Physical Access Control for Captured RFID Data

Travis Kriplean, Evan Welbourne, Nodira Khoussainova, Vibhor Rastogi,
Magda Balazinska, Gaetano Borriello, Tadayoshi Kohno, Dan Suciu
University of Washington, Seattle, WA
{travis, evan, nodira, vibhor, magda, gaetano, yoshi, suciu }@cs.washington.edu

1. Introduction

Radio Frequency Identification (RFID) technology has become popular as an effective and low-cost solution for tagging and wireless identification. Although early RFID deployments were focused primarily on industrial settings, successes have led to a boom in more personal, pervasive applications such as reminders [3] and elder-care [19]. RFID holds the promise of enhancing many everyday activities, but it also raises great challenges, in particular, with respect to privacy and security.

At the University of Washington, we have 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 [23]. It is a capture-and-access system: all data is streamed from the readers into a central database, where it is accessible to applications. Our goal with the RFID Ecosystem is to provide a laboratory for long-term research in privacy and security, as well as applications, data management, and systems issues for RFID-based, community-oriented pervasive computing.

RFID systems collect data as a stream of triples having the form: `(IDENTITY, LOCATION, TIME)`. Each triple captures the time and location where a reader sighted an RFID tag near one of its antennas. This fine-grained location information enables many new types of applications and services. For example, a reminder service can alert users when they forget to take an item with them as they go home for the day [3]. Alternatively, a personal digital diary can record the places a user visits, whom she has contact with, and what activities she is involved in so that she can later study trends in her use of time. The capture of this information over time, however, also raises significant privacy concerns. Imagine, for example, a system that allows your peers and superiors to query the length of your coffee breaks or how much time you spend socializing with colleagues. A fundamental design decision for a pervasive RFID system is thus how best to balance system utility with privacy.

While RFID security is now a vibrant research area and many protection mechanisms against unauthorized RFID cloning and reading attacks are emerging [13], little work has been done to address the complementary issue of privacy for RFID data *after* it has been captured and stored by an authorized system. In this article, we discuss the problem of privacy for personal RFID data and investigate one particular issue: peer-to-peer privacy. We assume a system with trusted owners and administrators, and focus on ways to constrain peers' access to information about one another. Our contribution is an access control policy called Physical Access Control (PAC). 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. Though it strictly limits information disclosure, the database view afforded by PAC is still useful because it augments users' memory of places, objects, and people. We posit that PAC is appropriate as a *default* level of access control because it models the physical boundaries found in everyday life. PAC is also a useful starting point for research in principled information disclosure–it can be carefully relaxed to increase utility without disclosing extensive private information. In this article, we focus on the privacy, utility, and security issues raised by its implementation in the RFID Ecosystem.

2. Privacy and Utility in Pervasive Architectures

Eighteenth-century legal philosopher Jeremy Bentham first described the perfect architecture for surveillance: the *panopticon*, a prison where cells are arranged 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 upon another's data. Under these conditions, the "state of conscious and permanent visibility [assures] the automatic functioning of power" [6] because the individual 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 fundamental conviction regarding privacy in the RFID Ecosystem is that privacy needs to be designed into the system from the ground up. The challenge in architecting privacy into a pervasive sensing system is in providing enough utility to support the desired suite of applications while carefully controlling what information should be disclosed, to whom, how, and under what conditions. To effectively make such decisions, the privacy and utility trade-off needs to be treated holistically [10]. In particular, perspectives and methods from computer security, databases, human-computer interaction, and the social sciences should be brought to bear on a proposed privacy mechanism. In this way, we can understand the mechanism's security vulnerabilities, how well it matches users' expectations of privacy, the ease with which it is understood, and the utility it affords.

