

Defining Strong Privacy for RFID

ARI JUELS
RSA Security
and
STEPHEN A. WEIS
Google

In this work, we consider privacy in Radio Frequency IDentification (RFID) systems. Our contribution is twofold: (i) We propose a simple, formal definition of strong privacy useful for basic analysis of RFID systems, as well as a different (weaker) definition applicable to multiverifier systems; (ii) We apply our definition to reveal vulnerabilities in several proposed privacy-enhancing RFID protocols; and (iii) We formally analyze and suggest improvements to hash-locks, one of the first privacy-enhancing RFID protocols in the literature.

Categories and Subject Descriptors: D.4.6 [**Security and Protection**]: Authentication; F.1.1 [**Models of Computation**]: Bounded-Action Devices

General Terms: Security, Theory

Additional Key Words and Phrases: RFID, privacy, security, EPC, proximity cards

ACM Reference Format:

Juels, A and Weis, S. A. 2009. Defining strong privacy for RFID. ACM Trans. Info. Syst. Sec. 13, 1, Article 7 (October 2009), 23 pages.

DOI = 10.1145/1609956.1609963 <http://doi.acm.org/10.1145/1609956.1609963>

1. INTRODUCTION

Radio Frequency IDentification (RFID) systems offer improved efficiency in inventory control, logistics, and supply chain management. As such, they are of great interest to enterprises dependent on efficient supply chains, particularly large retailers and consumer product manufacturers. The long-term goal of these organizations is to integrate RFID on the retail level. Without proper protection, widespread adoption of retail RFID could raise privacy concerns for everyday consumers.

Authors' addresses: A. Juels, RSA Laboratories, 11 Cambridge Center, Cambridge, MA 02142; email: ajuels@rsa.com; S. A. Weis, Google, 345 Spear Street, San Francisco, CA, 94105; email: sweis@google.com.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© 2009 ACM 1094-9224/2009/10-ART7 \$10.00
DOI 10.1145/1609956.1609963 <http://doi.acm.org/10.1145/1609956.1609963>

ACM Transactions on Information and System Security, Vol. 13, No. 1, Article 7, Publication date: October 2009.

Briefly, RFID systems consist of two main components: *tags* and *readers*. Tags are radio transponders attached to physical objects. Radio transceivers, or readers, query these tags for some (potentially unique) identifying information about objects which tags are attached to. Although readers may be treated as an intermediary to a independent backend database, for simplicity, we will treat a reader and a backend database as a single entity.

Historically, tags have been expensive, bulky devices used for tagging objects like shipping containers, munitions, automobile parts, or even live cattle. The high value of such objects justifies relatively expensive RFID devices that may perform complex computations, store ample data, and initiate their own communications. Since people do not typically carry cattle or shipping containers around with them, individual privacy is not an issue in these applications.

Concerns about privacy and security do arise, however, in new consumer applications of RFID. RFID technology is present in proximity cards or “prox cards” for physical access control, automobile keys, and in a variety of contactless payment systems. Many—if not the majority—consumers in industrialized countries today make regular use of such systems. One potential application will be electronic product code (EPC) tags. EPC tags will serve on consumer products as a successor to optical “UPC” bar codes.

In the case of EPC tags, private information might be read from tags attached to products like clothing, books, or prescription drugs. An insecure prox card payment system might leak information about the movements of its users. But even if the semantic meaning of information on tags is well protected, tags may still be recognizable between appearances, and thus subject to tracking. This could violate individual location privacy for consumers or could reveal strategic information about an enterprise or military supply chain.

Because of their intentionally simple design, EPC tags cannot support expensive, traditional cryptography and security functions—not even basic ciphers. Tight economic considerations suggest that this will remain the case for the foreseeable future. Next-generation EPC tags (such as the emerging, more expensive Class 2 type) and other RFID tags, though, may be capable of performing symmetric-key cryptography. While some of the RFID literature [Fishkin et al. 2004; Juels et al. 2003] addresses privacy in tags that cannot perform cryptography, *we focus in this article on privacy in RFID systems where symmetric-key cryptographic operations are possible.*

1.1 Privacy-Preserving RFID Protocols in the Literature

Fortunately, much attention has been devoted to RFID security and privacy in recent years. Juels [2006] offers a survey of much of the related literature in and Avoine [2006] maintains a current online bibliography. Rather than reviewing the literature here, we refer the reader to those sources. While many RFID security schemes have been proposed, there has been little formal analysis. Lacking formal security definitions, much existing work offers ad hoc notions of security.

1.2 Previous RFID-Privacy Definitions

Our proposed definition is similar in flavor to those of Avoine [2005], who proposes very general and flexible definitions of RFID privacy. The aim of Avoine, however, is to capture a range of adversarial abilities, while we seek to characterize a very strong adversary with a relatively simple definition. Thus, we aim for specificity and simplicity over flexibility. For example, the Avoine definitions distinguish between adversarial eavesdropping on the reader-to-tag channel, that is, reader transmissions to tags, and the tag-to-reader channel, that is, tag transmissions to the reader. (In practice, the former is easier to eavesdrop on than the latter, as readers emit more power than tags.) In contrast, we simply assume the ability of an adversary to eavesdrop on both channels.

On the other hand, our model is more general than the Avoine models in one important respect: Ours characterizes the privacy of systems in which tags carry *correlated secrets*, as in the proposed system of Molnar, Soppera, and Wagner (MSW) [Molnar and Wagner 2004; Molnar et al. 2005]. This feature is lacking in the Avoine models.

Consider the following scheme. Suppose that x is a secret value; let h represent a hash function (modeled as a random oracle) and E represent a symmetric-key cipher in the ideal-cipher model. At time step t , tag 1 outputs $(E_x[1, t], h(1, t, x))$; tag 2 similarly outputs $(E_x[2, t], h(2, t, x))$. Tag 3 outputs $(E_x[3, t], x \oplus h(1, t, x) \oplus h(2, t, x))$. Any pair of tags may be seen to offer strong privacy: An adversary cannot distinguish among different outputs of a given tag. But clearly an adversary capable of obtaining the outputs of the three tags at a given time t can completely break the privacy of this system by recovering x . The Avoine definitions, however, permit adversarial interaction with only two tags. They would classify this system as providing strong privacy (Existential-UNT).

While this three-tag construction is slightly artificial, it actually reflects a feature in realistic RFID architectures. In the MSW system [Molnar et al. 2005], for example, tags contain overlapping sets of secret keys derived from a tree structure.

In other work, Juels [2004] presents what he calls a “minimalist” model for tag privacy and security. This model is targeted at low-cost tags that cannot perform symmetric-key cryptography. The accompanying formal privacy definition may be regarded as a specialized narrowing of our definitions here. The Juels definition specifically assumes bounds on the number of queries that an adversary can make in mounting a man-in-the-middle attack. Thus, that definition does not aim to capture RFID privacy in a strong sense. Similarly, Golle et al. [2004] define privacy in a sense specific to a cryptographic primitive they propose called “universal re-encryption” of which they mention RFID as a possible application. Their definition excludes certain forms of active adversarial attack.

Organization

In Section 2, we present our formal definition for strong RFID privacy. In Section 3, we examine several different privacy-preserving RFID systems proposed

in the literature. We show how each of these systems falls short of our strong definition of privacy; we demonstrate practical attacks that undermine important privacy characteristics. Section 4 briefly examines a privacy-preserving RFID idea called hash-locks that meets our definitional requirements. We prove that one variant meets our privacy definition and propose a simple improvement that achieves privacy in a broader sense. We conclude in Section 5.

2. FORMAL PRIVACY DEFINITION

In this section, we give a formal definition of privacy in RFID systems that can model a variety of security protocols and attacks. We regard a system as comprising a single reader \mathcal{R} and a set of n tags $\mathcal{T}_1, \dots, \mathcal{T}_n$. Each party is a probabilistic interactive Turing machine with an independent source of randomness and unlimited internal storage. Tags and readers are modeled as ideal functionalities, similar to entities in Canetti’s Universal Composability paradigm [Canetti]. Functionalities have a well-defined interface, may receive messages, and may respond with messages of their own.

We do not use the term oracle to describe tags and readers, since it sometimes implies a deterministic function where inputs always have the same output. In our model, tags and readers may maintain internal state, are randomized, and may adaptively change their output. We also want to avoid confusion when we later use oracle in the context of the random oracle model.

