

Soft Blocking: Flexible Blocker Tags on the Cheap

Ari Juels and John Brainard

RSA Laboratories
Bedford, MA 01730, USA
e-mail: {ajuels,jbrainard}@rsasecurity.com

Abstract. A “blocker” tag is a form of privacy-enhancing radio-frequency identification (RFID) tag. It operates by interfering with the protocol in which a reader communicates individually with other RFID tags. While inexpensive to manufacture in quantity, blockers are nonetheless special-purpose devices, and thus introduce level of complexity that may pose an obstacle to their deployment.

We propose a variant on the blocker concept that we call *soft blocking*. This involves software (or firmware) modules that offer a different balance of characteristics than ordinary blockers. Soft blocking offers somewhat weaker privacy enforcement that is essentially voluntary or internally auditable (much like P3P). It has the significant advantage, however, of relying on standard (or very slightly modified) RFID tags. Additionally, soft blocking also offers the possibility of flexible privacy policies in which partial or scrubbed data is revealed about “private” tags, in lieu of the all-or-nothing policy enforced by a blocker.

We show, moreover, how the correct functioning of a soft-blocker system may be rendered externally auditable with minor modifications to the basic tag-reading protocol. We also briefly discuss the special, attractive approach of *unblocking*, a soft-blocking variant that permits an “opt-in” approach to consumer privacy.

Key words: ALOHA, blocker tags, privacy, RFID tags, tree-walking

1 Introduction

Deployment of RFID tags as a form of next-generation barcode seems imminent, thanks to the promise of enhanced efficiency and accuracy of tracking in industrial supply chains [7, 10]. Researchers, privacy advocates, and the popular press, however, have for some time recognized the possibility of privacy threats in the deployment of Radio-Frequency Identification (RFID) tags. Concretely, concern centers on: (1) The threat of physical tracking enabled by the clandestine reading of unique serial numbers in tags and (2) Surreptitious inventory-taking of the objects carried by individuals, since tags are likely to carry information about the items to which they are attached [4, 19, 21, 23]. The challenge in meeting these threats is that basic RFID tags, particularly the tags likely to see the widest

deployment, have severely constrained computational resources. Standard cryptographic algorithms, in particular, are likely to be well beyond their reach for some time to come [21].

With these constraints in sight, several recent papers have proposed privacy-enhancing techniques for RFID tags. Juels and Pappu [17] consider a purported plan by the European Central Bank to embed RFID tags in Euro banknotes [9]. They propose a privacy-protecting scheme in which RFID tags carry ciphertexts on the serial numbers of banknotes. These ciphertexts are subject to re-encryption by computational devices in shops, thereby rendering multiple appearances of a given RFID tag unlinkable. The Juels/Pappu scheme, however, assumes a single verifying entity – namely a law-enforcement organization – and is not obviously extensible to the multi-verifier systems likely in commercial and consumer environments.

A scheme of Golle, Jakobsson, Juels, and Syverson [15] builds on this idea with a primitive known as universal encryption, a special extension of the El Gamal cryptosystem [13] in which re-encryption is possible without knowledge of public keys. The Golle *et al.* approach is geared toward general consumer use, as it does not require a centralized verifying entity. It has the drawback, though, of requiring an infrastructure of agents capable of performing public-key-based re-encryption for privacy protection of RFID tags, and is thus of primarily academic interest at this point.

Weis, Sarma, Rivest, and Engels [24] also propose a collection of privacy-enforcement ideas for RFID tags in general environments. First, they identify the problem of attacks based on eavesdropping rather than active tag queries. Recognizing that transmission on the tag-to-reader channel is much weaker than that on the reader-to-tag channel, they propose protocols in which tag-identifying information is concealed on the stronger channel. Weis *et al.* also propose privacy-preserving schemes for active attacks. One scheme involves the use of a hash function to protect the key used for read-access to the tag. Another includes use of a pseudo-random number generator to protect tag identities. In a nutshell, their idea is for the tag to output the pair $(r, PRNG(ID, r))$, where r is a counter, ID is the secret tag identifier and $PRNG$ denotes a pseudo-random number generator. A verifier must perform an expensive brute-force lookup in order to extract the ID from such an output. The authors note that this drawback probably limits applicability of the idea to small systems. They also note that it is unclear how and when adequate pseudo-random number generators can be deployed on inexpensive RFID tags.

