Challenges for Pervasive RFID-based Infrastructures

Evan Welbourne, Magdalena Balazinska, Gaetano Borriello, and Waylon Brunette
University of Washington, Seattle, WA
{evan, magda, gaetano, wrb}@cs.washington.edu

Abstract

The success of RFID in supply chain management is leading many to consider more personal and pervasive deployments of this technology. Unlike industrial settings, however, deployments that involve humans raise new and critical problems related to privacy, security, uncertainty, and a more diverse and evolving set of applications.

At the University of Washington, we are deploying a building-wide RFID-based infrastructure with hundreds of antennas and thousands of tags. Our goal is to uncover the issues of pervasive RFID deployments and devise techniques for addressing these issues before such deployments become common place.

In this paper, we present the challenges encountered and lessons learned during a smaller-scale pilot deployment of the system. We show some preliminary results and, for each challenge, discuss how we addressed it or how we are planning on addressing it.

1. Introduction

In the last few years, Radio Frequency Identification (RFID) technology has gained increasing attention as a flexible, and relatively fast solution for tagging and wireless identification [16, 18]. Early successes in the asset tracking and supply-chain domains [15] coupled with the falling cost of tags have lead researchers to consider pervasive, public RFID deployments that support more user-oriented services. A number of investigations into personnel tracking and task automation using RFID [2, 11, 14] have shown the technology's potential to facilitate everyday life by seamlessly integrating the virtual and physical worlds. Unfortunately, the majority of such studies have been limited to technology and user evaluations over a short time in restricted scenarios (often in a laboratory). Furthermore, the publicity surrounding this work [12] has revealed an intense public concern with RFID privacy and policy issues that have gone largely unaddressed.

We believe that a more holistic approach is required to effectively design and evaluate RFID-based pervasive computing systems. To this end, we are deploying a long-term, building-wide RFID-based test-bed in our department's building that will involve hundreds of RFID readers and antennas and thousands of tags. Our intent with this "RFID Ecosystem" is to explore the benefits of pervasive RFID infrastructures while identifying and addressing their challenges before such systems are adopted widely in other public settings, where problems may have more serious implications.

Several properties distinguish RFID infrastructures for pervasive computing from those for supply-chain applications. First, pervasive RFID applications are likely to evolve and grow over time. We already see RFID in elder care and object finding [2, 11] applications, each of which requires a flexible infrastructure that facilitates provisioning. Supply-chain applications are typically less dynamic and apply the technology in a narrower capacity (mostly for inventory tracking). Second, because a pervasive application will typically track people and belongings rather than items in inventory, privacy issues must be considered much more carefully. Finally, people are less predictable than goods moving through established distribution patterns in a supply-chain. As such, we must develop fundamentally new ways to deal with the variable-rate, partial, and noisy data likely to be generated by human activity.

In this paper, we present the RFID Ecosystem architecture and discuss the lessons we have learned from a small-scale pilot deployment. The issues we faced included hardware installation mechanics, aesthetics, health regulations, system reliability, privacy concerns, and the flexibility and scalability of our data management infrastructure. We outline our approach for addressing each issue and, wherever possible, distill generally useful practices. It is clear from this study, however, that the two main issues facing pervasive RFID deployments are privacy and overall system reliability. These two challenges must be expressly met before pervasive RFID-based infrastructures become widespread.

The remainder of this paper is organized as follows. We motivate our work in Section 2, present the system architecture in Section 3, and describe the study in Section 4. In Sections 5 through 7, we discuss the main challenges and lessons learned. We present related work in Section 8 before concluding in Section 9.

