High-Power Proxies for Enhancing RFID Privacy and Utility

Ari Juels¹ and Paul Syverson² and Dan Bailey¹

¹ RSA Laboratories
Bedford, MA 01730, USA
e-mail: {ajuels,dbailey}@rsasecurity.com
² Naval Research Laboratory
Washington, D.C. 20375, USA
e-mail: syverson@itd.nrl.navy.mil

Abstract. A basic radio-frequency identification (RFID) tag is a small and inexpensive microchip that emits a static identifier in response to a query from a nearby reader. Basic tags of the "smart-label" variety are likely to serve as a next-generation replacement for barcodes. This would introduce a strong potential for various forms of privacy infringement, such as invasive physical tracking and inventorying of individuals. Researchers have proposed several types of external devices of moderate-to-high computational ability that interact with RFID devices with the

to-high computational ability that interact with RFID devices with the aim of protecting user privacy. In this paper, we propose a new design principle for a personal RFID-privacy device. We refer to such a device as a REP (RFID Enhancer Proxy).

Briefly stated, a REP assumes the identities of tags and simulates them by proxy. By merit of its greater computing power, the REP can enforce more sophisticated privacy policies than those available in tags. (As a side benefit, it can also provide more flexible and reliable communications in RFID systems.) Previous, similar systems have been vulnerable to a serious attack, namely malicious exchange of data between RFID tags. An important contribution of our proposal is a technique that helps prevent this attack, even when tags do not have access-control features.

1 Introduction

In this paper, we propose the design of a new type of device for protecting consumer privacy with respect to RFID. We refer to this device as a REP (*RFID Enhancer Proxy*). Before explaining the aims and functioning of a REP, we first review background on RFID and its associated privacy problems.

A passive radio-frequency identification (RFID) tag is a microchip that is capable of transmitting a static identifier or serial number for a short distance. It is typically activated by a query from a nearby reader, which also transmits power for the operation of the tag. Several varieties of RFID tag are already familiar in daily life. Examples include the ExxonMobil Speedpass TM payment device, the small plaques mounted on car windshields for the purpose of automated toll payment, and the proximity cards used to control physical access to buildings.

The cost of rudimentary RFID tags, sometimes called "smart labels," promises to drop to roughly 0.05/unit in the next several years [19]. Tags as small as 0.4mm \times 0.4mm, and thin enough to be embedded in paper are already commercially available [23]. Such improvements in cost and size will mean a rapid proliferation of RFID tags into many areas of use. The United States Department of Defense and major retailers such as Wal-mart have issued mandates to their top suppliers requiring RFID deployment starting in 2005.

One goal of RFID-tag development is to see RFID serve ubiquitously as a replacement for barcodes. The main industry consortium advancing this goal is EPCglobal, a joint venture between the UCC and EAN, the organizations overseeing barcode use in the United States and Europe. EPCglobal is the standards-setting body for a system of standardized "electronic product codes" (EPC) analogous to the printed barcode used universally on consumer products today. (See, e.g., [13] for a description of a 96-bit EPC standard.) Broadly speaking, the vision is for RFID tags to serve as unique identifiers. These identifiers may serve as pointers to database entries, thereby allowing the compilation of extensive automated histories for individual items. EPCglobal has recently ratified its Class 1 Generation 2 standard, which will likely dictate basic tag architectures for some time to come.

Improved supply-chain management is the initial goal of major RFID deployments in the next few years. Pallets of goods will carry RFID tags so as to automate inventory tracking.

The present cost of RFID tags is such that prevalent RFID-tagging of individual goods in retail environments will be impractical for some years. Pilots are already afoot, however, and with improved manufacturing processes and larger economies of scale, as well as technological innovations like plastic circuits, itemlevel tagging seems inevitable.

Basic item-level RFID tagging promises many benefits, allowing flexible and intelligent handling of consumer goods and devices. Among the intruiging possibilities are:

- Receiptless item returns: Retailers can record the purchase conditions of an item in a database entry for its attached RFID tag. This would permit customers to return items without receipts. With RFID tags used to record the full lifecycle of an item, retailers would benefit from the ability to track the source of item defects.
- "Smart" appliances: With RFID tagging of foodstuffs, refrigerators could alert consumers to the presence of expired or recalled comestibles, and also compile shopping lists automatically based on a scan of their contents. Washing machines could use RFID-tagged articles of apparel to select an appropriate wash cycle. Microwave ovens could use RFID tags on cartons of food to determine an appropriate power setting and cooking regime.
- Aids to the handicapped: Researchers at Intel are exploring ways in which RFID may furnish information to aid Alzheimer's patients in navigating their environments.

- Recycling: Sorting recyclables is a resource-intensive process. RFID tags could permit automated identification of different types of recyclable plastics and other materials.
- Smart phones: Mobile phone manufacturers have plans to embed RFID readers in their handsets [16]. Consumers could use such devices to scan movie posters for showtimes, to scan products so as to make price comparisons, and so forth.

