Faster and Shorter Password-Authenticated Key Exchange

Rosario Gennaro

IBM T.J. Watson Research Center
19 Skyline Drive,
Hawthorne 10532, USA
rosario@us.ibm.com

Abstract. This paper presents an improved password-based authenticated key exchange protocol in the common reference string model. Its security proof requires no idealized assumption (such as random oracles).

The protocol is based on the **GL** framework introduced by Gennaro and Lindell, which generalizes the **KOY** key exchange protocol of Katz et al. Both the **KOY** and the **GL** protocols use (one-time) signatures as a non-malleability tool in order to prevent a man-in-the-middle attack against the protocol. The efficiency of the resulting protocol is negatively affected, since if we use regular signatures, they require a large amount of computation (almost as much as the rest of the protocol) and further computational assumptions. If one-time signatures are used, they substantially increase the bandwidth requirement.

Our improvement avoids using digital signatures altogether, replacing them with faster and shorter message authentication codes. The crucial idea is to leverage as much as possible the non-malleability of the encryption scheme used in the protocol, by including various values into the ciphertexts as *labels*. As in the case of the **GL** framework, our protocol can be efficiently instantiated using either the DDH, Quadratic Residuosity or N -Residuosity Assumptions.

For typical security parameters our solution saves as much as 12 Kbytes of bandwidth if one-time signatures are implemented in **GL** with fast symmetric primitives. If we use number-theoretic signatures in the **GL** framework, our solution saves several large exponentiations (almost a third of the exponentiations computed in the **GL** protocol). The end result is that we bring provable security in the realm of password-authenticated key exchange one step closer to practical.

1 Introduction

The central problem of cryptography is to enable reliable and secure communication among parties in the presence of an adversary. In order to do this, parties must share a common secret key to secure communication using known techniques (e.g., applying encryption and message authentication codes to all messages).

A protocol that allows two parties to establish such a secret key is called a *key exchange* protocol. The key exchange problem was initially studied by Diffie

and Hellman [15] who considered a passive adversary that can eavesdrop on the honest parties' communication, but cannot actively modify it. In other words, parties are assumed to be connected by reliable, albeit non-private, channels. Many efficient and secure protocols are known for this scenario. The more realistic scenario however is that of a far more powerful adversary who can modify and delete messages sent between the parties, as well as insert messages of its own choice. This is the scenario we consider in this paper.

Once we allow such a powerful adversary it becomes clear that in order to securely exchange a key, any two parties (call them Alice and Bob) must hold some secret information. Otherwise, there is nothing preventing an adversary from pretending to be Bob while communicating with Alice (and vice versa). The most common type of secret information considered are (i) parties already share a high entropy secret key; (ii) each party holds a secret key matching an authenticated public key (i.e. a public key securely associated with his identity) and (iii) the case we consider in this paper: parties share only a low entropy *password* that can be remembered and typed in by human users.

Cryptography has long been concerned with cases (i) and (ii), while the scenario of low entropy passwords (arguably the most commonly used case) has only recently received attention. This paper proposes a new and improved family of protocols (a framework) for password-based key exchange in the face of a powerful, active adversary.

Password-based authenticated key-exchange. Our model consists of a group of parties, with each pair of them sharing a password chosen uniformly at random from some small dictionary (the assumption of uniformity is made for simplicity only). The parties communicate over a network in the presence of an active adversary who has full control over the communication lines. In other words all communication between parties is basically carried out through the adversary. Nevertheless, the goal of the parties is to generate session keys in order to secretly and reliably communicate with each other.

An immediate observation is that the small size of the dictionary implies a non-negligible probability that the attacker will succeed in impersonating one of the parties, since the adversary can always guess Bob's password and pretend to be him while communicating with Alice. This type of attack is called an **on-line guessing attack** and is inherent whenever security depends on low entropy passwords. The severity of on-line guessing attacks can be limited with other mechanisms (such as locking an account after a number of failed attempts). A more dangerous attack is the *off-line guessing attack*, in which the adversary obtains a transcript of an execution of the key exchange protocol and is then able to check guesses for Bob's password against this transcript *off-line*. The aim of password-based authenticated key exchange is to limit the adversary only to on-line guessing attacks, and rule out possible off-line ones.

Prior related work. Bellare and Merritt [3] proposed the first protocol for password-based session-key generation. Although the specific protocol of [3] can be attacked (see [28]), small modifications to the protocol prevent these attacks [28], and even allow one to prove it secure under the ideal cipher and random oracle

models [1]. Bellare and Merritt’s work was very influential and was followed by many protocols (e.g. [4,29,23,27,28,30]) which, however, have not been proven secure and their conjectured security is based on heuristic arguments.

