

Man-in-the-Middle Secure Authentication Schemes from LPN and Weak PRFs

Vadim Lyubashevsky*

Daniel Masny†

Abstract

We show how to construct, from any weak pseudorandom function, a 3-round symmetric-key authentication protocol that is secure against man-in-the-middle attacks. The construction is very efficient, requiring both the secret key and communication size to be only $3n$ bits long. Our techniques also extend to certain classes of randomized weak-PRFs, chiefly among which are those based on the classical LPN problem and its more efficient variants such as Toeplitz-LPN and Ring-LPN. Building a man-in-the-middle secure authentication scheme from any weak-PRF resolves a problem left open by Dodis et al. (Eurocrypt 2012), while building a man-in-the-middle secure scheme based on any variant of the LPN problem solves the main open question in a long line of research aimed at constructing a *practical* light-weight authentication scheme based on learning problems, which began with the work of Hopper and Blum (Asiacrypt 2001).

1 Introduction

The need for light-weight cryptography is increasing rapidly due to the growing deployment of low-cost devices, such as smart cards and RFID tags, in the real world. One of the most common cryptographic protocols required on these devices is a symmetric key authentication protocol in which the prover (usually referred to as the Tag) authenticates his identity to the verifier (usually referred to as the Reader). The most direct way in which this protocol can be constructed is by using a pseudorandom function f (e.g. AES) for which the Tag and the Reader share a common key. Then the authentication protocol simply consists of the Reader sending a challenge c to which the Tag replies with $f(c)$, and the Reader verifies that the received evaluation of c is indeed correct. The main problem with this approach is that the pseudorandom function, whether it is a “provably-secure” one based on some mathematical assumption or an “ad-hoc” block cipher like AES, is usually quite costly for light-weight devices. For this reason, researchers have worked on designing block ciphers specifically for low-cost devices (e.g. [LPPS07, BKL⁺07]). A different approach for addressing this problem is constructing authentication schemes from building blocks that have weaker security properties than block ciphers or pseudorandom functions. In the present work, we follow this latter avenue of research.

1.1 Authentication from LPN

The Learning Parity with Noise (LPN) problem was initially shown to have cryptographic applications by Blum et al. [BFKL93], and then used as a basis for authentication schemes by Hopper and Blum in their HB scheme [HB01]. In this latter paper, a simple LPN-based authentication scheme was proposed that was secure in the *passive* attack model. Later work by Juels and Weis [JW05],

*lyubash@di.ens.fr. INRIA and ENS, Paris

†daniel.masny@ruhr-uni-bochum.de. Ruhr-Universitat Bochum. Part of this work was done while visiting ENS.

and also by Katz and Shin [KS06], modified this protocol (the result was called HB^+) to be secure against *active* adversaries. Nevertheless, even these schemes had a serious security shortcoming. If the adversary were allowed to modify the communication between the Tag and the Reader and observe the response of the reader to verification queries, then, as shown by Gilbert et al. [GRS05], there exists a very simple attack that can recover the secret key in polynomial time. Because such a *man-in-the-middle* attack can be mounted with relatively small effort, schemes that fall to it cannot be considered secure enough for real-world applications that require some decent level of security. It was thus a major open problem to construct an efficient LPN-based authentication scheme that remains secure against man-in-the-middle attacks.

A notable advance was made by Gilbert et al. [GRS08b] who proposed a scheme (termed $\text{HB}^\#$) that was able to resist the attack from [GRS05] and was shown to be secure against restricted man-in-the-middle adversaries. A second contribution of this work was to offer a solution to another problematic feature of previous LPN-based protocols. All protocols that are based on LPN require either the key size or the communication complexity to be square in the security parameter. Thus either the key size or the communication complexity would have to be on the order of hundreds of thousands of bits. Since the main motivation for LPN-based protocols is low-cost hardware, this is clearly unacceptable. To this end, [GRS08b] proposed a protocol based on a related assumption, called Toeplitz-LPN (see Section 2.2 for definitions), where the communication complexity was small and the secret key had some structured form which allowed for compact representations. While there has been no known weakness caused by using the Toeplitz-LPN assumption, it did turn out that the restricted man-in-the-middle model introduced in [GRS08b] was not sufficient to prevent all practical attacks, and one such attack was shown by Ouafi et al. [OOV08].

There have been many other proposals, some without security proofs, others with claimed proofs that attempted to solve this problem, but all of these methods were ultimately shown to be flawed (see [GRS08a] for a small overview). A breakthrough finally came in a series of recent papers by Kiltz et al. [KPC⁺11] and Dodis et al. [DKPW12] who constructed relatively-efficient MACs based on the hardness of the LPN problem. Because MACs immediately give rise to man-in-the-middle secure authentication schemes, their work also resolved the problem of building such schemes from the LPN problem. This LPN MAC, however, suffered from the same drawback as other LPN-based schemes – the key size was prohibitively large. Thus in order to be useful in practice, the proof techniques would have to be adapted to work with more compact LPN-related assumptions, such as Toeplitz-LPN. But the constructions of [KPC⁺11] and [DKPW12] made use of certain algebraic structure of the LPN-problem, and the proofs turn out to be incompatible with other previously-considered versions of LPN.

1.2 Authentication from Weak-PRFs

Weak pseudo-random functions are keyed functions whose outputs on *random* inputs are indistinguishable from uniform. Weak PRFs are considered to be much “weaker” primitives than PRFs, and in particular, it is not known how to transform a weak-PRF into a PRF except by using tree techniques similar to the classical GGM construction [GGM86]. Additionally, it also appears to be much easier to build secure weak-PRFs than PRFs. For example, the function $f_a(x) = x^a \bmod p$ is a weak-PRF based on the DDH assumption, whereas the construction of a PRF based on DDH is much less efficient [NR97], requiring n multiplications in addition to the exponentiation in the weak-PRF. Similarly, the recent construction of lattice-based PRFs [BPR12] first builds a relatively efficient weak-PRF (which is just $f_a(x) = \text{Round}(x \cdot a \bmod p)$, where $x, a \in \mathbb{Z}_p^n$, and the $\text{Round}(\cdot)$

function drops some number of least-significant bits) and then converts it to a full PRF using techniques similar to [NR97, NRR02]. The resulting lattice-based PRF is both less efficient and requires a stronger computational assumption than the underlying weak-PRF.

Due to efficiency advantages and lower security requirements, there has been some research on constructions of cryptographic primitives such as symmetric encryption and stream ciphers built directly from weak-PRFs (e.g. [DN02, MS07, Pie09]). The work along this theme that is most related to ours is the aforementioned one of Dodis et al. [DKPW12], where it is shown how to build a 3-round authentication scheme secure against *active* attacks from any weak-PRF. As we mentioned earlier, the active security model, where the adversary is not allowed to send any verification queries to the Reader, is not considered strong enough for real-world applications. And so the problem of constructing man-in-the-middle secure authentication schemes from arbitrary weak-PRFs remained open.

1.3 Our Results

Our first result is a construction, from any weak pseudorandom function, of a 3-round symmetric-key authentication protocol that is secure against man-in-the-middle attacks. Our scheme has the exact same communication complexity as the actively-secure scheme of [DKPW12], and only has one extra key element. To be more precise, the secret keys in our scheme consist of the key of the weak-PRF plus the description of a pairwise-independent hash function, which requires an additional two elements whose size is the output length of the weak-PRF. So if we assume that both the domain and range of the weak-PRF is n bits, then the total key size is $3n$.

We then extend our construction of a weak-PRF scheme to *randomized* weak-PRFs. Randomized weak-PRFs are keyed functions that become computationally indistinguishable from uniform when their outputs are perturbed by some low-entropy noise. Noisy learning problems such as LPN and LWE [Reg09] can be equivalently viewed as problems of distinguishing the outputs of a randomized weak-PRF from the uniform distribution. To get a man-in-the-middle secure authentication scheme from a randomized weak-PRF, we require just one more secret key element than our weak-PRF based scheme.

Our constructions, and to some extent their security proofs as well, turn out to be surprisingly simple. The main insight is that one should embed the n -bit output of the (randomized) weak-PRF into a finite field of size 2^n . Then, in addition to the secret keys associated to the function, we also create secret keys in the field which end up being masked by the presumed indistinguishability from uniform of the (randomized) weak-PRF. We then show how the interplay in the field between the weak-PRF and the additional secret keys results in protocols that have the desired man-in-the-middle security.