1.1 The privacy problem

The impending ubiquity of RFID tags, however, also poses a potentially widespread threat to consumer privacy [14] and likewise to the privacy of corporate data. The initial RFID-chip designs proposed by EPCglobal are geared toward general corporate and consumer use. So as to permit inexpensive manufacture, tags of this kind carry only the most basic functionality, emitting a static, 96-to-256-bit identifier (EPC) on receiving a reader query [19]. Such a system would divulge a large amount of information about ordinary consumers. This threat is twofold: (1) Thanks to their unique identifiers, RFID tags could permit indiscriminate physical tracking of individuals, and (2) As RFID tags may carry product information (as in EPCglobal standards), they would permit surreptitious inventorying of their bearers and could facilitate corporate espionage. An attacker scanning the RFID tags contained in personal items could in principle gather information about a victim's clothing, medications, memberships and financial status (via RFID tags in wallet cards), and so forth. An attacker gaining access to RFID information in warehouses or store shelves can glean valuable corporate intelligence.

The privacy issues raised by RFID tags in the consumer domain have received considerable coverage in the popular press and attention from privacy advocates. Early fuel for these concerns included a purported plan by the European Central Bank to embed RFID tags in Euro banknotes. Public outcry has since forced the postponement or withdrawal of several retail RFID pilot projects. A number of states in the United States, including Utah, California, and Massachusetts have embarked upon legislation to address the problems of RFID privacy. (It should be noted that legislation in California was defeated this year by the California Assembly.) The risks that RFID poses to corporate data have been less well publicized, but have still received some attention [22].

1.2 Why "killing" is insufficient

EPCglobal chip designs address the privacy problem by permitting an RFID tag to be "killed." On receiving a short, specially designated PIN [18], a tag renders itself permanently inoperable. For example, a clothing shop might deploy RFID tags to facilitate tracking of shipments and monitoring of shelf stocks. To protect the privacy of customers, checkout stations might "kill" the tags of purchased goods. The concept is similar to the removal or deactivation of inventory-control tags as practiced today.

There will be many environments, however, in which simple measures like "kill" commands are unworkable or undesirable for privacy enforcement. The several examples above of beneficial consumer uses for RFID illustrate why consumers may not wish to have their tags killed. Likewise, "kill" commands will not protect privacy in cases where RFID tags are deployed to track borrowed items like library books. Libraries are already beginning to deploy RFID [21]. The same will be true for RFID-tagging of rented items, like DVDs. Killing cannot play a role in protecting consumer privacy in these cases.

In the corporate setting, of course, killing is unworkable, as it would negate the benefits of supply-chain visibility that RFID brings to begin with.

1.3 Why Faraday cages are insufficient

Another proposed tool for protecting RFID tags is known as a Faraday cage. This is a metal shield, e.g., a piece of alluminum foil, that is impenetrable by radio waves of certain frequencies, including those used by RFID systems. By enclosing an RFID tag in a Faraday cage, one can minimize its vulnerability to unwanted scanning.

In some cases, Faraday cages may indeed prove very effective. For example, to protect an RFID-enabled identity card when not in use, one might store it in a metal-lined case. For general consumer use, for example, the approach is unworkable. One could use a foil-lined bag, for instance, to protect groceries from scanning. As foil-lined bags can be used to evade inventory-control systems (i.e., theft detection systems), retail shops are unlikely to embrace their proliferation. Moreover, this approach will not work for items on one's person, including clothing, handbags, wristwatches, etc.

1.4 RFID-tag capabilities

Projections on the likely resources in several years of Class 1 RFID tags with cost in the vicinity of \$0.05 include several hundred bits of memory and somewhere between 500 to 5000 logical gates [26], of which a considerable fraction will be required for basic tag functions. Few gates will be available for security functionality. Thus such RFID tags may be expected to perform some basic computational operations, but not conventional cryptographic ones. Even hardware-efficient symmetric-key encryption algorithms like that recently proposed by Feldhofer et al. [2] are well beyond the reach of RFID tags of this kind. At best, low-cost RFID tags may include security functions involving static keys, such as keyed writes, i.e., essentially just PIN-controlled data accesses.

1.5 Our work

As explained above, we introduce in this paper a new type of RFID-privacy-protecting device known as a REP. The REP works by assuming the identities of RFID tags under its control. In particular, it loads their identifying information

and then simulates the tags in the presence of reading devices in order to enforce a privacy policy on behalf of the REP owner. These privacy policies may include the requirement for the reader and the REP to participate in an authentication protocol more sophisticated than an ordinary tag implements. When a tag is no longer to be simulated by a REP, it may have its identity re-implanted.

For consumer applications, we propose that the REP more-or-less continually rewrite the identifiers transmitted by tags under its control. The REP may write either ciphertexts or random pseudonyms to tags. This proposal is similar in flavor to those of Golle et al. [6] and Juels and Pappu [10], which propose re-encryption of ciphertexts on tag identifiers by computationally powerful and potentially untrusted external computing devices. In contrast to these proposals for tag re-encryption, however, we consider the REP as a trusted personal device. This removes the need for a reliance on public-key cryptography, and consequently leads to a different set of architectural choices, as we shall see.