2.1 Tag Functionalities

Each tag functionality \mathcal{T}_i stores an internal secret *key* and a session identifier *sid*. A tag can be assigned a new *key* via a SETKEY message. A tag responds to a SETKEY message by disgorging its current key. The caller may then send an arbitrary new key to replace the prior key. A tag SID can be set to a new value *sid* via the message (TAGINIT, *sid*). TAGINIT messages delete information associated with an existing *sid*. In other words, a tag may be involved in only one protocol session at a time.

A tag may respond to a protocol message or challenge, denoted c_j , with a response r_j . The tag functionality interface is illustrated in Figure 1. We explain it further in the following text.

2.2 Reader Functionality

The reader \mathcal{R} is initialized with private key material. For the purposes of this model, this key material is immutable and internal to the reader. Tag data may be thought of as residing in a backend database containing records of the tags owned by a particular reader. The reader functionality initializes a new session upon receipt of a message of the form READERINIT. When receiving a READERINIT message, \mathcal{R} generates a fresh session identifier, *sid*, and the first challenge of an interactive challenge-response protocol, c_0 .

For each READERINIT received, the reader creates a new internal entry of the form (*sid*, “open,” c_0). Any responses containing *sid* are appended to that entry, as well as subsequent challenges, or any other auxiliary data. This entry is marked as “closed” and becomes read-only when the reader ultimately accepts

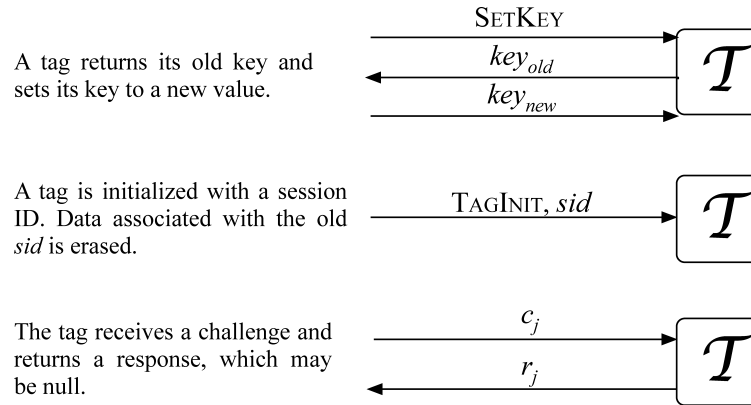


Fig. 1. Illustration of the tag functionality interface.




Fig. 2. Illustration of the reader functionality interface.

or rejects a session. The interface to a reader functionality \mathcal{R} is illustrated in Figure 2; further explanation will follow.

2.3 Functionality Interaction

The sid values generated by passing `READERINIT` messages to \mathcal{R} may be passed as a `TAGINIT` message to some \mathcal{T} . Once a tag is initialized with a sid , it can respond to challenge c_j with response r_j . This response can be a function of the key , sid , the $j - 1$ previous challenge-response pairs, and locally generated randomness. The tag may internally log both the latest received c_j and the response r_j . This protocol log and any other data except key are deleted by a tag upon receipt of a `TAGINIT` message. Thus, a tag's responses can only depend on that particular session.

Similarly, reader functionality \mathcal{R} responds to messages of the form (sid, r_j) by first searching for a $(sid, \text{"open"}, c_0, r_0, \dots, c_j)$ session entry, if any exists. At this point, \mathcal{R} performs a verification step by computing a function over its entire internal state, including all open and closed sessions and any internal key material. \mathcal{R} can output an "accept" or "reject" and mark the session entry as $(sid, \text{"closed"}, c_0, r_0, \dots, c_j, r_j)$. Alternatively, \mathcal{R} can compute the next challenge c_{j+1} and append it to the corresponding session entry.

2.4 Parameterized Adversaries

An adversary \mathcal{A} can issue its own `SETKEY`, `TAGINIT`, `READERINIT`, challenge, and response messages. We parameterize the number of `READERINIT` messages sent by \mathcal{A} with r , the number of computational steps¹ it performs by s , and the number of `TAGINIT` messages sent by t . We do not parameterize the number of explicit challenge and response messages an adversary may issue because it will be determined by the number of `TAGINIT` and `READERINIT` messages. That is, communication costs are entirely determined by r and t .

Nor do we parameterize the number of `SETKEY` messages sent by \mathcal{A} . An adversary able to issue `SETKEY` messages essentially models both the ability to corrupt a tag, since adversaries obtain its internal tag keys, and to manufacture arbitrary tags, since adversaries can then set tag's *key* to an arbitrary value. For this reason, we will consider any tag that receives a `SETKEY` message from the adversary as “corrupted.” We assume an adversary is able to send as many `SETKEY` messages as it likes (within its computational step bound s), at any time, to any set of at most $n - 2$ tags. The rationale for this $n - 2$ bound will be explained in the following section.

2.5 Privacy Experiment

We now present a privacy experiment reminiscent of the classic indistinguishability under chosen-plaintext attack (IND-CPA) and under chosen-ciphertext attack (IND-CCA) cryptosystem security experiments. The idea is that an RFID protocol may be considered private for some parameter values if no adversary has a significant advantage in this experiment. The goal of the adversary in this experiment is to distinguish between two different tags within the limits of its computational power and functionality-call bounds.

That is why at least two tags need to be uncorrupted. In the first phase of this experiment, an adversary \mathcal{A} will be able to issue any messages and perform any computation within its parameter bounds. The adversary \mathcal{A} controls the communication channel between \mathcal{R} and each \mathcal{T}_i , and may obtain or corrupt any *sid*, c_j , or r_j values at will.

The adversary selects two uncorrupted tags as challenge candidates. These challenge candidates are removed from the pool of tags at large. One of these challenge candidates is then randomly selected and presented to \mathcal{A} (effectively as a tag oracle). \mathcal{A} may then issue any messages and perform any computation it likes within its parameter bounds, *except* issuing a `SETKEY` command to the tag oracle. (For most protocols, that would make the experiment trivial, although later we briefly explore variants allowing corruption of challenge candidates.)

All key material is generated using a randomized function $\text{GEN}(1^k) \rightarrow (key_0, \dots, key_n)$. This function is considered part of the overall RFID system and thus privacy depends on its characteristics. This function is computable by an adversary. Obviously, its output in the experiment is kept secret from \mathcal{A} .

We denote our privacy experiment for an RFID system by $\text{Exp}_{\mathcal{A}, \mathcal{S}}^{\text{priv}}[k, n, r, s, t]$. Here, $\mathcal{S} = (\text{GEN}, \mathcal{R}, \{\mathcal{T}_i\})$, where $\{\mathcal{T}_i\}$ contains n tags. We let k be a security

¹We are using “computational steps” in the context of a traditional RAM model of computation.

Experiment $\mathbf{Exp}_{\mathcal{A},\mathcal{S}}^{priv}[k, n, r, s, t]$:

Setup:

- (1) $\text{GEN}(1^k) \rightarrow (key_0, \dots, key_n)$.
- (2) Initialize \mathcal{R} with (key_0, \dots, key_n) .
- (3) Set each \mathcal{T}_i 's key to key_i with a SETKEY call.

Phase 1 (Learning):

- (4) \mathcal{A} may do the following in any interleaved order:
 - (a) Make READERINIT calls, without exceeding r overall calls.
 - (b) Make TAGINIT calls, without exceeding t overall calls.
 - (c) Make arbitrary SETKEY calls to any $(n - 2)$ tags.
 - (d) Communicate and compute, without exceeding s overall steps.

Phase 2 (Challenge):

- (5) \mathcal{A} selects two tags \mathcal{T}_i and \mathcal{T}_j to which it did *not* send SETKEY messages.
 - (6) Let $\mathcal{T}_0^* = \mathcal{T}_i$ and $\mathcal{T}_1^* = \mathcal{T}_j$ and remove both of these from the current tag set.
 - (7) Let $b \xleftarrow{R} \{0, 1\}$ and provide \mathcal{A} access to \mathcal{T}_b^* .
 - (8) \mathcal{A} may do the following in any interleaved order:
 - (a) Make READERINIT calls, without exceeding r overall calls.
 - (b) Make TAGINIT calls, without exceeding t overall calls.
 - (c) Make arbitrary SETKEY calls to any tag in the current tag set *except* \mathcal{T}_b^* .
 - (d) Communicate and compute, without exceeding s overall steps.
 - (9) \mathcal{A} outputs a guess bit b' .
- Exp** succeeds if $b = b'$.

