cases such a procedure requires quite durable and resilient tags. But even if a certain percentage of tags were to be rendered defective in the process, the number of operable tags could be controlled by applying the necessary degree of redundancy.

If tags are uniformly distributed in a random manner over large areas, we can make assumptions about the average tag density and the coverage of the area. In some cases we can even randomly distribute tags and still maintain a certain regular distribution structure. By integrating RFID tags at regular intervals with string or the thread used for weaving carpets, for instance, it would be feasible to weave a complete carpet or produce carpet tiles that exhibited a regular RFID tag texture (e.g., forming a mesh of RFID tags). Even though we would not know the absolute positions of the tags after an RFID-augmented carpet had been laid out in a random fashion, we would still have relatively precise information about the distances between neighboring tags and about the overall tag density.

Although a random large-scale tag distribution enables a dense and, on average, uniformly distributed coverage comparatively cheaply and easily, it does pose some challenges. For instance, if the cost per RFID tag is too high even in mass production, a dense large-scale deployment may not be economically feasible. Further issues are the durability of RFID tags that are embedded in a carrier material, and the “bootstrapping” of super-distributed RFID infrastructures with respect to positioning and the provisioning of location-dependent context information.

3.1 Deployment of RFID Readers vs. Tags

Rather than tagging large areas with small, cheap RFID tags, it is also possible to distribute RFID reader antennas instead. By integrating an array of stationary RFID antennas into the floor, as described by [2], it is possible to detect tags that pass over particular readers, or even track certain tags if the output of several readers is combined and analyzed. Thus, by simply attaching a passive RFID tag to a mobile device, the latter is freed from the extra burden imposed by an energy-consuming RFID scanner.

However, such an approach has several disadvantages compared to the concept of SDRIs. First, integrating readers or antennas into the floor or into surfaces is resource-intensive, as quite a large quantity of expensive RFID equipment is needed to achieve a good resolution and coverage. For the continuous operation of these RFID readers and antennas, additional energized electronic devices are required. Furthermore, the deployment of the RFID equipment is complex and may require a considerable amount of construction work, not to mention the costly maintenance in the event that a device has to be replaced at a later point in time. Also, if the antennas are embedded into the environment and the mobile entities are tagged with passive RFID tags, the latter have only a limited means of controlling their degree of “visibility”. The mobile objects cannot easily prevent themselves from being detected or even tracked by the environment. If, on the other hand, the environment is being tagged and remains passive, the mobile entity itself is performing the sensing. Consequently, the mobile device remains in control of any interaction taking place with the environment, thus facilitating the implementation of a specific privacy policy, for example.

3.2 RFID Tag Distribution Patterns

In order to accomplish a large-scale distribution of RFID tags, there are a variety of possible tag distribution patterns to choose from. Typical distribution patterns include:

- *Random uniform distribution*: Tags are uniformly distributed over a certain area in a random manner.
- *Regular distribution*: Tags are distributed in a regular pattern, but usually with random tag identification numbers. Typical regular patterns are:
 - *Grid pattern*: Given a grid with edge length d , each non-border tag has four nearest adjacent neighbors at a distance d and four farther adjacent diagonal neighbors at a distance $\sqrt{2} * d$.
 - *Equilateral triangulation pattern*: Each tag has six equidistant neighbors at a distance d .

From the perspective of a mobile object, random tag distribution patterns have inherently different properties compared with regular patterns. One example is illustrated in Fig. 2. We expect there to be other generally advantageous patterns for the distribution of RFID tags, such as irregular but non-random distribution patterns, for example. This calls for the investigation of suitable tag distribution patterns and their respective properties as part of future research.

3.3 Sparse vs. Dense Tag Distributions

The density of the RFID tag distribution which can be achieved in an SDRI primarily depends on the properties of the underlying RFID technology. A secondary aspect is the required degree of resolution and the degree of tag redundancy, which can be expressed by the average number of tags that are within the range of the mobile reader antenna at an arbitrary location, for example.




Fig. 2. In an area of randomly and uniformly distributed tags, a mobile reader will come across another tag while moving on a straight path with probability $p > 0$ depending on the tag density and the distance traveled. In regular RFID tag distributions, such as a grid structure, it is possible that a mobile reader may choose a straight path where it will never encounter a single tag.