Juels [16] proposes a system known as “minimalist cryptography,” in which a tag carries multiple, pre-programmed pseudonyms. By releasing different pseudonyms during different reading sessions, the tag prevents tracking by unauthorized entities. An authorized entity, by contrast, has information about the linkage among pseudonyms. An additional feature of this proposal is a “throttling” mechanism, whereby the tag imposes a delay on the release of pseudonyms so as to prevent an adversary from harvesting multiple pseudonyms in rapid succession. In

the full-blown protocol, tags and readers authenticate to one another, and an authenticated reader may replace the set of pseudonyms on a tag.

Fishkin and Roy [11] have shown in preliminary experiments that the signal-to-noise ratio of the data received by a tag from a reader gives a rough indication of read distance. They propose as a privacy-enforcing measure that tags might be constructed to evince trust as a function of distance; in particular, when apparently scanned at a distance, tags might either release partial information or refuse to respond to queries. Follow-up work should provide a notion of how practical this approach is, and whether distance is indeed a suitable metric for trust.

A rather different, complementary perspective on privacy for RFID tags is that of Garfinkel [14], who elaborates a policy for consumer privacy-protection in the form of a proposed “RFID Bill of Rights.” Proposed there are: The right of the consumer to know what items possess RFID tags and the right to have tags removed or deactivated upon purchase of these items, the right of the consumer to access of the data associated with an RFID tag, the right to access of services without mandatory use of RFID tags, and finally the right to know to when, where, and why the data in RFID tags is used. A position statement issued by a number of civil libertarians and privacy advocates offers a similar perspective [8].

As explained, the point of departure for our proposals in this paper is the work of Juels, Rivest, and Szydlo [18]. They describe a privacy-protection tool they call a “blocker” tag. This is an RFID tag that can obstruct reading of tag identifiers within a certain numerical range by simulating the presence of RFID tags bearing *all* identifiers in that range. Such obstruction is accomplished through non-standard interaction with the “tree-walking” or ALOHA protocols employed in current tag-reading standards [5, 22]. So as not to serve as a purely disruptive mechanism, the blocker may be accompanied by a form of privacy “zoning,” according to which only the reading of a certain subset of identifiers is disrupted. Thus, tag identifiers may be “zoned” while in the possession of manufacturers and retailers such that their reading may take place without impediment. Before tags are placed in the hands of consumers, their “zoning” may be changed so that identifiers cannot be read in the presence of a “blocker,” thereby enforcing the privacy of consumers. At the same time, the “blocker” concept offers more flexible privacy options than tag disablement: By deactivating a “blocker” or removing it from the vicinity of her tags, a consumer can still make use of tags.

All of the proposed technical mechanisms described above for enforcing RFID-tag privacy carry the burden of supplementary resource requirements. For example: (1) The Weis *et al.* [24] proposal requires cryptographic functionality in tags, which is beyond their current capabilities; (2) The “minimalist cryptography” approach [16] involves enhanced memory in an RFID tag, as well as rewrite capabilities (in the full protocol); (3) The proposals of [15, 17] involve an infrastructure of re-encryption devices; (4) ‘Blocker’ tags [18] are special-purpose

devices enhanced with a non-compliant protocol variant – although we feel that they are perhaps the most practical of these proposals.

Soft blocking is a system that offers somewhat weaker privacy protection than full-blown blocking, but with the benefits of greater policy flexibility and of no alteration to either standard tags or readers.

Organization

In section 2, we briefly outline the soft-blocker system and illustrate the concept by means of examples. We provide further architectural details in section 3. In section 4, we describe a modification permitting an external auditor to verify that a soft-blocker system is compliant with a particular privacy policy. We conclude in section 5.

2 Overview and Examples

In a nutshell, soft blockers simply express the privacy preferences of their owners to RFID readers. This approach requires audit mechanisms to enforce reader respect for these preferences.

To give a more detailed explanation, the soft blocker proposal involves displacement of “blocker”-like functionality from the tag-reader protocol to a software (or firmware) module, which we refer to as a *tag privacy agent* (abbreviated TaPA). The role of the TaPA is to filter tag data output by a reader prior to their transmission to other parts of the RFID system, primarily back-end applications. The TaPA thus enforces privacy-policy compliance on an internal basis. As we show later in the paper, it is possible to render the TaPA externally auditable.

