

Physical Access Control for Captured RFID Data

To protect the privacy of RFID data after an authorized system captures it, this policy-based approach constrains the data users can access to system events that occurred when and where they were physically present.

Radio frequency identification technology has become popular as an effective, low-cost solution for tagging and wireless identification. Although early RFID deployments focused primarily on industrial settings, successes have led to a boom in more personal, pervasive applications such as reminders¹ and eldercare.² RFID promises to enhance many everyday activities but also raises great challenges—in particular, with respect to security and privacy.

At the University of Washington, we've deployed the RFID Ecosystem, a pervasive computing system based on a building-wide RFID infrastructure with 80 RFID readers, 300 antennas, tens of tagged people, and thousands of tagged objects.³ The RFID Ecosystem is a capture-and-access system that streams all data from the readers into a central database, where applications can access

it. Our goal is to provide a laboratory for long-term research in security and privacy, as well as applications, data management, and systems issues for RFID-based, community-oriented pervasive computing.

RFID security is a vibrant research area, with many protection mechanisms against unauthorized RFID cloning and reading attacks emerging.⁴ However, little work has yet addressed the complementary issue of protecting the privacy of RFID data after an authorized system has captured and stored it. We've investigated peer-to-

peer privacy for personal RFID data through an access-control policy called Physical Access Control. PAC protects privacy by constraining the data a user can obtain from the system to those events that occurred when and where that user was physically present. While strictly limiting information disclosure, PAC also affords a database view that augments users' memory of places, objects, and people. PAC is appropriate as a default level of access control because it models the physical boundaries in everyday life. Here, we focus on the privacy, utility, and security issues raised by its implementation in the RFID Ecosystem.

Privacy and utility in pervasive architectures

The 18th-century legal philosopher Jeremy Bentham first described the perfect architecture for surveillance: the *panopticon*, a prison that arranges its cells about a central tower from which a guard can monitor every cell while remaining invisible to the inmates. The architecture's innovation is that the guard's presence becomes unnecessary except for occasional public demonstrations of power. Many privacy concerns in pervasive computing stem from a similar potential for an unseen observer to access and act on data about someone else. Under these conditions, the "state of conscious and permanent visibility [assures] the automatic functioning of power"⁵ because individuals must constantly conform to the code of conduct their peers or superiors hold them to.

Just as surveillance can be built into an architecture, so can privacy assurances. Our funda-

Travis Kriplean, Evan Welbourne,
Nodira Khoussainova, Vibhor
Rastogi, Magdalena Balazinska,
Gaetano Borriello, Tadayoshi
Kohno, and Dan Suci
University of Washington

mental conviction regarding privacy in the RFID Ecosystem is that privacy must be designed into the system from the ground up. The challenge in architecting privacy into a pervasive-sensing system is to provide enough utility to support the desired applications suite while carefully controlling what information to disclose, to whom, how, and under what conditions. Effective decisions must holistically trade off privacy and utility.⁶ In particular, a proposed privacy mechanism must consider perspectives and methods from computer security, databases, human-computer interaction, and the social sciences. Only in this way can we understand the mechanism's security vulnerabilities, how well it matches users' expectations of privacy, how easy it is to understand, and its utility.

Most pervasive-sensing systems represent one of two architectural models: wearable or infrastructure. Generally speaking, each model makes different trade-offs between privacy and utility. The wearable model processes and stores sensors and data on devices that the user owns and wears. MyLifeBits embodies this model: users wear microphones, video cameras, and other sensors that continually record sensor data.⁷ Such systems can put the device wearer in control if they store the data locally and disclose no information without the user's explicit consent. These "perfect memory" systems, however, pose privacy concerns for others who encounter the user but don't consent to information capture.⁸ Plausible deniability is lost: although human memory is lossy, captured sensor data isn't.

In contrast, the infrastructure model has a central authority that manages sensor data on users' behalf. This model gives rise to the threat of permanent visibility, but we adopted it for the RFID Ecosystem because it enables much richer services through data and resource sharing. It's also less expensive because

cost is amortized over many users. Moreover, a central database allows a system to leverage database security and privacy techniques such as privacy-preserving data mining⁹ and k-anonymity.¹⁰ These techniques enable privacy-preserving statistical queries that complement and extend the utility of careful access control for point queries, such as "When did I see Bob today?" A principled, privacy-sensitive framework for managing data should employ both statistical and point queries—as in a Hippocratic database,¹¹ for example. However, our focus here is on point queries under one possible access control policy.