Fig. 3. The parameterized privacy experiment for the RFID system $(\text{GEN}, \mathcal{R}, \{\mathcal{T}_i\})$.

parameter. Adversary \mathcal{A} with parameters r, s , and t is denoted by $\mathcal{A}[r, s, t]$, where r, s and t are respective parameters for reader initializations, computation steps, and tag initializations. We detail our experiment in Figure 3.

2.6 Privacy Definition

A protocol run within an RFID system $\mathcal{S} = (\text{GEN}, \mathcal{R}, \{\mathcal{T}_i\})$ is defined to be private if no adversary $\mathcal{A}[r, s, t]$ has a nonnegligible advantage in successfully guessing b in the experiment in Figure 3.

We now define the notion of (r, s, t) -privacy for RFID systems. The variables r, s , and t can be functions of system parameters like k , if desired, as we shall see in our analyses. We use the notation $\text{poly}(k)$ to represent any polynomial function of k .

Definition 2.1 (RFID (r, s, t) -Privacy). A protocol initiated by \mathcal{R} in an RFID system $\mathcal{S} = (\text{GEN}, \mathcal{R}, \{\mathcal{T}_1, \dots, \mathcal{T}_n\})$ with security parameter k is (r, s, t) -private if:

$$\forall \mathcal{A}[r, s, t] \Pr [\mathbf{Exp}_{\mathcal{A},\mathcal{S}}^{priv}[k, n, r, s, t] \text{ succeeds in guessing } b.] \leq \frac{1}{2} + 1/\text{poly}(k).$$

2.7 Implications of (r, s, t) -Privacy

A few observations on our privacy definition are in order. Definition 2.1 requires that \mathcal{R} not differ significantly in its acceptance rate across tags. Otherwise, \mathcal{A} could just pick, for example, one tag that \mathcal{R} accepts frequently and one that it accepts less frequently as its challenge candidates. \mathcal{R} must essentially own all the tags in $\{\mathcal{T}_i\}$, or at least treat them all equally.

\mathcal{R} also cannot be trivially history-dependent nor can its acceptance be influenced significantly by \mathcal{A} . For example, an \mathcal{R} that accepts each tag only a fixed number m times cannot be private in the face of an \mathcal{A} able to make more than $m + 1$ `READERINIT` calls. Such an adversary could simply complete m legitimate queries to a tag Phase 1, and see if the reader accepts T_b^* in Phase 2. We discuss exactly this attack in Section 3.

Our definition treats protocol-level privacy issues only. In the real world, there are effectively many possible side channels. As Avoine and Oechslin [2005a] observe, for example, if tags emit a distinctive radio fingerprint, then no protocol-level countermeasures can prevent privacy infringement. And, of course, bearers of privacy-preserving RFID tags are vulnerable to tracking if they additionally carry wireless devices that lack privacy protections.

An important implication of Definition 2.1 is that the `GEN` function cannot output low-entropy or strongly correlated keys. Otherwise, \mathcal{A} might have a significant advantage in guessing keys given the output from `SETKEY` calls. This illustrates how privacy by our definition *depends on the entire RFID system* ($\text{GEN}, \mathcal{R}, \{T_i\}$). There are classes of \mathcal{R} and `GEN` that simply cannot be private under this definition.

2.7.0.1 Remarks. (i) A superficial difference between the Avoine model and our own is that the target tags in the Avoine model are determined not by the adversary, but by a Challenger. This Challenger has unspecified functionality, but presumably selects target tags uniformly at random. In an asymptotic sense, adversarial and random selection of target tags are equivalent: Random selection will yield any desired pair of target tags with nonnegligible probability.

(ii) One may consider a variant of $\text{Exp}_{\mathcal{A}, \mathcal{S}}^{\text{priv}}$ experiment with a stronger adversary \mathcal{A} that is allowed to corrupt T_b^* in Phase 2 of the experiment, that is, removing the words “except T_b^* ” from experiment step (8c). In most schemes, this will trivially violate privacy. It is worthwhile, however, to consider this definition of *forward* (r, s, t) -privacy. In a scheme with forward privacy of this kind, corrupting a tag does not link it to its past output. (Forward privacy is a goal of the OSK/AO scheme that we analyze later in the text.) Forward privacy is strictly stronger than (r, s, t) -privacy as defined in Definition 2.1.

2.8 Cross-Reader (r, s, t) -Privacy

A limitation of our definition is that it models a world in which a reader \mathcal{R} owns every tag, that is, a monolithic and closed system. While this model provides important insights into the privacy properties of an RFID system, different (weaker) definitions can be useful for analyzing real-world scenarios.

Consider an RFID system that contains two types of tags, type A and type B. The reader \mathcal{R} recognizes, that is, accepts valid tags of type A. But it rejects all tags of type B. Our definition suggests that such a system can *never* be privacy preserving. An attacker can break the privacy of the system by distinguishing between type A and type B tags.

In the real world, though, such scenarios frequently arise. There are many, independent RFID systems; people carry heterogeneous constellations of tags, such as SpeedPass tags, contactless credit cards, proximity cards, automobile immobilizer chips, and so forth. The very diversity of tag types leaks information. An unusual collection of tags could even be uniquely identifying. Thus, while our privacy definition is useful for characterizing privacy within a single RFID system, an extended or variant definition is needed to analyze multiverifier privacy—and capture its inherent limitations.

For this purpose, a different (weaker) privacy definition is useful. An attacker capable of mounting unrestricted man-in-the-middle attacks can determine whether or not a given tag is recognized by a given reader \mathcal{R} . Such an attacker can, of course, easily violate the privacy of an RFID-tag bearer by determining the presence of type A and type B tags on her person. For this reason—and also because man-in-the-middle attacks can be tricky to mount in RFID environments—we propose a variant of our privacy definition. This variant excludes man-in-the-middle attacks by means of a very simple restriction on the adversary in $\mathbf{Exp}_{\mathcal{A},\mathcal{S}}^{priv}[k, n, r, s, t]$: When \mathcal{A} emits a `READERINIT` call, the system first emits `TAGINIT` calls with random *sids* to all tags. We call the modified experiment $\mathbf{Exp}_{\mathcal{A},\mathcal{S}}^{cross-read-priv}[k, n, r, s, t]$, and characterize a system in the obvious way as *cross-reader* (r, s, t) -private.

Cross-reader (r, s, t) -privacy is useful in analyzing an environment with multiple, independent verifiers (RFID readers). We can model such a system as an environment with just one reader \mathcal{R} ; we simply assume type A and type B tags as described earlier. Again, type A tags are those that lie within the closed system for reader \mathcal{R} , and type B tags are those in other, independent systems.

What is especially interesting about cross-reader (r, s, t) -privacy is its intersection with tag authentication. As we show in the following text, in a system where \mathcal{R} recognizes all tags, that is, all tags are of type A, the original Weis et al. [2003] hash-lock system is (r, s, t) -private and cross-reader (r, s, t) -private for polynomial r, s, t . In a system with tags of types A and B, however, the Weis et al. system is not cross-reader private. This is true because the system does not provide strong authentication for tags. Without mounting a man-in-the-middle attack, an adversary can simulate a target RFID tag and determine whether it is of type A or B, thereby undermining privacy. As we show, however, by adding strong authentication to the Weis et al. system, we can confer cross-reader (r, s, t) -privacy for polynomial r, s, t in a multiverifier environment.

2.8.0.2 Remark. Even cross-reader privacy definition is unachievable when type A and type B tags may be distinguished outside the protocol layer. For example, if the two types operate at different frequencies, an attacker can distinguish between them without reference to any valid reader. In such cases, the best privacy that can be hoped for is indistinguishability, that is, cross-reader privacy among tags of a given physical type.

3. ANALYSIS OF SEVERAL PROPOSED RFID SCHEMES

We now examine several RFID privacy schemes proposed in the literature: (1) That of Ohkubo, Suzuki, and Kinoshita (OSK) [2004] and a variant proposed

by Avoine and Oechslin (AO) [2005b] (with a follow-up by Avoine, Dysli, and Oechslin [2005]); (2) That of Nohara, Inoue, Baba, and Yasuura (NIBY) [2005]; (3) The YA-TRAP scheme of Tsudik [2006b]; and (4) A recent zero-knowledge scheme of Engberg, Harning, and Jensen (EHJ) [2004].