Sparse Non-Overlay Tag Distribution. If the RFID technology used does not support collision detection and resolution, only a single tag should be within the range of the reader antenna at any given location. We call this *sparse* tag distribution. In this case, the maximum tolerable tag density of the SDRI is limited by the characteristics of the available reader antenna. The distributed tags should be spaced in such a way that each tag exclusively covers an area (typically larger than the scan range covered by the reader antenna). Otherwise, the tags cannot be reliably scanned due to frequently occurring collisions. But even if collision resolution is available, one might deliberately prefer a non-redundant RFID tag distribution in order to simplify or speed up tag processing.




Fig. 3. Examples of sparse non-overlapping tag distributions: a) random b) grid c) triangular. The circles enclosing the tags indicate the idealized area in which a specific tag can be detected.

A sparse, non-overlapping tag distribution has several disadvantages, however. First, it results in an inflexible partitioning of the area covered into coarse-grained cells whose dimensions are defined by the range of the mobile reader antenna. Secondly, the typical scan range of reader antennas is roughly circular or elliptical. As a consequence, a non-overlapping tag distribution would not cover the entire area, but would result in “blind spots” at the fringes where no tags could be detected at all (see Fig. 3). Thirdly, a non-overlapping tag distribution is not redundant, which means that tag failures cannot be compensated for. Even if blind spots at the fringes of non-overlapping detection ranges

can be regarded as negligible transition zones and are therefore acceptable, the failure of single tags results in larger areas where no tag can be detected at all.

Dense Overlay Tag Distribution. In order to establish an SDRI with overlapping tag scan ranges, so that multiple tags can be detected per scan and per location, the RFID system must support collision resolution. In this case, the tolerable tag density is limited firstly by the technically viable proximity of tags at which these tags still respond correctly during a scan (and are unaffected by *tag detuning* [9], for instance), and secondly by the maximum number of tags within antenna range that can be simultaneously detected by the anti-collision scheme in a reasonable time. Both factors depend on the particular characteristics of the RFID system.




Fig. 4. Examples of dense overlay tag distributions: a) random b) grid c) triangular.

Dense overlay tag distributions allow us to achieve a fine-grained and complete tag coverage without “blind spots” (see Fig. 4). Another important advantage is the introduction of redundancy with respect to the number of RFID tags that are detectable per scan at a given physical location. By using an anti-collision system and a sufficiently high scan range, we can detect several tags per location. A location could then be described by the set of all the tags that are detectable at the respective place. If we detect several tags, we can calculate the center point of their locations and use that value as a position approximation for the current location by using localization techniques as proposed in [5], for instance. This has the benefit of increased robustness with regard to tag failures: if a tag is destroyed or fails over time, we still detect the remaining functional tags at a location.

3.4 Tag Typing and Clustering

Even though each RFID tag in an SDRI has its own unique ID, in certain situations it may be sufficient to discern only particular categories or types of tags. For the large-scale deployment of RFID tags within a building, for example, one might want to use different types of tags in different sections of a building, such as a tag type *A* for corridors, a type *B* for public spaces, a type *C* for private areas, and another type *D* to mark stairways and elevators. So in addition to a unique tag ID, each tag could also be equipped with an additional data field containing a predefined type identifier. If the RFID tags supported physical (or virtual) write access, it would be possible to “impregnate” the desired type identifiers onto tags after they had been deployed. Alternatively,

tags of a particular type could be grouped and pre-packed together for efficient deployment.

There are many different fields of application for tag typing and clustering, such as marking potentially hazardous areas for visually impaired people who are equipped with a smart cane² (e.g., in order to warn the person about an approaching stairway), or defining different categories of area for mobile robots.

4 Current Prototypes and Future Work

So far we have implemented prototypes of two SDRI-based applications (using the Hitachi mu-chip [11] and LEGO Mindstorms [15]), which demonstrate the feasibility and versatility of our approach. The first prototype is a location-aware autonomous vacuum cleaner (equipped with a mobile RFID reader and antenna) which adjusts its behavior based on tags embedded in the floor space at its particular location, such as avoiding areas that are marked as off-limits or keeping within an area surrounded by a virtual barrier (consisting of tags that have been marked accordingly in the teaching mode of the mobile robot).