PAC: How it works

We could define many access-control policies for captured RFID data. For example, policies might allow a manager to track employee location during work hours, support staff to locate inventory objects, and individual users to grant conditional tracking permissions to their friends. However, these policies present problems when applied as a default. The managerial example raises surveillance concerns, the object finder can be abused to track people, and user-defined poli-

cies often lack foresight and vigilance.¹² On the other hand, a restrictive policy that lets users access only their own data precludes many interesting applications that might be safe.

Marc Langheinrich argues that proximity-based disclosure could limit the surveillance threat.¹³ PAC implements this proposal. It attempts to realize the spatial privacy features of wearable sys-

tems within an infrastructure architectural model. We propose it as a default access-control policy because it models spatial privacy in everyday life. It places upper and lower bounds on accessible information, restricting the information users can obtain to what they could have observed in person. Specifically, a user can "see" other users and her own tagged objects at any time and place when she was physically present, but can't see the other users' objects. For each user, the database stores a persistent record of all encountered persons and personal objects. PAC thus respects Yitao Duan and John Canny's *data discretion principle*,¹⁴ which states that users should have access to media recorded when they were physically present and shouldn't when they weren't present. We believe that PAC provides an intuitive, easily understandable flow of captured information.¹²

Although PAC is conservative in the information it reveals, its memory-like view of the data provides useful service primitives. For example, a user's query on the location of his lost object will return the location where he last saw it—a likely answer. Queries over a

A proposed privacy mechanism must consider perspectives and methods from computer security, databases, HCI, and the social sciences.

user's entire history of activities could enable other personal-memory applications—for example, a reminder service that alerts users when they forget to take an item with them as they go home for the day.¹ This balance between privacy and utility makes PAC a suitable default access-control policy for a pervasive RFID deployment. Moreover, privacy-preserving extensions and

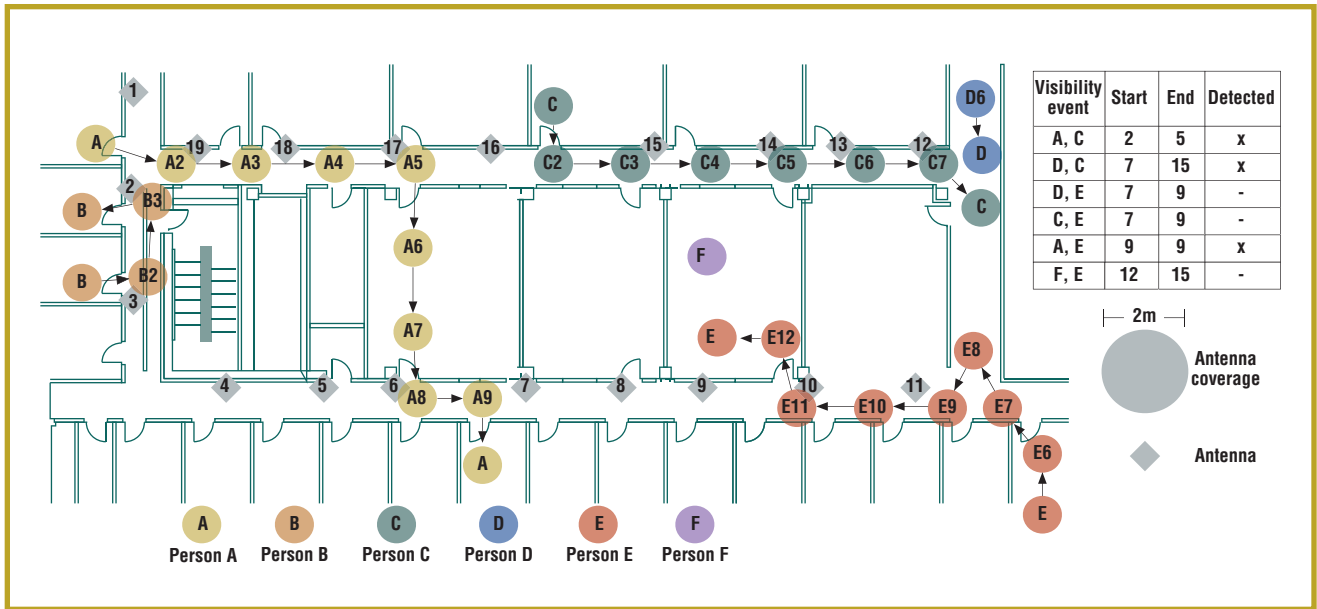


Figure 1. Six people moving over the course of 15 time steps. Each user’s location is annotated with the time stamp at which they move to the space. No time stamp indicates that they remain in that space. The table shows visibility events between users. The last column shows visibility events inferred by the system when the parameterizable time window Δ is 1 and antennas 1–3, 4–11, and 12–19 are mutually visible.

relaxations of PAC could enable an even greater range of applications. As such, rather than plunge a community into an information-promiscuous environment, our strategy is to start with information disclosure commensurate with everyday life and carefully extend it as needed for useful applications.