In our analyses of OSK/AO and NIBY, we demonstrate an adversarial algorithm that shows that the algorithms do not achieve strong privacy according to our definition. We believe that this adversarial algorithm is practical and that it reveals a privacy vulnerability that has gone unnoticed in the literature. The vulnerability is similar in both OSK/AO and NIBY. It involves an adversary interacting with a target tag and modifying its state. The aim of the modification is to mark the target tag so that it is recognizable by the adversary at a later time. The mark takes the form of state that is invalid in the view of a legitimate RFID reader. The adversary effectively plants a timestamp in a tag that is too far in the future to be treated as valid by the reader.

We believe that this vulnerability in tags that do not accept reader input has been overlooked in the literature because it bypasses the natural intuitive view of privacy modeling. For example, in the OSK/AO system, tags do not accept input from a querying reader (apart from a trivial read command). The only state change in tags is the incrementing of an internal counter. Thus, these systems appear to reduce an adversary to strictly passive attacks. Our formal definition of privacy, however, highlights the fact that an adversary can in fact mount an active attack by aggressively changing the internal state of tags. Tags in OSK/AO and NIBY systems are counter based and output-only. We also examine two systems in which readers modify tag state by means of timestamps.

The YA-TRAP system is a new proposal by Tsudik that has not yet received formal security analysis. Tsudik already notes that the system does not withstand the full range of possible active attacks. With these limitations in mind, Tsudik has designed YA-TRAP for settings in which RFID-tag information is batched, as in warehouses, rather than for consumer applications. Our analysis gives firmer shape to this design aim. We show a privacy limitation in which, as in OSK/AO, a tag may be marked by being invalidated. More interestingly, an extension of this attack can mark a tag with a unique timestamp that may be indirectly read by an adversary. Our analysis confirms Tsudik's view of the circumscribed settings for which YA-TRAP is appropriate.

The EHJ system is particularly interesting because it is under commercial development, with privacy being emphasized as its distinguishing feature. The vulnerability we highlight in the EHJ system is of a different nature than the others. In that system, an adversary can intercept a flow in the authentication protocol between a tag \mathcal{T}_i and the reader, and interrupt the associated session. By resuming the session at a later time using the intercepted flow, the adversary can uniquely identify \mathcal{T}_i .

There is one common thread across all of the vulnerabilities we explore. As our formal definition shows, even a single bit of information—namely whether a tag functions or malfunctions—has an important bearing on RFID-system privacy.

3.1 OSK/AO Schemes

OSK [2004] propose a simple scheme for privacy-preserving tag identification in the face of active attacks; the OSK scheme also provides forward security. Suppose that tag i is initialized with secret x_i ; let h_1 and h_2 be independent one-way functions (most conveniently modeled as random oracles). When queried, the tag updates its secret key by applying h_1 . The tag outputs $ID_{i,t} = h_2(h_1^{(t)}(x_i))$.

Readers in the OSK scheme share secrets with tags. On input $ID_{i,t}$, a reader determines i by brute-force search. OSK propose a fixed upper bound m on the number of time steps over which tags operate. (After m time steps, presumably a tag yields no output or random output while it awaits reinitialization or retirement.) OSK propose that readers precompute a giant table $T = \{[ID_{i,t}, (i, t)]\}$, where $1 \leq i \leq n$ and $1 \leq t \leq m$. When structured appropriately, T can support rapid look-up of a tag identifier. Kinoshita et al. [2005] have implemented the OSK protocol in an active RFID tag.

AO [Avoine and Oechslin 2005b; Avoine et al. 2005] propose use of Hellman tables as a special form of precomputation on T . Hellman [1980] originally studied the problem of breaking symmetric keys. He considered the resource requirements of an attacker seeking to recover secret key k from a ciphertext $C = e_k(M)$ on a predetermined message M . Of course, the attacker might simply perform a brute-force key search, investing computational effort $O(n)$, where n is the size of the full key space. Or she might precompute and store a table of size $O(n)$ consisting of $C = \{C_i = e_{k_i}(M)\}_{k_i \in K}$, and simply look up C in the table; this is what OSK originally proposed.

Hellman, however, identified a useful time-space trade-off. An attacker can pre-compute a table of size $O(N^{2/3})$, now known as a Hellman table, that permits successful key search with computational effort only $O(N^{2/3})$. Loosely speaking, the idea is to group sequences of ciphertexts in C into deterministically traversable chains of size $O(N^{2/3})$. Within the Hellman table are stored only the first and last elements of these chains. Starting with a ciphertext C , the attacker can traverse the induced chain until she locates the last element. By consulting the Hellman table, she can then determine the first element in the chain; on traversing the chain from this point, she can determine k_i .

AO cleverly observed that the problem of symmetric-key search for readers in privacy-preserving RFID systems is nearly identical with that of breaking keys. For a reader to determine which tag it is communicating with, the reader can crack the ciphertext C_i to determine k_i and thus T_i . AO apply a variant of the Hellman technique to the OSK system, yielding a scheme with considerably more practical look-up costs. (They in fact propose use of a variant called *rainbow chains* put forth by Oechslin [2003].) Rather than constructing a simple look-up table T , a reader can make use of a Hellman table \tilde{T} . The AO approach can thus render the original OSK more practical.²

²Note, however, that because a Hellman table \tilde{T} requires precomputation over a fixed number of entries, it is unclear how to modify the table to accommodate new entries. Insertion of new entries into a simple table T is more straightforward.

- (1) In Phase 1, \mathcal{A} selects a pair of distinct tags \mathcal{T}_i and \mathcal{T}_j uniformly at random.
- (2) \mathcal{A} sends m queries to \mathcal{T}_i (sends m TAGINITS).
- (3) \mathcal{A} submits \mathcal{T}_i and \mathcal{T}_j as its challenge candidates.
- (4) In Phase 2, \mathcal{A} initializes a protocol between \mathcal{T}_b^* and \mathcal{R} .
- (5) If \mathcal{R} accepts \mathcal{T}_b^* , \mathcal{A} guesses $b = 1$, i.e. $\mathcal{T}^* = \mathcal{T}_j$. Else \mathcal{A} guesses $b = 0$, i.e. $\mathcal{T}^* = \mathcal{T}_i$.

Fig. 4. Adversarial algorithm for OSK/AO.

3.1.0.3 Privacy Vulnerability in OSK/AO. Avoine analyzes the OSK (and, by extension, the AO) scheme for tag identification, and concludes that it affords the strongest form of privacy [Avoine 2005]. As we now explain, however, OSK is in fact vulnerable to an attack that undermines strong privacy in the system.³

Very simply, an adversary can exploit the upper bound m on the number of tag time steps accommodated in the look-up table T or \tilde{T} . The adversary queries a tag until its counter value exceeds m . In essence, it mounts a denial-of-service attack against the tag: The adversary exhausts the supply of valid output values stored in the tag.

This strategy depends on the ability of the adversary to obtain a single bit of information from a reader, namely whether or not it recognizes the output of a given tag as valid. For most natural RFID tag applications, such information will be easily obtainable. For example, a proximity card either opens or fails to open a door, a payment card is either accepted or rejected by a point-of-sale device, and so forth. The adversarial algorithm in Figure 4 demonstrates that OSK/AO does not achieve (r, s, t) -privacy for $t > m$ and $r \geq 1$ (and $s > m$). Indeed, for these parameters, $\mathbf{Exp}_{\mathcal{A}, S}^{priv}[n, k, r, s, t] = 1$.

A simple modification of the adversarial algorithm in Figure 4 allows an adversary to distinguish among multiple tags by desynchronizing them in different degrees. The adversary queries \mathcal{T}_i a total of q_i times to mark it; the adversary lets q_i be unique for each tag. On later sighting a target tag \mathcal{T} , the adversary queries \mathcal{T} a total of m times. Let r_1, r_2, \dots, r_m be the resulting outputs. The adversary feeds these outputs to the reader in reverse order r_m, r_{m-1}, \dots until the reader accepts some r_k . The adversary concludes that the counter value for the tag $q = m - k$. Assuming that no normal reading has taken place during the interval of time between the marking and the sighting of the tag, the adversary concludes that if $q = q_i$, the sighted tag is \mathcal{T}_i . Note that if the adversary staggers q_i values sufficiently, this strategy works even when some limited amount of normal reading has effected changes in tag counter values.