Most pervasive sensing systems represent one of two architectural models: wearable or infrastructure. Generally speaking, each model makes different tradeoffs with respect to privacy and utility. In the wearable model, sensors and data are processed and stored on devices owned and worn by the user. This model is embodied by MyLifeBits [8], in which users wear microphones, video cameras, and other sensors which continually record sensor data. Such systems can put the device wearer in control if data is stored locally and no information is disclosed without the user's explicit consent. However, these "perfect memory" systems pose privacy concerns for others who encounter the user but do not consent to information capture [12]. Plausible deniability is lost–although human memory is lossy, captured sensor data is not.

In contrast, the infrastructure model has a central authority which manages sensor data on behalf of users. This model gives rise to the threat of permanent visibility described earlier. Yet systems like the RFID Ecosystem and Aura [7] adopt the infrastructure model because it enables much richer services through data and resource sharing, and at less expense because cost is amortized over many users. Moreover, reliance on a central database allows a system to leverage database security and privacy techniques such as privacy-preserving data mining [1] and K-anonymity [22]. These techniques provide privacy-preserving statistical queries which nicely complement and extend the utility offered by careful access control for point queries (*e.g.*"When did I see Bob today?"). Both statistical and point queries should be employed in a principled, privacy-sensitive framework for managing data (*e.g.*a Hippocratic database [2]), however, our focus here is on point queries under one possible access control policy.

3. Physical Access Control

Many access control policies for captured RFID data could be defined. Policies might allow a manager to track the location of employees during work hours, support staff to locate objects in inventory, and individual users to grant conditional tracking permissions to their friends. However, these policies present problems when applied as a default. In particular, the managerial example raises surveillance concerns; the object finder can be abused to track people; and user-defined policies are usually plagued by lack of foresight and vigilance [17]. On the other hand, a restrictive policy that allows a user to access only their own data precludes many interesting applications that might be safe to provide.

We propose that a default access control policy should model spatial privacy in everyday life. To this end, we present PAC, an attempt to realize the spatial privacy features of wearable systems within the context of an infrastructural architecture. PAC restricts the information a user can obtain to that which the user could have observed in person. Specifically, a user can "see" his objects and other users when and where he was physically present, but not others' objects as they may be concealed from sight.

PAC places upper and lower bounds on accessible information: no more information is available than what a user could have observed in person, yet a persistent record of all encountered objects, persons, and events is kept. Langheinrich [16] argues that proximity-based disclosure could limit the threat of surveillance – PAC implements this proposal. We also argue that PAC provides an intuitive and easily understandable potential "flow" of captured information [17]. Furthermore, PAC's ethic respects Duan and Canny's [5] *data discretion principle* which states that users should have access to recorded media when they were physically present and should not if they were not present.

Although PAC is conservative in the information revealed, its memory-like view of the data provides useful service primitives. For example, a user's query on the location of her lost object will return the location where she last saw it, a likely answer. Similarly, queries over the history of a user's own activities could enable all the applications described in the introduction. We argue that this balance between privacy and utility makes PAC a suitable default access control policy for a pervasive RFID deployment. Moreover, privacy-preserving extensions and 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 as needed for useful applications.

3.1. Implementation

PAC defines a database view consisting of only those triples `(IDENTITY, LOCATION, TIME)` representing

a user's own location and the locations of users and objects he could *plausibly* have seen. User queries on collected data are always returned in reference to this PAC view rather than the complete data (*i.e.*PAC employs a Truman model [21]). For example, if a user asks "how many people were on the fourth floor yesterday", the system effectively responds "you saw 5 people on the fourth floor" instead of "you saw 5 out of the 11 total people on the fourth floor".

An implementation of PAC 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 called *mutually-visible* if they share an unobstructed line of sight. Every such instance of mutual-visibility is called a *visibility-event*. Consider the scenario presented in Figure 1. The figure shows a snapshot of six users going about their daily routines; all visibility-events are enumerated in the table.

This definition of mutual-visibility is an ideal that must be approximated using captured sensor data. In the RFID Ecosystem, the mutual-visibility computation must incorporate the spatio-temporal relationships between RFID tag reads by antennas. Before presenting a formal definition of mutual-visibility for the RFID Ecosystem, we formally describe the spatial and temporal relations between tag reads.