We prove security of our schemes in the sequential man-in-the-middle model, in which the adversary simultaneously interacts with one copy of the Tag and Reader (see Figure 1). There is a stronger notion of concurrent man-in-the-middle security where the adversary is allowed to simultaneously communicate with multiple clones of the Tag and Reader possessing the same secret key, but we do not prove security in this model.¹ While the concurrent model is theoretically stronger, we do not believe that it is practically relevant to the low-cost device setting considered in this paper. In particular, there is no reason that a low-cost Tag would have the need (or even

¹In Appendix A, we show that for certain instantiations of a randomized weak-PRF there indeed exists a concurrent man-in-the-middle attack.

Protocol	# r	Security			Complexity	
		assumption	active	MIM	key size	com.
weak-PRF [DKPW12]	3	weak-PRF	$\sqrt{\epsilon}$?	$2n$	$3n$
weak-PRF [this work]	3	weak-PRF	$q_v \cdot \sqrt{\epsilon}$		$3n$	$3n$
HB ⁺ [JW05, KS06]	3	LPN _{n,τ}	$\sqrt{\epsilon}$	X [GRS05]	$2n$	$2n^2$
Random-HB [#] [GRS08b]	3	LPN _{n,τ}	$\sqrt{\epsilon}$	X [OOV08]	$2n^2$	$3n$
HB [#] [GRS08b]	3	Toeplitz-LPN _{n,τ}	$\sqrt{\epsilon}$	X [OOV08]	$4n$	$3n$
MAC ₁ [KPC ⁺ 11]	2	LPN _{n,τ}	ϵ	$2^\lambda \cdot \epsilon$	$2n^2$	$4n$
MAC ₂ [KPC ⁺ 11]	2	LPN _{n,τ}	ϵ	$Q \cdot \epsilon$	λn^2	$4n$
Lapin [HKL ⁺ 12]	2	Ring-LPN _{n,τ}	ϵ	?	$2n$	$3n$
MAC ₁ + Lapin	2	Ring-LPN _{n,τ}	ϵ	$2^\lambda \cdot \epsilon$	$6n + 2\lambda$	$4n$
LPN-based [this work]	3	LPN _{n,τ}	$q_v \cdot \sqrt{\epsilon}$		n^2	$3n$
		Toeplitz-LPN _{n,τ}			$5n$	
		Ring-LPN _{n,τ}			$4n$	

Table 1: **Authentication Protocols Based on Weak-PRFs and the LPN-related Assumptions.** Listed is the amount of authentication rounds $\#r$, the security properties achieved by the protocol and its complexity (with lower order terms dropped) according to the key size and the communication. Let ϵ be the advantage in breaking the assumption, then the term depending on ϵ is proven to be the best possible advantage of breaking the protocol in the given model. Q is the amount of tag and verification queries whereas q_v is defined as the amount of verification queries, which is $q_v = 1$ in the active model. n parameterizes the hardness of the assumption and λ is the statistical security parameter. [DKPW12] gives an alternate construction of MAC₁ and MAC₂ with better computational complexity, but the rest of the properties are basically the same.

ability) to simultaneously participate in more than one authentication session. Furthermore, it also seems unlikely that in an ecosystem where one wants to have relatively strong security, secret keys would be shared among the Tags. Still, constructing an efficient authentication scheme from generic weak-PRFs that is secure in the concurrent man-in-the-middle model remains an interesting open problem. ²

1.4 Comparison to Other Works

Table 1 compares the results obtained in this paper with those of previous works. Compared to the protocols that only achieve active security, our scheme achieves the much stronger man-in-the-middle security at a fairly small cost. In the case of protocols based on a generic weak-PRF, we extend the security to the man-in-the middle model at the cost of only one extra secret key element and one extra field multiplication. We get similar results when comparing our protocol with actively-secure LPN-based ones.

It is also interesting to compare our LPN protocol to the MAC constructions in [KPC⁺11]. There are three advantages to the MAC constructions – they are only two rounds, they have slightly tighter reductions to LPN, and they are secure in the concurrent man-in-the-middle model, whereas our scheme is secure in the sequential man-in-the-middle model. The advantages of our construction are that the key sizes and the communication complexities are smaller.

²We note that our current scheme is still secure in the concurrent model against *active* attacks.

The above-listed differences between our LPN scheme and the MAC schemes are, in our opinion, fairly minor with several pluses and minuses on both sides. In practice, it makes almost no difference whether the authentication scheme is 2 or 3 rounds since the Tag is the one who starts the protocol – thus a 2-round protocol essentially becomes a 3-round one. And while security tightness is certainly a desirable property, it is very unclear what effects it has in practice. Similar public key authentication schemes, such as GQ [GQ88] and Schnorr [Sch91], have been studied for a long time, yet do not exhibit any weaknesses due to their non-tight reductions.

The major advantage of our construction is that it is *generic* and can be instantiated with virtually any version of the LPN function or a randomized weak-PRF satisfying a few mild properties (see Section 2.1). For example, our construction allows for authentication schemes based on the fairly well-studied Toeplitz-LPN assumption, which seems to provide a very good compromise between security and computational efficiency. The constructions of [KPC⁺11] and [DKPW12], on the other hand, can only construct MACs from functions with very “algebraic” properties.

The recent work of Heyse et al. [HKL⁺12] proposed a new LPN-type assumption, called Ring-LPN, to enable efficient constructions that are compatible with the MAC transformation in [KPC⁺11]. The assumption is relatively new, and its unclear at this point whether it has the same hardness as the more well-studied LPN and Toeplitz-LPN assumptions. Still, even if the Ring-LPN problem is hard, our LPN protocol can also be instantiated based on this assumption and is more efficient than the resulting MAC transformation.

2 Preliminaries and Notation

2.1 Function Families and their Properties

In this section we define the important classes of functions that will appear in the paper. As mentioned earlier, we will be considering embeddings of function outputs into a finite field. The embedding can be arbitrary, and the simplest one is to simply think of a function output string $s \in \{0, 1\}^n$ as a polynomial in a finite field $\mathbb{F} = (\mathbb{Z}_2^n, +, \times)$. Thus, without loss of generality, we will assume that all our functions output elements to some finite field \mathbb{F} .

Definition 2.1. *A function family $\mathcal{H} : \mathbb{D} \rightarrow \mathbb{F}$ is called pairwise-independent if for all $x_1 \neq x_2 \in \mathbb{D}$, $y_1, y_2 \in \mathbb{F}$,*

$$\Pr_{h \leftarrow \mathcal{H}} [h(x_1) = y_1 \wedge h(x_2) = y_2] = 1/|\mathbb{F}|^2$$

Definition 2.2. *A function family $\mathcal{F} : \mathbb{D} \rightarrow \mathbb{F}$ is said to be a weak-PRF family if for any polynomial-sized k , randomly-chosen $f \in \mathcal{F}$, and randomly-chosen $r_1, \dots, r_k \in \mathbb{D}$, the distribution of $(r_1, f(r_1)), \dots, (r_k, f(r_k))$ is computationally indistinguishable from the uniform distribution over $(\mathbb{D}, \mathbb{F})^k$.*

Even if $(r_1, f(r_1)), \dots, (r_k, f(r_k))$ can be distinguished from the uniform distribution over $(\mathbb{D}, \mathbb{F})^k$, it’s possible that the sequence can become indistinguishable if the outputs $f(r_i)$ were perturbed by some noise. Such function families are called randomized weak-PRFs. The noise perturbation can be anything, but in this paper we will only consider noise distributions with an eye towards LPN applications. In particular, both the noise and the output of $f(r_i)$ are group elements, and the perturbation consists of adding the two together. This is still consistent with our requirement

of being able to embed the output of all functions into a finite field \mathbb{F} since the group needed for LPN can simply be the underlying additive group of \mathbb{F} (see Section 2.2).

Definition 2.3. For a function $f(\cdot) : \mathbb{D} \rightarrow \mathbb{F}$ and a distribution χ over \mathbb{F} , we will write $f^\chi(r)$ to mean a randomized function that generates an element $e \in \mathbb{F}$ according to the distribution χ and outputs $f(r) + e$. A function family $\mathcal{F} : \mathbb{D} \rightarrow \mathbb{F}$ is said to be a randomized weak-PRF family with noise χ if for any polynomial-sized k , randomly-chosen $f \in \mathcal{F}$, and randomly-chosen $r_1, \dots, r_k \in \mathbb{D}$, the distribution of $(r_1, f^\chi(r_1)), \dots, (r_k, f^\chi(r_k))$ is computationally indistinguishable from the uniform distribution over $(\mathbb{D}, \mathbb{F})^k$.

