

Scalable RFID Systems: a Privacy-Preserving Protocol with Constant-Time Identification

Basel Alomair*, Andrew Clark*, Jorge Cuellar†, and Radha Poovendran*

*Network Security Lab (NSL), University of Washington, Seattle, Washington

†Siemens Corporate Technology, München, Germany

Email: {alomair,awclark,rp3}@uw.edu, jorge.cuellar@siemens.com

Abstract

In RFID literature, most “privacy-preserving” protocols require the reader to search all tags in the system in order to identify a single tag. In another class of protocols, the search complexity is reduced to be logarithmic in the number of tags, but it comes with two major drawbacks: it requires a large communication overhead over the fragile wireless channel, and the compromise of a tag in the system reveals secret information about other, uncompromised, tags in the same system. In this work, we take a different approach to address time-complexity of private identification in large-scale RFID systems. We utilize the special architecture of RFID systems to propose the first symmetric-key privacy-preserving authentication protocol for RFID systems with constant-time identification. Instead of increasing communication overhead, the existence of a large storage device in RFID systems, the database, is utilized for improving the time efficiency of tag identification.

1 Introduction and Related Work

The ability to trace RFID tags, and ultimately the individuals carrying them, is a major privacy concern in RFID systems. Privacy activists have been worried about the invasion of users’ privacy by RFID tags, calling for the delay or even the abandonment of their deployment. In extreme cases, companies have been forced to repudiate their plans for RFID deployment in response to the threat of being boycotted [9]. Consequently, significant effort has been made in the direction of designing RFID systems that preserve users’ privacy.

Two main objectives of typical RFID systems are identification and privacy. Identification, by itself, can be as straightforward as broadcasting tags’ identifiers in clear text. When combined with the privacy requirement, however, transmitting identifiers in clear text is obviously unacceptable. For RFID tags capable of performing asym-

metric cryptography, such as public-key encryption [4] or trapdoor functions [6], private identification can be achieved easily (for instance, by encrypting a randomized version of the tag’s ID with the reader’s public key). Public-key operations, however, are beyond the computational capabilities of low-cost tags. Hoping Moore’s law will eventually render tags capable of performing public-key operations, one might consider the computational limitations of RFID tags a temporary problem. The price of tags, however, will be a determining factor in the deployment of RFID systems. When RFID systems are to replace barcodes to identify tagged items, the price of tags will contribute to the product price. When retailers are to choose between tags that can perform sophisticated cryptographic operations and cheaper tags that cannot, it seems highly likely that the cheaper tags will prevail. Consequently, low-cost RFID systems are restricted to the use of symmetric-key cryptography in most applications.

Privacy-preserving symmetric-key protocols are faced with the following paradox. In one hand, a tag must encrypt its identity with its secret key so that only authorized readers can extract the identity. On the other hand, authorized readers will need to know the identity of the tag in order to determine which key is to be used for decryption. Therefore, given that tags’ responses are randomized (to protect users’ privacy), and that the length of tags’ responses is sufficiently long (so that easy to implement attacks such as random guessing and exhaustive search will have small probability of success), searching the database for those responses is a nontrivial task.

Most RFID protocols trade-off identification efficiency for the sake of privacy. That is, private identification is accomplished, but the reader is required to perform a linear search among all tags in the system in order to identify the tag being interrogated (see, e.g., [7], [8], [16]). In a typical protocol of this class, the reader interrogates a tag by sending a random nonce, r_1 . The tag generates another nonce,

r_2 , computes $h(ID, r_1, r_2)$, where h is a cryptographic hash function, and responds with $s = (r_2, h(ID, r_1, r_2))$. (Different protocols implement variants of this approach; but this is the main idea of this class of protocols.) Upon receiving the tag’s response, the reader performs a linear search of all the tags in the system, computing the hash of their identifiers with the transmitted nonces, until it finds a match. Obviously, unauthorized observers cannot correlate different responses of the same tag, as long as the nonce is never repeated.

Although protocols of this class have been shown to provide private identification, their practical implementation has a scalability issue. In a large-scale RFID system, performing a linear search for every identification run can be a cumbersome task (especially in applications requiring identification of multiple tags simultaneously), and can also lead to denial of service attacks. Hence, for an RFID system to be practical, one must aim for a scheme that can break the barrier of *linear-time* identification complexity.

A big step towards solving the scalability issue in RFID systems was introduced in [14]. This new approach traded-off computational and communication overhead on tags to speed up the identification process. The authors utilized a tree data structure, where each edge in the tree corresponds to a unique secret key, each leaf of the tree corresponds to a unique tag, and each tag carries the set of keys on the corresponding path from the root of the tree to its leaf. When a reader interrogates a tag, the tag responds with a message encrypted with its first key. By decrypting the tag’s response with the keys corresponding to all edges of the first level of the tree, the reader can determine to which edge the tag belongs. By traversing the tree from top to bottom, the tag can be identified in $O(\log N_T)$ time using $O(\log N_T)$ reader-tag interactions, where N_T is the number of tags in the system.

Arranging tags in a tree based on secret keys they possess, however, introduced a new security threat to the RFID system: every compromised tag will reveal the secret keys from the root of the tree to its leaf. Since these keys are shared by other tags in the system, compromising one tag will reveal secret information about all tags sharing a subset of those keys. In [5], the tree structure is analyzed showing that in a tree with a branching factor of two, compromising 20 tags in a system of 2^{20} tags leads to the identification of uncompromised tags with an average probability close to one. Researchers who believe that reducing identification complexity from $O(N_T)$ to $O(\log N_T)$ cannot be overlooked as a result of the vulnerability it introduced have been making significant effort to mitigate the tag compromise problem in tree based systems [12], [13], [17]. The idea shared by all such attempts is to employ a key updating mechanism to reduce the effect of tag capture. Other researchers, however, believe that

TABLE 1. Performance comparison as a function of the number of tags in the system, N_T . Class 1 represents protocols with linear-time identification, while Class 2 represents protocols with log-time identification. The overhead in the last column refers to computation and communication overhead on the tags’ side.