The adversarial algorithm we describe here would appear to be feasible in a practical setting. Avoine and Oechslin [2005b] propose $m = 1,024$ as an example parameter for a practically realizable AO system. As EPC tags (Class 1 Gen 2) have a theoretical read rate of up to 1,000 per second [Alien Technology 2005], the adversarial algorithm we describe could well be feasible in many systems. In the original OSK system, which is less efficient than the AO variant, m would have to be considerably smaller to permit practical deployment, and the vulnerability we describe more easily exploitable.

³The vulnerability we demonstrate appears to exist even within the Avoine privacy model. The analysis in Avoine [2005] may simply represent an oversight.

3.1.0.4 *Remarks.* (i) The adversarial strategy we describe for OSK/AO here induces desynchronization between tags and the reader. It is tempting to try to evade this problem with a protocol modification in which the reader rejects input from a \mathcal{T}_i if the current tag counter q_i differs by more than d from the previous counter perceived by the reader. This does not solve the privacy problem; however, it is equivalent to setting $m = d$. Moreover, if d is small, then the system becomes vulnerable to denial-of-service attacks. Another possible countermeasure is simply to cause tags to stop yielding output after the m^{th} query. Again, this creates a vulnerability to denial-of-service attacks.

In a natural setting, the adversary might not always have access to fresh tags, that is, tags whose counters are equal to 0. A simple modification of the adversarial algorithm in Figure 4, however, can still distinguish between tags. By querying a target tag a sufficient number of times, the adversary can mark it. For example, suppose that $m = 1,024$, and suppose that prior to the testing phase of our definition experiment, \mathcal{T}_i and \mathcal{T}_j have each been queried q_i and q_j times, where $q_i, q_j \in_U [0, 1, 2, \dots, 50]$. Then, the adversary merely has to submit 51 queries to \mathcal{T}_i to mark it such that it may be distinguished from \mathcal{T}_j in the guessing Phase 2.

One possible countermeasure to such attacks is throttling, that is, deliberate slowing of the rate of tag response. A rather different privacy model for tags would be needed to accommodate the effect of this countermeasure, however [Juels 2004].

(ii) MSW carries a similar vulnerability in principle, as tags can support only a pre-determined number of queries, as observed by Molnar et al. [2005]. Under realistic parameterization, however, a privacy-infringing attack would require an infeasibly large number of queries. In a practical sense, therefore, MSW is not vulnerable.

3.2 Unlinkable ID-Matching Scheme

At the 2005 Workshop on Privacy in the Electronic Society, Nohara, Inoue, Baba, and Yasuura (NIBY) [Nohara et al. 2005] proposed a scheme similar in flavor to MSW. Their basic protocol permits leaves of potentially nonuniform depth within a tree T of maximum depth d . A tag is presumed to output random values when depth-first-search queries exceed the depth of its corresponding leaf.

An adversary can undermine strong privacy in this scheme in much the same manner as OSK/AO; nonuniformity in the NIBY scheme, however, renders it more vulnerable. In particular, for a given tag \mathcal{T}_i , an adversary can learn the depth d_i of the corresponding leaf in the identifier tree with overwhelming probability using the following algorithm:

Algorithm Probe(d, \mathcal{T}_i).

- (1) Send d queries to \mathcal{T}_i (send d TAGINITS.); let $Z_e = \{z_e^{(i)}\}_{i=1}^d$ denote the sequence of d outputs yielded by the e^{th} query, for $1 \leq e \leq d$.
- (2) For every e , $1 \leq e \leq d$, replace $z_e^{(d)}, z_e^{(d-1)}, \dots, z_e^{(e)}$ with random values.
- (3) Feed Z_e, Z_{e-1}, \dots to \mathcal{R} until \mathcal{R} accepts some $Z_{e'}$.
- (4) Output e' .

- (1) In Phase 1, \mathcal{A} selects a pair of distinct tags \mathcal{T}_i and \mathcal{T}_j uniformly at random.
- (2) \mathcal{A} computes $d_i = \text{Probe}(d, \mathcal{T}_i)$ and $d_j = \text{Probe}(d, \mathcal{T}_j)$.
- (3) \mathcal{A} submits \mathcal{T}_i and \mathcal{T}_j as its challenge candidates.
- (4) In Phase 2, \mathcal{A} computes $d' = \text{Probe}(d, \mathcal{T}_b^*)$.
- (5) If $d' = d_j$, \mathcal{A} guesses $b = 1$, i.e. $\mathcal{T}^* = \mathcal{T}_j$. Else, \mathcal{A} guesses $b = 0$, i.e. $\mathcal{T}^* = \mathcal{T}_i$.

Fig. 5. Adversarial algorithm for NIBY.

Suppose that the probability that two leaves in the tree T differ in depth is non-negligible. With use of the Probe algorithm as a subroutine, we can construct an adversarial algorithm that breaks (r, s, t) -privacy in NIBY for $r, t > 3d$ (and $s > 2d$). This algorithm is depicted in Figure 5.

A reader can ostensibly detect this type of attack. In practice, however, the natural noisiness of RFID systems might make such detection difficult.

The authors analyze the security of a restricted case of their scheme in which all leaves are presumed to have equal depth. The adversarial algorithm we describe here is not applicable in that case. Our analysis, therefore, suggests that use of a uniform tree for NIBY is essential to privacy.

3.3 YA-TRAP: Yet Another Trivial RFID Authentication Protocol

Tsudik [2006b] describes an RFID authentication protocol that he calls YA-TRAP (yet another trivial RFID authentication protocol). As discussed earlier, Tsudik aims this protocol at environments in which tag information is processed in batches, rather than for more fine-grained applications like access control and tagging of individual consumer items. Tsudik [2006a] has not yet formally analyzed the security properties of YA-TRAP; such analysis is forthcoming. In the meantime, our investigation here confirms the importance of confining YA-TRAP to the limited settings for which it is designed.

In YA-TRAP, the reader shares a unique key x_i with tag \mathcal{T}_i . A tag \mathcal{T}_i also stores an internal timestamp p_i namely the time at which it was last interrogated by a reader.⁴ To interrogate a tag, a reader transmits the current time $p_{\mathcal{R}}$. The tag compares $p_{\mathcal{R}}$ with p_i . If $p_{\mathcal{R}}$ is stale with respect to p_i , that is, $p_{\mathcal{R}} \leq p_i$, then the tag outputs a random response. Otherwise, the tag outputs $R = H_{x_i}[p_{\mathcal{R}}]$, where H_{x_i} represents an HMAC computed with secret key x_i ; the tag also sets $p_i \leftarrow p_{\mathcal{R}}$, i.e., updates its timestamp.⁵ To validate the response of a tag, a reader checks whether $R = H_{x_j}[p_{\mathcal{R}}]$ for any secret key x_j in its database. If so, the reader accepts the tag, otherwise it rejects it.

We consider two adversarial strategies against the YA-TRAP protocol. The first is based on denial of service. An adversary queries a tag with p_{max} to mark it, where p_{max} is some time in the far future. The tag sets its internal time value $p_i \leftarrow p_{max}$, and outputs random values in response to all future queries. Thus, a reader rejects the tag in all future sessions, permitting the adversary

⁴The original article uses t to denote timestamps. We replace this with p to avoid confusion with the security parameter in the (r, s, t) privacy definition.

⁵The premise is that HMAC output is indistinguishable from a random distribution, so successful or failed tag responses are indistinguishable to an adversary. While this premise is not true of HMAC in general, it may be achieved with an appropriate underlying hash function.

- (1) In Phase 1, \mathcal{A} selects an arbitrary pair of distinct tags \mathcal{T}_i and \mathcal{T}_j .
- (2) \mathcal{A} sends a TAGINIT to \mathcal{T}_i and sends challenge $c_0 = p_{max}$.
- (3) \mathcal{A} submits \mathcal{T}_i and \mathcal{T}_j as its challenge candidates.
- (4) In Phase 2, \mathcal{A} initializes a query by sending a READERINIT to \mathcal{R} and a TAGINIT to \mathcal{T}_b^* .
- (5) \mathcal{A} relays \mathcal{R} 's challenge $c_0 = p_{\mathcal{R}}$ to \mathcal{T}_b^* and the response r_0 back to \mathcal{R} .
- (6) If \mathcal{R} accepts r_0 , \mathcal{A} guesses $b = 0$, i.e. $\mathcal{T}^* = \mathcal{T}_i$, and otherwise guesses $b = 1$.