Implementation

RFID systems collect data as a stream of triples having the form (Identity, Location, Time). PAC defines a database view consisting of only those triples representing a user’s own location and the locations of users and objects he or she could plausibly have seen. A database respecting the PAC policy always responds to a user’s queries through her view, rather than the complete data—in other words, PAC employs a Truman model.¹⁵ For example, if a user asks, “How many people were on the fourth floor yesterday?” the system effectively responds, “You saw five people on the fourth floor,” instead of “You saw five out of the 11 total people on the fourth floor.”

A PAC implementation thus requires

a procedure for inferring when a user could plausibly have seen another user or object. Our implementation relies on a notion of mutual visibility for this procedure. Two users or a user and an object are *mutually visible* if they share an unobstructed line of sight. Every such instance of mutual visibility is called a *visibility event*.

Consider the scenario in figure 1. It presents a snapshot of six users, A–F, going about their daily routines and a table enumerating all visibility events.

This definition of mutual visibility is an ideal that the system must approximate using captured sensor data. In the RFID Ecosystem, the mutual-visibility computation incorporates the spatio-temporal relationships between the RFID tag reads that antennas collect. The RFID Ecosystem’s formal definition of mutual visibility depends on these relations between tag reads:

- *Spatial*. Determining the unobstructed line of sight between two tag reads poses two challenges. First, a sighted tag’s exact location is unknown; in-

stead, the antenna’s location serves as a proxy for the tag’s location. Second, two tags might be mutually visible yet read by two different antennas, which motivates the definition of *mutually visible antennas*: pairings of antennas A_1 and A_2 such that the system can interpret a tag read at A_1 as mutually visible with a tag read at A_2 .

- *Temporal*. By protocol, each antenna reads tags rapidly and in sequence, so two tags are rarely read at exactly the same time. We therefore use a parameterizable time window, Δ , that defines how close in time two tag reads must occur for the tags to be considered mutually visible.

We can now express a formal definition of mutual visibility in terms of the data captured by the system: Two tags X and Y are mutually visible if X is read by antenna A_1 at time t_X and Y is read by antenna A_2 at time t_Y , such that $|t_X - t_Y| \leq \Delta$ and A_1 and A_2 are mutually visible.

This definition yields the visibility events marked in figure 1. In this sce-

nario, antennas 18 and 15 are mutually visible, so the system will correctly interpret A and C as mutually visible at $t = 3$. In contrast, antennas 11 and 12 are not mutually visible, so E and C won't be considered mutually visible. Note how varying Δ tunes mutual visibility's strictness. In figure 1, $\Delta = 1$, so the system never detects D and A as mutually visible; however, if $\Delta = 2$, then they are mutually visible during time steps 5 to 7.

Users, objects, and ownership

We distinguish between *user tags* and *object tags*. A user can see an object's location only when the user and object are mutually visible and the user owns that object. The ownership restriction is required because RFID tags are readable through opaque materials, such as backpacks. X-ray vision is not part of the PAC information ethic. In our current model, ownership is simple: each object is singly owned. However, our goal is to study community-oriented systems, so future work will need to investigate how PAC can operate with shared objects.

Measuring mutual visibility

Given a pair of antennas A_i, A_j , we would like to label them as mutually visible or not in a way that minimizes false visibility events. Our approach has been to label each antenna pair with the probability that two tags in these coverage areas share a line of sight. The motivation is to give system administrators a way to systematically reason about potential information leakage.