	Search time	Key size	Database size	Overhead
Class 1	$O(N_T)$	$O(1)$	$O(N_T)$	$O(1)$
Class 2	$O(\lg N_T)$	$O(\lg N_T)$	$O(N_T)$	$O(\lg N_T)$
Proposed	$O(1)$	$O(1)$	$O(N_T)$	$O(1)$

the new threat outweighs the reduction in identification complexity, thus, proceeding with the linear-time class of protocols and trying to improve on its performance (see, e.g., [5], [7], [8], [16]).

Another major drawback of the tree based class of protocols is the increase in communication and computation overhead on tags. In a typical RFID system, the reader interrogates multiple tags simultaneously. Consequently, even in the linear-time identification protocols, where communication overhead is $O(1)$, collision avoidance and medium access control are among the most challenging problems in the design of efficient RFID systems [10], [11]. Increasing the communication overhead to $O(\log N_T)$ can only complicate access control further. Extra computation overhead can also be problematic for passive tags as it leads to more energy consumption.

In this paper, we address the private identification problem in large-scale RFID systems. We propose a protocol that, in addition to being *resilient to tag compromise attacks*, allows *constant-time identification*, without imposing extra *communication or computation overhead* on the resource limited tags. The main drive behind devising our protocol is the intuition that, in order to overcome the problems in both linear and log time identification classes, one must aim for a solution that is fundamentally different than both of them. We do not resort to tree structure, nor do we incur more communication overhead. Instead, we utilize resources that are already available in RFID systems to improve identification efficiency. That is, since in any RFID system there is a database, to store information about tags in the system, and since storage is relatively cheap in today’s technology, we tradeoff storage for the sake of better identification efficiency. To the best of our knowledge, the proposed protocol is the first symmetric-key privacy-preserving protocol that allows constant-time tag identification. Table 1 compares our protocol to the class of linear-time identification protocols, Class 1, and to the class of log-time identification protocols, Class 2.

2 Model Assumptions

2.1 System Model

RFID systems are typically composed of three main components: tags, readers, and a database. In our model,

the tag is assumed to have limited computing power: hash computations are the most expensive operations tags can perform. The reader is a computationally powerful device with the ability to perform sophisticated cryptographic operations. The database is a storage resource at which information about tags in the system is stored. Readers-database communications are assumed to be secure.

2.2 Adversarial Model

We assume adversaries with complete control over the communication channel. Adversaries can observe all exchanged messages, modify exchanged messages, block exchanged messages and replay them later, and generate messages of their own. We do not consider an adversary whose only goal is to jam the communication channel. Distinguishing tags by the physical fingerprints of their transmissions requires sophisticated devices and cannot be solved using cryptographic solutions. It is out of the scope of this work as in the majority of similar proposals.

The adversary \mathcal{A} is modeled as a polynomial-time algorithm. Given a tag, T , and a reader, R , we assume \mathcal{A} has access to the following oracles:

- *Query* (T, m_1, x_2, m_3): \mathcal{A} sends m_1 as the first message to T ; receives a response, x_2 ; and then sends the message $m_3 = f(m_1, x_2)$. This oracle models the adversary's ability to interrogate tags in the system.
- *Send* (R, x_1, m_2, x_3): \mathcal{A} receives x_1 from the reader R ; replies with $m_2 = f(x_1)$; and receives the reader's response x_3 . This oracle models the adversary's ability to act as a tag in the system.
- *Execute* (T, R): The tag, T , and the reader, R , execute an instance of the protocol. \mathcal{A} eavesdrops on the channel, and can also tamper with the messages exchanged between T and R . This oracle models the adversary's ability to actively monitor the channel between tag and reader.
- *Block* (\cdot): \mathcal{A} blocks any part of the protocol. This query models the adversary's ability to launch a denial of service attack.
- *Reveal* (T): This query models the exposure of the tags' secret parameters to \mathcal{A} . The oracle simulates the adversary's ability to physically capture the tag and obtain its secret information.

\mathcal{A} can call the oracles *Query*, *Send*, *Execute*, and *Block* any polynomial number of times. The *Reveal* oracle can be called only once (on the same tag), at which the tag is considered compromised and, thus, there is no point of calling the *Reveal* oracle on the same tag multiple times. To model tag compromise attacks, however, the adversary is allowed to call other oracles after the *Reveal* oracle on the same tag; detailed discussion about this is provided in Section 7.

2.3 Security Model

The security model presented in this section does not consider the adversary's ability to perform pre-processing before engaging in the games. In Section 7, however, we will modify the security model to give the adversary such ability to perform pre-processing that involves calling the *Reveal* oracle on tags in the system. The main purpose of this modification is to allow modeling tag compromise attacks.

The two main security goals of our protocol are tags' privacy and tag-reader mutual authentication. Privacy is measured by the adversary's ability to trace tags by means of their responses in different protocol runs. We define three notions of untraceability, *universal*, *forward*, and *existential*.

Definition 1 (Universal Untraceability): In an RFID system, tags are said to be universally untraceable if an adversary cannot track a tag based on information gained before the tag's last authentication with a valid reader. In other words, there is no correlation between a tag's responses before and after completing a protocol run with a valid reader.

Universal untraceability is modeled by the following game between the challenger C (an RFID system) and a polynomial time adversary \mathcal{A} .

- 1) C selects two tags, T_0 and T_1 , and a valid reader, R .
- 2) \mathcal{A} makes queries on T_0 , T_1 , and R using the *Query*, *Send*, *Execute*, and *Block* oracles for a number of times of its choice.
- 3) \mathcal{A} stops calling the oracles and notifies C .
- 4) C carries out an instance of the protocol with T_0 and T_1 , during which mutual authentication of both tags with R is achieved.
- 5) C selects a random bit, b , and sets $T = T_b$.
- 6) \mathcal{A} makes queries of T and R using the *Query*, *Send*, *Execute*, and *Block* oracles.
- 7) \mathcal{A} outputs a bit, b' , and wins the game if $b' = b$.

The second notion of privacy, forward untraceability, is defined as follows.

Definition 2 (Forward Untraceability): In an RFID system with forward untraceability, an adversary capturing the tag's secret information cannot correlate the tag with its responses before the last complete protocol run with a valid reader.