Fig. 6. Adversarial DoS algorithm for YA-TRAP.

to distinguish the tag from one that is unmarked. We outline this adversarial strategy in Figure 6.

As the adversarial algorithm in Figure 6 demonstrates, the YA-TRAP protocol does not achieve $(r, s, t) = (1, O(k), 2)$ -privacy. It is worth noting that the Phase 1 portion of the attack does not involve \mathcal{R} and can be mounted by an adversary with access to \mathcal{T}_i alone. The Phase 2 portion does not require a man-in-the-middle setting. A real-life adversary could wait until \mathcal{R} tries to legitimately query \mathcal{T}_b^* , and then eavesdrop, that is, play a purely passive role.

Building on this first strategy, we arrive at a second and more potent adversarial strategy that allows an adversary to mark a tag uniquely—not simply to invalidate it. To mark a tag \mathcal{T}_i in this way, the adversary selects a future time u_i . The adversary queries tag \mathcal{T}_i with it, causing $p_i \leftarrow u_i$. The time u_i serves as a uniquely individuating mark for tag \mathcal{T}_i . Starting at any time prior to u_i , an attacker can test whether a given, unidentified tag \mathcal{T} is in fact \mathcal{T}_i . The process involves two stages.

- (1) Probing \mathcal{T} . The adversary selects two slightly spaced times u_i^{before} and u_i^{after} such that $u_i^{before} \leq u_i < u_i^{after}$. The adversary then queries the tag \mathcal{T} with u_i^{before} and u_i^{after} , obtaining responses r_{before} and r_{after} .
- (2) Testing the results. The adversary interacts with \mathcal{R} at time u_i^{before} and u_i^{after} and replays responses r_{before} and r_{after} . If the reader accepts r_{before} and rejects r_{after} , then $\mathcal{T} = \mathcal{T}_i$ with negligible error.

Note that this strategy does not require an active man in the middle. \mathcal{A} can interact with the tag and the reader at separate times. This adversarial algorithm can be extended, of course, to mark multiple tags uniquely. In the real world, an attacker could identify an individual on the basis of a single marked tag \mathcal{T}_i —regardless of how many other tags the consumer is carrying.

In the setting for which YA-TRAP was designed, where a reader batches tag information exclusively for processing by a backend system that does not reveal information to an adversary about the results of its identification of tags, YA-TRAP is not subject to the attacks we describe.

3.3.0.5 Remark. If the granularity of time in the system is small, then it may be difficult for \mathcal{A} to interact with \mathcal{R} at exactly time u_i^{before} or u_i^{after} . To increase its likelihood of success, \mathcal{A} can query the tag at multiple times around u_i^{before} and u_i^{after} , that is, harvest a staggered set of responses from the target tag.

By querying each of n tags, an adversary can set the time value of \mathcal{T}_i to a unique time value p_{2i} . These time values are interleaved among a set of marker values of the form p_{2i-1} such that $p_1 > p_2 > \dots > p_{2n-1} > p_{2n} > p_{2n+1}$.

Instead of just being able to identify a tag among two challenges \mathcal{T}_i and \mathcal{T}_j like in the (r, s, t) -privacy experiment, an adversary will now be able to uniquely identify some \mathcal{T}^* from any of these n possibilities. To do so, an adversary \mathcal{A} sends `READERINIT` messages to have \mathcal{R} repeatedly query \mathcal{T}^* . A legitimate \mathcal{R} 's time value will increase monotonically. To identify \mathcal{T}^* , \mathcal{A} will wait for the first time value t^* that \mathcal{T}^* is accepted. If $p_{2i-1} < t^* < p_{2i+1}$, then \mathcal{A} will identify \mathcal{T}^* as \mathcal{T}_i .

This attack could even identify a particular tag from a set of heterogeneous unknown \mathcal{T}^* tags. In the real world, this could identify an individual carrying a tag marked with a specific time value t^* , regardless of how many other heterogeneous RFID systems are present.

3.4 A Zero-Knowledge RFID Protocol

Engberg, Harning, and Jensen [2004] (EHJ) have proposed a privacy-preserving RFID protocol that they call zero-knowledge device authentication. (The protocol, as they note, is not zero-knowledge in the received sense of the term.) Media reports suggest that this protocol is serving as the basis of a commercial offering by a company called RFIDSec [Kahn 2005]. As we show, their protocol does not provide strong privacy according to our definition.

The authors do not fully explain a core aspect of their system. Their authentication protocol involves reader-side initiation, but they do not indicate how a reader is to determine which tag it is communicating with (and thus which key to use). Therefore, we assume that the number of tags in the vicinity of the reader is small, such that it can try all keys in attempting to communicate with a tag. We also gloss over the question of how a reader singulates a tag, as this has little bearing on our analysis.

The protocol involves two passes. To use the notation of EHJ, DT is a timestamp, and RSK is a random nonce selected by the reader. The value $SSDK$ is a tag-specific secret key. We let \oplus denote the XOR operation:

- (1) $\mathcal{R} \rightarrow \mathcal{T}: [DT, (RSK \oplus h(DT \oplus SSDK)), h(RSK \oplus SSDK)]$
- (2) $\mathcal{T} \rightarrow \mathcal{R}: h(RSK \oplus SSDK \oplus DT)$.

A tag stores the timestamp DT' that it received in the last successful protocol execution. When it receives a new authentication message based on some DT , it only responds if the message is valid (with respect to $SSDK$) and if $DT > DT'$. Otherwise, the tag outputs no response.

3.4.0.6 Privacy Vulnerability in EHJ. A privacy vulnerability in this system lies in the way that a tag decides whether or not to respond to a first-flow message. The first flow of the protocol is meant to permit the reader to authenticate to the tag. This flow, however, is keyed according to the tag with which the reader communicates. Suppose that adversary intercepts and stores a first flow message M directed to \mathcal{T}_i . If the adversary transmits M to \mathcal{T}_i (while M is still fresh), then \mathcal{T}_i will output a response. Any other tag will disregard M .

Verification of $DT > DT'$ is meant to serve against replay attacks. It is easy to see, however, that this protocol feature causes a tag to reveal 1 bit of information, namely whether or not DT' is fresh in the view of the tag. As

- (1) In Phase 1, \mathcal{A} selects a pair of distinct tags \mathcal{T}_i and \mathcal{T}_j uniformly at random.
- (2) The attacker causes the reader to issue an authentication message M to \mathcal{T}_i at time DT ; \mathcal{A} stores M and does not allow it to reach the tag.
- (3) In Phase 2, the attacker transmits M to tag \mathcal{T}_b^* .
- (4) If \mathcal{T}_b^* outputs a response, \mathcal{A} guesses $b = 0$, i.e. $\mathcal{T}^* = \mathcal{T}_i$. Else \mathcal{A} guesses $b = 1$, i.e. $\mathcal{T}^* = \mathcal{T}_j$.

Fig. 7. Adversarial algorithm for EHJ.

we show, an attacker can exploit this information to undermine the privacy protection of the system.

Suppose that we execute our privacy experiment with two tags, \mathcal{T}_0 and \mathcal{T}_1 . In Phase 1, the attacker causes the reader to authenticate to \mathcal{T}_0 at time DT_0 ; the attacker further causes the reader to issue an authentication message M to the tag again at time $DT_0' > DT_0$, but intercepts this message (i.e., stores it and does not allow it to reach \mathcal{T}_0). Finally, the attacker causes the reader to issue an authentication message to \mathcal{T}_1 and \mathcal{T}_1 at time DT_1 , where $DT_1 > DT_0'$.

At this point, observe that message M carries a timestamp that is fresh with respect to the stored timestamp DT_0 in \mathcal{T}_0 , but stale with respect to the stored timestamp DT_1 in \mathcal{T}_1 . Thus, in the second phase of our experiment now, the attacker merely has to transmit M to tag \mathcal{T}_b^* . If the tag accepts the message and transmits a response, the attacker concludes that $b = 0$; otherwise, the tag ignores M , and the attacker concludes that $b = 1$.

In Figure 7, we specify an adversarial algorithm that shows that EHJ is not (r, s, t) -private for any $r \geq 1$ and $t \geq 2$. Of course, an adversary can extend this strategy to discriminate easily among multiple tags, $\mathcal{T}_0, \mathcal{T}_1, \dots, \mathcal{T}_{n-1}$. The adversary need simply harvest first-flow messages M_0, \dots, M_{n-1} for each of the tags. To identify a given tag, the adversary tests each of its harvested messages.