Fig. 5. Mobile model vehicle built using LEGO Mindstorms technology. The vehicle motor and sensors are operated by the LEGO RCX unit, which receives control commands from or sends status information back to the LEGO infrared transmitter which is connected to a laptop computer. The RFID reader is also connected to the laptop on which the software for reading tags and calculating the tag positions is executed. The bottom view of the vehicle shows the RFID antenna, which is mounted approx. two centimeters above the floor, and the two rotation sensors measuring the revolutions of the back wheels. At the front of the vehicle, a bumper is connected to a pressure sensor to detect collisions while the vehicle is in motion.

The second is a system for collaborative map making where independent mobile vehicles (each again equipped with a mobile RFID reader and antenna) explore a previ-

² Goto et al. have developed a smart cane with an integrated RFID reader, for instance. It responds to RFID tags which serve as “data-carriers” embedded in the floor at places of interest [10]. Here an SDRI could provide a fine-grained and complete distribution of data-carriers over large areas, as opposed to the cumbersome process of deploying such data-carriers selectively.

ously unknown area (see Fig. 5). Starting from a known position, each vehicle chooses a random path through the area and thereby keeps track of the tags encountered and the relative inter-tag distances. The separate tag observations are subsequently merged by means of an efficient least-squares coordinate transformation algorithm, yielding a global map containing the absolute positions of known RFID tags. This map can then be used by other mobile devices for navigation and self-positioning purposes. Based on this prototype, we plan to implement an SDRI-based positioning system to be integrated with a modular probabilistic positioning system [4] which we have already developed. For the map-making trials, a sparse random tag distribution on a $0.5 \text{ m} \times 0.5 \text{ m}$ floor area has been used. The mobile reader antenna covers an area of about $9 \text{ cm} \times 6 \text{ cm}$ at a distance of about 1 cm above the floor. Using mu-chips equipped with a 4 cm long film antenna (inlet), a density of about 120 tags/m^2 has proven to be sufficient for this type of application.

Further experiments are underway to explore different tag distribution patterns and tag densities in practice and to determine their influence on scalability, efficiency, and robustness. The potential of autonomous vehicles with two or more reader antennas will also be explored. Apart from building specific SDRI-based applications for demonstration and more systematic evaluation purposes, we are also looking into ways of developing general middleware which will provide an efficient and reliable means of accessing an underlying physical SDRI, including fault-tolerant read/write operations, an automated maintenance procedure for the “hot” integration of newly distributed tags during operation, and high-level services such as location management, self-positioning, and local data sharing. We are particularly interested in the issue of robustness and the degree of fault-tolerance that can be achieved through massive redundancy of “super-distributed” RFID tags.

Acknowledgements. We wish to acknowledge Svetlana Domnitcheva, Julio Perez, and Matthias Sala for implementing the location-aware autonomous smart vacuum cleaner. We would also like to thank Marco Bär for his work on the collaborative map-making prototype and Hitachi SDL for providing us with mu-chips and tag readers.

References

1. Harold Abelson, Don Allen, Daniel Coore, Chris Hanson, George Homsy, Thomas F. Knight, Jr., Radhika Nagpal, Erik Rauch, Gerald Jay Sussman, and Ron Weiss. Amorphous Computing. *Communications of the ACM*, 43(5):74–82, March 2000.
2. Gregory D. Abowd, Agathe Ballestini, and Thomas O’Connell. The Location Service: A framework for handling multiple location sensing technologies. Available at www.cc.gatech.edu/fce/ahri/publications/location_service.pdf, 2002.
3. I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci. Wireless Sensor Networks: A Survey. *Computer Networks*, 38(4):393–422, March 2002.
4. Jürgen Bohn and Harald Vogt. Robust probabilistic positioning based on high-level sensor-fusion and map knowledge. Technical Report 421, Institute for Pervasive Computing, Dept. of Computer Science, ETH Zurich, Switzerland, April 2003.
5. Nirupama Bulusu, John Heidemann, and Deborah Estrin. GPS-less low cost outdoor localization for very small devices. *IEEE Personal Communications Magazine*, 7(5):28–34, October 2000.