Forward untraceability is modeled by the following game between C and \mathcal{A} .

- 1) C selects two tags, T_0 and T_1 , and a valid reader, R .
- 2) \mathcal{A} makes queries of T_0 , T_1 , and R using the *Query*, *Send*, *Execute*, and *Block* oracles for a number of times of its choice.
- 3) \mathcal{A} stops calling the oracles and notifies C .
- 4) C carries out an instance of the protocol with T_0 and T_1 , during which mutual authentication of both tags

with R is achieved.

- 5) C selects a random bit, b , and sets $T = T_b$.
- 6) \mathcal{A} calls the oracle *Reveal* (T).
- 7) \mathcal{A} outputs a bit, b' , and wins the game if $b' = b$.

Finally, the third notion of privacy, existential untraceability, is defined as follows.

Definition 3 (Existential Untraceability): Tags in an RFID system are said to be existentially untraceable if an active adversary cannot track a tag based on its responses to multiple interrogation, even if the tag has not been able to accomplish mutual authentication with an authorized reader.

Existential untraceability is modeled by the following game between C and \mathcal{A} .

- 1) C selects two tags, T_0 and T_1 .
- 2) \mathcal{A} makes queries of T_0 and T_1 using the *Query* oracle for at most $C-1$ number of times for each tag, where C is a pre-specified system security parameter.
- 3) \mathcal{A} stops calling the oracles and notifies C .
- 4) C selects a random bit, b , and sets $T = T_b$.
- 5) \mathcal{A} makes a query of T using the *Query* oracle.
- 6) \mathcal{A} outputs a bit, b' , and wins the game if $b' = b$.

To quantify the adversary's ability to trace RFID tags, we define the adversary's advantage of successfully identifying the tag in the previous games as

$$Adv_{\mathcal{A}} = 2\left(\Pr[b' = b] - \frac{1}{2}\right). \quad (1)$$

If the adversary cannot do any better than a random guess, then $\Pr(b' = b) = 1/2$. Consequently, the adversary's advantage, $Adv_{\mathcal{A}}$, is zero, at which point we say that tags are untraceable.

The other security goal of our protocol is mutual authentication. An *honest protocol run* is defined as follows [3]: A mutual authentication protocol run in the symmetric key setup is said to be honest if the parties involved in the protocol run use their shared key to exchange messages, and the messages exchanged in the protocol run have been relayed faithfully (without modification).

Another term that will be used for the remainder of the paper is the definition of negligible functions: A function $\gamma : \mathbb{N} \rightarrow \mathbb{R}$ is said to be negligible if for any nonzero polynomial φ , there exists N_0 such that for all $N > N_0$, $|\gamma(N)| < (1/|\varphi(N)|)$. That is, the function is said to be negligible if it converges to zero faster than the reciprocal of any polynomial function.

We now give the formal definition of secure mutual authentication for RFID systems as appeared in [3].

Definition 4 (Secure Mutual Authentication): A mutual authentication protocol for RFID systems is said to be secure if and only if it satisfies all the following conditions: 1. No information about the secret parameters of an RFID tag is revealed by messages exchanged in protocol runs.

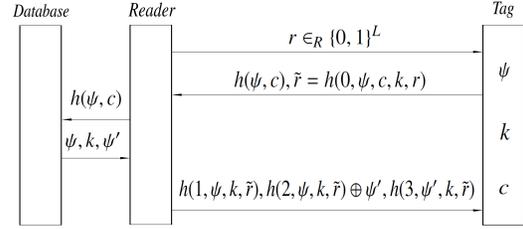


Fig. 1. A schematic of one instance of the protocol.

2. **Authentication \Rightarrow Honest protocol:** the probability of authentication when the protocol run is not honest is negligible in the security parameter.

3. **Honest protocol \Rightarrow Authentication:** if the protocol run is honest, the tag-reader pair must authenticate each other with probability one.

To model the adversary's attempt to authenticate herself to a reader (tag), we propose the following game between the challenger C and adversary \mathcal{A} .

- 1) C chooses a tag, T , at random, and a reader, R .
- 2) \mathcal{A} calls the oracles *Query*, *Send*, *Execute*, and *Block* using T and R for a number of times of its choice.
- 3) \mathcal{A} decides to stop and notifies C .
- 4) \mathcal{A} calls the oracle *Send* (*Query*) to impersonate a tag (reader) in the system.
- 5) If \mathcal{A} is authenticated as a valid tag (reader), \mathcal{A} wins the game.

Definition 4 implies that the protocol achieves secure mutual authentication only if the adversary's probability of winning the previous game is negligible.

3 System Description

3.1 Protocol Overview

In our system, each tag has an internal counter, c , and is preloaded with a unique *secret* pseudonym, ψ , and a secret key, k . The secret key and the secret pseudonym are updated whenever mutual authentication with a valid reader is accomplished, while the counter is incremented every time authentication fails.

When an RFID reader is to identify and authenticate a tag within its range, it generates a random nonce, $r \in_R \{0, 1\}^L$, and transmits it to the tag. Upon receiving r , the tag computes $h(\psi, c)$ and $\tilde{r} := h(0, \psi, c, k, r)$, where ψ is the tag's current pseudonym, k is the tag's current secret key, c is the tag's internal counter, and r is the received nonce. The tag then increments its counter, $c \leftarrow c+1$. With $h(\psi, c)$, the reader accesses the database to identify the tag and obtain its information, including its pseudonym, ψ , its secret key, k , and a new pseudonym, ψ' , to update the tag. With \tilde{r} , the reader authenticates the tag by confirming its knowledge of the secret key, k , obtained from the database.

Once the tag has been identified and authenticated, the reader responds with $h(1, \psi, k, \tilde{r})$, $h(2, \psi, k, \tilde{r}) \oplus \psi'$, and

$h(3, \psi', k, \tilde{r})$. With $h(1, \psi, k, \tilde{r})$, the tag authenticates the reader (by verifying its knowledge of its secret key, k). If the reader is authenticated, the tag uses $h(2, \psi, k, \tilde{r}) \oplus \psi'$ to extract its new pseudonym, ψ' . Once the new pseudonym has been computed, the tag verifies its integrity using $h(3, \psi', k, \tilde{r})$. The tag and the reader then update the tag's secret key to $k' = h(k)$ truncated to the required length, ℓ . Figure 1 depicts a single protocol run between an RFID reader-tag pair.