In our proposed system, tags have associated *classifications*. These are labels or cues made available to the TaPA to determine what privacy policy should be implemented during a reading session. For example, a tag classified as a “blocker” might cause a TaPA to filter out sensitive tag data. A TaPA may determine classifications on the basis of tag identifiers themselves. E.g., a leading bit-pair of the form ‘11’ in a tag identifier might indicate that the tag is a blocker. Alternatively, the TaPA may learn the classification of a tag by performing a database or directory lookup.

Soft blocking is perhaps best illustrated by means of a series of examples. Our first example implements the essential functionality of the original blocker proposal of Juels, Rivest, and Szydlo [18].

Example 1. Tags have one of three classifications: “blocker,” “private,” or “public.” When a TaPA receives a collection of identifiers, it does the following: If there is a “blocker” tag present, then it returns the data released by “public” tags only. Otherwise, the TaPA returns the data of both “public” and “private” tags.

As a deployment example, we might consider a retail shop, such as a supermarket. Here, objects would be initially classified as “public.” When an item is

purchased, its classification is changed to “private” (upon use of an appropriate, access-controlled operation like PIN transmission). This action would provide the consumer with the option of using a soft blocker tag to protect the item from undesired reading once outside the shop.

Of course, blocking may be more refined than this. For instance, policy might rely on a combined view of several classifications of tags. Indeed, different TaPAs may respond in different ways to classification data. Another example illustrates this:

Example 2. A medical “unblocker” tag informs a TaPA that if the TaPA belongs to a certified medical device, then it may scan all tags of the “medical” classification – even if a blocker is present. Under other circumstances, a blocker tag indicates that “medical” tags should be treated as private.

Another possible refinement is for blocking policy to apply just to certain data fields in tags, as is the following example.

Example 3. An “enviroblocker” tag causes a TaPA to block all data fields on “private” tags except the field containing the recycling number on plastic containers.

Of course, privacy policies here may be arbitrarily general. Another possibility includes privacy policies that pre-process data so as to hide individual data elements while harvesting data for the purpose of computing aggregate statistics [3]. Yet another is to mix, i.e., randomly permute data so as to dissociate data from tag identities or from other linked data elements, essentially as a mix network or single mix server [6].

2.1 Unblocking

Unblocking, as in Example 2 above, is unachievable using the original blocking techniques. As such, it exemplifies the flexibility of soft blocking. Unblocking is particularly potent when considered as the basis for an “opt-in” approach to consumer privacy. This approach deserves special emphasis as a way of offering consumer privacy protection as a system default. It can also be useful in addressing tag reliability problems of the sort we mention in section 5.4. We illustrate the idea in our next example, which should be contrasted with Example 1 above.

Example 4. Tags have one of three classifications: “unblocker,” “private,” or “public.” When a TaPA receives a collection of identifiers, it does the following: If there is no “unblocker” tag present, then it returns the data released by “public” tags only. Only if a “unblocker” is present does a TaPA return the data of both “public” and “private” tags.

As in Example 1, objects in retail environments might be initially classified as “public.” When items are purchased, their classification is changed to “private.”

Consumers might have unblockers attached to their appliances at home so as to enable RFID use there. Additionally, the readers used for item returns in shops might have unblockers attached to them. (This limited use might be enforced using the audit techniques described below.) In all other circumstances, “private” RFID tags would be unreadable by default.

2.2 Reader licensing

Another strategy for simplification is to program readers to behave as though a soft blocker were always present. Viewed another way, readers may be subject to classification or “licensing”: particular readers may have special restrictions determining what tags they are permitted to scan and when, and these restrictions may be classified according to “licenses” corresponding to soft blocker types. To enhance its auditability, a reader might broadcast a specification of its license type before scanning tags in its vicinity, i.e., the reader may give an initial indication of what soft blocker (or blockers) it is simulating proximity to. An auditor can then determine whether a given reader is correctly licensed by checking whether the indicated license is consistent with the ownership and function of the reader. Using the techniques discussed in section 4, an auditor can further determine whether the reader is compliant with the policy specified by its claimed zoning. As an alternative, licensing may be determined by having soft blocker tags physically affixed to readers or their antennae; such tags might even provide a visual indication of their associated policy.

This approach is somewhat less flexible than the direct possession of soft blockers by consumers. On the other hand, the enforcement of policy directly on readers also entirely removes from consumers the onus of blocker-tag management, as in the following example.