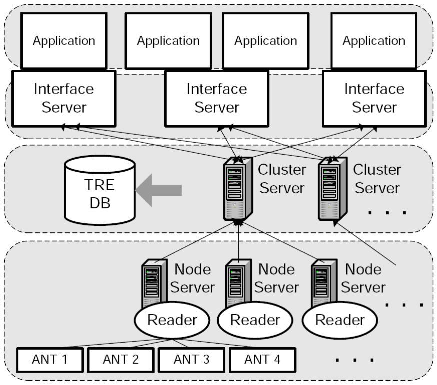

Figure 1. System architecture

2. Applications

The deployment of a pervasive RFID-based infrastructure in an everyday environment holds the promise of enabling new classes of applications that go beyond tracking and monitoring. Such a system could, for example, support logging and analysis of individual tag movements over time, allowing a user to ask questions such as "how often do I get interrupted in my office on an average day?". Current and historical data on groups of tags could also be used to identify and analyze aggregate phenomena such as the impact of seminars on improving communication between researchers. Real-time streams of tag reads could be used in "find my object" applications, reminding services [2], and to actuate devices.

Enabling these classes of application, however, presents a significant challenge to the system design. First, the system must be able to consistently and accurately read tags and it must do so at a granularity sufficient for the intended applications. The system must support archiving and retrieval of tag reads and real-time reporting of new reads. To enable the monitoring of large spaces with a possibly large number of tags, the system must scale to handle high-volume streams of tag reads. Finally, because such a system will manage large amounts of (potentially sensitive) personal data from multiple users, it must be secure and support an appropriate privacy model. In the following section, we present our preliminary system architecture and discuss how these application requirements affected its design.

3. System Architecture

A pervasive RFID-based infrastructure must manage three types of information. First, it must manage streams of tag read events (TREs) generated by antennas as they detect tags in their vicinity. Each TRE contains a tag ID, the ID of the antenna that detected the tag (*i.e.*, reader ID and antenna number), and a timestamp. Second, the infrastructure must store tag metadata (TMD) which includes tag ID, the name of the tagged object, the name of that object's owner, privacy parameters, and possibly other information. Finally, the system must manage reader metadata (RMD), that is the location of each reader and its antennas.

The RFID Ecosystem is designed to support the collection, storage, transport, and sharing of TREs, TMD, and RMD across a set of applications in a reliable, scalable, secure, and privacy-oriented fashion.

As illustrated in Figure 1, the system architecture is divided into four layers. The bottom layer consists of the RFID readers, their antennas and, for each reader, a piece of software called the Node Server that simply collects data from the reader. Each Node Server streams its TREs to a nearby Cluster Server, which stores them in a centralized database (which also holds TMD and RMD) and forwards them, as requested, to one or more interface servers. Interface servers are application-specific RFID Ecosystem clients. They can do as little as manage network connections and forward streams of TREs or they can perform sophisticated event detection [21] and continuous stream processing [1] on behalf of a set of applications. The fourth and top layer of the system is the application layer.

With the exception of the interface server and the application which can run on the same machine, all components exist on physically distinct machines. To ensure high-availability, each component should be replicated, although we do not yet use replication in our prototype deployment. We further discuss system reliability in Sections 5 and 6.

This layered and partitioned design ensures scalability. Node servers filter and smooth data [9]. Cluster Servers only handle data generated by small groups of Node Servers and can be added incrementally as the system grows. Interface servers handle their own load and are provisioned as the number of applications and users grow. Finally, the centralized database can be partitioned and distributed.

The layered design also helps support privacy and security requirements as we discuss in Section 7.

4. Pilot Study

To evaluate the RFID Ecosystem and gain insights on practical issues, we conducted a two-week pilot study of a prototype deployment. We also ran a series of laboratory RFID benchmarks to supplement the pilot.