3.2 Database Overview

As mentioned above, the tag is identified by its randomized response, $h(\psi, c)$, which is an L -bit long string. Since security requires that L is sufficiently long, it is infeasible to construct a physical storage that can accommodate all possible 2^L responses, for direct addressing. (This is the reason why previous schemes resorted to linear search amongst all tags in the system to identify a response.) For ease of presentation, the structure of the database is divided into three logical parts, M-I, M-II, and M-III.

To allow for constant-time identification, with feasible storage, we truncate the L -bit identifiers to their s most significant bits, where s is small enough so that a storage of size 2^s is feasible. Of course, many identifiers will share the same s most significant bits (to be exact, 2^{L-s} possible identifiers will share the same truncated value). M-I is a table of size $O(2^s)$, with addresses ranging from 0 to $2^s - 1$, and each table entry contains a pointer to an entry in M-II (similar to a hashtable data structure, with truncation instead of hashing). All identifiers with the same s most significant bits will be stored in a smaller table in M-II, and the pointer at address s in M-I will point to the head of this smaller table. Finally, actual information about tags in the system is stored in M-III. Detailed construction of the database and description of the identification process will be the focus of the remainder of this section.

The proposed protocol can be broken into four main phases: parameters selection phase, system initialization phase, tag identification phase, and identity randomization and system update phase. Each phase is detailed below.

3.3 Parameters Selection

During this phase, the database is initialized and each tag is loaded with secret information. The secret information includes the tag's secret key, which the tag and reader use to authenticate one another, and the tag's pseudonym, which is used for tag identification.

Given the total number of tags the RFID system is suppose to handle, N_T , and predefined security and performance requirements (more about this later), the system designer chooses the following parameters to start the initialization phase:

TABLE 2. A list of parameters and used notations.

Symbol	Definition
N_T	The total number of tags in the system
N	The total number of pseudonyms in the system
ψ_i	The pseudonym corresponding to the i^{th} tag
C	The maximum counter value
ℓ	The length of the secret parameter in bits
h	Cryptographic hash function
L	The output length of the used hash function
n	The length of the truncated hash values
$\Psi_{i,c}$	A tag identifier, $\Psi_{i,c} := h(\psi_i, c)$
$\Psi_{i,c}^n$	The n most significant bits of $\Psi_{i,c}$

- The total number of pseudonyms, N . Since pseudonyms will be used as unique tag identifiers, there must be at least one pseudonym for every tag in the system. Furthermore, since tags are assigned new identifiers following every successful mutual authentication process with an authorized reader, the total number of pseudonyms must be greater than the total number of tags in the system, i.e., $N > N_T$.
- The maximum counter value, C . The counter is used by RFID tags to mitigate traceability by active adversaries; the larger the counter is, the more difficult it will be for active adversaries to track the tag; on the downside, the size of the database will grow linearly with the counter (the database size is $O(NC)$). Therefore, the size of the counter is a tradeoff between tags' privacy and system complexity.
- The length, ℓ , in bits, of the tags' secret parameters (pseudonyms and keys). As in any symmetric key cryptosystem, ℓ should be chosen properly to prevent easy-to-implement attacks, such as exhaustive search and random guessing. Obviously, ℓ must be long enough to generate N distinct pseudonyms, i.e., $\ell \geq \lceil \log_2 N \rceil$. In practice, however, ℓ will be much longer.
- The hash function, h . In particular, the output length of the hash values, L , is of special importance. The length must be chosen large enough so that there are no collisions during database initialization, which is described below.
- The length, n , of the truncated hashes. The size of n is the key for constant-time identification and practicality of the system. It will be determined in Section 4.

Table 2 summarizes the list of system parameters and used notations.

3.4 System Initialization

Once the system parameters have been chosen, the initialization phase can start. The initialization phase can be summarized in the following steps.

- 1) Given the number of pseudonyms, N , and the length of each pseudonym, ℓ , the system designer draws, *without*

$h(\psi_1, 0)$	$h(\psi_1, 1)$	\dots	$h(\psi_1, C-1)$
$h(\psi_2, 0)$	$h(\psi_2, 1)$	\dots	$h(\psi_2, C-1)$
\vdots	\vdots		\vdots
$h(\psi_N, 0)$	$h(\psi_N, 1)$	\dots	$h(\psi_N, C-1)$

Fig. 2. During database initialization, all values of $h(\psi, c)$ are computed.

replacement, N pseudonyms randomly from the set of all possible ℓ -bit strings. That is, N distinct pseudonyms, $\psi_1, \psi_2, \dots, \psi_N$, are chosen at random from $\{0, 1\}^\ell$. Each tag is given a unique pseudonym and a secret key, and each tag's counter is initially set to zero. We emphasize that the drawn pseudonyms are not publicly known; otherwise, tags' privacy can be breached.

2) For each pseudonym, ψ_i , the hash value $h(\psi_i, c)$ is computed for all $i = 1, \dots, N$ and all $c = 0, \dots, C-1$. That is, a total of NC hash operations must be performed, as depicted in Figure 2. Each row of the table in Figure 2 corresponds to the same pseudonym. Therefore, all entries in the i^{th} row must point to the same memory address carrying information about the tag identified by the pseudonym ψ_i .

In order for tags to be identified uniquely, the hash values in the table of Figure 2 must be distinct. This can be achieved by choosing the hash function, h , to be an expansion function, as opposed to the usual use of hash functions as compression functions, so that collision will occur with small probability.¹ We will assume that the output of the hash function has length L bits, which must be at least equal to $\lceil \log_2 NC \rceil$ so that the table in Figure 2, which is of size NC , can be constructed without collisions (L will be much larger in practice). If a pseudonym that causes a collision in Figure 2 is found, the pseudonym is replaced by another one that does not cause a collision. (Observe that the pool of possible pseudonyms is of size 2^ℓ , which is much larger than the required number of pseudonyms N , giving the system designer a sufficient degree of freedom in constructing the system.) With the appropriate choice of the hash function, a table of hash values with no collisions can be constructed. Note that this operation is performed only once during the initialization phase, thus, it does not undermine the performance of the system.