Example 5. As a policy requirement, readers used at point-of-sale stations in retail stores might incorporate soft-blockers that forbid them to scan tags licensed as “private.” Such readers might further emit a “point-of-sale” code prior to scanning tags. This would permit an auditor to ensure that a reader carries a valid soft blocker.

Readers used for item returns in retail stores, in contrast, might be permitted to scan “private” tags. The power level of such readers, however, could be regulated. In particular, one can imagine that such readers might be permitted to emit only enough power to read tags at a very short distance when scanning in a privacy zone.

It would make sense for RFID readers used in the home to carry no restriction on their scanning abilities.

Given this regulatory environment for readers, and good policy enforcement, a consumer could be assured that her tags are only scanned at home or when held in proximity to readers designated for item returns in retail environment.

3 TaPA Architecture

We assume for simplicity in our description that all tags emit unique identifiers (or pseudonyms) in response to reader queries. Of course, our techniques are equally applicable to other forms of tag data. A TaPA then comprises three components:

- **Tag database DB :** This is a set $\{T_i, S_i\}_{i=1}^n$ of tag identifiers T_i and associated data S_i . S_i might include such data as tag classifications, history, “kill” PINs, and so forth. Naturally, DB may include various forms of access control, a topic we do not touch on here. DB might be a private information store or, alternatively, a public directory supported by, e.g., the Object-Name Service (ONS) [1].
- **Classification engine CE :** This is an algorithm that takes as input a set of tag identifiers $\{T_i\}_{i=1}^m$ of both ordinary tags and blockers, and also associated auxiliary data $\{D_i\}_{i=1}^m$ (e.g., tag manufacturer number, tag type, etc.) released by tags upon query. From some pre-established (possibly standardized) set Γ , it outputs tag classifications $\{\gamma_i\}$. CE may access DB in computing γ_i values. CE might in some cases consist of a simple algebraic function $f : \{0, 1\}^k \rightarrow \Gamma$ on k -bit identifiers.
- **Data filter DF :** This algorithm takes as input a set of tag identifiers $\{T_i\}_{i=1}^m$ and associated auxiliary data $\{D_i\}_{i=1}^m$ released by tags upon query. DF calls upon CE to obtain the classifications of tags it processes. It may additionally access DB while performing its computation. A DF emits filtered tag-associated data according to some policy P . (As P may be quite general, and dependent on external factors, e.g., time-of-day, we do not provide notation for it here.)

The data filter DF is the core of the TaPA. It is here that the privacy policy P of the TaPA is expressed and implemented. As an example, here is one way in which we might implement the basic system outlined in Example 4:

Example 6. A natural classification set here would be $\Gamma = \{\text{“blocker”}, \text{“private”}, \text{“public”}\}$. A very simple implementation would involve no DB . The classification engine CE might consist of a simple function f on the first two bits of a tag identifier that maps ‘00’ and ‘01’ to “blocker,” ‘10’ to “private,” and ‘11’ to “public.” In this case, any tag with an identifier possessing a leading ‘0’ would serve as a blocker. On input consisting of a set $\{T_i\}_{i=1}^m$, the data filter DF would perform the following steps:

1. Call on CE to compute classifications $G = \{\gamma_i = f(T_i)\}_{i=1}^m$;
2. If “blocker” $\in G$, then output $\{T_i : 1 \leq i \leq m, \gamma_i = \text{“public”}\}$; otherwise, output $\{T_i\}_{i=1}^m$.

As an alternative, it would be straightforward to implement this system using a database DB , rather than a classification function f . This database would associate with each tag identifier T_i its classification γ_i in Γ . The classification engine CE would then perform a simple lookup to determine the classification of a given tag.

The policy P in DF can be arbitrarily sophisticated. It might dictate not only which tags are visible to a TaPA, but also which associated information in DB is made available during a given scanning session. It can also incorporate information external to a TaPA. For example, a TaPA in a restaurant might be permitted to read “private” tags outside of business hours for the personal use of staff. It might determine the time of day from an internal clock or even a remote call to a Web site.

Furthermore, it is important to recognize that P might govern not just the logical environment of a reader, but its physical parameters as well. In Example 5, for instance, we describe a soft blocker that requires a reader to broadcast at low power – and thus to scan at only short distances – when reading “private” tags.

Tag classification changes: While the classification of a tag may be readily changed through modification of an associated database entry in DB , such a change may be local in nature, unless DB is a globally accessible database (like the ONS).