One method is to sample a large number of points from each antenna's expected coverage area and calculate the fraction of point pairs that share a line of sight. Let C_i and C_j be two sets of points uniformly drawn from A_i 's and A_j 's respective coverage areas and let $\text{visiblePoints}(p_a, p_b)$ be 1 when points p_a and p_b share a line of sight, and 0 otherwise. Then

$$\text{visibility}(A_i, A_j) = \frac{\sum_{p_a \in C_i, p_b \in C_j} \text{visiblePoints}(p_a, p_b)}{|C_i||C_j|}$$

We then consider A_i, A_j mutually visible if $\text{visibility}(A_i, A_j) \geq \tau$, where τ sets the lower bound on the fraction of point pairs that must share a line of sight for two antennas to be considered mutually visible. With $\tau = 1.0$, all points must share a line of sight, providing the highest privacy by minimizing falsely detected, mutually visible tags. However,

$\tau = 1.0$ will likely miss many actual visibility events, thus degrading utility. After studying our deployment, we labeled the antennas we thought should be mutually visible, yielding $\tau = 0.84$.

This measure has limitations. First, τ is an approximation because true antenna coverage varies over time and with environmental conditions. Second, antennas can sometimes read through opaque surfaces (for example, an interior laboratory window with curtains drawn). Further techniques are necessary to accurately model antenna behavior.

PAC feasibility

Our definition of PAC assumes a well-behaved, lossless model for our RFID equipment. We therefore wanted to determine how PAC performed in a real deployment with antennas that might not behave as expected. Erroneous behavior can have adverse effects on both privacy and utility. For example, privacy violations can occur if an antenna reads a tag beyond its expected range (possibly causing false visibility events); likewise, utility can be degraded when an antenna fails

to read a tag. Here, we discuss our experiment to evaluate PAC in practice.

Experimental setup and methodology

We evaluated PAC over a set of user scenarios that cover some ways visibility events could occur. For each scenario, we enacted multiple trials and collected the corresponding stream of raw tag reads that the antennas captured on each trial. For each trial, we also captured ground-truth location data. A simulator then processed the ground truth data,

The motivation is to give system administrators a way to systematically reason about potential information leakage.

producing a stream of simulated tag reads that models the case of a well-behaved, lossless RFID deployment.

Representative scenarios. A visibility event can occur in many ways. We've defined four scenarios that, while not exhaustive, represent common types of visibility events:

- In the *personal-objects* scenario, a single user walked around the halls carrying six tagged objects on various parts of the body and inside a duffel bag.
- In the *glimpses* scenario, one user stood at one end of a hallway while another entered the opposite end from an office and quickly walked around the corner.
- In the *walking-together* scenario, two users walked around the hallways together.
- In the *passing-by* scenario, two users passed one another while walking in opposite directions.

Data collection and ground truth. The scenarios gave a rough script for exper-

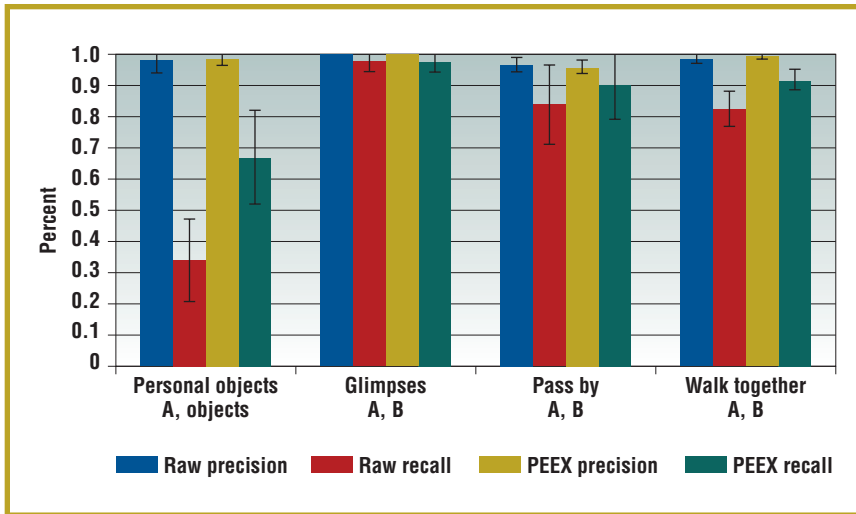


Figure 2. Average precision and recall for visibility events in each scenario. Each group of experimental results represents a single visibility event between two tags, except for the first, which is the average of A's visibility events with their objects.

imenters to follow. To collect ground truth accurately for each trial, experimenters used tablet PCs with a map-based data-collection tool.¹⁶ Using the tool's stylus, experimenters could track their current location as they enacted the scenarios by moving a cross-hair to the corresponding location on a map. In this way, each trial produced an XML trace of time-stamped latitude and longitude coordinates. This trace is fed to the simulator to produce the simulated tag reads. On the basis of experience with our RFID deployment, we set the antenna coverage area as a circle with a 2-meter radius about the antenna.