6. Vlad Coroamă and Felix Röthenbacher. The Chatty Environment – providing everyday independence to the visually impaired. Workshop on Ubiquitous Computing for Pervasive Healthcare Applications at Ubicomp 2003, October 2003.
7. Alberto Elfes. Using occupancy grids for mobile robot perception and navigation. *IEEE Computer*, 22(6):46–57, June 1989.
8. Richard Fairhurst, Ben Jameson, Andrew Bolt, Robert Brown, and Matthew Slattery. Geowiki - a map which you can annotate. Homepage at www.geowiki.com.
9. Christian Flörkemeier and Matthias Lampe. Issues with RFID usage in ubiquitous computing applications. In *Proc. PERSASIVE 2004*, number 3001 in LNCS, pages 188–193. Springer-Verlag, April 2004.
10. K. Goto, Matsubara H., and S. Myojo. A mobile guide system for visually disabled persons. In *Proc. 4th Int. Symposium on Autonomous Decentralized Systems*, pages 12–17, March 1999.
11. Hitachi Ltd. Hitachi mu-chip - the World's smallest RFID IC. Hitachi Mu-Solutions. Homepage at www.hitachi.co.jp/Prod/mu-chip/, 2004.
12. Tim Kindberg. Implementing physical hyperlinks using ubiquitous identifier resolution. In *Proc. 11th Int. Conference on World Wide Web*, pages 191–199. ACM Press, 2002.
13. O. Kubitz, M.O. Berger, M. Perlick, and R. Dumoulin. Application of radio frequency identification devices to support navigation of autonomous mobile robots. In *IEEE 47th Vehicular Technology Conference*, volume 1, pages 126–130, May 1997.
14. Matthias Lampe and Christian Flörkemeier. The Smart Box application model. In Alois Ferscha, Horst Hörtnner, and Gabriele Kotsis, editors, *Advances in Pervasive Computing*, pages 351–356. Austrian Computer Society (OCG), April 2004.
15. LEGO Mindstorms. Homepage at <http://mindstorms.lego.com>, 2004.
16. L.M. Ni, Yunhao Liu, Yiu Cho Lau, and A.P. Patil. Landmarc: indoor location sensing using active RFID. In *Proc. 1st IEEE Int. Conference on Pervasive Computing and Communications (PerCom 2003)*, pages 407–415, March 2003.
17. M. Philipose, K.P. Fishkin, D. Fox, D. Hahnel, and W. Burgard. Mapping and localization with RFID technology. Technical Report IRS-TR-03-014, Intel Research, December 2003.
18. Michael Rohs and Jürgen Bohn. Entry Points into a Smart Campus Environment – Overview of the ETHOC System. In *Proc. 23rd Int. Conference on Distributed Computing Systems - Int. Workshop on Smart Appliances and Wearable Computing (IWSAWC)*, Providence, Rhode Island, USA, May 2003.
19. Sanjay E. Sarma. Towards the five-cent tag. Technical Report MIT-AUTOID-WH-006, MIT Auto-ID Center, 2001. www.autoidcenter.org/research/MIT-AUTOID-WH-006.pdf.
20. Super Distributed Objects Domain Special Interest Group (SDO DSIG). Request for Proposal: PIM and PSM for Super Distributed Objects. OMG Document sdo/02-01-04, Object Management Group (OMG), www.omg.org, 2002.
21. Marc Smith, Duncan Davenport, Howard Hwa, and Lik Mui. The Annotated Planet: A mobile platform for object and location annotation. In *Proc. 1st Int. Workshop on Ubiquitous Systems for Supporting Social Interaction and Face-to-Face Communication in Public Spaces at UbiComp 2003, Seattle, Washington, USA*. Microsoft Research, October 2003.
22. J.R. Tuttle. Traditional and emerging technologies and applications in the radio frequency identification (RFID) industry. In *IEEE Radio Frequency Integrated Circuits (RFIC) Symposium*, pages 5–8, June 1997.
23. Roy Want, Kenneth P. Fishkin, Anuj Gujar, and Beverly L. Harrison. Bridging physical and virtual worlds with electronic tags. In *Proc. SIGCHI Conference on Human Factors in Computing Systems*, pages 370–377. ACM Press, 1999.
24. Kho Hao Yuan, Ang Chip Hong, M. Ang, and Goi Sio Peng. Unmanned library: an intelligent robotic books retrieval & return system utilizing RFID tags. In *Proc. IEEE Int. Conference on Systems, Man and Cybernetics (SMC '02)*, volume 4, October 2002.