The study had 6 participants, 2 female and 4 male (3 students and 3 faculty), and included the first three authors as well as other members of our research group. The prototype deployment used 11 RFID readers with 34 fixed-position antennas, each with a range of about 2 meters. Antennas were deployed in the hallways of the top three floors of our building. Passive tags (900 MHz UHF EPC1) were used to tag mobile objects. Each reader had an associated node server and there was one cluster server per floor. Participants were asked to wear a "person ID" tag and to tag any

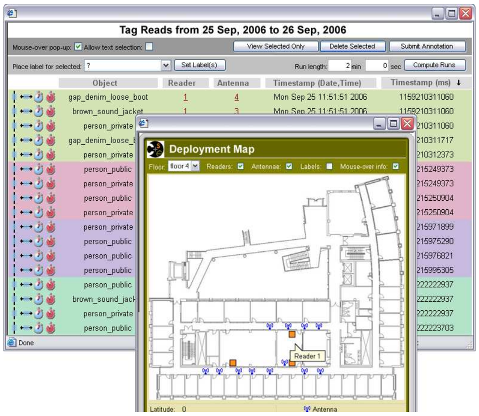
Tag Reads from 25 Sep, 2006 to 26 Sep, 2006

Object	Reader	Antenna	Timestamp (Date,Time)	Timestamp (ms)
gap_denim_loose_boot	1	4	Mon Sep 25 11:51:51 2006	1159210311060
brown_sound_jacket	1	3	Mon Sep 25 11:51:51 2006	1159210311060
person_private				210311060
gap_denim_loose_t				210311717
person_private				210312373
person_public				215249373
person_private				215249373
person_public				215250904
person_private				215250904
person_private				215971899
person_public				215975290
person_public				215976821
person_public				215995305
person_public				222222937
brown_sound_jack				222222937
person_private				222222937
person_public				222223703

Figure 2. Screenshot of the web application.

personal objects they wished, 54 tags were used in the study. A Javascript-enabled web application allowed users to issue preset queries on the collected data such as: "where is object X?"; "where is person Y?"; and "how much time have I spent in the building this week?". The interface also allowed users to view their own raw data and to delete any TREs at any time. A map-based interface was used to display antenna and TRE locations. Figure 2 shows a screenshot of the application.

To collect ground truth to compare against TRE data, users kept a daily web diary of their movements (and the movements of their tagged objects) throughout the building. Participants also regularly discussed their experiences and any interesting system-related phenomenon.

5. Deployment Challenges

The physical installation of an RFID-based infrastructure in a public environment raises several issues ranging from health considerations, to reliable tag detection, and even aesthetics. In this section, we use results from our pilot study to present a checklist of what we found to be some of the most important points to consider when performing such a deployment. For two of them, we provide some preliminary data. We also list four others that turned out to be more important than we originally thought.

Consider the relative positions of tags and antennas. It is widely known [4, 18] and our own microbenchmarks confirm that the probability with which a tag is read depends greatly on (a) how the tag is mounted on an object, ideally tags should be mounted such that normal use of the object results in the proper orientation of the tag relative to the antenna; (b) the orientation of the antenna itself, which may help ensure a large number of tags will be perpendicular to it; and (c) the object's material properties. Figure 3(a) shows the read rates resulting from a microbenchmark where an experimenter walked by a single overhead antenna ten times in each direction while holding a particular object at her side. The figure shows great variance depending on the material properties of the object: *e.g.*, a purse with no metal objects is almost always detected, but a purse holding a PDA or cell phone is not.

In a public setting, constraints (a) and (b) translate into considering how tagged objects will be used or carried (*e.g.*, books in a bag tend to be oriented perpendicular to the floor). Constraints (a) through (c) can also be addressed by constraining where to place tags on objects and people. As shown in Figure 4, a tag hung loosely around the neck results in a better read rate than a tag clipped tightly to clothing. This is likely due to the high water content of human bodies. It is, however, difficult to predict or control what users will do with their objects and their tags. In the pilot study, we found the read rates to be much lower than in microbenchmarks and greatly variable between objects and people. Figures 3(b) and 3(c) show the read rates for various tags throughout the pilot; these rates were computed as the number of "antenna crossings" while in transit as reported by the TRE data, divided by the actual number of antenna crossings as inferred from the web diary.