In order for randomized weak-PRFs to be useful for cryptographic constructions, the range \mathbb{F} and the error distribution should have certain characteristics. For example, the weak-PRFs would be of very little use if the error distribution χ was just the uniform distribution over \mathbb{F} . In this paper we will assume that the additive group of the field \mathbb{F} and the error distribution χ satisfy the following three properties:

1. There exists a weight function $\|\cdot\| : \mathbb{F} \rightarrow \mathbb{R}^+$ such that the additive group that underlies the field \mathbb{F} satisfies the triangle inequality – that is for all $a, b \in \mathbb{F}$, $\|a \pm b\| \leq \|a\| + \|b\|$. Additionally, $\|a\| = 0$ if and only if $a = 0$.
2. There exists a positive real $\tau' \in \mathbb{R}$ such that $\Pr_{e \sim \chi}[\|e\| \leq \tau'] = 1 - n^{-\omega(1)}$.³
3. For a positive real α , let $\beta(\alpha) = \{z \in \mathbb{F} : \|z\| \leq \alpha\}$. We will assume that $|\beta(2\tau')|/|\mathbb{F}| = n^{-\omega(1)}$.

The first property essentially makes sure that the randomness in the randomized weak-PRF behaves “nicely” via the triangular inequality.⁴ The second property determines the completeness of our protocol. Additionally, because of the way our security proof works, the completeness of the protocol also plays a role in the soundness of the protocol.⁵ Thus this value should be very close to 1. The third property determines the soundness of the protocol. Intuitively, it is related to the probability that an adversary can randomly guess a response and be accepted by the verifier.

Due to their similarity, we will be presenting our authentication scheme and its proof based on weak-PRFs together with the ones based on randomized weak-PRFs. Since a weak-PRF is just a randomized weak-PRF whose error distribution χ has its support entirely on 0, it’s easy to see that it can trivially be made to satisfy the above three properties. We can define the weight function as $\|x\| = 1$ for all $x \neq 0$ and set $\tau' = 0$. Thus for weak-PRFs we have $\Pr_{e \sim \chi}[\|e\| \leq \tau'] = 1$ (and so the protocol will have perfect completeness) and $|\beta(2\tau')|/|\mathbb{F}| = |\{0\}|/|\mathbb{F}| = 1/|\mathbb{F}|$.

2.2 Randomized Weak-PRFs from the LPN Problem and its Variants

The classical decisional $\text{LPN}_{n,\tau}$ assumption states that the uniform distribution over $\mathbb{Z}_2^n \times \mathbb{Z}_2$ is computationally-indistinguishable from the following distribution: for a fixed randomly-chosen vector $s \in \mathbb{Z}_2^n$, output $(r, r \cdot s + e)$ where r is chosen uniformly random from \mathbb{Z}_2^n and e is a Bernoulli

³More formally, τ' is a function of n , $\tau'(n)$, but we will omit the n throughout the paper.

⁴Even though we are using the standard notation for “norm”, the weight function $\|\cdot\|$ is not quite a norm because it’s not true that for all integers α , $\alpha\|a\| = \|\alpha a\|$ (since we are working over a finite field).

⁵This seems to be a common feature of protocols that have man-in-the-middle security because the simulator replies to the adversary under the assumption that properly-formed responses by the Tag are accepted by the Reader. Even though it is not stated in [KPC⁺11, DKPW12], the soundness of their protocols also depends on their completeness in exactly the same way as in this work.

random variable that is 1 with probability τ . By the hybrid argument, it is easy to see that if the fixed secret is now a matrix $S \in \mathbb{Z}_2^{m \times n}$ then the distribution $(r, Sr + e)$, where r is chosen as before and e is a vector each of whose coefficients is 1 with probability τ , is also computationally-indistinguishable from the uniform distribution over $\mathbb{Z}_2^n \times \mathbb{Z}_2^m$ (with a loss of a factor m in the reduction). We now formulate this latter statement in terms of the randomized weak-PRF notation from the previous subsection.

Let Ber_τ^m be a distribution over \mathbb{Z}_2^m where every coordinate is independently chosen to be 1 with probability τ and 0 with probability $1 - \tau$.

Definition 2.4 (LPN). *Let $\mathcal{F} : \mathbb{Z}_2^n \rightarrow \mathbb{Z}_2^m$ be a function family indexed by matrices $S \in \mathbb{Z}_2^{m \times n}$. For a function $f_S \in \mathcal{F}$ and a vector $r \in \mathbb{Z}_2^n$, define $f_S(r) := Sr$. Then the $LPN_{n,\tau}$ assumption implies that \mathcal{F} is a randomized weak-PRF family with noise Ber_τ^m .*

In the above definition, the domain \mathbb{D} of \mathcal{F} is \mathbb{Z}_2^n . Because we insisted in Definition 2.3 that the range of the function family \mathcal{F} be a finite field (this will be used in our protocol) and the LPN problem only requires an additive group structure, we have some freedom as to how to define this field. The LPN assumption requires the range to have the group structure $(\mathbb{Z}_2^m, +)$, thus \mathbb{F} can be any finite field that has $(\mathbb{Z}_2^m, +)$ as its underlying additive group. The most natural definition is $\mathbb{F} = \mathbb{Z}_2[x]/(g(x))$ where $g(x)$ is a polynomial of degree m that is irreducible over \mathbb{Z}_2 , and addition and multiplication are just standard polynomial addition and multiplications modulo 2 and $g(x)$. Thus addition in $(\mathbb{F}, +, \times)$ exactly corresponds to addition in $(\mathbb{Z}_2^m, +)$.

The randomized weak-PRF based on LPN can also quite naturally be made to satisfy the three properties after Definition 2.3. The weight function $\|\cdot\|$ can be defined to be the Hamming weight. That is, for any element $a \in \mathbb{Z}_2^m$, $\|a\|$ is the number of 1's in a . With this definition of the weight function, one can compute, via the Chernoff bound, a τ' such that any element e chosen according to Ber_τ^m satisfies $\|e\| \leq \tau'$ with overwhelming probability. To satisfy the third property, we would need that $|\beta(2\tau')|/|\mathbb{F}| = n^{-\omega(1)}$, which is equivalent to the condition that $\left(\sum_{i=0}^{\lfloor 2\tau' \rfloor} \binom{m}{i}\right) / 2^m = n^{-\omega(1)}$. The above conditions are identical to those in other authentication protocols, such as [KS06, KSS10, GRS08b], and so the LPN parameters needed to make those schemes secure, also carry over to ours.

Because the LPN problem yields rather inefficient schemes, Gilbert et al. [GRS08b] proposed protocols based on the hardness of the Toeplitz-LPN problem, which is just like the LPN problem except that the secret matrix S is a Toeplitz matrix.

Definition 2.5 (Toeplitz-LPN). *Let $\mathcal{F} : \mathbb{Z}_2^n \rightarrow \mathbb{Z}_2^m$ be a function family indexed by Toeplitz matrices $S \in \mathbb{Z}_2^{m \times n}$. For a function $f_S \in \mathcal{F}$ and a vector $r \in \mathbb{Z}_2^n$, define $f_S(r) := Sr$. Then the $Toeplitz-LPN_{n,\tau}$ assumption implies that \mathcal{F} is a randomized weak-PRF family with noise Ber_τ^m .*

Heyse et al. [HKL⁺12] recently introduced the Ring-LPN problem, which also results in more efficient protocols. While the Ring-LPN problem has not been well-studied, it does have some resemblance to the better-studied Ring-LWE problem [LPR10] in lattice cryptography, and so there are some reasons to believe that it might be secure.

Definition 2.6 (Ring-LPN). *Let $g(x)$ be a polynomial of degree n in $\mathbb{Z}_2[x]$ irreducible over \mathbb{Z}_2 and define the field \mathbb{F} to be $\mathbb{F} = \mathbb{Z}_2[x]/(g(x))$. Let $\mathcal{F} : \mathbb{F} \rightarrow \mathbb{F}$ be a function family indexed by polynomials $s \in \mathbb{F}$. For a function $f_s \in \mathcal{F}$ and a polynomial $r \in \mathbb{F}$, define $f_s(r) := sr$. Then the $Ring-LPN_{n,\tau}$ assumption implies that \mathcal{F} is a randomized weak-PRF family with noise Ber_τ^n .*




Figure 1: **Man-in-the-Middle Attack Model.**

2.3 Security Models

All authentication schemes are protocols in which the Tag and the Reader possess some secret key sk and then perform an interaction in which the Tag must convince the Reader of his identity. The difference in the security models depends on the strength that we give the adversary. The three most natural security models are *passive*, *active*, and *man-in-the-middle*. All three models consist of two stages. In the first stage, depending on the model, the Adversary is allowed to have some interaction with the Tag and the Reader. In the second stage, in all three models, he loses the interaction with the Tag and must interact with the Reader in hopes of getting the latter to accept the interaction. We now briefly describe the the passive and active attack models, and then describe in detail the man-in-the-middle model.