A first rigorous treatment of the problem was provided by Halevi and Krawczyk [22]. They consider an *asymmetric* model in which some parties (called servers) hold certified public keys available to the all parties, including the clients who instead hold only passwords. In this model (which requires a public-key infrastructure) Halevi and Krawczyk provide a secure password-based key exchange. The first (and only currently known) protocol to achieve security without *any* additional setup is that of Goldreich and Lindell [20]. Their protocol is based on general assumptions (i.e., the existence of trapdoor permutations) and constitutes a proof that password-based authenticated key exchange can actually be obtained. Unfortunately, the protocol of [20] is not very efficient and thus cannot be used in practice.

Katz, Ostrovsky and Yung (KOY) [24] present an efficient and practical protocol for the problem of password-authenticated key-exchange in the common reference string model. In this model, the extra setup assumption is that all parties have access to some public parameters, chosen by some trusted third party. This assumption is clearly weaker than assuming a public-key infrastructure, and there are settings in which it can be implemented safely and efficiently (such as a corporation wanting to provide secure password login for its employees, and thus can be trusted to choose and distribute the common reference string). The KOY protocol is based on the security against chosen-ciphertext attack [16] of the original Cramer-Shoup encryption scheme [11]. This in turn can be reduced to the Decisional Diffie-Hellman (DDH) assumption. The complexity of the KOY protocol is only 5–8 times the complexity of a Diffie-Hellman unauthenticated key-exchange protocol.

The KOY protocol was generalized by Gennaro and Lindell [19], using generic building blocks instead of specific number-theoretic assumptions. More specifically, they use the notion of *projective hash functions* and the CCA-secure encryption schemes defined in [12]. The resulting protocol GL has a much more intuitive proof of security and can be proven secure under a variety of computational assumptions (such as Quadratic Residuosity and N -Residuosity).

We note that there are password-authenticated key-exchange protocols which are more efficient than KOY and GL, but whose proof holds in an idealized model of computation such as the ideal cipher and random oracle models [1,6]. The common interpretation of such results is that security is *likely* to hold even if the random oracle is replaced by a (“reasonable”) concrete function known explicitly to all parties (e.g., SHA-1). However, it has been shown that it is impossible to replace the random oracle in a generic manner with any concrete function [7]. Thus, the proofs of security of these protocols are actually heuristic in nature.

1.1 Our Contributions

We improve on the both the GL and KOY protocols, in particular by reducing the communication bandwidth required by the protocol.

Both the KOY and the GL protocol use (one-time) signatures as a non-malleability tool in order to prevent a man-in-the-middle attack against the protocol. This negatively affects the efficiency of the resulting protocol. Indeed in order to preserve provable security without use of the random oracle an implementation of the KOY or GL protocol is presented with two choices.

One-time signature schemes (i.e. signature schemes which are secure if the key is used to sign only one message) can be implemented from fast symmetric key primitives (such as one-way functions). However the length of the resulting keys and signatures is problematic and causes a substantial increase in the required bandwidth.

One could use “regular” signature schemes (i.e. secure for many messages) but then, if we require provable security in the standard model, the amount of computation would substantially increase. Moreover if we want to use the most efficient provably secure signature schemes in the literature (e.g. [18,14,25]) we would introduce new computational assumptions such as the Strong RSA assumption, on top of the ones required by the GL protocol.

Our improvement avoids using digital signatures altogether, replacing them with faster and shorter message authentication codes. The crucial idea is to leverage as much as possible the non-malleability of the encryption scheme used in the protocol, by including various values into the encryption as *labels*. For typical security parameters our improvement saves as much as 12 Kbytes of bandwidth in a protocol execution.

As in the case of the GL framework, our protocol can be efficiently instantiated using either the DDH, Quadratic Residuosity or N -Residuosity Assumption.

1.2 Our Construction in a Nutshell

Let us describe our construction informally. We start by first describing the tools that we are going to use and then describing the protocol.

Chosen-Ciphertext Secure Public-Key Encryption [16]: We use an encryption scheme $\mathcal{E}$ which is secure against chosen-ciphertext attack. The common reference string for our password protocol is simply the public key PK for such an encryption scheme. We stress that the corresponding secret key does not have to be known by any party¹.