Since the length of $h(\psi_i, c)$ (the tags' identifiers), L , is large to avoid collision, it would be infeasible to have a

1. For example, this can be accomplished by concatenating multiple hash functions, i.e., $h(x) = h_1(x) \parallel \dots \parallel h_m(x)$, so that $h(x)$ has the required length.

physical storage that can accommodate all possible L -bit strings (for direct addressing). For example, if $L = 128$, a database of size in the order of 4×10^{28} Gigabyte will be required. Previously proposed privacy-preserving schemes solve this problem in one of two approaches. The first approach requires $O(N_T)$ memory space to store information about each tag in the system, and requires the reader to perform a linear search among tags in the system to identify tags' responses; thus requiring $O(N_T)$ space and $O(N_T)$ time for identification. The other method identifies tags based on their key information and requires the reader to perform logarithmic search to identify tags' responses; thus requiring $O(N_T)$ space and $O(\log N_T)$ time for identification.

3) For ease of presentation, we will divide the database into three logical parts, M-I, M-II, and M-III. The first part, M-I, consists of a single table of size $O(2^n)$. The second part, M-II, consists of multiple smaller tables; the total size of all the tables in M-II is $O(NC)$. Finally, the last part, M-III, is of size $O(N)$.

The table in M-I is a table of pointers. The addresses of M-I range from 0^n to 1^n ; each entry in the table points to the head of one of the mini tables in M-II (according to a specific relation explained below).

Each entry of M-II contains two fields. In the first field, the hash values obtained in the table of Figure 2 are stored (i.e., $h(\psi_i, c)$ for all $i = 1, \dots, N$ and all $c = 0, \dots, C-1$). M-II is organized based on the hash values stored in the first field. We say that two hash values $h(\psi_1, c_1)$ and $h(\psi_2, c_2)$ are in the same position, b , if their n most significant bits are the same (recall that the output length of the hash function is $L > n$). All hash values that have the same position, i.e., share the n most significant bits, are stored in the same mini table in M-II (e.g., the hash values with $b = s$ in Figure 3). Hash values with distinct positions are stored in different tables (e.g., hash values with $b = 0^n, s, 1^n$ in Figure 3). (Recall that Figure 2 contains the computed hash values; hence, table M-II can be viewed as a reorganized version of the two-dimensional table in Figure 2 into a one-dimensional table of size $O(NC)$.) The second field of each entry of M-II stores a pointer to an entry in M-III containing information about a tag in the system (depending on the value of the first field). For example, if the value stored in the first field is $h(\psi_i, c)$, then the value in the second field will be a pointer to the data entry in M-III where information about the tag with pseudonym ψ_i can be found.

After M-II has been constructed, the pointers at M-I are chosen to satisfy the following: the pointer stored at address a in M-I must point to the mini table in M-II that stores identifiers with position a . In other words, each pointer in M-I must point to the identifiers with position equal to the address of the pointer.

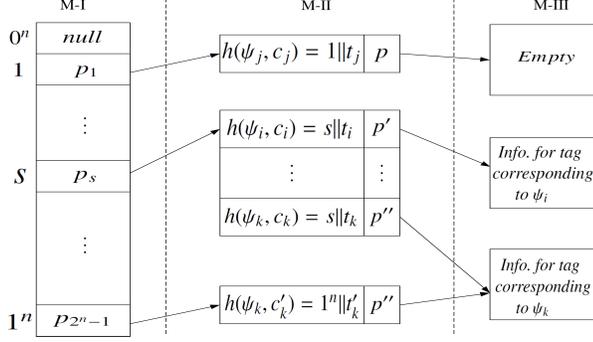


Fig. 3. The three-tier architecture of the database.

Finally, M-III is the actual storage where tags' information is stored. Figure 3 depicts the architecture of the database with the three logical partitions. The identification phase below will further illustrate the structure of the database.

3.5 Tag Identification

Tags in a protocol run of the system are identified by the hash of their pseudonyms concatenated with their internal counters. Denote by $\Psi_{i,c}$ the hash value of the i^{th} pseudonym concatenated with a counter c ; that is, $\Psi_{i,c} := h(\psi_i, c)$. Furthermore, we will denote by $\Psi_{i,c}^n$ the truncated value of $\Psi_{i,c}$; more precisely, $\Psi_{i,c}^n$ represents the n most significant bits of $\Psi_{i,c}$ (i.e., the position of $\Psi_{i,c}$).

Once $\Psi_{i,c}$ has been received, the reader accesses the data entry at address $\Psi_{i,c}^n$ in M-I. This table entry is actually a pointer, p , to one of the tables in M-II. There are three possible scenarios here:

- The value at address $\Psi_{i,c}^n$ in M-I is a *null*. This implies that, during the construction of the table in Figure 2, no identifier with position $\Psi_{i,c}$ is constructed. Therefore, either the tag is not a valid one or the tag's response has been modified. In the example of Figure 3, if the n most significant bits of the received $\Psi_{i,c}$ are *zeros*, then no valid tag matches this response.
- The pointer, p , at address $\Psi_{i,c}^n$ points to a table in M-II with exactly one entry. In this scenario, the first field of the entry pointed at by p must be the entire (untruncated) $\Psi_{i,c}$; the value at the second field will be a pointer to the entry in M-III that contains information about the interrogated tag. In the example of Figure 3, if the n most significant bits of the received $\Psi_{i,c}$ are *ones*, then the pointer at address 1^n in M-I will point to the entry at M-II at which $\Psi_{k,c'_k} = 1^n || t'_k$ and the pointer, p'' , are stored. In turn, p'' will point to the entry at M-III where information about the tag with pseudonym ψ_k is stored.
- The pointer at address $\Psi_{i,c}^n$ of M-I points to a table in M-II with more than one entry. In this scenario, the reader searches the first fields of the mini table in M-II until it reaches the entry that matches the complete (untruncated)

received identifier, $\Psi_{i,c}$; and then follows the pointer (in the corresponding second field) to get the tag's information. In the example of Figure 3, if the received identifier is $\Psi_{k,c_k} = s || t_k$, the reader will follow the pointer at address s of M-I. The pointer, however, points to a table in M-II with more than one entry. Therefore, the reader must search until it reaches the last entry of the table to find a match for the received $\Psi_{k,c_k} = s || t_k$. Once the match is found, the reader can follow the pointer, p'' , to the entry in M-III containing information about the tag with pseudonym ψ_k .