A particular concern is a change in classification of a tag from “public” to “private,” as needed in Example 4 above. If Shop A flags tag T in its database as having the classification “private,” Shop B will not necessarily have any awareness of this, and therefore will not respect the classification. A better approach in such cases is to change a classification indicator on the tag itself.

This may be implemented straightforwardly by having a bit in a tag that indicates its classification as either “public” or “private.” Any operation that changes the bit should be PIN-protected so as to prevent malicious alteration. (PIN protection, of course, is already present in EPCglobal specifications for the “kill” function.) The classification bit can be rewriteable, but a write-once bit may be sufficient for retail use, as this would be sufficient to ensure that tags are classified as “private” whenever they are owned by consumers.

4 Auditability

The soft-blocker system as described above may be thought of as loosely analogous to the Platform for Privacy Preferences (P3P) [2]. Users (or more precisely, their tags) specify a privacy policy, but have no means of actively enforcing or of auditing compliance with it. Indeed, the only way to verify compliance at all is to inspect a TaPA and ensure that it is properly configured within a system. With rigorous enforcement of privacy policies – perhaps supported by legislation and appropriate reader firmware defaults – software verification of this kind may be adequate.

It is possible to do much better, though. We can obtain external auditability in our proposal with a minor modification to the basic scheme combined with a slight change in the RFID-reader protocol. The key idea here is to perform *selective scanning*. In other words, the reader should touch only those portions of the tag space to which it is entitled; an auditor may determine reader compliance

by examining its scanning patterns then. To support this approach, we introduce a fourth component into the TaPA architecture that may reside either on the reader or in application software:

- **Reader driver RD :** This component interacts dynamically with a reader, indicating what portions of the tag-identifier space the reader should scan.

For the reader driver to be able to operate meaningfully, we require that tag classifications be determined through a partitioning of the tag-identifier space (or some portion thereof). Let $Z = \{z_0, z_1, \dots, z_m\}$ be a collection of non-overlapping identifier sets, which we call *zones*. Let tags specifying a privacy policy have identifiers residing in the zone z_0 . We propose that tag scanning, then, involve a two-phase process. First, a reader scans identifiers in z_0 . The identifiers (and other tag data) retrieved from this process determine the privacy policy to be enforced and, in particular, what subset $z \in Z$ the reader is permitted to scan. In the second reading phase, the reader scans the zones in Z . (Of course, variants with more than two phases are possible.)

The tree-walking algorithm for tag singulation makes it easily possible for an auditing device to determine whether or not a given reader is adhering to a particular privacy policy f . In the standard tree-walking algorithm, at each stage of the reading process, the reader specifies a prefix. Only tags whose identifiers have this prefix participate in the communication protocol; we refer to these as “communicating” tags. If the reader ever specifies a prefix that corresponds to identifiers in a zone that the reader should not scan, then it may be determined to be in breach of the privacy policy. Similarly, the ALOHA algorithm, when implemented according to EPC specifications [5], has provisions for identifier-prefix specification. A similar approach to auditing may be adopted in this case.

Example 7. We describe a simple, auditable implementation of the policy given in Example 4. Here, z_0 consists of all identifiers with a leading ‘0’ bit, z_1 consists of all identifiers with the leading bit pair ‘10,’ and z_2 consists of all identifiers with the leading bit pair ‘11.’ Tags in z_0 are “blockers.” Those in z_1 are “private” and those in z_2 are “public.” The policy P , then, is such that if any tag is detected in z_0 , i.e., any tag with a leading ‘0’ bits, then only identifiers in z_2 should be returned. Otherwise, identifiers from z_1 and z_2 , i.e., all identifiers with leading ‘1’ bits are returned.¹

This system may easily be audited as follows. In the first reading phase, the reader should specify a ‘0’ prefix for all communicating tags. In the second reading phase, if a blocker is present, then the reader should specify the prefix ‘11’ for all communicating tags. If the reader specifies in the first phase a prefix that permits communication by any tag with a leading ‘1’ bit, or if, in the

¹ The fact of learning whether there are any tags in z_1 at all constitutes a minor privacy violation. EPC tags permit a reader, however, to specify a mask, i.e., a leading substring of bits that tags must possess in their identifiers in order to respond to reader queries. Thus a reader can scan z_2 without learning whether there are any tags in z_1 .

presence of a blocker, the reader specifies in the second phase a prefix that permits communication by any tag with a leading ‘11’ bit pair, then the reader may be deemed to violate the privacy policy.