Comparing visibility events. For each pair of tags X and Y in a given trial, we compared the visibility events detected in the raw and simulated data. Let S and R denote the set of all visibility events for X and Y in the simulated and raw data, respectively. By definition, a visibility event v occurs during a window of time (t_X, t_Y) such that $|t_X - t_Y| \leq \Delta$. We define the *visibility-event time stamp* of v as $\mu_v = (t_Y + t_X)/2$. A visibility event s in S occurs in R if there exists an r in R such that $|\mu_r - \mu_s| \leq \epsilon$. So, a visibility event in S also occurs in R when a visibility event in R has a time stamp within ϵ of the visibility-event's time stamp in S . In our experiments, $\epsilon = 1$ second.

To measure privacy and utility, we

compared the recall and precision of the visibility events detected in the raw and simulated data. Recall measures utility according to the fraction of simulated visibility events that also occurred in the raw data. A recall of 1 indicates that the raw data accurately captured all visibility events in the simulated data. A recall less than 1 indicates that the simulation missed some visibility events because of missed tag reads. In contrast, precision measures privacy according to the fraction of detected visibility events that also occurred in the simulated data. A precision less than 1 indicates privacy loss because the system detects false visibility events.

Computing results. We performed 10 trials of each scenario. After each trial, we calculated precision and recall for the visibility events of every tag pair. We then computed the mean and standard deviations of precision and recall across all the trials for these visibility events. In the personal-objects scenario, visibility events occur between the experimenter and each of his or her objects. In this case, we give the average and standard deviation across all these pairings.

Results

Figure 2 shows the precision and recall for all four scenarios. The experimental outcome is encouraging and indicates that PAC can indeed operate effectively

in a lossy environment to provide both privacy and utility.

Privacy. The high precision demonstrates that nearly all the visibility events detected by our RFID deployment also occurred in the simulated data, suggesting that little information leaked. This result also verifies the integrity of the data-collection procedure because high precision depends on correct ground truth input.

Utility. Recall suffers when antennas fail to read tags and so miss visibility events. This can happen for various reasons, such as the properties of the material to which the tag is affixed and the tag's orientation with respect to the antennas. The tags hung from the experimenters' shirt or pants. This resulted in high read rates for the user tags (recall between 90 and 95 percent) in all the experiments and correspondingly high detection of visibility events between user tags (more than 80 percent).

In the personal-objects scenario, we observed lower recall because the antennas couldn't consistently read the tags in pockets or the duffel bag. However, several algorithms and tools could ameliorate this problem by cleaning RFID data. We evaluated whether such tools could improve PAC performance by comparing the precision and recall of the raw data stream against a third set of tag reads produced by the PEEEX (Probabilistic Event Extractor for RFID Data) research prototype.¹⁷ PEEEX uses integrity constraints to correct raw RFID data. We ran PEEEX with a few integrity constraints that capture logical, physical relations between objects and people (for example, an object can't

move by itself). Figure 2 shows that PEEEX significantly improves recall in the personal-objects and walking-together scenarios (t-test with $p < 0.001$), without affecting precision.

Overall, our results show the practical feasibility of employing PAC. Despite the noisy, lossy, and inaccurate nature of real RFID data, PAC effectively limits information disclosure while providing good system utility. It captures user-user interactions quite well. Inherent RFID unreliabilities hamper the capture of user-object interactions, but even simple cleaning tools such as PEEEX significantly improve performance in this area.

“Misplaced” user tags

Our PAC implementation assumes that users are always wearing their user tags. We must, however, anticipate users who accidentally or intentionally “misbehave.” For example, Alice might accidentally forget her user tag in Bob’s office, or she might maliciously place it in Bob’s backpack. In both scenarios, the system would incorrectly believe that Alice is in Bob’s proximity and would grant her access to his data. (Such errors aren’t possible with object tags because they can only be mutually visible with their owner.)

We’re developing several mechanisms for addressing “misplaced” user tags. We focus here on users who intentionally misbehave; mechanisms that defend against malicious parties will also account for accidental misuse. Our defensive techniques fall under the principle of security risk management. While an adversary might still be able to circumvent our security mechanisms, the cost of mounting an attack against users’ privacy should outweigh the benefits to the adversary.

Detection

The threat of detection can deter malicious activities because it might lead to social sanctions and punitive measures for

the offending party. We’re exploring two classes of detection mechanisms. First, the RFID Ecosystem could automatically detect anomalies in a user tag’s movements. The system could trigger an alert if Alice’s user tag is always mutually visible with Bob or one of Bob’s objects, such as his backpack, or if Alice’s user tag has been in an unusual location for too long.