The identification process allows for unique identification of tags in the system. This is due to the requirement that, in the initialization phase, the values in the table of Figure 2 are distinct. Consequently, the entries in M-II are distinct, allowing for the unique identification of tags.

Remark 1: Recall that the pseudonyms drawn in the initialization are not publicly known. If the pseudonyms were published, an adversary can, in principle, construct her own system and identify tags in constant-time. Further discussion about the adversary's ability to expose secret pseudonyms is provided in Section 7.

3.6 System Update

Once a tag has been authenticated, the reader draws one of the unoccupied pseudonyms generated in the initialization phase. (Recall that the number of pseudonyms is greater than the number of tags in the system; consequently, there will always be unused pseudonyms available for identity randomization.) Once an unoccupied pseudonym has been chosen, it is to be transmitted to the tag in a secret and authenticated way.

To allow for correct identification of a tag after its pseudonym has been updated, the database must be updated accordingly. A straightforward way of updating the database is by updating the pointers corresponding to the outdated and updated pseudonyms. For example, if the tag's outdated pseudonym is ψ_i and its updated pseudonym is ψ_k , then all pointers in M-II corresponding to entries $\Psi_{i,0}, \Psi_{i,1}, \dots, \Psi_{i,C-1}$ must point to a null; and all pointers in M-II corresponding to entries $\Psi_{k,0}, \Psi_{k,1}, \dots, \Psi_{k,C-1}$ must point to the entry in M-III containing information about the tag. This method, however, requires $O(C)$ updates.

An alternative method that allows a faster update is depicted in Figure 4. Instead of updating the pointers as in the previous method, the tag's information is moved to the entry in M-III pointed at by the pointers corresponding to the updated pseudonym in M-II. The only price to pay for this method over the previous one is that the size of M-III will increase from $O(N_T)$ to $O(N)$ (asymptotically, N and N_T are of the same size). In the example of Figure 4, instead of changing all entries in M-II with pointer p' to p , and changing entries with pointer p to *null*, the tag's information is moved to the entry in M-III pointed at by

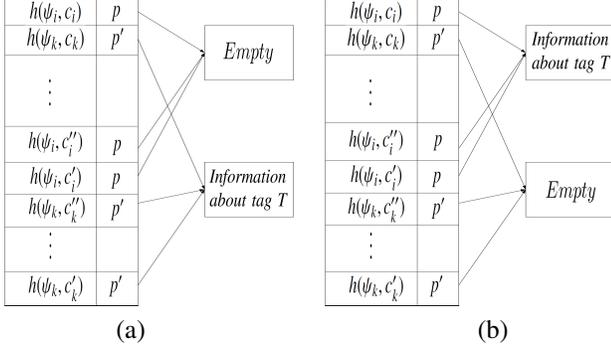


Fig. 4. (a) Before (b) After; an illustration of database update.

p' and the entry pointed at by p is emptied.

4 Performance Analysis

For the proposed scheme to be practical, we must show that a set of parameters can be chosen such that our claim of constant-time identification can be achieved with feasible resources (namely, feasible database size). This section is devoted to showing that, with a set of appropriately chosen parameters, the proposed technique can achieve constant-time identification with a database of size $O(N_T)$.

Assuming that the $\Psi_{i,c}$'s are uniformly distributed, the probability that the truncated version $\Psi_{i,c}^m$ takes a specific value, s , is $\alpha = \Pr(\Psi_{i,c}^m = s) = 2^{-n}$, for any $s \in \{0, 1\}^n$. Let $M := NC$ and define $m := \log_2 M$, where N is the total number of pseudonyms and C is the maximum counter value. Then, out of the M values of $\Psi_{i,c}$'s, the probability that exactly k of them share the same truncation value (i.e., exactly k of them have the same n most significant bits) is

$$\Pr[\mathbf{k} = k] = \binom{M}{k} \alpha^k (1 - \alpha)^{M-k}, \quad (2)$$

where \mathbf{k} is the random variable representing the number of $\Psi_{i,c}^m$ sharing the same value, s , for any $s \in \{0, 1\}^n$. Then, for $k \ll M$,

$$\binom{M}{k} = \frac{M!}{k!(M-k)!} \approx \frac{M^k}{k!}. \quad (3)$$

Using the facts that $\lim_{n \rightarrow \infty} (1 - \frac{1}{n})^n = 1/e$, $M = 2^m$, and $\alpha = 2^{-n}$ we get

$$(1 - \alpha)^{M-k} \approx (1 - \alpha)^M = (1 - 2^{-n})^{2^m \cdot 2^{m-n}} \approx e^{-2^{m-n}}. \quad (4)$$

Substituting equations (3) and (4) into (2) yields,

$$\Pr[\mathbf{k} = k] \approx \frac{M^k}{k!} \cdot \alpha^k \cdot e^{-2^{m-n}} = \frac{1}{k!} \cdot \beta^k \cdot e^{-\beta}, \quad (5)$$

where $\beta = 2^{m-n}$. Choosing $m = n$ yields $\beta = 1$ and equation (5) can be reduced to

$$\Pr[\mathbf{k} = k] \approx \frac{1}{k!} \cdot e^{-1} \text{ for } k = 0, 1, \dots \quad (6)$$

(One can show that equation (6) is a valid probability mass function by verifying that $\sum_{k=0}^{\infty} \Pr[\mathbf{k} = k] = 1$.)