Spatial. There are two challenges in determining unobstructed line of sight between two tag reads. First, the exact location of a sighted tag is unknown; the location of the antenna is used as a proxy for the tag's location. Second, two tags may be mutually-visible, yet read by two different antennas. This motivates the definition of *mutually-visible antennas* – pairings of antennas A_1 and A_2, such that a tag read at A_1 can be interpreted as mutually-visible with a tag read at A_2.

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

A formal definition of mutual-visibility can now be expressed 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.

The above definition yields the visibility-events marked in Figure 1. In this scenario, antennas 18 and 15 are mutually-visible to one another, and therefore the system will correctly interpret A and C as mutually-visible at $t = 3$. In contrast, antennas 11 and 12 are not mutually-visible and therefore E and C will not be considered mutually-visible. Note how varying Δ tunes the strictness of mutual-visibility. In Figure 1, $\Delta = 1$ and therefore D and A are never detected as mutually-visible; however, if $\Delta = 2$, then they are mutually-visible during time $5 - 7$.

3.2. Users, objects, and ownership

We make a distinction between *user-tags* and *object-tags*. A user can see the location of an object only when the user and object are mutually-visible *and* the user owns that object. The ownership restriction is required because RFID tags may be read 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, as our goal is to study community-oriented systems, future work will need to investigate how PAC can operate with shared objects.

3.3. 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 manner that minimizes false visibility-events. Our approach has been to label each pair of antennas with the probability that two tags in the two antennas' respective areas of coverage 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 pairs of points 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:

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

A_i, A_j are then considered mutually visible if $\text{visibility}(A_i, A_j) \geq \tau$, where τ sets the lower bound on the fraction of pairs of points that need to 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 because it minimizes 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 $\tau = .84$.

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

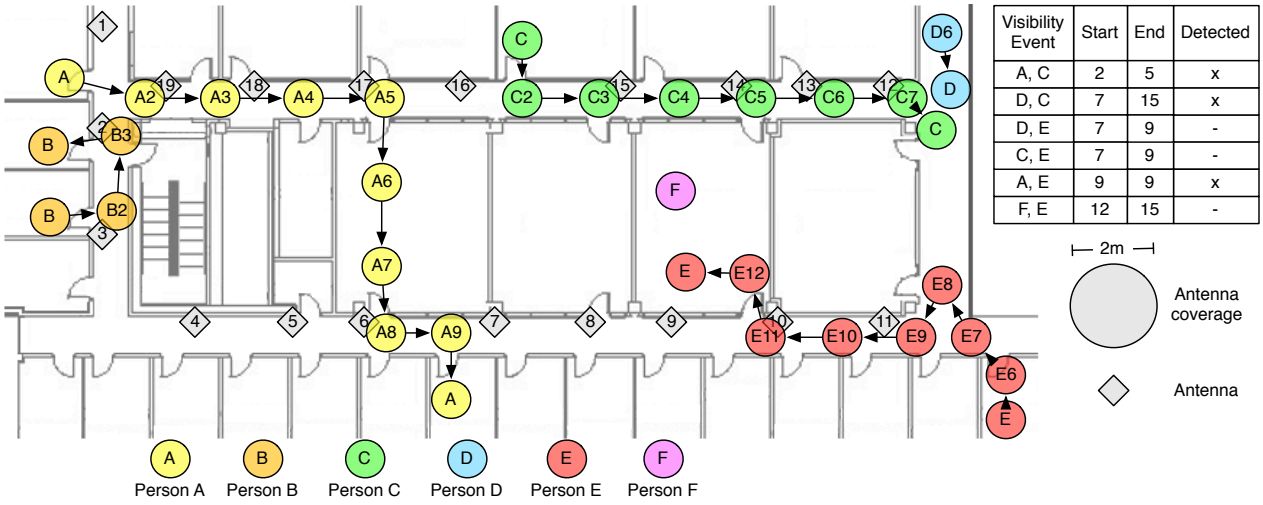

Visibility Event	Start	End	Detected
A, C	2	5	x
D, C	7	15	x
D, E	7	9	-
C, E	7	9	-
A, E	9	9	x
F, E	12	15	-