A REP must perform four different operations:

- 1. Tag acquisition: When the owner of a REP and RFID tag wishes the REP to simulate the tag, the REP must acquire all of the necessary tag information and place the tag in a state permitting the REP to act as its proxy. The main technical challenge occurs when the tag has associated secrets, like PINs for access control or "killing," that must be transferred securely.
- 2. Tag relabeling (or re-encryption): The REP changes the identifiers on tags in its control so as to prevent surveillance of these tags. (Tags could instead be put in a "sleep" mode, but this has drawbacks that we discuss later.) Relabeling introduces various integrity problems, particularly the need to prevent adversarial re-writing of tags. Indeed, one of our contributions is a simple technique for preventing an adversary from swapping the identities of two different tags, e.g., swapping the identifiers on two medications with differing dosages. Previous proposals [6] are vulnerable to this type of attack or require special physical prevention mechanisms [10]. The technique we propose, which involves random input from the tag in the creation of pseudonyms, works even when tags do not have access-control features.
- 3. Tag simulation: The REP simulates tags in interaction with readers. The REP may also simulate spurious tags to prevent leakage of information about the number of tags carried by its owner. As a REP is presumed to be a powerful device, it can enforce more-or-less any privacy policy desired by its owner. We do not therefore specify simulation policies in this paper. We note, however, that these could include robust public-key based authentication schemes.
- 4. Tag release: When the owner of a REP wishes it no longer to simulate a tag, the REP must release its control and reimprint the tag with its original identity.

We note that blocker devices as proposed by Juels, Rivest, and Szydlo [11] and the variant proposed in [9] can serve as alternatives to REPs for consumer privacy protection. REPs, however, have a couple of features that make them

an attractive alternative to blockers: (1) If a tag temporarily exits the broadcast range of a blocker, it is subject to complete compromise; by contrast, a tag under the control of a REP will merely go without an identity change during this period, and (2) Blocker tags can only be effective as a universal standard implemented on both tags and readers, while a REP requires only tag-based support, and can be compatible with any reading system.

In fact, though, the idea of blocking can be viewed as complementary to the tag-simulation aspect of our REP proposal: A blocker could, for instance, act as a REP under certain circumstances. It might, for instance, simulate tags that it is protecting in its "private" space so as to allow reader access to these tags when policy permits. This "block-and-simulate" approach is conceptually simple, and an attractive alternative to ideas we describe here.

1.6 Organization

In section 2, we briefly describe previous work relevant to our proposal here. We outline our REP proposal in section 3, delineating ideas for the functions of tag acquisition, tag simulation, and re-implantation of tag identities. We conclude in section 4.

2 Previous Work

Researchers have from the outset recognized the limitations of the "killing" approach, and the consequent possibility of privacy threats from physical tracking in the deployment of RFID tags [18]. Several recent papers have proposed ways of addressing the problem. As explained above, the major challenge is that inexpensive RFID tags, the type likely to be deployed most widely, may well be incapable of performing even the most basic cryptographic operations, and also have little memory (just a few hundred bits).

Weis, Sarma, Rivest, and Engels [26] 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. They 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 pseudorandom number generator to protect tag identities. In a nutshell, their idea is for the tag to output the pair (r, PRNG(ID, r)), where 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 and Pappu [10] consider a plan by the European Central Bank to embed RFID tags in Euro banknotes. 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. Thus tags themselves perform no cryptographic operations. Verification of correct behavior by re-encryption agents in the Juels and Pappu system may be performed by any entity with optical access to banknotes, e.g., shops and banks. Thus, while their scheme involves changes in the identities of RFID tags, they require optical contact for the purpose of authentication, which our scheme does not.

Juels, Rivest, and Szydlo [11] propose a special form of RFID tag called a "blocker." This tag disrupts the protocol used by the reader to establish communications with individual tags among a set of tags. By targeting this disruption selectively, the "blocker" tag aims to protect consumer privacy while permitting normal inventory-control processes to proceed normally. A "blocker" could be a more sophisticated device than an RFID tag, e.g., a mobile phone.

Juels [12] proposes the concept of "minimalist cryptography." This involves a scheme in which RFID tags store a small set of unlinkable pseudonyms. They rotate through these as a privacy-protection measure. To ensure against an attacker exhausting the set of pseudonyms, Juels proposes a form of "throttling," i.e., timed delay on pseudonym changes. The full-blown scheme here includes use of one-time padding to enforce privacy and authentity of tags, and is accompanied by a formal model and analysis.

Molnar and Wagner [15] examine RFID privacy in the special setting of libraries, where tag deactivation is naturally infeasible. They propose a range of schemes. Some of these do not require symmetric-key cryptography on tags; for example, they consider the idea of having tags transmit random strings to readers for use in protecting communications on the stronger reader-to-tag link and the idea of relabelling tags with new identifiers at the time of check-out. As libraries may be in a position to purchase relatively high-cost RFID tags, Molnar and Wagner also consider some schemes that involve pseudo-random number generation on tags.

Garfinkel [5] proposes a different approach based on an "RFID Bill of Rights," which consists of five articles proposed as a voluntary framework for commercial deployment of RFID tags. Included are: (1) the right of the consumer to know what items possess RFID tags, (2) the right to have tags removed or deactivated upon purchase of these items, (3) the right of the consumer to access of the data associated with an RFID tag, (4) the right to access of services without mandatory use of RFID tags, and finally (5) the right to know to when, where, and why the data in RFID tags is accessed. In a similar vein, Floerkemeier et al. consider ways of harmonizing RFID use with the Fair Information Principles of the OECD [4]. They also propose the concept of a "Watchdog Tag," a high-powered device that monitors policy compliance.