Using the fact that $e = \sum_{k=0}^{\infty} \frac{1}{k!}$, the expected number of truncated $\Psi_{i,c}$'s with the same value is

$$E[\mathbf{k}] = \sum_{k=0}^{\infty} k \cdot \Pr[\mathbf{k} = k] = e^{-1} \sum_{k=0}^{\infty} \frac{1}{k!} = 1. \quad (7)$$

Recall that identifiers $\Psi_{i,c}$ with the same truncated value $\Psi_{i,c}^m$ will be in the same table in M-II; and when the reader receives one of these identifiers it will have to search the table to be able to identify the tag. Equation (7), however, implies that the expected size of the tables in M-II is *one*. Therefore, upon receiving a tag identifier $\Psi_{i,c}$, the reader goes to the table entry in M-I at address $\Psi_{i,c}^m$, follows the pointer p_1 stored at that address, searches the table in M-II pointed at by p_1 for the received $\Psi_{i,c}$ (on average there will be only one entry by (7)), and then follows a pointer p_2 to information about the tag. Indeed, the search time is independent of the number of tags in the system (on average).

Since the database consists of three parts, M-I, M-II, and M-III; and since the size of M-I is $O(2^n)$, the size of M-II is $O(NC)$, and the size of M-III is $O(N)$, the only concern is the size of M-I. The above analysis shows that, by choosing $n = \lceil \log_2 NC \rceil$, the system achieves the constant-time identification claim. Therefore, the size of M-I is $O(NC)$ and, consequently, the total size of the database is $O(NC)$. However, C is a constant, independent from the number of tags in the system; and N is $O(N_T)$. Therefore, with the proposed system, the required size of the database for constant-time identification to be achieved is $O(N_T)$.

5 Case Study

Since big O analysis can be misleading by absorbing big constants, we give here a numerical example of the practicality of our system. Assume an enterprise with one billion items to be tagged, i.e., $N_T = 10^9$. Assume further that the total number of pseudonyms is two billions, i.e., $N = 2N_T$ and $C = 1,000$. Then, the truncated identifiers are $n = \lceil \log_2 NC \rceil = 41$ -bit long. Therefore, M-I can be constructed with a storage smaller than 12 terabyte; a practical storage even for personal usage.²

Therefore, an active adversary must interrogate a tag more than 1,000 consecutive times, not separated by a protocol run with a valid reader, in order to correlate its responses. Observe that, unlike security for general computer communications, 1,000 consecutive interrogations is an unlikely scenario for RFID systems. A web server, for instance, is always online. In a typical RFID systems,

² Western Digital has already released 8-TB hard drives for personal use [1].

however, adversaries must be in close proximity to tags in order to interrogate them. Furthermore, an adversary who is always in the vicinity of a tag can track it down visually without interrogation. Therefore, limiting the number of consecutive tag interrogations is a reasonable relaxation in RFID models.

6 Security Analysis

Due to lack of space, we give below informal sketches illustrating the main ideas behind the proofs. For formal security analysis, please refer to the full version of this paper [2].

An important observation is that blocking the last message (from the reader to the tag) will lead the reader to update the tag's pseudonym while the tag has not,³ i.e., a desynchronization attack. Fortunately, however, this can be solved by storing both the updated and the outdated pseudonyms in the database (the database must be designed accordingly).

The interactive protocol used in our system implements an instance of the challenge-response class of protocols, a well-studied class in cryptography. The main concept behind the proofs is that the used hash function is a one-way function and that a computationally bounded adversary cannot invert the hash function in polynomial time. Given these two assumptions, an adversary observing a sequence of hash values cannot extract the inputs to the hash function. Therefore, the adversary's probability of generating a valid response that will be authenticated by a valid reader (tag) is negligible in the length of the response. This is the basic concept behind the mutual authentication proof (please refer to [2] for a formal proof).

Using the same fact, that the adversary cannot recover tags' secret parameters from the hash values, universal untraceability follows from the fact that the adversary does not know the tag's outdated nor updated pseudonyms. Forward untraceability can be proven using the fact that, even if an adversary has broken into the tag and captured its secret parameter, the tag's future responses cannot be correlated to the tag's responses prior to parameter exposure. This is due to the fact that outdated pseudonyms are independent from the updated pseudonyms. Likewise, existential untraceability follows from the random oracle assumption. That is, given that the counter is incremented following failed protocol runs, the outputs cannot be correlated by an adversary without the knowledge of the secret pseudonym. Therefore, an adversary interrogating the tag less than the maximum number of times allowed cannot trace the tag using its responses.

3. This is an inherited problem shared by all interactive protocols. The fundamental problem here is that the sender of the last message has no means of confirming that the message has been successfully delivered.

7 Tag Compromise Vulnerability

7.1 The Compromise attack

Each tag in the proposed protocol has two pieces of secret information, its pseudonym and its key. Since tags' pseudonyms and keys are designed to be statistically independent for different tags, compromising some tags in the system does not affect the security of other, uncompromised tags. An adversary, however, can compromise a tag in the system and attempt to harvest as many pseudonyms as possible by performing multiple protocol runs with a valid reader.

The adversarial model of Section 2 can be modified to capture the tag compromise attack. Let an adversary calling the *Reveal* (T) oracle, thus capturing the tag T , have the ability to perform multiple protocol runs with the system. Let q be the number of protocol runs an adversary has performed with the system using compromised tags. The number of interest here is how many distinct pseudonyms the adversary has collected, after q protocol runs. This is known in the literature of probability theory as the "coupon collecting problem" [15]. Given there are N distinct pseudonyms and the adversary has performed q protocol runs, assuming each pseudonym is equally likely to be selected, the expected number of distinct pseudonyms collected by the adversary is [15]:

$$N \left(1 - \left(\frac{N-1}{N} \right)^q \right). \quad (8)$$

Assume an adversary has built a system, similar to our construction, with the collected pseudonyms. The adversary's advantage of distinguishing between two tags, given by equation (1), will be greater than zero if at least one of the two tags' pseudonyms is in the constructed table. Thus, given the adversary has performed q protocol runs with a system of N pseudonyms, the probability of distinguishing between two tags is:

$$1 - \left(\frac{N-1}{N} \right)^{2q}. \quad (9)$$

Consider the numbers given in Section 5, i.e., $N = 2 \times 10^9$. To have a 0.001 probability of distinguishing between two tags, an adversary needs to compromise a tag and complete more than a million protocol runs with the system. Figure 5 shows the adversary's probability of having an advantage greater than zero as a function of the number of protocol runs performed with the system using compromised tags.