Passive Adversary. The weakest adversary is a *passive* one. In the first stage of the security game, the Adversary simply observes the interaction between the Tag and the Reader, and in the second stage he interacts with the Reader in hopes of convincing the latter to accept the interaction. Such a scheme is very simple to construct from any weak-PRF. The secret key is a random weak-PRF $f \in \mathcal{F}$, and the protocol consists of the Reader sending a random $r \in \mathbb{D}$ and the Tag replying with $z = f(r)$. The Reader accepts the interaction iff $z = f(r)$. To the passive adversary, the pairs $(r, f(r))$ look uniformly-random, and it's easy to finish the proof by showing that an adversary who is able to get the reader to accept can be used to distinguish the weak-PRF outputs from uniform. The same idea can also be used to build a scheme based on a randomized weak-PRF, where the reader sends a random $r \in \mathbb{D}$, the tag replies with $z = f^x(r)$, and the reader accepts iff $\|f(r) - z\| \leq \tau'$. This idea was used by Hopper and Blum to construct a passive authentication scheme based on the LPN problem [HB01].

Active Adversary. A somewhat stronger adversary is one who, in addition to just watching the interaction between the Tag and the Reader in the first stage, can also interact with the Tag (although he is not able to interact with the Reader, and in particular, not allowed to make any verification queries). After the first stage, the Adversary loses access to the Tag and interacts with the Reader in hopes of being accepted. The first actively-secure protocol based on the LPN problem was built by Juels and Weis [JW05], and then simplified and improved by Katz and Shin [KS06]. The LPN construction of [KS06] was generalized by Dodis et al. to any weak-PRF [DKPW12].

Man-in-the-Middle Adversary. The strongest type of Adversary is one who in the first stage is able to simultaneously interact with the Tag and the Reader and make *verification queries* to the Reader. In the second stage, the Adversary again loses access to the Tag, and interacts with the Reader hoping to make the latter accept. In this paper, the protocols we will be constructing will be sigma protocols (i.e. have three rounds usually referred to as *commit*, *challenge*, and *response*) and will use a model that is simpler to describe and is at least as secure as the man-in-the-middle one. We now describe the security game and the Adversary’s condition for winning it:

Setup: Generate a secret key and give it to the Tag \mathcal{T} and the Reader \mathcal{R} .

Attack: Invoke the Adversary \mathcal{A} who has access to \mathcal{T} and \mathcal{R} and let him interact with them t times. Each of the interactions is as follows (see Figure 1):

\mathcal{A} receives a commitment r_i from \mathcal{T} and sends a commitment $r_i + r'_i$ to \mathcal{R} . \mathcal{R} responds with a challenge c_i and \mathcal{A} sends a challenge $c_i + c'_i$ to \mathcal{T} . \mathcal{T} answers with a valid response z_i . \mathcal{A} can now send his response $z_i + z'_i$ for verification to \mathcal{R} . \mathcal{R} answers with *accept*, if $(r_i + r'_i, c'_i, z_i + z'_i)$ is valid, otherwise he answers with *reject*.

Winning Condition: We say that the Adversary \mathcal{A} wins the game if at some point he makes a query to \mathcal{R} such that $(r'_i, c'_i, z'_i) \neq (0, 0, 0)$ and the Reader \mathcal{R} sends *reply = accept*.

Notice that if there is an Adversary who can win the two stage Man-in-the-Middle game (i.e. where he loses access to the Tag in the second stage and must get the reader to accept), then he can also win the game described above since he can simply ignore the messages sent by the Tag in the second stage. Thus security in the model that we will be using in this paper implies security in the “more natural” two stage model.

3 Construction Based on a (randomized) weak-PRF

In this section we present our main construction, an authentication protocol secure against man-in-the-middle attacks from any weak-PRF or a randomized weak-PRF that satisfies the three properties stated after Definition 2.3. The protocol based on a weak-PRF is very similar to the one based on a randomized weak-PRF, and so we present them together in Figure 2. The security proofs are also very similar, and we also present them together in the next section.

The underlying building blocks of the protocol in Figure 2 are a pairwise-independent function family \mathcal{H} and a family \mathcal{F} of randomized weak-PRFs with noise χ . If \mathcal{F} is a family regular (non-randomized) weak-PRFs, then it’s the same as a randomized weak-PRF with noise χ , where χ has all of its support on 0 – thus for all $f \in \mathcal{F}$, $f^\chi(\cdot) = f(\cdot)$. The secret keys of the authentication scheme are randomly chosen $f \in \mathcal{F}$, $h \in \mathcal{H}$, and $s \in \mathbb{F}$. In the case that \mathcal{F} is a regular weak-PRF family, we do not need the extra key s , and in the protocol we can assume that $s = 1$. In the case that \mathcal{F} is a randomized weak-PRF family, we assume that it satisfies the three properties after Definition 2.3. Thus there is an associated weight function $\|\cdot\|$ and a value τ' such that the error e chosen from χ satisfies $\|e\| \leq \tau'$ with overwhelming probability.

In the first step of the protocol, the Tag picks a random element $r \in \mathbb{D}$ and sends it to the Reader. The reader chooses a random $c \in \mathbb{F}$ and sends it to the Tag. In its turn, the Tag evaluates $f^\chi(r)$ and $h(r)$, and sends $z = f^\chi(r)s + h(r)c$ back to the Reader, where all addition and multiplication operations take place in the field \mathbb{F} . In the case that \mathcal{F} is a regular weak-PRF family, the response




Figure 2: **Authentication Protocol Based on a (randomized) weak-PRF.** If the weak-PRF is not randomized, (i.e. the support of the distribution χ is 0 and $\tau' = 0$), then we can set $s = 1$. In this case, the condition $\|f(r) - s^{-1}(z - h(r)c)\| \leq \tau'$ simplifies to $f(r) = z - h(r)c$.

of the Tag is simply $z = f(r) + h(r)c$. The Reader accepts the Tag if $\|f(r) - s^{-1}(z - h(r)c)\| \leq \tau'$. In case of a regular weak-PRF family without noise, this condition is equivalent to $f(r) = z - h(r)c$.

Efficiency Note. As in most 3-round authentication schemes, it is only necessary for the challenge c to have as much entropy as the security of the scheme. Thus if the the field \mathbb{F} is very large, it is fine to restrict the domain of c to some arbitrary, large enough, subset of \mathbb{F} . This would slightly reduce the communication complexity of the protocol and possibly make the multiplication $h(r)c$ more efficient. Similarly, the secret key s need not be chosen uniformly from all of \mathbb{F} either. The only time the randomness of s is used in the proof is in Lemma 4.3, and there it is enough for s to also be chosen from a large enough subset of \mathbb{F} . Restricting the domain of s would make the multiplication by it more efficient and also slightly decrease the key size of all our LPN-based instantiations. So for example, the Toeplitz-LPN and Ring-LPN instances in Table 1 would have key sizes $4n + \lambda$ and $3n + \lambda$ instead of $5n$ and $4n$. For simplicity, in this paper we just pretend that s and c get chosen from all of \mathbb{F} .

Example Instantiation. We now give an example instantiation of the protocol using the $\text{LPN}_{n,\tau}$ assumption from Definition 2.4. The noise distribution χ is Ber_τ^m and to choose the secret key f , a random $S \in \mathbb{Z}_2^{n \times m}$ is picked and $f^\chi(r) := Sr + e$ where $e \sim \text{Ber}_\tau^m$. Thus f maps the domain \mathbb{Z}_2^n to \mathbb{Z}_2^m . As in the discussion following Definition 2.4, the field \mathbb{F} is defined to be $\mathbb{Z}_2[x]/(g(x))$ where $g(x)$ is any irreducible polynomial of degree m . The simplest definition of a pairwise independent function family that maps \mathbb{Z}_2^n to \mathbb{F} is to index the family by two polynomials in \mathbb{F} . To pick a random element of the family, one randomly picks $a_1, a_2 \in \mathbb{F}$ and defines $h(r) = a_1r + a_2$, where r is treated like a polynomial in \mathbb{F} and multiplication and addition is performed over \mathbb{F} .⁶ The final secret key is a random polynomial $s \in \mathbb{F}$. Thus the secret keys are (S, a_1, a_2, s) .

⁶To be able to treat r as an element of \mathbb{F} , it is important that $m \geq n$. If $m < n$, then one can define the pairwise-independent function differently (e.g. $h(r) = a_1r_1 + \dots + a_kr_k + a_{k+1}$, where $r = r_1 | \dots | r_k$).