In their work on mix networks, Golle et al. [6] and also Danezis and Lysyan-skaya [1] independently propose a cryptographic tool known as universal re-

encryption. Universal re-encryption permits a (semantically secure) public-key ciphertext to be re-encrypted by an entity without knowledge of the associated public key. Taking advantage of this property, Golle et al. briefly propose a system in which an RFID tag stores a public-key ciphertext of its unique identifier in a form subject to universal re-encryption. In order to change the appearance of this ciphertext, it is necessary to permit its re-encryption by external agents, namely RFID reader/writers with adequate computational power to perform cryptographic operations.

The Golle et al. approach is similar in flavor to that of Juels and Pappu. The major difference, however, is that in the case of banknotes, a single public key may be used for the complete system. In contrast, in a consumer environment, it is likely that many public keys will be employed. Every consumer, for example, may wish to possess an individual key to permit direct management of his or her privacy. Thus, in such an environment, it is important that no public key be involved in the process of re-encryption: The public key itself could otherwise serve as a privacy-compromising identifier. Universal re-encryption provides exactly this feature of public-key concealment, and thereby permits unlimited privacy-preserving re-encryption of the ciphertexts carried by tags. With this approach, one can imagine special privacy-enhancing readers scattered throughout a city to re-encrypt ciphertexts on behalf of the owners of tags.

Golle et al. also propose the idea of having privacy-concerned users of RFID-tags carry personal re-encryption devices with them. This is similar in flavor to our REP proposal in this paper. As we shall see, however, by exploiting the fact that a REP is a trusted device, we are able to solve an important integrity problem present in the Golle et al. proposal, namely the problem of attackers swapping identifiers between tags.

More generally, the idea of small devices communicating through more powerful proxy devices has already proven of value in a number of computing systems. This is a means by which small, embedded computational devices in the Oxygen project at MIT, for example, enforce privacy for users [25], and by which some privacy-preserving systems have operated [17]. Our main contribution in this paper is the application of the idea in the face of the special challenges that the limited computational capabilities, high mobility, and sensitive nature of RFID devices pose.

The "RFID Guardian" project [24], an effort contemporaneous with our own research, aims shortly to build a device similar in flavor to a REP. That project does not at present treat the issues of fine-grained control such as tag acquisition and ownership transfer.

3 How a REP Works

A REP, as we have explained, functions as a proxy for RFID tags. As such, it is able to simulate these tags and therefore enforce privacy policies of more-or-less arbitrary sophistication. Additionally – and quite importantly for many applications – a REP, being a powered device, can serve as a much more reliable

interface for transmitting RFID data than an RFID tag. Stated more generally, a REP can serve as a more trustworthy conduit for RFID data than the tags it controls. This can be particularly valuable in, e.g., environments in which there are physical impediments to RFID scanning. Metals and liquids can both interfere with RFID scanning, for instance; manufacturers have already confronted challenges in scanning such items as cans of drinking soda.

We now offer details on the four processes involved in REP management of tags: Tag acquisition, Tag relabeling, Tag simulation, and Tag release.

3.1 Tag acquisition

Acquisition of a tag by a REP involves transfer of the complete set of tag data. For tags that simply broadcast identifiers and other public information, this is a straightforward matter: The REP need merely scan the tag. Where it becomes more complicated is when a tag has associated secrets, particularly PINs required to implement secure tag operations such as writing and "killing." The transfer of these data may take place in one of two ways:

- 1. The tag data may be transferred directly to the REP from a trusted higher powered device such as a reader. In all cases, care should be taken of course to protect the privacy and integrity of data during this transfer.
 - At checkout from a shop, for example, private data associated with the tags on purchased products might be communicated by the checkout register directly to the REP via, e.g., a Bluetooth link.
 - In a supply chain, before shipment to a supply-chain partner, a pallet of tagged items might itself be tagged with a REP. The REP is programmed with private data about the tags in its pallet from a reader. This data transfer may take place using the RFID data transport or another physical layer such as Bluetooth, ZigBee, or IrDA.
- 2. The tag data may be released by the tag on suitable out-of-band authentication of the REP to the tag, or this data transfer may take place in an environment with adequate compensating controls. There are several channels by which the RFID reader might authenticate itself to the tag as a trusted device. If tags bear printed keys, then optical scanning of these keys might serve this function [10,26]. A more convenient alternative might be release of tag data upon physical contact or proximity between the REP and the tag in accordance with the "resurrecting duckling" paradigm of [20]. Indeed, researchers have demonstrated methods by which tags may be able to ascertain (very roughly) whether a reader is in close proximity [3].

3.2 Tag relabeling

As explained above, we advocate relabeling of tags by the REP as a means of protecting against privacy compromise thorugh direct tag scanning. One way to accomplish this is to have the REP re-encrypt a public-key ciphertext carried by a tag, as proposed in previous work. The setting we consider, however, in which

the REP serves as a proxy permits a simpler approach involving the assignment of changing pseudonyms to tags. In particular, for timeslot t, the REP can assign a k-bit pseudonym $p_{t,i}$ to tag i. (Time here would be maintained by the REP alone, as it is infeasible for tags to keep time.) This pseudonym may be generated uniformly at random by the REP and stored in a table in association with the tag identity. Alternatively, it could be computed as a k-bit symmetric-key ciphertext based on a master key σ held by the REP. In particular, we might simply compute $p_{t,i} = E_{\sigma}[t,i]$.