Exploit redundancy. In our study, we deployed multiple antennas in each hallway. As a result, users had to pass a sequence of up to three antennas on their way to or from their offices. Figure 4 shows the probability that at least one, at least two, or all three antennas in the sequence detected a tag. These measures were computed by counting antenna crossings in the TRE data for each trip past a sequence of three antennas where at least one TRE occurred for the tag in question. For the objects in our study, the redundancy increased read rates by 50% to 400%. A possible reason for this improvement is that the position and orientation of a tag changes as a participant walks along a path. Instead of deploying antennas in more locations, another technique for exploiting redundancy would have been to mount multiple antennas at different angles covering the same area, preferably the region around a doorway or a hallway intersection. In general, both options work and it is not yet clear if one deployment outperforms the other.

Remember health regulations. An important constraint for any RFID deployment is that FCC regulation limits the amount of power a person may absorb from an antenna. For our equipment this translates into a 9 inch standoff between an antenna and any occupiable space (*e.g.*, hallways, offices). To comply with the regulation while achieving a reasonably good read rate, we hung antennas from 8-foot high hallway ceilings (pointing down). This mounting position allowed for a read range that covered most tags above a participant's waist.

Aesthetics matter. The large amount of equipment in the hallways raised concerns and objections regarding the inaesthetic appearance of RFID readers and antennas. We found, however, that no one noticed antennas attached to the ceiling as long as the cables were hidden. Other use-

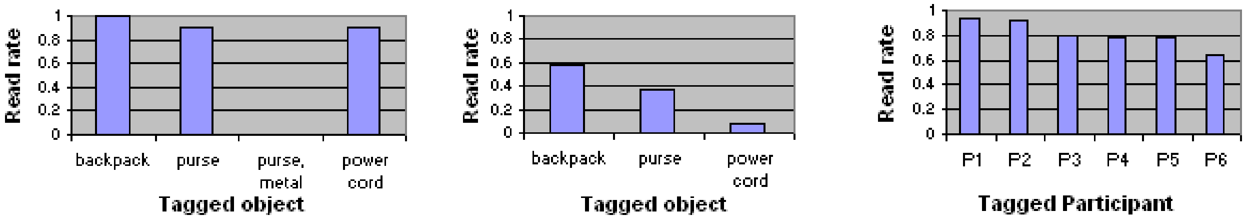

(a) Microbenchmark (b) Pilot study (c) Pilot study

Figure 3. Read rates for tags in a laboratory benchmark and in the pilot study

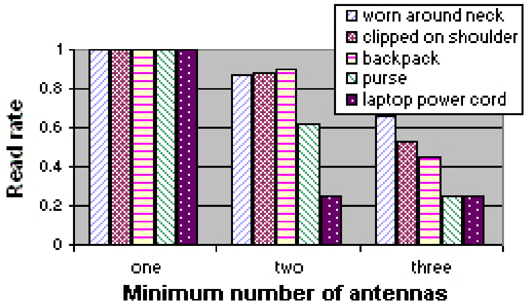

Figure 4. The effect of antenna redundancy on read rate for various types of tagged objects.

ful camouflage included non-metal ceiling tiles, glass windows, and non-metallic paint.

6. Systems Challenges

The systems challenges that arose during the pilot had to do primarily with reliability and the lack of predictability in the generated data streams.

System failures. As with any distributed system, the RFID Ecosystem faces the risk of node and network failures due to system updates, network maintenance, software bugs, etc. Since we started the deployment 2 months ago, our system experienced over twelve interruptions outside of our control (*i.e.*, excluding software bugs and upgrades), showing that fault-tolerance is critical in real deployments.

Input data errors. As discussed in the previous section, an RFID-deployment also suffers from the inherent imprecision of TREs. In our study, we saw many missed TREs but also a few duplicate TREs where the same tag was detected by adjacent antennas. These errors hurt the usefulness of the system because they propagate to applications. Cleaning RFID data is currently an active area of research [9, 10].