In the protocol, the Tag chooses an $r \in \mathbb{Z}_2^n$ and sends it to the Reader, who replies with a randomly-chosen $c \in \mathbb{F}$. The Tag receives the c computes $f^\chi(r) = Sr + e \in \mathbb{Z}_2^m$, and treats the result as a polynomial in \mathbb{F} . He then multiplies it by s and adds it to $h(r)c = (a_1r + a_2)c$, and sends the resulting $z = f^\chi(r)s + h(r)c$ to the Reader. The reader computes $f(r) = Sr$ and $s^{-1}(z - h(r)c)$, and accepts if the weight of $f(r) - s^{-1}(z - h(r)c)$ is less than or equal to τ' .

Notice that the protocol would be exactly the same for the Toeplitz-LPN $_{n,\tau}$ problem, with the only difference being how S is defined. By having S be a Toeplitz matrix, the key storage space shrinks from $mn + 3m$ to $n + 4m$, and the matrix-vector multiplication Sr can be computed more efficiently. The Ring-LPN $_{n,\tau}$ protocol would also work in essentially the same way. In this case, we set $m = n$ and have $\mathbb{D} = \mathbb{F}$. The secret key S will just be a random polynomial in \mathbb{F} just like s, a_1 , and a_2 . Thus Sr will simply be a multiplication of two polynomials in the field \mathbb{F} .

Lemma 3.1. *The completeness of the authentication protocol is $\Pr_{e \sim \chi}[\|e\| \leq \tau']$. And in particular, if the weak-PRF is not randomized, the completeness is 1.*

Proof. The Tag sets $z \leftarrow f^\chi(r)s + h(r)c = (f(r) + e)s + h(r)c$, where $e \sim \chi$. Thus $f(r) - s^{-1}(z - h(r)c) = e$, and so the Reader accepts whenever $\|f(r) - s^{-1}(z - h(r)c)\| = \|e\| \leq \tau'$. \square

4 Security of the Authentication Scheme

Theorem 4.1. *Suppose that the authentication protocol in Figure 2 has completeness κ and there is a man-in-the-middle adversary who successfully breaks this scheme with probability ϵ while making at most q_v verification queries. Then there exists an algorithm which, in the same amount of time, has advantage $\frac{1}{2} \left(\kappa^{q_v-1} (\epsilon/q_v - 1/|\mathbb{F}|)^2 - \beta(2\tau')/|\mathbb{F}| \right)$ in breaking the uniformity assumption of the (randomized) weak-PRF family \mathcal{F} .*

Proof. If an adversary making q_v verification queries wins the game, then one of these q_v queries can be thought of as the “winning query”. By “winning query”, we mean that it is the *first* accepted query such that $(r'_i, c'_i, z'_i) \neq (0, 0, 0)$ (where r'_i, c'_i, z'_i are as in Figure 1). Once the Adversary sends such a query, he wins the game. If the Adversary has an ϵ success probability of winning the MIM-game, then by an averaging argument there must be some integer $i^* \leq q_v$ such that the probability that the Adversary wins the game and query number i^* is the “winning query” is at least ϵ/q_v . For the rest of the proof, we will assume that we know this i^* (which can be determined a priori by running the adversary on known inputs.)

The Challenger gives us ordered pairs $(r_i, y_i) \in \mathbb{D} \times \mathbb{F}$ where the r_i are uniformly random in \mathbb{D} and the y_i are either uniformly random in \mathbb{F} or equal to $f^\chi(r_i)$ (where f is a randomly-chosen function from the (randomized) weak-PRF family \mathcal{F} with noise χ). We will show how to use the adversary who breaks the authentication protocol with the i^* th winning query to decide which of the two distributions the Challenger is outputting.

We now proceed to show how to simulate the Tag and the Reader before the Adversary’s i^* th verification query (see Figure 3). We pick a random $s \in \mathbb{F}$ and $h \in \mathcal{H}$ as the secret keys, and upon receiving a pair (r_i, y_i) from the Challenger, we send r_i to the Adversary. The Adversary can then modify this and forward $r_i + r'_i$ to the Reader. The Reader picks a random $c_i \in \mathbb{F}$, sends it to the adversary, who then sends the possibly modified challenge $c_i + c'_i$ to the Tag. The Simulator playing as the Tag computes $h(r_i)(c_i + c'_i)$ using his secret key h , and then uses the y_i received from the challenger together with his other secret key s , to send $z_i = y_i s + h(r_i)(c_i + c'_i)$. After receiving z_i , the Adversary may send $z_i + z'_i$ to the verifier and make a verification query.




Figure 3: **Simulating the Tag and the Reader before the Adversary’s i^* th verification query.** If the weak-PRF is not randomized, then we set the secret key $s=1$ instead of choosing it at random from \mathbb{F} .

Case 1: The Challenger’s distribution is not random. Notice that if the Challenger sends $(r_i, y_i = f^\chi(r_i))$, then the responses of the Tag are exactly what they should be if the secret key were (f, s, h) . Thus if $(r'_i, c'_i, z'_i) = (0, 0, 0)$, the Reader who always sends “accept” is correct with probability κ (the completeness of the protocol). And if $(r'_i, c'_i, z'_i) \neq (0, 0, 0)$, the response of “reject” is also correct since the i^* th verification query has not yet been reached. Because the simulator has faithfully simulated the valid Tag and Reader up to this point with probability κ^{q_v-1} , the Adversary’s i^* th query will be the “winning one” (i.e. $(r', c', z') \neq (0, 0, 0)$ and $\|f(r + r') - s^{-1}((z + z') - h(r + r')c)\| \leq \tau'$) with probability $\kappa^{q_v-1}\epsilon/q_v$. In the next section we show that if the Adversary sends this winning query, then the simulator will be able to detect that it is indeed correct, and thus reply to the Challenger that the distribution is *not random*. The way that the simulator does the detection depends on whether $r' = 0$ or $r' \neq 0$. If $r' = 0$, then the detection is performed as in Figure 4 and the proof that the detection works (with probability 1) is in Lemma 4.3. On the other hand, if $r' \neq 0$, then we perform the detection as in Figure 5. For this part of the proof, we will need to rewind the adversary and have him respond correctly for two *different* challenges, which will happen with probability $\kappa^{q_v-1}(\epsilon/q_v - 1/|\mathbb{F}|)^2$ by Lemma B.1. If the Adversary does successfully respond for two different queries, then in Lemma 4.5 we show that we can detect this (again with probability 1) and thus correctly reply that the Challenger’s distribution is not random. We summarize the preceding paragraph with the following Lemma:

Lemma 4.1. *If the Challenger sends non-random ordered pairs $(r_i, y_i = f^\chi(r_i))$, then the simulator responds “not random” with probability at least $\kappa^{q_v-1}(\epsilon/q_v - 1/|\mathbb{F}|)^2$.*

Case 2: The Challenger’s distribution is uniformly random. We now move to the case that the Challenger sends uniformly random pairs $(r_i, y_i) \in (\mathbb{D}, \mathbb{F})$. We will show that in this case, even an all-powerful adversary cannot make the simulator reply “not random” to the challenger except with negligible probability. Notice that in the case that the y_i are uniform and independent of the r_i , the secret keys h, s chosen by the simulator are information-theoretically hidden throughout




Figure 4: Answering the Challenger after the Adversary’s i^* th verification query if $r'=0$.

the interaction in Figure 3. In the next section, we use this fact as well as the field properties of \mathbb{F} and the pairwise-independence of the function h to show that even an all-powerful adversary cannot fool the simulator into replying “not random” to the Challenger with probability greater than $\beta(2\tau')/|\mathbb{F}|$ (which is $1/|\mathbb{F}|$ in case \mathcal{F} is not randomized). The preceding is proven in Lemmas 4.4 and 4.6, for the cases where $r' = 0$ and $r' \neq 0$, respectively. This result is summarized in the following Lemma:

Lemma 4.2. *If the Challenger sends uniformly random ordered pairs $(r_i, y_i) \in (\mathbb{D}, \mathbb{F})$, then the simulator responds “not random” with probability at most $\beta(2\tau')/|\mathbb{F}|$.*

The statement of the Theorem follows by combining Lemmas 4.1 and 4.2. □

4.1 Answering the Challenger when $r' = 0$

We now proceed to explain how to use the Adversary’s i^* th verification query to construct a response for the Challenger. In this subsection, we deal with the case that when the Adversary makes this query, he does not change the commitment r – in other words $r' = 0$. If $r' = 0$, then it means that the value of $f(r)$ will be included in two places in the case that the Challenger sends pairs of the form $(r, y = f^x(r))$. First, the Challenger’s output $y = f^x(r)$ contains something close to $f(r)$, and second the winning query of the Adversary contains $f(r)$ since it must be that $\|f(r) - s^{-1}(z + z' - h(r)c)\| \leq \tau'$ in order for the query to be accepted. In Lemma 4.3, we show that by simply using the triangular inequality (Property 1 following Definition 2.3), it must be the case that $\|s^{-1}(z' + h(r)c')\| \leq 2\tau'$.