The approach of re-encryption or more generally, re-naming of tags, however, introduces a serious security problem, that of *data integrity*. Because tags are computationally too weak to authenticate re-writing entities, it is hard to enforce write-control permissions on tags that preclude adversarial tampering. This means that an attacker can corrupt tag data.

Writing of tag data is typically a PIN-protected process in RFID tags. This mitigates the risk of malicious corruption of tag data, but does not eliminate it. An adversary can potentially intercept REP-to-tag communications and thus learn the write PIN for the tag. Alternatively, if attacking at sufficiently close range, the adversary can hijack a write session between the REP and tag. Provided that k is sufficiently large, i.e., pseudonyms are long enough, an attacker has very little chance of being able to forge a pseudonym existentially.

More serious is the possibility of a *swapping attack*, in which an adversary exchanges the ciphertexts $p_{t,i}$ and $p_{t,j}$ between two tags i and j. This can have very serious consequences. It suffices to consider the possibility of an attacker exchanging ciphertexts associated with two medications or two spare aircraft parts. Previous proposals involving re-encryption of ciphertexts have been unable to address this attack, and have indeed left its resolution as an open problem.

In the case where the PIN associated with a tag is locked, i.e., not subject to alteration, the PIN itself can serve as a kind of authenticator for the tag [7]. Thus, a PIN can be used as a mechanism to defend against swapping attacks: If a tag is discovered to carry a pseudonym that does not match its PIN, then it may be presumed that a swapping attack has occurred.

The use of PINs to defend against swapping attacks, however, is twofold. First, as noted above, a frequently-used PIN is subject to compromise. And a compromised PIN is effectively a kind of static identifier. An attacker capable of testing the correctness of a PIN can use it to track a tag. Of course, if a PIN is not used to authenticate the operation of identifier-writing, but only to test periodically for swapping, then the risk of PIN compromise is diminished. A second, more serious problem is the basic one of PIN management. We have already noted that tag acquisition may need to involve out-of-band transfer of tag secrets. In general, management of tag PINs is like the general problem of key management in data-security systems. It is conceptually simple, but operationally thorny. Hence, it seems very likely that consumers will carry RFID tags that do not have associated PINs, or will not know the associated PINs of their tags!

Happily, we are able to provide a simple defense against identifier swapping that works even when write access to tags is universal.

The idea is for a tag i to participate itself in the generation of a given pseudonym $p_{t,i}$. In principle, if the tag itself could perform symmetric-key encryption under an appropriate cipher E, then the data-integrity problem would be solved: The tag would not need to have its pseudonyms updated by the REP. Cryptography of this kind, however, as we have explained, is well beyond the reach of low-cost tag capabilities.

Tags can, however, generate a certain amount of randomness. We might therefore consider a protocol in which a tag generates a new pseudonym $p_{t,i}$ for a counter t maintained (internally) on the tag. If it receives an "update" command from the REP, along with a valid write key, the tag transmits $p_{t,i}$ to the reader and adopts $p_{t,i}$ as its new pseudonym. (In order to prevent desynchronization due to an interrupted session, a tag might await a final "ack" from the reader before effecting the update.) This approach would render swapping attacks infeasible, as the REP – and thus an adversary – would be unable to dictate tag pseudonyms.

In practice, tags are capable of generating only a limited number of random bits in the course of a given session. Moreover, much of the randomness that a tag generates is already bespoke by other protocol requirements. (See the remark below.) A tag may therefore be unable to generate a full-length random pseudonym in each session.

Even partial generation of a pseudonym by a tag, however, can help alleviate the risk of swapping attacks. In particular, tag might emit a random nonce r of length k' < k before accepting the writing of a new pseudonym. The tag then only accepts a new pseudonym if it "matches" this nonce, e.g., if the last bits of the pseudonym are equal to r. (As an alternative, a tag might simply "declare" the last bits of its pseudonym to be r and accept only the other bits from the reader.) In other words, a tag can participate partially in the generation of its pseudonyms. An adversary attempting to swap pseudonyms, then, will be unable to do so unless it can locate a pair of tags simultaneously emitting the same nonces.

The probability of successful attack by an adversary, then, is a function of the number of tags N managed by a REP, the number of timeslots s available to the adversary for its attack, and the bit-length k'. Consider, for instance, a pallet carrying some 100 tags relabelled every minute, and seeking protection against attacks lasting up to one day (1440 minutes), and employing tags that generate 32-bit nonces. The probability that a given tag shares a random pseudonym with any of the 99 others may be crudely bounded above by $99/2^{32}$. Thus the probability of a successful swapping attack in this case is easily seen to be less than $(1-(1-99/2^{32})) \times 1440 < 0.000034$.

Denial-of-service: Even if an attacker cannot successfully initiate a swapping attack, corruption of tag data has a second effect: Denial of service. If an attacker is able to implant a pseudonym in a tag, the tag effectively becomes desynchronized with the REP: The REP no longer recognizes the tag's pseudonym. If the

REP were consequently to halt rotation of new pseudonyms into a tag, a breach of privacy could result, since tag identifier would remain static. A REP might alert a user to unexpected de-synchronization events of this kind by emitting a warning tone, for instance. (Alternatively, a REP might continue to relabel tags even if it does not know their true underlying identifiers; the REP can, of course, simply generate temporary identifiers for tags it does not recognize. This might have an undesirable spillover effect if a REP relabels tags that do not belong to its owner!)