An auditor could use a special-purpose device to simulate a set of tags (of various classifications) and record all values broadcast by a reader. This would permit the detection of breaches of a particular privacy policy. Given a widely adopted set of privacy policies, it would be possible to manufacture small devices that would function automatically, illuminating an LED or otherwise informing an auditor when a non-compliant reader is encountered.

5 The R_xA Pharmacy Soft-Blocker Demonstration

RSA Laboratories constructed a prototype soft-blocking system for a demonstration environment called the “R_xA Pharmacy.” Exhibited at the RSA Conference 2004, the R_xA Pharmacy illustrated both the consumer benefits and potential privacy risks that RFID might bring to a pharmacy of the future, and how (soft) blocker tags might preserve the benefits of RFID tags while helping to prevent their abuse.

5.1 Equipment

The pharmacy employed Tag-It™ high-frequency (HF) RFID tags manufactured by Texas Instruments. These tags measure approximately 55mm × 110 mm; they have the appearance of blank white adhesive labels that, when held up to a light source, reveal an internal antenna and circuitry. HF tags, which operate at a frequency of 13.56 Mhz, possess three properties especially suitable for demonstration environments: (1) Thanks to their short operational range of some 40cm or so, they present a minimized risk of cross-contamination, i.e., inadvertent tag reading; (2) HF tags do not suffer from substantial reading interference in the presence of liquids and other substances, a problem with higher-frequency tag varieties; and (3) HF tags, unlike some other RFID-tag types, are compliant with the RF spectrum allocation regulations of most nations.

Readers in the demonstration were the ThingMagic™ Mercury 3 model. The application software was implemented in C++, using the Microsoft(R) Foundation Classes (MFC) for the user interface. Communication between the application and the reader was managed by a subclass of the MFC CSocket class called a CRFIDReaderSocket. The CRFIDReaderSocket object maintained the network connection to the reader and periodically queried the reader, using simple SQL queries. The readers’ responses were parsed and stored in a list of tags currently “visible” to the reader. The application received tag information from the CRFIDReaderSocket, performed the classification algorithm, then updated its display accordingly. A Microsoft(R) Access database, shared among multiple applications, was used to store the status of each tag. The database was updated with user-supplied data after the completion of each transaction.

5.2 The demonstration experience

The R_xA Pharmacy stocked a large collection of ordinary pharmacy bottles filled with colorful pills (otherwise known as jellybeans). Each bottle bore an RFID tag hidden beneath a label that read “R_xA Pharmacy. Contents: 1 oz. jellybeans. To be taken orally as needed. Warning: Side effects of RFID tag may include loss of user’s privacy! RSA(R) Blocker Tag is designed for protection.” Apart from the unique serial numbers in their RFID tags, all bottles were identical.

“Pharmacists” in the demonstration asked visitors to fill out a mock prescription form requiring a name (or pseudonym) and a medication preference. The R_xA Pharmacy offered three different types of “medications”: “happiness” pills, “wisdom” pills, and “tranquility” pills. A collection of laptops with RFID readers and external monitors as peripherals served as point-of-sale terminals for the demonstration.

On receiving a visitor’s prescription form, a pharmacist entered the provided information into one of the terminals, thereby creating a record in a central database (implemented in Microsoft(R) Access). The pharmacist then scanned one of the RFID-enabled medication bottles. This established a binding in the database between the visitor record and the RFID serial number associated with the medication bottle. The point-of-sale screen displayed the unique serial number of the bottle as well as the associated visitor information. The pharmacist then marked the bottle in the database as having been paid for, i.e., as now being situated in the privacy zone for the system.

To demonstrate the potential benefits of RFID to the consumer, pharmacists explained that the scanning of bottles could support automated verification of correct medication dispensing. Pharmacists further explained the value to consumers of retaining live RFID tags on leaving the pharmacy. In the future, these tags could support home applications, such as “smart” medicine cabinets capable of monitoring medication compliance, detecting expired medications, and alerting customers to potentially dangerous drug interactions [12]. Further, RFID-tagged medication bottles could be returned to a pharmacy to ensure quick and accurate refills.