One might think that if a tag does not respond to a read request that no privacy issue arises. This idea is misleading for two reasons. First of all, the absence of a tag can be conspicuous. For example, if the tag is a disabled proximity card, the fact of an individual being unable to enter a building is an easily observable event. More importantly, though, in the EHJ scheme, a tag must provide some indication of its existence at the outset. If it does not, then the reader cannot determine that it should initiate the first flow. Any practical embodiment of EHJ therefore would probably need at a minimum to involve a liveness signal from tags.

3.4.0.7 Vulnerability in EHJ with Shared Keys. Even if SSDK is shared across tags (which would make the protocol more scalable), a privacy-infringing attack is still possible. In particular, the adversary causes the reader to set timestamps DT_0 and DT_2 in tags \mathcal{T}_i and \mathcal{T}_j , respectively, by causing two authentication sessions to run to completion. The adversary also causes the reader to output a first-flow message M at time DT_1 , where $DT_0 < DT_1 < DT_2$. The message M will then be fresh with respect to \mathcal{T}_i , but stale with respect to \mathcal{T}_j , permitting the adversary to distinguish between them. This adversarial strategy too can be extended in an obvious manner to distinguish among multiple tags.

3.4.0.8 *Repairing EHJ.* The EHJ protocol can easily be modified to alleviate the privacy vulnerability we present here. Rather than ignoring an authentication message with a stale timestamp, a tag can output a random or pseudorandom response (computationally indistinguishable from a legitimate one). In the absence of a carefully specified protocol, however, we do not make any formal claims about the security of this variant. Indeed, the inevitable presence of clock skew and the naturally slow response times of RFID tags make timestamp-based protocols like this one tricky to analyze.

One might think that tag silence confers privacy protection by concealing the presence of RFID tags on a given person or in a given environment. We note, however, that tag silence probably confers minimal privacy protection. Even if a tag does not reply to a reader query, its presence is still likely to be detectable through simple power analysis of reflected RF reader emissions. Recent work by Oren and Shamir [O'Connor 2006] suggests the potency of such attacks. Logical-layer silence does not necessarily imply RF silence.

3.5 Summary of Minimum Bounds for Privacy Breaks

This table summarizes the minimum bounds necessary to break privacy in the schemes discussed in this article. Definitions of the security parameters r , s , and t appear in Section 2.4.

Scheme	r	s	t	Note
OSK/AO	1	m	m	m is the maximum times a tag may operate
NIBY	$3d$	$2d$	$3d$	d is the depth of the tree used in NIBY
YA-TRAP	1	$O(k)$	2	k is the security parameter
EHJ	1	$O(1)$	2	

4. APPENDIX: HASH-LOCKS

Weis, Sarma, Rivest, and Engels [2003] (WSRE) first advanced the general approach of key search for RFID-tag identification. They proposed two basic schemes dubbed hash-locks and randomized hash-locks. We now show that deterministic hash-locks are not private by our definition. We will show, however, that randomized hash locks are (r, s, t) -private for any r , s , and t polynomial in k . We propose a new improved randomized hash-lock that is both private and resistant to replay attacks, and also satisfies our cross-reader privacy definition.

4.1 Deterministic Hash-Locks

Hash-locks are an access control mechanism designed for efficiency. The general idea is as follows. A tag normally resides in a locked state in which it only responds to reader queries with a temporary *metaID*, rather than some permanent identifier. This *metaID* is the hash of some random nonce selected by \mathcal{R} and programmed into the tag. \mathcal{R} stores both the *metaID* and its preimage for its later interaction with the tag.

\mathcal{R} may unlock a tag by sending it the preimage of its *metaID*. The tag verifies that this value hashes to the *metaID*. Alternatively, the tag need not use a hash function at all. It can instead be locked with some secret *metaKey*. The tag is programmed with a $(metaID, metaKey)$ pair when locked, and unlocked with

the *metaKey*. The hash-function-based scheme may be justified, however, if the hardware cost of implementing a hash function is less than EEPROM costs of *metaKey* storage or if the *metaID* needs to be assigned in the clear.

Using the notation of this article, the session identifier *sid* will play the role of the *metaID*. Let $h : \{0, 1\}^* \rightarrow \{0, 1\}^{2k}$ be a random oracle. Hash-locks work as follows.

Setup:

- (1) GEN outputs a uniformly chosen vector of k -bit keys (key_1, \dots, key_n) .
- (2) \mathcal{R} is initialized with GEN's output.
- (3) Each \mathcal{T}_i is initialized by a SETKEY call with a unique key_i .
- (4) \mathcal{R} makes a TAGINIT call to each \mathcal{T}_i with $sid = h(nonce)$ for a random k -bit *nonce*.
- (5) \mathcal{R} stores each $(sid, nonce)$ pair internally.

System execution:

- (6) \mathcal{R} 's challenge c_0 to a tag \mathcal{T} is empty.
- (7) \mathcal{T} 's response $r_0 = sid$.
- (8) \mathcal{R} 's challenge $c_1 = nonce$.
- (9) If $h(c_1) = sid$, then \mathcal{T} 's r_1 is some identifier determined by *key*.
- (10) Otherwise \mathcal{T} 's r_1 is empty.

Clearly, this protocol is not private by our definition, since a tag's response r_0 is always its current *sid*. Or \mathcal{A} can simply eavesdrop on the protocol between \mathcal{T}_b^* and \mathcal{R} to obtain some identifier. Thus, deterministic hash-locks are not (r, s, t) -private for any nontrivial (r, s, t) values. In fact, any protocol where tags respond deterministically is necessarily not private by our definition.

4.2 Randomized Hash-Locks

In contrast, WRSE's randomized hash-lock scheme is a mechanism for private tag identification that is in fact (r, s, t) -private in the random oracle model. We present a short proof of this fact. Letting $h : \{0, 1\}^* \rightarrow \{0, 1\}^{2k}$ be a random oracle, randomized hash-locks work as follows:

Setup:

- (1) GEN outputs a uniformly chosen vector of k -bit keys (key_1, \dots, key_n) .
- (2) \mathcal{R} is initialized with GEN's output.
- (3) Each \mathcal{T}_i is initialized by a SETKEY call with a unique key_i .
- (4) Tags do not store any session state or *sid*.

System execution:

- (5) \mathcal{R} 's challenges c_0 are empty.
- (6) \mathcal{T} 's response $r_0 = (nonce, h(nonce||key))$, for a random k -bit *nonce*.
- (7) \mathcal{R} searches for key_i among GEN's output such that r_0 contains $h(nonce||key_i)$.
- (8) \mathcal{R} accepts if such a key_i exists and rejects otherwise.

A simple illustration is given in Figure 8.

4.3 Improved Randomized Hash-Locks

Since r_0 does not depend on any reader input, tag outputs are vulnerable to replay. Consequently, as explained in Section 2.8, the WRSE system does not provide privacy in a multiverifier model. In particular, it does not provide cross-reader (r, s, t) -privacy.

Fortunately, this is an easily addressable flaw. Instead of sending an empty c_0 , \mathcal{R} can instead send $c_0 = \text{nonce}_{\mathcal{R}}$, where $\text{nonce}_{\mathcal{R}}$ is generated uniformly at random. (Alternatively, sid could contain this nonce value.) A tag will generate its own nonce $\text{nonce}_{\mathcal{T}}$ and send the value $r_0 = (\text{nonce}_{\mathcal{T}}, h(\text{nonce}_{\mathcal{R}} || \text{nonce}_{\mathcal{T}} || \text{key}))$. The reader \mathcal{R} will then search for a key among GEN's output that would generate r_0 .

Adversaries cannot replay a previously used r_0 to a \mathcal{R} , since with high probability, it will not match the $\text{nonce}_{\mathcal{R}}$ value generated by \mathcal{R} for that session. Unlike the OSK/AO or NIBY protocol, the tag does not store internal state based on $\text{nonce}_{\mathcal{R}}$. Thus, there is no counter or timestamp that can be manipulated by an adversary. Unlike EHJ, the tag always yields an output value.

We do not formally analyze this variant protocol here.

5. CONCLUSION: FURTHER RESEARCH DIRECTIONS

We have proposed new privacy definitions strong enough to capture highly stringent application-level design requirements for RFID systems in the real world. We believe that the relative simplicity of these definitions is useful for the design and analysis of privacy-preserving RFID protocols. In this article, we have examined several published systems that fail to fulfill our privacy definition. In so doing, we believe that we have highlighted potential design flaws.