Figure 1. Six people moving over the course of 15 time steps. Each user's location is annotated with the timestamp at which they move to the space. No timestamp indicates that they remain in that space. Visibility-events between users are shown in the table. The detected column of the table shows the visibility-events inferred by the system when $\Delta = 1$ and antennas $1-3$, $4-11$, and $12-19$ are defined to be mutually-visible.

4. Feasibility of PAC

In adapting PAC for use in the RFID Ecosystem we have assumed a well-behaved and lossless model for our RFID equipment. We would like to determine how PAC performs in a real deployment with antennas which may not behave as expected–the uncertainty of RFID in practice could have adverse effects on both privacy and utility. For example, privacy violations could occur if an antenna reads a tag beyond its expected range (possibly causing false visibility-events); utility is degraded when visibility-events are either missed or labeled with the wrong timestamp. Here we discuss our experiment to evaluate PAC in practice.

4.1. Experimental setup and methodology

We evaluate PAC over a set of user scenarios which cover some ways visibility-events can occur. For each scenario, we enact multiple trials and collect the corresponding stream of *raw* tag reads captured by antennas on each trial. For each trial, we also capture ground truth location data which we subsequently process using a simulator that models a well-behaved, lossless RFID deployment and produces a stream of *simulated* tag reads.

4.1.1 Representative scenarios

There are many ways in which a visibility-event may occur. While not exhaustive, we have defined four scenarios that represent common types of visibility-events. In the *personal objects* scenario, a single user walks around the halls carrying six tagged objects on various parts of his body, including inside a duffel bag. The second is the *glimpses* scenario in which one user stands at the end of a hallway while another user B enters the opposite end of the hallway from an office and quickly walks around the corner. Third, in the *walking together* scenario, two users walk around the hallways together. Finally, two users pass one another while walking in opposite directions for the *passing by* scenario.

4.1.2 Data collection and ground truth

The scenarios gave a rough script for experimenters to follow. To accurately collect ground truth for each trial, experimenters used tablet PCs with a map-based data collection tool [18]. The tool allows current location to be captured by moving a cross-hair to the current location on a map using the stylus. Each trial with the tool produced an XML trace of timestamped latitude and longitude coordinates which could be fed to the simulator to produce the simulated tag reads. Based on experience with our RFID deployment, we set the antenna coverage area to be a circle with a radius of two meters (six feet) about the antenna.

4.1.3 Comparing visibility-events

For each pair of tags X and Y involved in a given trial, we compare the visibility-events detected in the raw and simulated data. Let S and R denote the set of all visibility-events of 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 timestamp of v as $\mu_v = \frac{t_Y + t_X}{2}$. A visibility-event s in S is said to *occur* in R if there exists an r in R such that $|\mu_r - \mu_s| \leq \epsilon$. A visibility-event in S thus also occurs in R when there is a visibility-event in R whose timestamp is within ϵ of the timestamp of the visibility-event in S. In our experiments, $\epsilon = 1$ second.

To gain a measure of the privacy and utility achieved, we compare the visibility-events detected in the raw and simulated data in terms of recall and precision. Recall measures utility. It is the fraction of simulated visibility-events that also occurred in the raw data. A recall equal to 1 indicates that all visibility-events in the simulated data were accurately captured by the raw data. A recall below one indicates that some visibility-events are missed due to missed tag reads. In contrast, precision is a measure of privacy. Precision is the fraction of detected visibility-events that also occurred in the simulated data. A precision below one indicates privacy loss because the system detects false visibility-events.

4.1.4 Computing results

Each scenario is performed ten times. After each trial, we calculate precision and recall for the visibility-events of every pair of tags. We then compute the mean and standard deviations of the precision and recall across all the trials for these visibility-events. In the personal objects scenario, there are visibility-events between the experimenter and each of his objects. In this case, we give the average and standard deviation across all of these pairings.