At the same time, pharmacists explained to visitors the potential privacy risks of item-level tagging. As the RFID tags in the bottles contain unique serial numbers, any organization with access to a network of RFID readers could potentially perform physical tracking of consumers carrying medications. In the EPCglobal standard, the RFID tags on bottles would indicate their medication types, an additional source of worry. As an example, a consumer carrying, say, Oxycontin(R), a pain-relief medication with a high black-market value, might be vulnerable to thieves employing RFID readers. Additionally, an unscrupulous pharmacy could share its database information with other organizations so as to permit widespread association between the RFID-tag serial number and the name and personal information of the individual carrying the medication.

To protect against such abuses, R_xA pharmacists provided visitors with ordinary pharmacy bags equipped with soft blocker tags. Pharmacists demonstrated that when placed in these bags, medication bottles were not visible to R_xA Phar-

macy software. Point-of-sale screens displayed only the word “Blocked” when such scanning took place. Pharmacists also explained that bottles not marked as having been paid for, i.e., stolen merchandise, would remain visible to the R_xA system, as blocker tags serve to protect privacy, not to abet theft.

In the RSA Conference demonstration, pharmacists stapled visitors’ prescription forms to the bags containing their medication bottles (as in a typical U.S. pharmacy), and gave these as take-aways.

5.3 Tag classification

The classification-engine in this demonstration relied on a classification scheme in which blocker tags were distinguished by their serial numbers. In particular, blocker tags contained serial numbers residing in a specially designated partition z_0 of the full identifier space for tags. The zoning, i.e., public or private status of the tags for bottles, was determined by reference to the database. This allowed quick identification of blocker tags without requiring modification of tag identifiers to change status.

To draw on our architectural notation from above to describe the R_xA demo, the classification engine CE took as input the serial number n of a tag. If $n \in z_0$, then the tag was classified as a blocker; otherwise, CE consulted DB to determine whether the tag was registered “private” or “public.” The policy of the filter DF was then simple: If the set of tags presented to the filter included a blocker, then the filter output only the serial numbers of tags marked as “public”; otherwise DF output all tag serial numbers. In effect, then, the R_xA demo involved a TaPA like that in Example 1 above. Because of the use of a local database to record tag classifications, of course this particular architecture would not protect privacy outside the demonstration environment.

5.4 Tag reliability

The prototype system generally functioned as designed, but sensitivity to tag orientation led to occasional failures. The power available to a tag on scanning depends on the area it presents perpendicular to the field emitted by the antenna. Thus, when a blocker was aligned nearly parallel to this field (and perpendicular to the plane of the antenna), it would sometimes fail. We also found that when attached too high on pharmacy bags, and thus at a distance from the RFID tag on the bottle inside the bag, the blocker would sometimes fail to disrupt scanning of the bag contents. Finally, a small proportion of tags appeared to suffer from manufacturing defects, and could not be scanned reliably. On the other hand, some results of the demonstration were encouraging. For example, when blockers were attached on the lower portion of bags, they proved very effective at most angles of presentation – indeed, even when the bags were scanned with the blocker positioned away from the reader antenna relative to the contents of the bag.

The scanning problems revealed in the R_xA demo suggest the need for careful tag and/or system engineering to support soft blocking in its basic form. Passive

blocker tags may require careful placement on consumer bags or some form of antenna enhancement to achieve their full potential. Alternatively, an “opt-in” approach involving unblockers may prove attractive as a way of ensuring privacy as a default in the advent of system failures. With this approach, a scanning failure would cause private information to remain undisclosed (even when a unblocker is in fact present to authorize disclosure).

More active forms of blocking offer another avenue of exploration. Given industry plans to implement RFID readers in mobile phones [20], the most appealing option may ultimately be for mobile phones and similar devices to implement blocking, whether of the soft or full-blown variety. This would involve simulation of tag functionality, rather than reader functionality. (Note that current proposals for mobile phones involve a range of only a few centimeters, but this could well change.)

The physical effectiveness and reliability of blockers is, in short, a subject of study whose evolution must continue alongside that of RFID technology more generally.

6 Conclusion

We have introduced the notion of soft blocking, a cheap and flexible alternative to the original blocker-tag scheme. A soft blocker tag does not provide the same strong privacy guarantees as a full-blown blocker. Soft blockers, however, have the advantage of being almost identical to conventional RFID tags in terms of their hardware and functioning. They also support a wider range of privacy policies, as the notion of unblocking illustrates. Finally, although loosely evocative of the P3P standard for Web browsers, soft blockers have the potential advantage of supporting external audit procedures.