Unpredictable input streams. In our deployment, unlike in supply-chain management, the characteristics of the TRE stream are often unpredictable. Because antennas are mounted in hallways, TREs are typically generated in short bursts (about 3 TREs per object as a person passes by an antenna). Several times, however, an item was left in front of an antenna and generated thousands of TREs within a few minutes. Managing such unpredictable data rates will require transforming low-level tag reads into higher-level tag events (*e.g.*, "tag X is at location Y", "tag X has left location Y"). This process, however, will likely require application-specific parameters and will likely not be able to detect all events with certainty.

7. Privacy Challenges

Pervasive RFID-based deployments raise privacy concerns because they can enable the tracking of people and personal objects by parties that would otherwise be unable or unauthorized to do so. These concerns involve: the physical security of the communication between tags and readers, the security of the data stored in and processed by the system, and controlled access to the data. In this study, we focused on the latter problem and studied in-situ many of the privacy concerns experienced by the participants.

Chief among the participants' concerns was the perceived ease with which one's activities could be inferred from the data (*e.g.*, time of day, direction of movement, and set of tags seen). We were able to validate this concern by writing a simple script that could detect lunch breaks with better than 75% accuracy for three of the participants, showing that participant P1 took 29 minute lunch breaks on average, P2 took about 32 minutes, and P3 took the longest (40 minute) breaks. Similarly, it was easy to infer potentially more sensitive information such as when and how many times a participant used the restroom in a day.

Our initial approach to addressing these privacy concerns was to allow display and deletion of one's personal data via the web interface. The limitation of this approach is that a user can still see another's data before that data is deleted. A more appropriate option would be for users to specify high-level rules that describe which TREs should be accessible to which users and which TREs should be dropped automatically (*e.g.*, all trips from my office to the restroom shorter than 2 minutes). However, rule-based access control may have limited expressiveness and can limit the utility of the system. To this end, techniques for limiting queries or anonymizing and perturbing query results could also be used [3, 17]. For example, a query on a colleague's location could return approximate information (*e.g.*, 4th floor) by default, or more exact information (*e.g.*, room 490) only a few times per day. We are currently exploring the suitability of such techniques for various applications.

Finally, to protect the privacy of non-participants, each Node Server automatically discards any tag reads for a tag that is not registered in our database.

8. Related Work

There is a large body of work on RFID middleware design [5, 6, 13]. However, the proposed architectures are oriented toward the supply chain and asset tracking domain. As such, these designs give little attention to privacy and security and typically support only a narrow range of enterprise applications. By contrast, privacy, security, scalability, and extensibility are fundamental design principles of the RFID Ecosystem.

Past work with RFID in pervasive computing [14, 16, 20] has demonstrated the promise of more user-oriented RFID-based applications. Unfortunately, the majority of these studies have been conducted in settings that are too restricted to bring forth the real challenges and subtleties of pervasive RFID in everyday life. A key goal of the RFID Ecosystem is to enable such research in a microcosm of the real-world where real-life concerns can be adequately modeled and addressed. A few related real-world deployments exist, including a dedicated RFID test-bed at WINMEC [13] and a number of pervasive computing deployments [8, 19]. However, none of these deployments focuses on the problem of a private, scalable, extensible RFID infrastructure for pervasive computing.

Compared to other indoor location tracking systems (*e.g.*, IR based, IEEE 802.11 based) [7], the RFID Ecosystem offers some compelling trade-offs. First, though the one-time cost of installing RFID infrastructure could be an order of magnitude greater than installing IR beacons or wifi access points, the cost of each new RFID tag is at least an order of magnitude less than a new IR or wifi enabled device. To this end, the RFID Ecosystem can scale to track many more objects at a fraction of the cost. Furthermore, the burden of maintaining an RFID system is less because there are no batteries to be replaced or complex, on-going radio signal strength calibration tasks required.