4.2. Results

Figure 2 shows the precision and recall for all four scenarios. The outcome of the experiment 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 there was little information leakage. This also verifies the integrity of our data collection procedure, as such high precision requires the generation of correct ground truth input.

Utility. Recall suffers when visibility-events are not detected because antennas fail to read tags. This can happen for a variety of reasons, such as the properties of the material to which the tag is affixed, as well as the tag's orientation with respect to the antennas. Experimenters wore the tags hanging from their shirt or pants. This resulted in high read rates for the user-tags (recall between 90-95%) in all of the experiments and correspondingly high detection of visibility-events between user-tags ($> 80\%$).

In the personal objects scenario, we observe lower recall because the antennas couldn't consistently read the tags which were in pockets or the bag. There exist, however, several algorithms and tools which 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 a research prototype called PEEX [15]. PEEX corrects raw RFID data through the use of integrity constraints. We ran PEEX with a few integrity constraints which capture logical, physical relations between objects and people (*e.g.*an object cannot move by itself). Figure 2 shows that PEEX significantly improves the recall in the personal objects and walking together scenario (t-test with $p < 0.001$), without affecting the precision.

Overall, our results show that it is feasible in practice to employ PAC. Despite the noisy, lossy, and inaccurate nature of real RFID data, PAC effectively limits information disclosure while providing good system utility. User-user interactions are captured quite well while user-object interactions are hampered by the inherent unreliability of RFID. However, even simple applications of cleaning tools such as PEEX suffice to significantly improve capture of user-object interactions.

5. "Misplaced" user-tags

Our implementation of PAC 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 the proximity of Bob (and would grant her access to his data). It should be noted that such errors are not possible with object-tags because they can only be mutually visible with their owner.

We are developing several mechanisms for addressing "misplaced" user-tags. Our discussion focuses 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 to an adversary of mounting an attack against users' privacy should outweigh the benefits to the adversary.

Detection. The threat of detection can deter malicious activities because detection might lead to social sanction and/or punitive measures for the offending party. We are exploring two classes of detection mechanisms. First, the RFID Ecosystem could automatically detect anomalies in the movement of a user tag. The system could trigger an alert if Alice's user-tag is always mutually-visible with Bob or one of Bob's objects, like his backpack, or if Alice's user-tag has been in an unusual location for too long.