Second, because nonvisually impaired users generally know the ground truth about the people (or at least the number of people) in their immediate vicinity, we can explicitly involve users in anomaly detection. For example, Bob could detect Alice’s maliciously planted RFID user tag if he’s in an elevator alone but the elevator’s front panel says that there are two occupants.

Prevention

We’re exploring two classes of prevention mechanisms: making attacks too costly or inconvenient for adversaries and periodically verifying that the RFID user tags are in the appropriate user’s possession.

One method for increasing the cost to an adversary is to combine user tags with expensive or essential devices, such as

Other attack vectors

Alice could share her legitimate observations of Bob (as accessed through the PAC system) with Charlie, thereby revealing to Charlie information about Bob’s location that Charlie couldn’t have observed. Direct attacks on the RFID hardware (for example, cloning tags) are also possible, but we don’t consider them here. (Ari Juels surveys the state of the art in preventing such attacks.⁴)

Future work

We’re focusing our future work in three areas.

Principled PAC relaxations

We’re looking at other access-control mechanisms to implement alongside PAC to provide additional information when socially appropriate. For example, user-defined access-control rules are important for applications that rely on shared context between users, such as location-shared buddy lists.¹⁸ Because users explicitly grant permission to make their information available when they opt in to such applications, physical proximity is not a necessary access requirement.

Socially situated events offer another

The threat of detection can deter malicious activities because it might lead to social sanctions and punitive measures for the offending party.

cell phones or employee badges. A second method is to stop capturing reads for user tags that become separated from their legitimate owners. For example, we could consider a user tag “capturable” for n time units whenever the RFID Ecosystem detects that the legitimate user is actually in possession of that tag—say, by detecting when the tag enters that user’s office.

potential relaxation. For example, an augmented calendar system might let a user query the location of a meeting’s invitees during the scheduled meeting time. Such a relaxation would give users socially acceptable information at particular times.

A third class of relaxations could involve mediation with the system. For example, a lost object’s owner might

request its location. The system could choose to reveal the object's location to the requester, or it could send an email to the person most recently detected to have moved it. By involving the system or an administrator, this relaxation could prevent abuse by not revealing sensitive information while still supporting useful actions.

Finally, an opportunistic access-control scheme¹⁹ could let users access private information in rare circumstances such as emergencies. Administrators would log and investigate these actions to decide if they were legitimate. Determining the conditions and frequency under which to use this access mechanism is an open problem.

User studies

We still need to empirically validate our assertion that PAC is an intuitive policy that will match users' expectations of privacy in everyday life. Moreover, a number of studies have examined user privacy expectations for captured audio and video data (for example, Giovanni Iachello and his colleagues²⁰) and disclosure of information to others across potentially great distances (for example, Sunny Consolvo and her colleagues²¹), but few studies examine how physical space factors into people's expectations of privacy for captured RFID data. Apu Kapadia and his colleagues have begun to explore the use of spatial metaphors,²² but further work is necessary. The PAC definition and implementation has helped in obtaining the Institutional Review Board "minimal risk" approval for these user studies.

Probabilistic data

We've demonstrated that using PEEEX to clean data enhances data utility, but PEEEX can also produce probabilistic data. In this case, each tuple has an associated probability that represents the system's confidence about its validity.

Implementing PAC in a probabilistic

context leads to a challenging problem. Suppose user *A* asks, "Is user *B* currently at location *L*?" If *A* is at *L*, then PAC allows the correct answer. If *A* isn't at *L*, then PAC refuses to reveal *B*'s location. However, cleaned data would assign probabilities p_A and p_B to the chance that *A* is at *L* and *B* is at *L*, respectively. In the probabilistic context, the correct answer is no longer yes or no but p_B . Yet the system can't return p_B when *A* isn't at *L*.

One approach is to return $p_A \cdot p_B$, the probability that both *A* and *B* are at *L*. However, this reveals too much. Even if p_A is small (*A* is not likely to be at *L*), *A* can still compute p_B . More generally, the requirement is that if *A* is likely to be at *L* (p_A is large), then the system should reveal p_B . Otherwise, the system should hide this information. One ad hoc strategy is to return $\min(p_A, p_B)$. We plan to explore more principled approaches that fulfill this requirement.