A secondary effect of a corruption attack is that the REP cannot properly release tags: If it does not recognize their pseudonyms, it cannot manage them properly. Thus, one of two approaches might be needed for tag restoration: (1) If the REP possesses PINs for the tags in its control, it can try to match tags to PINs via exhaustive search or (2) Some kind of manual intervention on the part of the user might be necessary, e.g., the user might have to key in a printed product code from items that the REP has "lost." Given the current and probably persistent imperfections in RFID, we expect some level of back-up identifer recovery and manual intervention to occur regularly. Since denial-of-service attacks would likely be a rarity, anyway, they would probably constitute little more than a nuisance in our system.

Remarks: Communications on the reader-to-tag (or REP-to-tag) channel, which is often called the *forward channel*, are typically transmitted at a higher power than on the tag-to-reader channel, which is often called the *back channel*. One way to achieve privacy protection of REP-to-tag communications, therefore, is to have the tag generate a random value R and send it on the back channel. The REP can then protect transmission of a message on the forward channel, as the REP can then transmit the write PIN XORed with R. Techniques such as these can in principle prevent compromise of write PINs for tags via long-range eavesdropping, and thus reduce the overall threat of data corruption. They do not, however, address the problems of short-range eavesdropping and hijacking.

An entirely different and stronger approach to data integrity is possible using somewhat more heavyweight techniques. For example, the "minimalist cryptography" concept in [12] could be used to establish shared secrets between tags and the REP. On top of this might be layered a kind of lightweight message authentication code (MAC) as in [8]. Under the modeling assumptions of [12], this combination of techniques would permit the REP and tag to authenticate new pseudonyms.

3.3 Tag simulation

Once the REP has acquired a tag, it can, of course, simulate it as desired in the presence of an RFID reader. As explained above, this has the benefit of making tag reading more reliable: The REP, as a higher-powered device can transmit information to a reader more reliably than a tag. The REP might essentially enforce the kind of data filtering envisioned in the "soft blocking" approach to tag privacy. "Soft blocking," however, relies upon a universal set of

policy conventions. A REP, by contrast, can achieve a wholly personalized set of privacy policies. Additionally, a REP can enforce these policies in the presence of any reader – even a malicious one. We give two examples of REP capabilities unavailable in previously proposed approaches:

- Geographical conditioning: A REP may make decisions about whether to release information based on its geographical location. For example, a REP might release information about a pallet's RFID tags only when the pallet arrives at its destination. There is a variety of channels by which the REP might determine whether or not it is present at its destination, e.g.: (1) A built-in GPS unit; (2) Authenticated transmissions from readers; or (3) An authenticated notification from another protocol such as Bluetooth.
- Object simulation: To deceive attackers, a REP may simulate RFID tags associated with objects that the user does not possess. Here are two examples:
 - 1. A consumer can "carry" information about an object by simulating it. When the owner of a refrigerator wants to purchase a new handle of the correct type, or the owner of a stereo system wants to know which speakers are appropriate for her home theater system, she can simulate the associated RFID tags in order to acquire, carry, and convey this information conveniently.
 - 2. The owner of a Patek Philippe watch might program her REP to simulate the Patek Philippe RFID tag when she is present in upscale shops (so as to improve her level of customer service), but to mask her watch (or simulate a cheap one) when she is walking the streets.

Additionally, there are other scenarios in which a REP can enforce privacy policies. For example, jewelry retailers typically perform nightly inventories of their stock, given the high value of individual items. One can imagine that they would find RFID-tagging of their stock useful in this process. Such tagging, however, would make it possible for a competitor to scan a jewelry case quickly and in secrecy, and thereby learn the rate of stock turnover. A REP might simulate non-existent jewels to render this more difficult. This approach would, similarly, be very useful in military environments.

Finally, we note that a REP can transmit tag information to devices other than RFID readers. A REP might, for instance transmit tag data via WiFi, thereby serving as a bridge between RFID and other wireless systems.

3.4 Tag release

When a REP is to release an RFID tag, it must restore the tag's original identity. This process is straightforward if the REP has unrestricted write access to the tag. The technique we introduce in section 3.2 for preventing swapping attacks introduces a problem here, however, as its aim is precisely to restrict the identifiers that may be written to a tag. We propose, therefore, that on release of a tag, the randomly assigned portion of its identifier be retained, and that the rest of its identifier be restored to its original state. For example, the identifier

on an EPC tag has two segments, roughly speaking: (1) A (numerical) identifier segment that specifies the object the tag is attached to, e.g., says, "This is a 100g tablette of Valhrona chocolate" and (2) A unique numerical segment, effectively a serial number. During the period in which a tag is simulated by a REP, segment (1) can be effaced or overwritten by the REP, while segment (2) (or a portion thereof) is generated at random by the tag. When the tag is released, the randomness in (2) is retained, while (1) is restored. Effectively, then, a tag gets a new serial number at the time it is released.