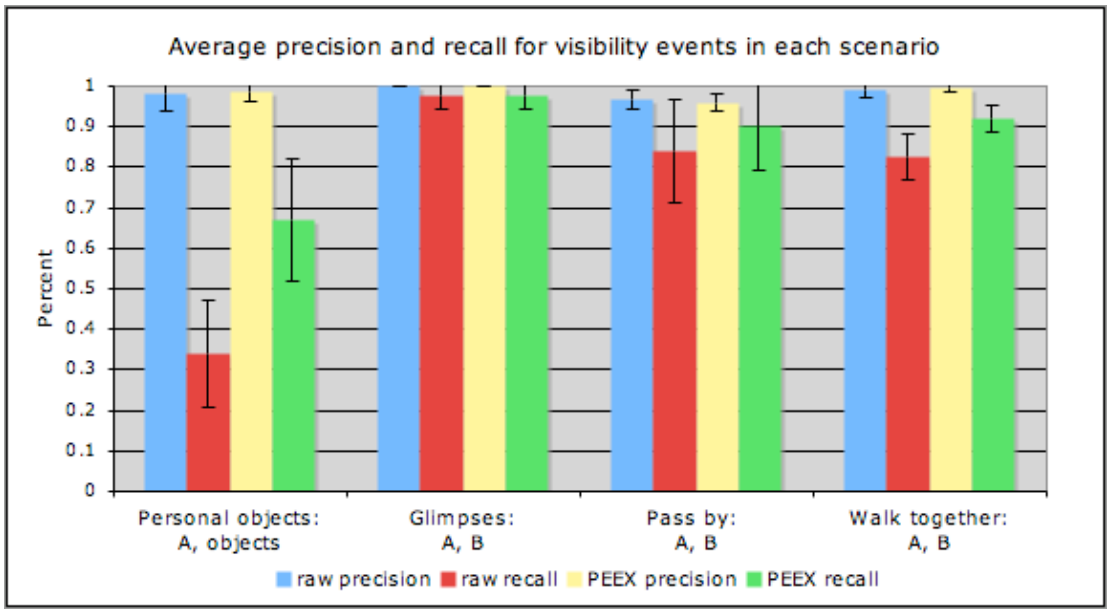

Figure 2. Experimental results. Each group of results represents a single visibility-event between two tags, except for the first, which is the average of A's visibility-events with their objects.

Second, since non-visually 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 the user in anomaly detection. For example, Bob could detect Alice's maliciously planted RFID user-tag if he is in an elevator alone but the front panel of the elevator says that there are two occupants.

Prevention. We are exploring two classes of prevention mechanisms: making attacks too costly or inconvenient for the adversary, and periodically verifying that the RFID user-tags are actually in the possession of the appropriate user.

One method for increasing the cost to an adversary is to combine user-tags with expensive or essential devices, such as cell phones or employee badges. A second method could be to stop capturing reads for a user-tag if that tag becomes separated from its legitimate owner. 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, e.g., by detecting when the tag enters the legitimate user's office.

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 could not have observed. Direct attacks on the RFID hardware (*e.g.*cloning tags) are also possible, but not considered here. See Juels [13] for a survey of the state of the art in preventing such attacks.

6. Future work

In this section, we briefly present a number of other areas of future work.

Principled relaxations of PAC. Other access control mechanisms can be defined on top of PAC to provide additional information when appropriate.

First, user-defined access control rules are important for enabling applications that rely on shared context between users, such as location-shared buddy lists [9]. Making information available without physical proximity would be justified because because users opt-in by explicitly granting permission.

Another related relaxation is through socially-situated events. For example, there might be an augmented calendar system that would allow a user to query for the location of a meeting's missing attendees during the time of the scheduled meeting. Such a relaxation would allow users to gain useful information at particular times that would be socially acceptable.

A third class of relaxations may involve mediation with the system. For example, an owner of a lost object might request the location of an object. The system could choose to reveal the object's location to the requester; alternately, it may send an email conveying that the owner is looking for their object to the person most recently detected to have moved the object. By involving the system or an administrator, this relaxation may prevent abuse by not revealing sensitive information, while allowing useful actions to be taken.

Finally, an opportunistic access control scheme [20] might be employed to allow users to access private information in rare circumstances such as emergencies. These actions are logged and investigated by administrators in order to decide if they were legitimate. Determining the conditions and the frequency under which this access mechanism might be used is an open problem.

User studies. Our claim is that PAC is an intuitive policy which will match users' expectations of privacy in everyday life. This assertion needs to be empirically validated through user studies. Moreover, while there has been a number of studies examining user privacy expectations for captured audio and video data (*e.g.* [11]) and disclosure of information to others across potentially great distances (*e.g.* [4]), there have been few studies examining how physical space factors into people's expectations of privacy for captured RFID data. Kapadia, et al. [14] have begun to explore the use of spatial metaphors, but further work is necessary. It should be noted that the definition and implementation of PAC has proven very useful in obtaining "minimal risk" Institutional Review Board approval for user studies.

Probabilistic Data. As demonstrated earlier, the PEEX system cleans data to enhance utility, but PEEX can also produce probabilistic data. In this case, each tuple has an associated probability which represents the system's confidence about its validity. Implementing PAC in a probabilistic context leads to a challenging problem which we illustrate below.