The RFID Ecosystem project aims to research socially appropriate RFID systems to provide the community (including businesses and policy makers) examples of effective methods for balancing utility with privacy. PAC is a first step in this direction because, as a default access control policy, it provides an upper and lower bound on accessible information that models human experience. It also leaves open the possibility of utility-enhancing extensions. Our experiments show that PAC is a practical solution that works well in a real-world RFID deployment where sensors are unreliable. ■

ACKNOWLEDGMENTS

We thank Garret Cole, Patricia Lee, Robert Spies, and Jordan Walke, with special thanks to Caitlin Lustig for her work on the simulator we used. We also acknowledge the University of Washington

College of Engineering. The US National Science Foundation funded this research under its Computing Research Initiative grants 0454394, IIS-0428168, and IIS-0415193.

REFERENCES

1. G. Borriello et al., "Reminding ABOUT Tagged Objects Using Passive RFIDs," *Proc. Ubiquitous Computing 6th Int'l Conf. (UbiComp 04)*, LNCS 3205, Springer, 2004, pp. 36–53.
2. D. Patterson et al., "Fine-Grained Activity Recognition by Aggregating Abstract Object Usage," *Proc. 9th Int'l Symp. Wearable Computers (ISWC 05)*, IEEE CS Press, 2005, pp. 44–51.
3. E. Welbourne et al., "Challenges for Pervasive RFID-Based Infrastructures," *Proc. 5th Ann. IEEE Int'l Conf. Pervasive Computing and Communications Workshops (Perccomp 07)*, IEEE CS Press, 2007, pp. 388–394.
4. A. Juels, "RFID Security and Privacy: A Research Survey," *IEEE J. Selected Areas in Communications*, Feb. 2006, pp. 381–395.
5. M. Foucault, *Discipline and Punish*, Random House, 1975.
6. G. Iachello et al., "Privacy and Proportionality: Adapting Legal Evaluation Techniques to Inform Design in Ubiquitous Computing," *Proc. SIGCHI Conf. Human Factors in Computing Systems*, ACM Press, 2005, pp. 91–100.
7. J. Gemmell et al., "Passive Capture and Ensuing Issues for a Personal Lifetime Store," *Proc. 1st ACM Workshop on Continuous Archival and Retrieval of Personal Experiences (CARPE 04)*, ACM Press, 2004, pp. 48–55.
8. S. Intille et al., *New Challenges for Privacy Law: Wearable Computers that Create Electronic Digital Diaries*, tech. report, MIT Dept. of Architecture House_n, Sept. 2003.
9. R. Agrawal et al., "Privacy-Preserving Data Mining," *ACM SIGMOD Record*, vol. 29, no. 2, 2000, pp. 439–450.
10. L. Sweeney, "K-Anonymity: A Model for Protecting Privacy," *Int'l J. Uncertainty, Fuzziness and Knowledge-based Systems*, vol. 10, no. 5, 2002, pp. 557–570.

11. R. Agrawal et al., "Hippocratic Databases," *Proc. 28th Int'l Conf. Very Large Databases (VLDB 02)*, Morgan Kaufmann, 2002, pp. 143–154.
12. S. Lederer et al., "Personal Privacy through Understanding and Action," *Personal Ubiquitous Computing*, vol. 8, no. 6, 2004, pp. 440–454.
13. M. Langheinrich, "Privacy by Design: Principles of Privacy-Aware Ubiquitous Systems," *Proc. Ubiquitous Computing 3rd Int'l Conf. (UbiComp 01)*, LNCS 2201, Springer, 2001, pp. 273–291.
14. Y. Duan and J. Canny, "Protecting User Data in Ubiquitous Computing," *Privacy Enhancing Technologies*, LNCS 3424, Springer, 2004, pp. 273–291.
15. S. Rizvi et al., "Extending Query Rewriting Techniques for Fine-Grained Access Control," *Proc. SIGMOD*, ACM Press, 2004, pp. 551–562.
16. Y. Li et al., "Design and Experimental Analysis of Continuous Location Tracking Techniques for Wizard of Oz Testing," *Proc. SIGCHI Conf. Human Factors in Computing Systems*, ACM Press, 2006, pp. 1019–1022.
17. N. Khossainova et al., "Probabilistic RFID Data Management," tech. report UW-CSE-07-03-01, Univ. of Washington, Computer Science and Engineering Dept., Mar. 2007.
18. J. Hong et al., "An Architecture for Privacy-Sensitive Ubiquitous Computing," *Proc. Mobisys*, ACM Press, 2004, pp. 177–189.
19. D. Povey, "Optimistic Security," *Proc. 1999 Workshop on New Security Paradigms*, ACM Press, 1999, pp. 40–45.
20. G. Iachello et al., "Prototyping and Sampling Experience to Evaluate Ubiquitous Computing Privacy in the Real World," *Proc. SIGCHI Conf. Human Factors in Computing Systems*, ACM Press, 2006, pp. 1009–1018.
21. S. Consolvo et al., "Location Disclosure to Social Relations," *Proc. SIGCHI Conf. Human Factors in Computing Systems*, ACM Press, 2005, pp. 81–90.
22. A. Kapadia et al., "Virtual Walls," *Proc. Pervasive*, LNCS 4480, Springer, May 2007, pp. 162–179.