On the other hand, if the queries sent by the Challenger are uniformly random, then, as we already observed in the previous section, the view of the Adversary is independent of the secret keys s and h . Therefore the Adversary’s behavior will be exactly the same as in the case where s and h are chosen *after* he outputs his i^* th query. In Lemma 4.4, we use this to show that even an all-powerful adversary cannot produce a query $z + z'$ such that $\|s^{-1}(z' + h(r)c')\| \leq 2\tau'$, except with probability $\beta(2\tau')/|\mathbb{F}|$.

Lemma 4.3. *If the challenger sent a valid pair, i.e. $(r, y = f^x(r))$, and the adversary's response is valid, i.e. $\|f(r) - s^{-1}(z + z' - h(r)c)\| \leq \tau'$, then $\|s^{-1}(z' + h(r)c')\| \leq 2\tau'$.*

Proof. Because we set z to $f^x(r)s + h(r)(c + c')$, we know that $\|f(r) - s^{-1}(z - h(r)(c + c'))\| \leq \tau'$. Since the adversary's response is valid, we also have that $\|f(r) - s^{-1}(z + z' - h(r)c)\| \leq \tau'$. Thus subtracting the first value from the second and using the triangular inequality, we obtain $\|s^{-1}(z' + h(r)c')\| \leq 2\tau'$. \square

Lemma 4.4. *If the ordered pair (r, y) sent by the challenger is uniformly random in $\mathbb{D} \times \mathbb{F}$, then the probability that even an all-powerful adversary can output $(c', z') \neq (0, 0)$ such that $\|s^{-1}(z' + h(r)c')\| \leq 2\tau'$ is at most $|\beta(2\tau')|/|\mathbb{F}|$.*

Proof. We first handle the case where f is a weak-PRF without any noise (i.e. the support of the distribution χ is 0 and $\tau' = 0$). In this case, the extra random key s is not necessary in the protocol (i.e. $s = 1$) and so the condition $\|s^{-1}(z' + h(r)c')\| \leq 2\tau'$ becomes $0 = z' + h(r)c'$. Since y is uniformly random in \mathbb{F} and independent of everything else, the value z that the adversary receives is also uniformly random and independent of the pairwise independent hash function h . Thus the adversary will behave in the same way if the function h were chosen *after* the adversary chooses c' and z' . Notice that the adversary must set $c' \neq 0$ because otherwise z' is also necessarily 0. Thus,

$$\forall r \in \mathbb{D}, z' \in \mathbb{F}, c' \in \mathbb{F} \setminus \{0\}, \Pr_h[0 = z' + h(r)c'] = \Pr_h[h(r) = -z'c'^{-1}] = 1/|\mathbb{F}|.$$

We now move to the case where the support of χ is not restricted to 0, and thus the secret key s is chosen uniformly at random from \mathbb{F} . Again, since y is uniformly random in \mathbb{F} , the distribution of $z = ys + h(r)(c + c')$ is uniform and independent of s and h . Thus the adversary would behave the same way if the values of s and h are chosen *after* the adversary chooses c' and z' . By definition, there are $|\beta(2\tau')|$ elements $t \in \mathbb{F}$ such that $\|t\| \leq 2\tau'$. If $t \neq 0$, we have

$$\forall r \in \mathbb{D}, z', c' \in \mathbb{F}, t \in \beta(2\tau') \setminus \{0\}, \Pr_s[s^{-1}(z' + h(r)c') = t] = \Pr_s[s = t^{-1}(z' + h(r)c')] = 1/|\mathbb{F}|.$$

If $t = 0$, then we necessarily have $c' \neq 0$ (otherwise $z' = 0$ as above) and so

$$\forall r \in \mathbb{D}, z' \in \mathbb{F}, c' \in \mathbb{F} \setminus \{0\}, \Pr_h[s^{-1}(z' + h(r)c') = 0] = \Pr_h[h(r) = -z'c'^{-1}] = 1/|\mathbb{F}|.$$

Therefore, we have that

$$\Pr_{s,h}[\|s^{-1}(z' + h(r)c')\| \leq 2\tau'] = \Pr_{s,h}[s^{-1}(z' + h(r)c') \in \beta(2\tau')] = |\beta(2\tau')|/|\mathbb{F}|.$$

\square

4.2 Answering the Challenger when $r' \neq 0$

In this subsection, we deal with the case that when the Adversary makes the i 'th verification query, he changes the commitment r to $r + r'$ where $r' \neq 0$. As in the previous subsection where $r' = 0$, in order to detect that the Adversary indeed answered the query correctly, we need to obtain the value of $f(r + r')$ twice. In the previous section, we had this because r' was 0. If $r' \neq 0$, then we can obtain this value twice by rewinding the Adversary and thus having him provide two winning queries that use the same commitment $r + r'$ (see Figure 5). Thus we will have two equations




Figure 5: **Answering the Challenger after the Adversary's i^* th verification query if $r' \neq 0$.**

containing $f(r + r')$ that we will be able to subtract from each other so that the result has no dependence on f .

When the Adversary sends the commitment $r + r'$, we first send the challenge c_0 , which the Adversary possibly modifies and forwards to the Tag as $c_0 + c'_0$. The simulator playing as the tag sends $z_0 = ys + h(r)(c_0 + c'_0)$ to the Adversary, who then possibly modifies it and sends $z_0 + z'_0$ to the Reader. We then rewind the adversary to the point after he sent $r + r'$ and send him a random challenge c_1 , which he forwards to the Tag as $c_1 + c'_1$, receives the response $z_1 = ys + h(r)(c_1 + c'_1)$ and forwards it to the Reader as $z_1 + z'_1$. If both of the Adversary's responses are correct (which happens with probability $(\epsilon/q_v - 1/|\mathbb{F}|)^2$ by Lemma B.1 because all the randomness except the challenge is exactly the same in both runs), then, as shown in Lemma 4.5, it will be always true that

$$\|s^{-1}((z_0 + z'_0) - (z_1 + z'_1) - h(r + r')(c_0 - c_1))\| \leq 2\tau'.$$

On the other hand, if the Challenger's outputs were uniformly random in (\mathbb{D}, \mathbb{F}) , then as we previously noted, until the i^* th verification query, the secret keys s, h are information-theoretically hidden from the adversary. Because of the rewinding in the i^* th step, some information about h does get revealed - in particular, because we send both $z_0 = ys + h(r)(c_0 + c'_0)$ and $z_1 = ys + h(r)(c_1 + c'_1)$, we are in fact committing to the value $h(r)$. But because h is pairwise-independent, nothing is known about $h(r + r')$ for a non-zero r' , and this is used in Lemma 4.6 to conclude that even an all-powerful adversary cannot fool the simulator except with probability $\beta(2\tau')/|\mathbb{F}|$.

Lemma 4.5. *If both of the adversary's responses are valid, i.e. $\|f(r + r') - s^{-1}(z_0 + z'_0 - h(r + r')c_0)\| \leq \tau'$ and $\|f(r + r') - s^{-1}(z_1 + z'_1 - h(r + r')c_1)\| \leq \tau'$, then $\|s^{-1}((z_0 + z'_0) - (z_1 + z'_1) - h(r + r')(c_0 - c_1))\| \leq 2\tau'$.*

Proof. The claim is obtained by subtracting the first two equations from each other and using the triangle inequality. \square

Lemma 4.6. *If the ordered pair (r, y) sent by the challenger is uniformly random in $\mathbb{D} \times \mathbb{F}$, and $r' \neq 0$, and $c_0 \neq c_1$, then the probability that even an all-powerful adversary can output z'_0 and z'_1 such that $\|s^{-1}((z_0 + z'_0) - (z_1 + z'_1) - h(r + r')(c_0 - c_1))\| \leq 2\tau'$ is at most $|\beta(2\tau')|/|\mathbb{F}|$.*