7.2 Countermeasures

Remember, however, that the database is a powerful device. Therefore, designing the database to record timing information about the tag's past protocol runs can mitigate this threat. For example, the database can store information

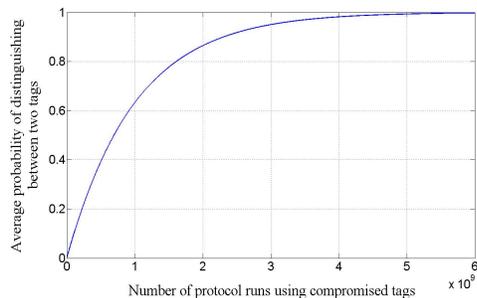


Fig. 5. The adversary's average probability of distinguishing between two tags vs. the number of protocol runs using a compromised tag, in a system with 2×10^9 pseudonyms.

about the tag's last five protocol runs (this can be stored as part of the tag's information, i.e., in M-III). If the adversary tries to harvest different pseudonyms by performing multiple protocol runs with the system, the tag performing the attack can be detected. Therefore, to harvest enough pseudonyms, the adversary will need to compromise more than one tag, depending on the system's parameters and the required probability of success.

Furthermore, the database can periodically update the system by replacing vacant pseudonyms with new pseudonyms (recall that the number of pseudonyms in the database, N , is only a small fraction of the number of all possible pseudonyms, 2^ℓ). This pseudonym update procedure is performed offline by the database, thus, not affecting identification time. Moreover, as a result of the independence of secret parameters amongst tags, the updating procedure is independent of tags.

With the periodic update described earlier, the space of possible pseudonyms will increase to all possible ℓ -bit long strings, as opposed to the predefined *smaller* number N . Therefore, for a bounded adversary, any polynomial number of collected pseudonyms is negligible in the security parameter ℓ . (Recall that the size of the actual database is still proportional to N ; only from the adversary's point of view the size is proportional to 2^ℓ .) Consequently, the adversary's probability of breaking the privacy of the system is negligible in ℓ , provided the periodic update of the database.

8 Conclusion

In this paper, we addressed the problem of individual tag identification in large-scale RFID systems. We proposed a protocol that enables the private identification of tags in the system with constant-time complexity. By utilizing the existence of a large storage device in the system, the constant-time identification is achieved by performing the necessary time consuming computations offline (independent of the reader-tag interactions). As opposed to tree based protocols, the proposed protocol does

not further complicate the already challenging problems in RFID systems, namely, collision avoidance and medium access control. Furthermore, tag compromise threats can be mitigated by periodically updating the database which, due to independence of secret parameters amongst tags, can be performed independent of any tag-reader interaction. To the best of our knowledge, this is the first symmetric-key, constant-time identification protocol in the literature of RFID.

References

- [1] <http://www.westerndigital.com/en/>.
- [2] B. Alomair, A. Clark, J. Cuellar, and R. Poovendran. Scalable RFID Systems: A Privacy-Preserving Protocol with Constant-Time Identification. *Network Security Lab (NSL) Technical Report #006*, available at <http://www.ee.washington.edu/research/nsl/faculty/radhaypublicationsmain.html>, 2009.
- [3] B. Alomair, L. Lazos, and R. Poovendran. Securing Low-cost RFID Systems: an Unconditionally Secure Approach. *2010 Workshop on RFID Security-RFIDsec'10 Asia*, 2010.
- [4] G. Avoine. Privacy issues in RFID banknote protection schemes. In *International Conference on Smart Card Research and Advanced Applications-CARDIS'04*. IFIP, 2004.
- [5] G. Avoine, E. Dysli, and P. Oechslin. Reducing time complexity in RFID systems. *Selected Areas in Cryptography-SAC*, 3897:291–306, 2005.
- [6] M. Burmester, B. De Medeiros, and R. Motta. Anonymous RFID authentication supporting constant-cost key-lookup against active adversaries. *International Journal of Applied Cryptography*, 1(2):79–90, 2008.
- [7] H.-Y. Chien. SASI: A New Ultralightweight RFID Authentication Protocol Providing Strong Authentication and Strong Integrity. *IEEE Transactions on Dependable and Secure Computing*, 2007.
- [8] T. Dimitriou. A Lightweight RFID Protocol to protect against Traceability and Cloning attacks. In *Conference on Security and Privacy for Emerging Areas in Communication Networks – SecureComm*, 2005.
- [9] S. Garfinkel, A. Juels, and R. Pappu. RFID Privacy: An Overview of Problems and Proposed Solutions. *IEEE SECURITY & PRIVACY*, 2005.
- [10] G. Khandelwal, K. Lee, A. Yener, and S. Serbetli. ASAP: a MAC protocol for dense and time-constrained RFID systems. *EURASIP Journal on Wireless Communications and Networking*, 2007.
- [11] M. Kodialam and T. Nandagopal. Fast and reliable estimation schemes in RFID systems. In *12th annual international conference on Mobile computing and networking, MobiCom*, 2006.
- [12] L. Lu, J. Han, L. Hu, Y. Liu, and L. Ni. Dynamic Key-Updating: Privacy-Preserving Authentication for RFID Systems. *International Conference on Pervasive Computing and Communications*, 2007.
- [13] L. Lu, J. Han, R. Xiao, and Y. Liu. ACTION: Breaking the Privacy Barrier for RFID Systems. *INFOCOM 2009. The 28th IEEE Conference on Computer Communications*.
- [14] D. Molnar and D. Wagner. Privacy and security in library RFID: issues, practices, and architectures. *11th ACM CCS*, 2004.
- [15] S. Ross. *A First Course in Probability*. Prentice Hall, 2002.
- [16] B. Song and C. J. Mitchell. RFID Authentication Protocol for Low-cost Tags. In *ACM WiSec*, 2008.
- [17] W. Wang, Y. Li, L. Hu, and L. Lu. Storage-Awareness: RFID Private Authentication based on Sparse Tree. In *Privacy and Trust in Pervasive and Ubiquitous Computing, SECPeU 2007.*, 2007.