For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.



Travis Kriplean is a doctoral student in computer science and engineering at the University of Washington. His primary research interests lie in human-computer interaction, specifically in computer-supported cooperative work. Aside from working on privacy issues for the RFID Ecosystem, he works on supporting public deliberation in the urban-planning domain through the UrbanSim project. He received his BS in computer science and sociology from the University of Wisconsin. Contact him at 101 Paul G. Allen Center, CSE Dept., Univ. of Washington, Box 352350, Seattle, WA 98195-2350; travis@cs.washington.edu.



Evan Welbourne is a doctoral student in computer science at the University of Washington. His research interests lie at the intersection of pervasive computing, databases, and human-computer interaction. He's the lead graduate student on the RFID Ecosystem project and has worked on a variety of pervasive computing projects at the University of Washington, Intel Research, and Microsoft Research. He received his MS in computer science from the University of Washington. Contact him at 101 Paul G. Allen Center, CSE Dept., Univ. of Washington, Box 352350, Seattle, WA 98195-2350; evan@cs.washington.edu.



Nodira Khossainova is a doctoral student in the University of Washington's database group. Her research interests include the management of uncertain data from sensors such as RFID antennas. More specifically, she focuses on the areas of extracting high-level events from raw data and of data privacy. She received her Bachelor's degree in computer science from the University of Auckland in New Zealand. Contact her at 101 Paul G. Allen Center, CSE Dept., Univ. of Washington, Box 352350, Seattle, WA 98195-2350; nodira@cs.washington.edu.




Vibhor Rastogi is a doctoral student in the University of Washington's database group. His research interests include database privacy, security, and probabilistic databases. He obtained his master's in computer science from the University of Washington. Contact him at 101 Paul G. Allen Center, CSE Dept., Univ. of Washington, Box 352350, Seattle, WA 98195-2350; vibhor@cs.washington.edu.



Magdalena Balazinska is an assistant professor of computer science and engineering at the University of Washington. Her research interests focus on data management systems for distributed data streams, stream data archives, and dirty, uncertain sensor data. She received her PhD in computer science from the Massachusetts Institute of Technology. She's a member of the IEEE, the ACM, and ACM SIGMOD. Contact her at 101 Paul G. Allen Center, CSE Dept., Univ. of Washington, Box 352350, Seattle, WA 98195-2350; magda@cs.washington.edu.



Gaetano Borriello is a professor of computer science and engineering at the University of Washington. He also founded Intel Research Seattle, where he launched the lab on applications of ubiquitous computing technology to healthcare and elder care, in particular. His research interests include location-based systems, sensor-based inferencing, and tagging objects. He received his PhD in computer science from University of California Berkeley. He's an associate editor in chief of *IEEE Pervasive Computing*. Contact him at 101 Paul G. Allen Center, CSE Dept., Univ. of Washington, Box 352350, Seattle, WA 98195-2350; gaetano@cs.washington.edu.



Tadayoshi Kohno is an assistant professor of computer science and engineering at the University of Washington. His research interests are computer security, privacy, and cryptography. He received his PhD in computer science at the University of California, San Diego. He is a member of the ACM, the International Association for Cryptological Research, and the IEEE. Contact him at 101 Paul G. Allen Center, CSE Dept., Univ. of Washington, Box 352350, Seattle, WA 98195-2350; yoshi@cs.washington.edu.



Dan Suciu is a professor of computer science and engineering at the University of Washington. His research is data management, emphasizing topics arising from sharing data on the Internet. He received his PhD in computer science from the University of Pennsylvania. Contact him at 101 Paul G. Allen Center, CSE Dept., Univ. of Washington, Box 352350, Seattle, WA 98195-2350; suciu@cs.washington.edu.