9. Conclusion

In this paper, we motivated the benefits of a pervasive RFID-based infrastructure by outlining some of the applications that such an infrastructure enables. We presented the architecture of our RFID Ecosystem, a system designed to provide data management services for building-scale, pervasive RFID deployments. We used our system to conduct a small-scale pilot study and presented the deployment, systems, and privacy challenges that we encountered. By far, the two key issues were the security and privacy of the data and the overall system reliability. The latter is challenging to achieve not only because of system failures, but also because of the intrinsic unreliability of the RFID technology and the unpredictability of a pervasive environment. We are currently completing our building-wide deployment and are planning to conduct longer studies with a larger number of participants.

Acknowledgments. We would like to thank Justin Cutter, Alex May, Jonathan Scoles, and Robert Spies for technical support. This work is supported by the University of Washington CoE and the NSF under Grant No. 0454394.

References

[1] Abadi et. al. The design of the Borealis stream processing engine. In *Proc. of the CIDR Conf.*, Jan. 2005.

[2] Borriello, G. et. al. Reminding about Tagged Objects using Passive RFIDs. In *Proc. of the 8th Ubicomp Conf.*, Sept. 2006.

[3] Dinur, I. and Nissim K. Revealing information while preserving privacy. In *Proc. of the 22nd PODS Conf.*, pages 202–210, June 2003.

[4] Floerkemeier, C. and Lampe, M. Issues with RFID usage in ubiquitous computing applications. In *Proc. of the 2nd Pervasive Conf.*, Apr. 2004.

[5] Floerkemeier, C. and Lampe, M. RFID middleware design - addressing application requirements and RFID constraints. In *sOc-EUSAI 05*, Oct. 2005.

[6] Franklin, M. J. et. al. Design considerations for high fan-in systems: The hifi approach. In *Proc. of the CIDR Conf.*, Jan. 2005.

[7] Hightower, J. and Borriello, G. A survey and taxonomy of location sensing systems for ubiquitous computing. UW CSE 01-08-03, Seattle, WA, Aug. 2001.

[8] Intille, S. S. et. al. A living laboratory for the design and evaluation of ubiquitous computing interfaces. Apr. 2005.

[9] Jeffery, S. et. al. Adaptive cleaning for RFID data streams. In *Proc. of the 32nd VLDB Conf.*, Sept. 2006.

[10] N. Khoussainova, M. Balazinska, and D. Suciu. Towards correcting input data errors probabilistically using integrity constraints. In *Proc. of Fifth MobiDE Workshop*, June 2006.

[11] Liu, X. et. al. Ferret: RFID Localization for Pervasive Multimedia. In *Proc. of the 8th Ubicomp Conf.*, Sept. 2006.

[12] McCullagh, D. Perspective: RFID tags: Big brother in small packages. C—NET News.com., Jan. 2003.

[13] Prabhu, B. S. et. al. *WinRFID - A Middleware for the enablement of Radio Frequency Identification (RFID) based Applications, In Mobile, Wireless and Sensor Networks: Technology, Applications, and Future*. Wiley 2005, 2005.

[14] Smith, J. R. RFID-based techniques for human-activity detection. *Communications of the ACM*, 48(9), Sept. 2005.

[15] Songini, M. L. Wal-Mart details its RFID journey. ComputerWorld, Mar. 2006.

[16] Stanford, V. Pervasive computing goes the last hundred feet with RFID systems. *IEEE Pervasive Computing*, 2(2), Apr. 2003.

[17] Sweeney, L. k-anonymity: A model for protecting privacy. *IJUFKS*, 10(5):557–570, Oct. 2002.

[18] Want, R. The magic of RFID. *ACM Queue*, 2(7), Oct. 2004.

[19] Want, R. et. al. An overview of the PARCTAB ubiquitous computing experiment. *IEEE Personal Communications*, 2(6):28–33, Dec 1995.

[20] Want, R. et. al. Bridging physical and virtual worlds with electronic tags. In *CHI*, pages 370–377, 1999.

[21] E. Wu, Y. Diao, and S. Rizvi. High-performance complex event processing over streams. In *Proc. of the 2006 SIGMOD Conf.*, June 2006.