Proof. For simplicity, we will define $w = (z_0 + z'_0) - (z_1 + z'_1)$. The information given to the adversary (in the two rewindings) by the simulator playing as the tag is $z_0 = ys + h(r)(c_0 + c'_0)$ and $z_1 = ys + h(r)(c_1 + c'_1)$. This is exactly the same as receiving z_0 and $\tilde{z} = z_0 - z_1 = h(r)(c_0 + c'_0 - (c_1 + c'_1))$. Notice that since z_0 contains the term ys , the value of z_0 is uniform and independent of the function h . The value of \tilde{z} , on the other hand, does depend on $h(r)$. So the behavior of the adversary would be unchanged if we chose z_0 uniformly at random, chose a random element u for $h(r)$ and set $\tilde{z} = h(r)(c_0 + c'_0 - (c_1 + c'_1))$, and then *after* the adversary picks z'_0, z'_1 , we finally choose h (conditioned on the already set value of $h(r)$). Thus we have that $\forall t \in \beta(2\tau'), c_0 \neq c_1 \in \mathbb{F}, r \in \mathbb{D}, r' \neq 0, w, s, u \in \mathbb{F}$

$$\Pr_h[s^{-1}(w - h(r + r')(c_0 - c_1)) = t \mid h(r) = u] = \Pr_h[h(r + r') = (w - st)(c_0 - c_1)^{-1} \mid h(r) = u] = 1/|\mathbb{F}|,$$

where $(c_0 - c_1)^{-1}$ exists since we assumed $c_0 \neq c_1$ and the last equality is true because h is a pairwise-independent function and $r' \neq 0$. \square

Acknowledgements

We are very grateful to Eike Kiltz and Krzysztof Pietrzak for numerous discussions pertaining to their work on the LPN problem.

References

- [BFKL93] Avrim Blum, Merrick L. Furst, Michael J. Kearns, and Richard J. Lipton. Cryptographic primitives based on hard learning problems. In *CRYPTO*, pages 278–291, 1993.
- [BKL⁺07] Andrey Bogdanov, Lars R. Knudsen, Gregor Leander, Christof Paar, Axel Poschmann, Matthew J. B. Robshaw, Yannick Seurin, and C. Vikkelsoe. Present: An ultralightweight block cipher. In *CHES*, pages 450–466, 2007.
- [BP02] Mihir Bellare and Adriana Palacio. Gq and schnorr identification schemes: Proofs of security against impersonation under active and concurrent attacks. In *CRYPTO*, pages 162–177, 2002.
- [BPR12] Abhishek Banerjee, Chris Peikert, and Alon Rosen. Pseudorandom functions and lattices. In *EUROCRYPT*, pages 719–737, 2012.
- [DKPW12] Yevgeniy Dodis, Eike Kiltz, Krzysztof Pietrzak, and Daniel Wichs. Message authentication, revisited. In *EUROCRYPT*, pages 355–374, 2012.

- [DN02] Ivan Damgård and Jesper Buus Nielsen. Expanding pseudorandom functions; or: From known-plaintext security to chosen-plaintext security. In *CRYPTO*, pages 449–464, 2002.
- [GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, 1986.
- [GQ88] Louis C. Guillou and Jean-Jacques Quisquater. A ”paradoxical” indentity-based signature scheme resulting from zero-knowledge. In *CRYPTO*, pages 216–231, 1988.
- [GRS05] Henri Gilbert, Matt Robshaw, and Herve Sibert. An active attack against hb^+ - a provably secure lightweight authentication protocol. Cryptology ePrint Archive, Report 2005/237, 2005.
- [GRS08a] Henri Gilbert, Matthew J. B. Robshaw, and Yannick Seurin. Good variants of hb^+ are hard to find. In *Financial Cryptography*, pages 156–170, 2008.
- [GRS08b] Henri Gilbert, Matthew J. B. Robshaw, and Yannick Seurin. $Hb^\#$: Increasing the security and efficiency of hb^+ . In *EUROCRYPT*, pages 361–378, 2008.
- [HB01] Nicholas J. Hopper and Manuel Blum. Secure human identification protocols. In *ASIACRYPT*, pages 52–66, 2001.
- [HKL⁺12] Stefan Heyse, Eike Kiltz, Vadim Lyubashevsky, Christof Paar, and Krzysztof Pietrzak. Lapin: An efficient authentication protocol based on ring-lpn. In *FSE*, pages 346–365, 2012.
- [JW05] Ari Juels and Stephen A. Weis. Authenticating pervasive devices with human protocols. In *CRYPTO*, pages 293–308, 2005.
- [KPC⁺11] Eike Kiltz, Krzysztof Pietrzak, David Cash, Abhishek Jain, and Daniele Venturi. Efficient authentication from hard learning problems. In *EUROCRYPT*, pages 7–26, 2011.
- [KS06] Jonathan Katz and Ji Sun Shin. Parallel and concurrent security of the hb and hb^+ protocols. In *EUROCRYPT*, pages 73–87, 2006.
- [KSS10] Jonathan Katz, Ji Sun Shin, and Adam Smith. Parallel and concurrent security of the hb and hb^+ protocols. *J. Cryptology*, 23(3):402–421, 2010.
- [LPPS07] Gregor Leander, Christof Paar, Axel Poschmann, and Kai Schramm. New lightweight des variants. In *FSE*, pages 196–210, 2007.
- [LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In *EUROCRYPT*, pages 1–23, 2010.
- [MS07] Ueli M. Maurer and Johan Sjödin. A fast and key-efficient reduction of chosen-ciphertext to known-plaintext security. In *EUROCRYPT*, pages 498–516, 2007.
- [NR97] Moni Naor and Omer Reingold. Number-theoretic constructions of efficient pseudorandom functions. In *FOCS*, pages 458–467, 1997.

- [NRR02] Moni Naor, Omer Reingold, and Alon Rosen. Pseudorandom functions and factoring. *SIAM J. Comput.*, 31(5):1383–1404, 2002.
- [OOV08] Khaled Ouafi, Raphael Overbeck, and Serge Vaudenay. On the security of hb# against a man-in-the-middle attack. In *ASIACRYPT*, pages 108–124, 2008.
- [Pie09] Krzysztof Pietrzak. A leakage-resilient mode of operation. In *EUROCRYPT*, pages 462–482, 2009.
- [Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009.
- [Sch91] Claus-Peter Schnorr. Efficient signature generation by smart cards. *J. Cryptology*, 4(3):161–174, 1991.

A A Concurrent Attack Against Randomized Weak-PRFs

The security proof of Lemma 4.6 gives a hint as to why our scheme may not be secure in the concurrent model against man-in-the-middle attacks. When rewinding the adversary, we also in a sense rewind the Tag, and as a consequence end up committing to one value of the pairwise-independent hash function h . In the proof, it was crucial that h was still free to be set to any value for any other input. If the adversary were to have access to multiple copies of the Tag, he could query all of them after the rewinding, and we would end up committing to multiple positions in h , which would in essence define the whole function.

We will show a possible attack in the concurrent model for function families which satisfy a certain homomorphic property, like LPN. Suppose that $\mathcal{F} : \mathbb{D} \rightarrow \mathbb{F}$ is a family of randomized weak-PRF's that is homomorphic under addition. In other words, for $f \in \mathcal{F}$ we have $f(r_1 + r_2) = f(r_1) + f(r_2)$. Also, without loss of generality, suppose that a part of the secret key is a pairwise independent hash function h defined as $h(r) = s_1 r + s_2$. Then the Adversary's man-in-the-middle attack proceeds as follows (see Figure 6):

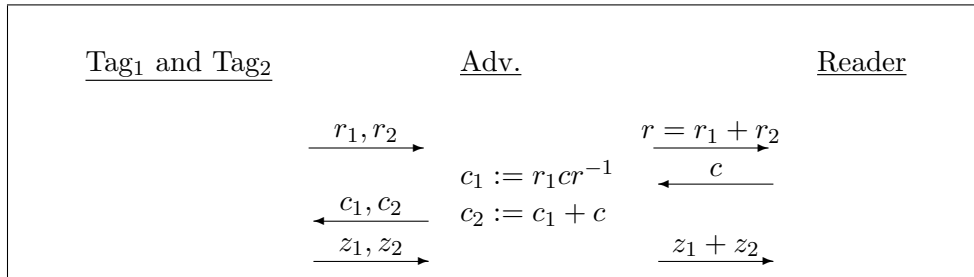


Figure 6: **Concurrent Attack Against Randomized Weak-PRFs.**

The adversary instantiates two copies of the Tag whose secret keys are (f, s, s_1, s_2) and receives two commitments r_1, r_2 from it. It sends the sum of their commitments $r = r_1 + r_2$ to the Reader. The reader sends the challenge c , and the adversary forwards $c_1 = r_1 c r^{-1}$ to the first copy of the Tag, and $c_2 = c_1 + c$ to the second copy. He receives responses $z_1 = (f(r_1) + e_1)s + h(r_1)c_1$ and $z_2 = (f(r_2) + e_2)s + h(r_2)c_2$ from the Tags. We will now show that if the error $e_1 + e_2$ is small,