Smooth projective hashing [12]: Let X be a set and $L \subset X$ a language. Loosely speaking, a hash function H_k that maps X to some set is **projective** if there exists a projection key that defines the action of H_k over the subset L of the domain X . That is, there exists a projection function $\alpha(\cdot)$ that maps

¹ In the GL protocol the requirement is actually weaker, as all is needed is a non-interactive non-malleable (with respect to many commitments) commitment which in the common reference string can be built out of CCA-Secure Encryption. For simplicity we describe our protocol using encryption, but in the final version we show that we can also use commitments. In practice this does not make much difference since CCA encryption is the most efficient known implementation of for this type of non-malleable commitments.

keys k into their projections $s = \alpha(k)$. The projection key s is such that for every $x \in L$ it holds that the value of $H_k(x)$ is uniquely determined by s and x . In contrast, nothing is guaranteed for $x \notin L$, and it may not be possible to compute $H_k(x)$ from s and x . A **smooth** projective hash function has the additional property that for $x \notin L$, the projection key s actually says *nothing* about the value of $H_k(x)$. More specifically, given x and $s = \alpha(k)$, the value $H_k(x)$ is uniformly distributed (or statistically close) to a random element in the range of H_k .

What makes smooth projective hashing a powerful tool (in both our application and the original one in [12]) is that if L is an NP-language, then for every $x \in L$ it is possible to efficiently compute $H_k(x)$ using the projection key $s = \alpha(k)$ and a witness of the fact that $x \in L$. Alternatively, given k itself, it is possible to efficiently compute $H_k(x)$ even without knowing a witness. Gennaro and Lindell [19] also prove another important property of smooth projective hash functions that holds when L is a hard-on-the-average NP-language. For a random $x \in_R L$, given x and $s = \alpha(k)$ the value $H_k(x)$ is *computationally indistinguishable* from a random value in the range of $H_k(x)$. Thus, even if $x \in L$, the value $H_k(x)$ is pseudorandom, unless a witness is known.

The basic idea behind the KOY and GL protocols is to have the parties exchange non-malleable encryptions of the joint password. The session key is then computed as the result of applying smooth projective hash functions to these encryptions (in this case the hard-on-the-average NP language consists of correct ciphertext/message pairs). Figure 1 shows the basic layout of the protocol.

The basic problem with the protocol described in Figure 1 is that the projective hash function themselves can be malleable, and an adversary could manage to get information about the session key by playing man-in-the-middle. In order to avoid this attack, the GL and KOY protocols add a signature step. A verification key is chosen by party A in the first message and *bound* together with the first encryption, by including it as a *label*². Then A signs the whole transcript in the third message. Party B accepts only if the signature is correct. Since the verification key cannot be changed (being protected by the non-malleability of the encryption in the first step), the adversary cannot modify the projection keys, unless it is able to produce a forgery.

In our protocol we expand the use of encryption labels. We protect the first projection s , by including it as a label in the second ciphertext³ c' . Now the adversary is left with the possibility of manipulating the second projection key s' . But at this point the master key computed by party A is already pseudorandom

² A *label* is a public string that accompanies a ciphertext and is an integral part of it. It must be submitted together with the ciphertext in order to obtain a decryption and the adversary should not be able to modify it. See Section 2.1 for details.

³ Interestingly this is already done in the KOY protocol, but it is not used in the proof in any significant way. Indeed the GL proof shows that the use of digital signatures make this step unnecessary. We reinstate it exactly because we want to avoid using signatures.

Insecure Password-Authenticated Key Exchange

- **Common reference string:** The public key PK for a chosen-ciphertext secure encryption scheme $\mathcal{E}$. A description of a smooth projective hashing family H_k over the set X of ciphertext/password pairs (c, w) .
- **Common input:** a shared (low-entropy) password w .
- **The protocol:**
 1. Party A computes an encryption $c = \mathcal{E}_{PK}(w)$ and sends it to party B .
 2. Party B chooses a key k for the smooth projective hash function, and computes its projection $s = \alpha(k)$. Also B computes the projective hash over (c, w) , i.e. $sk_B = H_k(c, w)$.
Finally B computes another encryption of the password i.e., $c' = \mathcal{E}_{PK}(w)$.
 B sends s, c' to party A .
 3. Party A chooses another key k' for the smooth projective hash function, and computes its projection $s' = \alpha(k')$. Also A computes the projective hash over (c', w) , i.e. $sk_A = H_{k'}(c', w)$.
 A sends s', t to party B .
- **Session Key Definition:**
 1. Party B computes sk_A using the projection s' and its knowledge of a witness for the fact that c' is an encryption to the password w (it knows a witness because it generated c') and outputs the session key $sk = sk_A \oplus sk_B$.
 2. Party A also computes sk_B using the projection s and its knowledge of a witness for the fact that c is an encryption to the string w (it knows a witness because it generated c) and outputs $sk = sk_A \oplus sk_B$.