Acknowledgements

The authors wish to thank Dan Bailey, Burt Kaliski, and Ron Rivest for many stimulating discussions and suggestions regarding soft blocking and RFID security and privacy more generally.

References

1. Auto-ID Object Name Service (ONS) 1.0, 12 Aug. 2003. Auto-ID Working Draft. M. Mealling, editor. Available to members at develop.autoidcenter.org/TR/ons-1.0.pdf.
2. Platform for privacy preferences (P3P) project, 2003. World-Wide Web Consortium. Available at <http://www.w3.org/P3P/>.
3. R. Agrawal and R. Srikant. Privacy-preserving data mining. In *Proc. of the ACM SIGMOD Conference on Management of Data*, pages 439–450. ACM Press, 2000.
4. Benetton undecided on use of ‘smart tags’. *Associated Press*, 8 April 2003.

5. AutoID Center. 13.56 MHz ISM band class 1 radio frequency identification tag interference specification: Candidate recommendation, version 1.0.0. Technical Report MIT-AUTOID-WH-002, MIT Auto ID Center, 2003. Available from <http://www.epcglobalinc.org>.
6. D. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–88, 1981.
7. J. Collins. The cost of Wal-Mart’s RFID edict. *RFID Journal*, 10 Sept. 2003.
8. American Civil Liberties Union (ACLU) et al. Rfid position statement of consumer privacy and civil liberties organizations, 20 November 2003. Referenced 2004 at <http://www.privacyrights.org/ar/RFIDposition.htm>.
9. Security technology: Where’s the smart money? *The Economist*, pages 69–70. 9 February 2002.
10. D.M. Ewatt and M. Hayes. Gillette razors get new edge: RFID tags. *Information Week*, 13 January 2003. Available at <http://www.informationweek.com/story/IWK20030110S0028>.
11. K.P. Fishkin and S. Roy. Enhancing RFID privacy via antenna energy analysis. Technical Report Technical Memo IRS-TR-03-012, Intel Research Seattle, 2003. Presented at the MIT RFID Privacy Workshop, November 2003.
12. K.P. Fishkin and S. Roy. A ubiquitous system for medication monitoring. In *Pervasive 2004*, 2004. To appear. Also available as Intel Research Seattle Technical Memo IRS-TR-03-012.
13. T. El Gamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory*, 31:469–472, 1985.
14. S. Garfinkel. An RFID Bill of Rights. *Technology Review*, page 35, October 2002.
15. P. Golle, M. Jakobsson, A. Juels, and P. Syverson. Universal re-encryption for mixnets. In T. Okamoto, editor, *RSA Conference - Cryptographers’ Track (CT-RSA)*. Springer-Verlag, 2004. To appear.
16. A. Juels. Minimalist cryptography for low-cost RFID tags, 2003. In submission.
17. A. Juels and R. Pappu. Squealing Euros: Privacy protection in RFID-enabled banknotes. In R. Wright, editor, *Financial Cryptography ’03*, pages 103–121. Springer-Verlag, 2003. LNCS no. 2742.
18. A. Juels, R.L. Rivest, and M. Szydlo. The blocker tag: Selective blocking of RFID tags for consumer privacy. In V. Atluri, editor, *8th ACM Conference on Computer and Communications Security*, pages 103–111. ACM Press, 2003.
19. D. McCullagh. RFID tags: Big Brother in small packages. *CNet*, 13 January 2003. Available at <http://news.com.com/2010-1069-980325.html>.
20. Nokia unveils rfid phone reader. *RFID Journal*, 17 March 2004. Referenced 2004 at www.rfidjournal.com.
21. S. E. Sarma, S. A. Weis, and D.W. Engels. Radio-frequency identification systems. In Burton S. Kaliski Jr., Çetin Kaya Koç, and Christof Paar, editors, *CHES ’02*, pages 454–469. Springer-Verlag, 2002. LNCS no. 2523.
22. S.E. Sarma. Towards the five-cent tag. Technical Report MIT-AUTOID-WH-006, MIT Auto ID Center, 2001. Available from <http://www.autoidcenter.org>.
23. R. Shim. Benetton to track clothing with ID chips. *CNET*, 11 March 2003. URL: <http://news.com.com/2100-1019-992131.html>.
24. S. A. Weis, S. Sarma, R. Rivest, and D. Engels. Security and privacy aspects of low-cost radio frequency identification systems. In *First International Conference on Security in Pervasive Computing*, 2003. To appear.