Suppose a user A asks: "Is user B currently at location L?". If A is at L, then PAC allows the correct answer. If A is not 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 in fact p_B. Yet the system cannot return p_B in the case that A is not at L. One approach is to return $p_A \cdot p_B$, the probability that both A and B are at L. This reveals too much, however – 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 fulfil this requirement.

7. Conclusion

The goal of the RFID Ecosystem is to enable research that provides the community (including businesses and policy makers) with examples of effective RFID systems that balance utility with privacy by design. PAC is a first step in this direction because it allows further experiments to be performed in a privacy respecting way while remaining amenable to utility-enhancing extensions. Our experiments show that PAC is a practical solution that works well in a real-world unreliable sensor architecture.

8. Acknowledgements

We would like to thank Garret Cole, Patricia Lee, Caitlin Lustig, Robert Spies, and Jordan Walke, with special thanks to Caitlin for her work on the simulator we used. We would also like to acknowledge the University of Washington CoE. This research has been funded by the NSF CRI 0454394, IIS-0428168 and IIS-0415193 grants.

References

[1] Agrawal, R. et al. Privacy-preserving data mining. *ACM SIGMOD Record*, 29, 2000.
[2] Agrawal, R. et al. Hippocratic databases. In *Proc. of VLDB*, 2002.
[3] Borriello, G. et al. Reminding about tagged objects using passive RFIDs. In *Proc. of Ubicomp*, volume 3205. Springer, 2004.
[4] Consolvo, S. et al. Location disclosure to social relations. In *Proc. of CHI*. ACM Press, 2005.
[5] Duan, Y. et al. Protecting user data in ubiquitous computing. In *Privacy Enhancing Technologies*, volume 3424. Springer, 2004.
[6] M. Foucault. *Discipline and Punish*. Random House, 1975.
[7] Garlan, D. et al. Project aura. *IEEE Pervasive Computing*, 1(2), Apr. 2002.
[8] Gemmell, J. et al. Passive capture and ensuing issues for a personal lifetime store. In *CARPE04*, 2004.
[9] Hong, J. et al. An architecture for privacy–sensitive ubiquitous computing. In *Proc. of Mobisys*, June 2004.
[10] Iachello, G. et al. Privacy and proportionality. In *Proc. of CHI*. ACM Press, 2005.
[11] Iachello, G. et al. Prototyping and sampling experience to evaluate ubiquitous computing privacy in the real world. In *Proc. of CHI*. ACM Press, 2006.
[12] Intille, S. et al. New challenges for privacy law: Wearable computers that create electronic digital diaries. Technical report, MIT House n, Sept. 2003.
[13] Juels, A. RFID security and privacy: a research survey. *IEEE Journal on Selected Areas in Communications*, 24, Feb. 2006.
[14] Kapadia, A. et al. Virtual walls. In *Proc. of Pervasive*. Springer-Verlag, May 2007.
[15] Khoussainova, N. et al. Probabilistic RFID data management. *UW-CSE-07-03-01*, Mar. 2007.
[16] M. Langheinrich. Privacy by design. In *Proc. of UbiComp*. Springer-Verlag, 2001.
[17] Lederer, S. et al. Personal privacy through understanding and action. *Personal Ubiquitous Comput.*, 8(6), 2004.
[18] Li, Y. et al. Design and experimental analysis of continuous location tracking techniques for wizard of oz testing. In *Proc. of CHI Notes*, May 2006.
[19] Patterson, D. et al. Fine-grained activity recognition by aggregating abstract object usage. In *Proc. of ISWC*, Oct. 2005.
[20] Povey, D. Optimistic security. In *NSPW99*, 1999.
[21] Rizvi, S. et al. Extending query rewriting techniques for fine-grained access control. In *Proc. of SIGMOD*. ACM Press, 2004.
[22] Sweeney, L. k-anonymity: A model for protecting privacy. *IJUFKS02*, 10(5), 2002.
[23] Welbourne, E. et al. Challenges for pervasive rfid-based infrastructures. In *Proc. of PERTEC 2007 Workshop*, Mar. 2007.