This change in segment (2) could be problematic in some cases. For example, if a user has a warranty associated with an item that is referenced by its initial serial number, the user would like to retain that serial number. We note, however, that the REP can help provided serial-number translation as desired. For example, if a carton of milk has had its serial number changed through relabelling, the REP can transfer the old serial number to a "smart" refrigerator when a consumer puts the milk away.

In the case where a tag has an associated PIN, of course, the PIN may be used to place the tag in a special state in which its serial number may be completely rewritten. Alternatively, physical mechanisms like reader proximity, as that described in [3], might trigger restoration of original tag state.

An important logistical question is how the REP is to determine when to release a tag. This process may in many cases be controlled by the user or performed automatically based on external environmental cues, e.g., when a user's home network informs the REP that the user has entered her home. Some experiments suggest that as tags enter the limit of range of a reading device, their response rate degrades [3]. Based on such information, a REP might be able to detect the removal of a tag from its vicinity and restore its initial state automatically. (To achieve early detection of impending tag departure from its read radius, a REP might periodically reduce its power level.) The opposite is alternatively possible: Release of a tag might be effected by bringing the REP into close proximity or actual physical contact with the tag. This latter case has a useful feature: Physical proximity is effectively a kind of authentication, and might serve as the basis for full restoration of a tag identifier, thereby bypassing the problem of serial-number-changes discussed above. For library books and similar items, this could be especially useful.

3.5 Putting tags to sleep

In principle, a REP can put tags to sleep while it is simulating them, and then wake them for identity re-implantation, thereby obviating the need for tag relabeling. ("Sleeping" is not supported by the EPC Class 1 Generation 2 standard, but could in principle be incorporated into inexpensive tags.) The process of waking, however, can be a problematic one. For logical access control, sleep/wake commands must be keyed with PINs so as to prevent malicious alteration of tag behavior. A problem then arises: Unless a tag identifies itself, a reader (or REP) cannot know which waking key to transmit; a sleeping tag cannot, of course, identify itself. Trial-and-error transmission of PINs to tags would be possible,

but cumbersome. Alternatively, it is possible for a REP to broadcast a waking PIN to all tags in its vicinity, but given the likely movement of tags in and out of the field of control of the REP over the course of time, this approach seems impractical for consumer applications. This is particularly the case if a REP wishes to transfer control of a tag to a different device: the secondary device must be able to identify the tag of which it is taking control.

Putting tags to sleep *might* present a more feasible approach to access control if waking involves some form of physical access control. For example, it would be possible to touch an RFID device to a tag in order to wake it. We expect, however, that it would be cumbersome for consumers to have to engage in a fine-grained physical process to control the state of their tags.

In supply chains, where rigorous logistical controls are available, sleep/wake patterns may be more managable. In such settings it may make sense for a REP to put tags to sleep while simulating them. As maintenance of live tags is the technically more challenging option, it is the approach of having a REP relabel tags on a regular basis that we primarily explore in this paper. In supply chains, where rigorous logistical controls are available, sleep/wake patterns may be more managable. In such settings it may make sense for a REP to put tags to sleep while simulating them.

3.6 REPs and the EPCglobal standard

We can make the basic REP approach work particularly effectively with Class 1 Generation 2 tags in commercial settings – if they have writeable IDs. (Many memory technologies such as EEPROM impose limitations on the number of times memory cells may be rewritten, but some thousands of rewrite operations should be supportable.) PIN management would also be essential to prevent swapping attacks, as EPC tags do not, of course, support our idea of taggenerated randomness in identifiers. Alternatively, the "block-and-simulate" approach would be workable here.

Unlike the application of REPs to provide personal privacy, REPs in commercial settings would manage a limited population of tags. The goal in this situation is to transport a container of items from one trusted environment such as a factory to another such as a distribution center. This transportation of goods has been a major source of loss for manufacturers and retailers alike.

Some have noted that the management of unique individual tag PINs would require a new data communications infrastructure to be built for this purpose among supply chain participants. By having the tags managed by a REP, one can reduce this key distribution problem to the authentication of a reader to the REP. REPs by their virtue of relaxed cost constraints could be Class 3 tags capable of public-key cryptography. Issuance of digital certificates to REPs and readers would eliminate the need for a new secret-key distribution infrastructure.

Upon arrival of a pallet at its destination, the reader and REP would perform public-key-based mutual authentication using their digital certificates. On completion of the protocol, the reader would issue a special command causing the REP to unconceal and relabel all the tags in the pallet with their true identities.

4 Conclusion

We have proposed the idea of a REP, a device that serves as proxy for basic RFID tags, such as those of the Class 1 Gen 2 variety. A REP renders RFID tags dormant and then simulates them to other devices, e.g., RFID readers. Thanks to its moderate-to-high computational ability, a REP can enforce sophisticated privacy and security policies, even taking factors like time and location into account. We have shown how REPs can protect the privacy of consumers and the sensitive information of industrial RFID systems, and can withstand attacks against the integrity of tag identifiers, e.g., swapping attacks.

As we have noted, in addition to restricting information access where appropriate, a REP can also facilitate communications in RFID systems. As RFID tags are passive devices and therefore not wholly reliable communicators, a REP can improve the reliability of communications with an RFID system by acting as a proxy for RFID tags. A REP can also communicate via protocols other than RFID, e.g., Bluetooth, thereby acting as a bridge between divergent communication systems. We think that REPs are a powerful and practical notion and believe that Class 3 EPCglobal devices might serve as REPs in some degree, perhaps along some of the lines we have proposed here.