then $z_1 + z_2$ could be a valid response. This could be problematic for the protocol's security since it will reveal something about the size of the errors in z_1 and z_2 . In fact, similar information was used in the attack of [OOV08] against the scheme in [GRS08b] to recover the secret key. We now show that if $\|e_1 + e_2\| \leq \tau'$, then the Reader will accept the adversary's response of $z_1 + z_2$. The reader will accept iff $\|f(r) - s^{-1}((z_1 + z_2) - (s_1r + s_2)c)\| \leq \tau'$, where

$$\begin{aligned}
& \|f(r) - s^{-1}((z_1 + z_2) - (s_1r + s_2)c)\| \\
&= \|f(r) - (f(r_1) + f(r_2)) - (e_1 + e_2) - s^{-1}(s_2(c_1 + c_2) + s_1(r_1c_1 + r_2c_2) - (s_1r + s_2)c)\| \\
&= \|- (e_1 + e_2) - s^{-1}(s_1(r_1c_1 + (r - r_1)(c_1 + c)) - s_1rc)\| \\
&= \|- (e_1 + e_2) - s^{-1}s_1(rc_1 + rc - r_1c - rc)\| = \|- (e_1 + e_2) - s^{-1}s_1(rr_1cr^{-1} - r_1c)\| \\
&= \|- (e_1 + e_2)\|.
\end{aligned}$$

B (Slightly) Extended Reset Lemma

For the rewinding step of the proof, it is necessary to show that we can use an adversary to extract two different valid responses. This issue was already elegantly addressed by Bellare and Palacio in their Reset Lemma [BP02, Lemma 3.1] whenever the only difference in the randomness given to an adversary during the rewinding is the challenge itself. We have to slightly extend their lemma to handle the MIM case. Our proof very closely follows theirs.

First we describe all of the used algorithms as deterministic algorithms, receiving their randomness as auxiliary inputs or parameters. Now we're able to guarantee that all the randomness beside the challenge is for every algorithm in both rewinding cases exactly the same by fixating their inputs.

Further, we want to modularize the behavior of the Adversary \mathcal{A} during the protocol allowing us to analyze it more easily. A very common way to do this is to explicitly use the state St of \mathcal{A} . \mathcal{A} has auxiliary input a , Tag \mathcal{T} t and Reader \mathcal{R} r .

Lemma B.1 (extended Reset Lemma). *Given the random tape R , the deterministic algorithms $Com_{\mathcal{T}}(t)$, $Rsp_{\mathcal{T}}(t, CMT, CH)$, $Dec_{\mathcal{R}}(r, CMT, CH, RSP)$ and challenge set $ChSet_{\mathcal{R}}$ representing Tag $\mathcal{T}(t)$ and Reader $\mathcal{R}(r)$. Then $acc(a, t, r)$ is the probability for $d = 1$ for an algorithm $\mathcal{A}(a, R)$ in the following experiment:*

$$\begin{aligned}
& St \leftarrow \mathcal{A}(a, R); \quad CMT \leftarrow Com_{\mathcal{T}}(t); \quad (CMT', St) \leftarrow \mathcal{A}(CMT, St); \\
& CH \stackrel{\$}{\leftarrow} ChSet_{\mathcal{R}}; \quad (CH', St) \leftarrow \mathcal{A}(CH, St); \quad RSP \leftarrow Rsp_{\mathcal{T}}(t, CMT, CH'); \quad RSP' \leftarrow \mathcal{A}(RSP, St); \\
& d \leftarrow Dec_{\mathcal{R}}(r, CMT', CH, RSP')
\end{aligned}$$

$res(a, t, r)$ is the probability d being $d = 1$ in this experiment:

$$\begin{aligned}
& St \leftarrow \mathcal{A}(a, R); \quad CMT \leftarrow Com_{\mathcal{T}}(t); \quad (CMT', St) \leftarrow \mathcal{A}(CMT, St); \\
& CH_1 \stackrel{\$}{\leftarrow} ChSet_{\mathcal{R}}; \quad (CH'_1, St_1) \leftarrow \mathcal{A}(CH_1, St); \quad RSP_1 \leftarrow Rsp_{\mathcal{T}}(t, CMT, CH'_1); \quad RSP'_1 \leftarrow \mathcal{A}(RSP_1, St_1); \\
& CH_2 \stackrel{\$}{\leftarrow} ChSet_{\mathcal{R}}; \quad (CH'_2, St_2) \leftarrow \mathcal{A}(CH_2, St); \quad RSP_2 \leftarrow Rsp_{\mathcal{T}}(t, CMT, CH'_2); \quad RSP'_2 \leftarrow \mathcal{A}(RSP_2, St_2); \\
& d_1 \leftarrow Dec_{\mathcal{R}}(r, CMT', CH_1, RSP'_1); \quad d_2 \leftarrow Dec_{\mathcal{R}}(r, CMT', CH_2, RSP'_2); \quad d = (d_1 \wedge d_2 \wedge (CH_1 \neq CH_2))
\end{aligned}$$

Then the following statement holds:

$$acc(a, t, r) \leq \frac{1}{|ChSet_{\mathcal{R}}|} + \sqrt{res(a, t, r)}$$

Proof. Under normal circumstances Adversary \mathcal{A} will have advantage $acc(a, t, r)$ for a random challenge CH_1 . Now we're able to run him twice to receive another valid response for a challenge CH_2 . Since CH_1 was chosen independently uniform exactly as CH_2 , the advantage of \mathcal{A} is the same in both cases. We want to show now, that the probability to get two different, valid results, as it is necessary for the rewinding, is lower bounded by $acc(a, t, r)$.

The output behavior of $Dec_{\mathcal{R}}$ depends on the random tape R , the challenge CH and the behavior of \mathcal{T} and \mathcal{R} which are determinate by t and r . $CMT'(a, R, t)$ and $RSP'(a, R, CH, t)$ are the determinate outputs of \mathcal{A} depending only on a, t, R and CH .

Now we define the functions: $X : \{0, 1\}^l \rightarrow \{0, 1\}$ and $Y : \{0, 1\}^l \rightarrow \{0, 1\}$ for $R \in \{0, 1\}^l$:

$$\begin{aligned} X(R) &:= \Pr[Dec_{\mathcal{R}}(r, CMT'(a, R, t), CH, RSP'(a, R, CH, t)) = 1] \\ Y(R) &:= \Pr \left[\begin{array}{l} Dec_{\mathcal{R}}(r, CMT'(a, R, t), CH_1, RSP'(a, R, CH_1, t)) = 1 \wedge \\ Dec_{\mathcal{R}}(r, CMT'(a, R, t), CH_2, RSP'(a, R, CH_2, t)) = 1 \wedge \\ CH_1 \neq CH_2 \end{array} \right] \end{aligned}$$

and for every $R \in \{0, 1\}^r$

$$Y(R) = X(R) \cdot X(R) \cdot \left(1 - \frac{1}{|ChSet_{\mathcal{R}}|}\right) \geq X(R) \cdot \left(X(R) - \frac{1}{|ChSet_{\mathcal{R}}|}\right),$$

since r, a, R and t are fixed and only the challenge is different. The last step of the proof is exactly the same as in the proof of the Reset Lemma. We define $p := 1/|ChSet_{\mathcal{R}}|$ and relate the expectation values of $X(R)$ and $Y(R)$ for an uniform random $R \xleftarrow{\$} \{0, 1\}^l$:

$$\begin{aligned} res(a, t, r) &= E(Y) \\ &\geq E(X \cdot (X - p)) \\ &= E(X^2) - p \cdot E(X) \\ &\geq E(X)^2 - p \cdot E(X) = acc(a, t, r)^2 - p \cdot acc(a, t, r) \end{aligned}$$

We're completing the squares:

$$\begin{aligned} acc(a, t, r)^2 - p \cdot acc(a, t, r) &\leq res(a, t, r) \\ \Leftrightarrow \left(acc(a, t, r) - \frac{p}{2} \right)^2 &\leq res(a, t, r) + \frac{p^2}{4} \\ \Leftrightarrow acc(a, t, r) &\leq \sqrt{res(a, t, r) + \frac{p^2}{4}} + \frac{p}{2} \leq \sqrt{res(a, t, r)} + p \end{aligned}$$

This holds, since $\sqrt{res(a, t, r) + \frac{p^2}{4}} \leq \sqrt{res(a, t, r)} + \sqrt{\frac{p^2}{4}}$. □