That said, our proposed privacy definitions are just a starting point. They certainly do not capture the full spectrum of real-world needs. Stronger definitions, which could perhaps model side-channel attacks as well as weaker privacy definitions, like cross-reader privacy, both could be useful in RFID protocol design. With this in mind, we propose two important areas for further work.

5.1 Stronger Definitions

In one sense, even our definition of strong privacy is not strong enough, as it does not capture the effects of side information. Side-channel attacks are pertinent to newly emerging protocol proposals. For example, in independent and contemporaneous work, Burmester et al. [2006] propose a new protocol called O-TRAP that is similar in flavor to hash locks. They propose a definition of privacy and security in the Universal Composability framework and prove that O-TRAP fulfills their definition. We believe that O-TRAP is also provably secure with respect to our definition of strong privacy.

O-TRAP, however, is not secure according to a variant of our definition in which the reader functionality outputs timing information. The reader (or back-end server) in the O-TRAP protocol performs a single table look-up in the optimistic case where a tag is synchronized with a reader; otherwise, it performs

a search of the full key space for the system. If an adversary can distinguish between these two cases based on timing—which seems a serious, practical threat—she can mark tags by desynchronizing them, much as we do in the attacks we describe previously for OSK/AO and NIBY. (The reader in O-TRAP could process tags in uniform time, but then the system would be no more efficient than brute-force key search, like the improved hash-lock scheme we proposed earlier.) Thus, a small extension of our model in which the reader outputs information about computation time can shed critical light on protocol design. Modeling other forms of side information may prove equally valuable.

5.2 Weaker Definitions

On the other hand, some aspects of our definitions (even the cross-reader one) may actually be too strong. For RFID tags capable of only symmetric-key cryptography, we believe that our definition may require the reader to perform brute-force search to identify tags—at least in cases of desynchronization. Such a system scales poorly. Obviously public-key cryptography would alleviate this burden, but is impractical for the vast majority of RFID devices. (Indeed, many classes of RFID tags will be incapable of executing even standard symmetric-key protocols for some years.) While the cryptographic literature often aims at increasingly strong definitions and protocols, we emphasize like Juels [2006] that a fertile and essential area of investigation is definitions and protocols for RFID privacy that are weaker, but more practical and useful.

ACKNOWLEDGMENTS

Thanks to Gildas Avoine, Burt Kaliski, Shingo Kinoshita, Miyako Ohkubo, Koutarou Suzuki, and Gene Tsudik for their very helpful comments on this work.

REFERENCES

- ALIEN TECHNOLOGY. 2005. Alien Technology Corporation achieves another step toward pervasive, economic RFID with announcement of 12.9 cent RFID labels. Alien Technology Press release. <http://www.alientechnology.com>.
- AVOINE, G. 2005. Adversarial model for radio frequency identification. Cryptology ePrint Archive. Report 2005/049. <http://eprint.iacr.org>
- AVOINE, G. 2006. Security and privacy in RFID systems. <http://lasecwww.ep.ch/figavoine/rfid/>.
- AVOINE, G., DYSLI, E., AND OECHSLIN, P. 2005. Reducing time complexity in RFID systems. In *Proceedings of the 12th Annual Workshop on Selected Areas in Cryptography (SAC'05)*. Springer-Verlag, Berlin.
- AVOINE, G. AND OECHSLIN, P. 2005a. RFID traceability: A multilayer problem. In *Proceedings of the 9th International Conference on Financial Cryptography and Data Security (FC'05)*. Springer-Verlag, Berlin, 125–140.
- AVOINE, G. AND OECHSLIN, P. 2005b. A scalable and provably secure hash based RFID protocol. In *Proceedings of the 2nd IEEE International Workshop on Pervasive Computing and Communication Security (PerSec'05)*. IEEE, Los Alamitos, CA, 110–114.
- BURMESTER, M., VAN LE, T., AND DE MEDEIROS, B. 2006. Provably secure ubiquitous systems: Universally composable RFID authentication protocols. <http://eprint.iacr.org/2006/131.pdf>.
- CANETTI, R. Universally composable security: A new paradigm for cryptographic protocols. IACR ePrint Report 2000/067. <http://eprint.iacr.org/2000/067>

- ENGBERG, S., HARNING, M., AND JENSEN, C. 2004. Zero-knowledge device authentication: Privacy and security enhanced RFID preserving business value and consumer convenience. In *Proceedings of the 2nd Annual Conference on Privacy, Security, and Trust*. IEEE, Los Alamitos, CA.
- FISHKIN, K. P., ROY, S., AND JIANG, B. 2004. Some methods for privacy in RFID communication. In *Proceedings of the 1st European Workshop on Security in Ad-Hoc and Sensor Networks (ESAS'04)*. Springer, Berlin.
- GOLLE, P., JAKOBSSON, M., JUELS, A., AND SYVERSON, P. 2004. Universal re-encryption for mixnets. In *Proceedings of the Cryptographers' Track RSA Conference (CT-RSA)*. Springer, Berlin, 163–178.
- HELLMAN, M. 1980. A cryptanalytic time-memory tradeoff. *IEEE Trans. Inf. Theor.* 26, 401–406.
- JUELS, A. 2004. Minimalist cryptography for low-cost RFID tags. In *Proceedings of the 4th International Conference on Security in Communication Networks (SCN'04)*. Springer-Verlag, Berlin, 149–164.
- JUELS, A. 2006. RFID security and privacy: A research survey. *IEEE J. Sel. Areas Comm.* 24, 2.
- JUELS, A., RIVEST, R., AND SZYDLO, M. 2003. The blocker tag: Selective blocking of RFID tags for consumer privacy. In *Proceedings of the 8th ACM Conference on Computer and Communications Security*. ACM, New York, 103–111.
- KAHN, F. 2005. Can zero-knowledge tags protect privacy? *RFID J.* <http://www.rfidjournal.com/article/articleview/1891/1/1/>.
- KINOSHITA, A., OHKUBO, M., HOSHINO, F., MOROHASHI, G., SHIONOIRI, O., AND KANAI, A. 2005. Privacy enhanced active RFID tag. In *Proceedings of the International Workshop on Exploiting Context Histories in Smart Environments*. Springer-Verlag, Berlin.
- MOLNAR, D., SOPPERA, A., AND WAGNER, D. 2005. A scalable, delegatable pseudonym protocol enabling ownership transfer of RFID tags. In *Proceedings of the 12th Annual Workshop on Selected Areas in Cryptography (SAC'05)*. Springer-Verlag, Berlin.
- MOLNAR, D. AND WAGNER, D. 2004. Privacy and security in library RFID: Issues, practices, and architectures. In *Proceedings of the ACM Conference on Communications and Computer Security*. ACM, New York, 210–219.
- NOHARA, Y., INOUE, S., BABA, K., AND YASUURA, H. 2005. Quantitative evaluation of unlinkable ID matching schemes. In *Proceedings of the Workshop on Privacy in the Electronic Society (WPES'05)*. ACM, New York.
- O'CONNOR, M. 2006. EPC tags subject to phone attacks. *RFID J.* <http://www1.rfidjournal.com/article/articleview/2167/1/1/>.
- OECHSLIN, P. 2003. Making a faster cryptanalytic time-memory trade-off. In *Proceedings of the International Conference on the Theory and Application of Cryptographic Techniques*. Springer-Verlag, Berlin, 617–630.
- OHKUBO, M., SUZUKI, K., AND KINOSHITA, S. 2004. Efficient hash-chain-based RFID privacy protection scheme. In *Proceedings of the International Conference on Ubiquitous Computing*. Springer-Verlag, Berlin.
- TSUDIK, G. 2006a. Personal communication.
- TSUDIK, G. 2006b. YA-TRAP: Yet another trivial RFID authentication protocol. In *Proceedings of the 4th Annual Conference on Pervasive Computing and Communications (PerCom'06)*. IEEE, Los Alamitos, CA.
- WEIS, S., SARMA, S., RIVEST, R., AND ENGELS, D. 2003. Security and privacy aspects of low-cost radio frequency identification systems. In *Proceedings of the International Conference on Security in Pervasive Computing (SPC'03)*. Springer-Verlag, Berlin, 454–469.

Received April 2007; revised November 2008; accepted January 2009