References

- 1. G. Danezis, 2003. Personal communications.
- M. Feldhofer, S. Dominikus, and J. Wolkerstorfer. Strong authentication for RFID systems using the AES algorithm. In M. Joye and J.-J. Quisquater, editors, Cryptographic Hardware and Embedded Systems (CHES), pages 357–370. Springer-Verlag, 2004. LNCS no. 3156.
- 3. K. P. Fishkin, S. Roy, and B. Jiang. Some methods for privacy in RFID communication. In 1st European Workshop on Security in Ad-Hoc and Sensor Networks (ESAS 2004), pages 42–53, 2004.
- 4. C. Floerkemeier, R. Schneider, and M. Langheinrich. Scanning with a purpose supporting the Fair Information Principles in RFID protocols. In 2nd International Symposium on Ubiquitous Computing Systems (UCS 2004), November 2004. Available at http://www.vs.inf.ethz.ch/publ/?author=floerkem.
- 5. S. Garfinkel. An RFID Bill of Rights. Technology Review, page 35, October 2002.
- P. Golle, M. Jakobsson, A. Juels, and P. Syverson. Universal re-encryption for mixnets. In T. Okamoto, editor, RSA Conference - Cryptographers' Track (CT-RSA), pages 163–178. Springer-Verlag, 2004.
- 7. A. Juels. Strengthening EPC tags against cloning, 2004. In submission. Referenced at rfid-security.com.
- A. Juels. 'Yoking-proofs' for RFID tags. In PerCom Workshops 2004, pages 138– 143. IEEE Computer Society, 2004.
- 9. A. Juels and J. Brainard. Soft blocking: Flexible blocker tags on the cheap. In S. De Capitani di Vimercatiand P. Syverson, editor, Wireless Privacy in the Electronic Society (WPES 04), pages 1–8. ACM Press, 2004.

- 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.
- 11. 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.
- Ari Juels. Minimalist Cryptography for RFID Tags. In C. Blundo and S. Cimato, editors, Security in Communication Networks, pages 149–164. Springer-Verlag, 2004.
- 13. AutoID Labs. 860 MHz-960 Mhz class 1 radio frequency identification tag radio frequency and logical communication interface specification recommended standard, version 1.0.0. Technical Report MIT-AUTOID-WH-007, Auto-ID Labs, 2002. Referenced in 2005 at http://www.autoidlabs.com.
- 14. D. McCullagh. RFID tags: Big Brother in small packages. *CNet*, 13 January 2003. Available at http://news.com.com/2010-1069-980325.html.
- 15. D. Molnar and D. Wagner. Privacy and security in library RFID : Issues, practices, and architectures. In B. Pfitzmann and P. McDaniel, editors, ACM CCS, pages $210-219,\ 2004.$
- 16. Nokia unveils RFID phone reader. *RFID Journal*, 17 March 2004. Available at http://www.rfidjournal.com/article/view/834.
- 17. Michael G. Reed, Paul F. Syverson, and David M. Goldschlag. Protocols using anonymous connections: Mobile applications. In Bruce Christianson, Bruno Crispo, Mark Lomas, and Michael Roe, editors, Security Protocols, 5th International Workshop, pages 13–23. Springer-Verlag, LNCS 1361, April 1997. Available at http://chacs.nrl.navy.mil/publications/CHACS/1997/.
- 18. 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.
- S.E. Sarma. Towards the five-cent tag. Technical Report MIT-AUTOID-WH-006, MIT Auto ID Center, 2001. Available from http://www.epcglobalinc.org.
- 20. F. Stajano and R. Anderson. The resurrecting duckling: Security issues for adhoc wireless networks. In 7th International Workshop on Security Protocols, pages 172–194. Springer-Verlag, 1999. LNCS no. 1796.
- 21. J. Stanley. Chip away at privacy: Library tracking system spawns Big Brother ire. San Francisco Chronicle, 2 July 2004.
- 22. R. Stapleton-Gray. Would Macy's scan Gimbels? competitive intelligence and RFID. Technical report, Stapleton-Gray & Associates, Inc., 2003. Available at http://www.stapleton-gray.com/papers/ci-20031027.PDF.
- K. Takaragi, M. Usami, R. Imura, R. Itsuki, and T. Satoh. An ultra small individual recognition security chip. *IEEE Micro*, 21(6):43–49, 2001.
- 24. A. Tanenbaum, G. Gaydadjiev, B. Crispo, M. Rieback, D. Stafylarakis, and C. Zhang. The RFID Guardian project. URL: http://www.cs.vu.nl/~melanie/rfid_guardian/people.html.
- 25. R. Tuchinda. Security and privacy in the intelligent room. Master's thesis, M.I.T., 15 May 2002.
- 26. S. A. Weis, S. Sarma, R. Rivest, and D. Engels. Security and privacy aspects of low-cost radio frequency identification systems. In Hutter et al., editor, First International Conference on Security in Pervasive Computing (SPC), pages 201– 212. Springer-Verlag, 2003. LNCS no. 2802.