Fig. 1. Common skeleton of KOY, GL and our protocol

for the adversary and thus it can be used as a key to MAC the projection key s' in order to prevent $\mathcal{A}$ from changing it.

A technical issue arises here, as party B has to use the same key that party A uses to compute the MAC, but party B has to yet finish the protocol and compute such key. Moreover the adversary can make B compute a different key from A , by modifying the projection s' . This issue can be solved by using sk_B as a MAC key since B already knows it. However the explicit use of only one component of the session key would allow an off-line attack from $\mathcal{A}$ (see Section 4).

The final solution is to MAC the transcript with sk_B and then use sk_A to “mask” the value of the MAC from the adversary. In the proof if B accepts after $\mathcal{A}$ modified s' he will be able to retrieve a forgery on the MAC keyed with sk_B . The protocol in full details is shown in Figure 2.

1.3 Efficiency Gains

Using symmetric primitives. In terms of computation, the most efficient implementation of the KOY or GL protocols uses one-time signatures based on symmetric primitives, such as one-way functions. One example of such a signature is the Lamport signature [26]: to sign a single bit b the public key consists of two values y_0, y_1 and the secret key is x_0, x_1 where $y_i = F(x_i)$ for a one-way function F . To sign bit b the signer reveals x_b .

Assuming a security parameter of 128 (e.g. a one-way function applied to 128 bits input, and messages hashed to 256 bits using a collision resistant function), we have that transmitting the key and the signature requires about 12 KBytes. Other solutions exist that create shorter signatures at the expense of an increase in computation time (see a survey of possible one-time signatures in e.g. [10]). In contrast our solution requires only 256 bits for the MAC.

Number-Theoretic Signatures. Of course one could implement the signature step in the KOY or GL protocol using provably secure signature schemes such as Gennaro-Halevi-Rabin [18] or Cramer-Shoup [14] which are based on the Strong RSA Assumption: they not only introduce another computational assumption for the security of the scheme, but require several modular exponentiations and about 4 Kbit of bandwidth to transmit keys and signatures. A shorter alternative would be the Boneh-Boyen [5] which requires only 160 bit for the signature, but it would still require 2 Kbits to send the verification key. Moreover signature verification in the Boneh-Boyen scheme is particularly expensive since it requires the computation of a bilinear map.

The above signatures are secure against many messages. There are more efficient number-theoretic one-time signatures such as the one obtained by a chain of length two in the GMR scheme [21], or the one recently proposed in [10] based on chameleon hashing with two trapdoors. Still because of the computation of modular exponentiations and the transmission of verification key and signature, these options are much more expensive than sending a simple MAC. It is not hard to see that for each one of these options the reduction in the number of exponentiations is at least a third.

1.4 Organization

We first recall the cryptographic tools that we need in Section 2: chosen-ciphertext secure public-key encryption, and message authentication codes. In Section 3 we review the notions of smooth projective hash functions (mostly lifted *verbatim*, with permission, from [19]). The protocol is then presented in Section 4 with an intuitive informal proof. Some concluding remarks are presented in Section 5.

For lack of space we refer the reader to [19] for the formal definition of password-authenticated key exchange. Also the formal proof of our protocol can be found in the expanded version of this paper [17].

2 Cryptographic Tools

We denote by n the security parameter.

If S is a set, with $|S|$ we denote its cardinality. $|m|$ denotes the bit length of m , if m is a string or a number.

If $A(\cdot, \cdot, \dots)$ is a probabilistic algorithm, then $x \in_R A(x_1, x_2, \dots)$ denotes the experiment of running A on input $x_1, x_2, \dots$ with x being the outcome. If S is a set, $x \in_R S$ denotes the experiment of choosing $x \in S$ uniformly at random. If X is a probability distribution over S then $x \in_R X$ denotes the experiment of choosing $x \in S$ according to the distribution X .

Finally, we denote statistical closeness of probability ensembles by $\stackrel{s}{\equiv}$, and computational indistinguishability (with respect to non-uniform polynomial-time machines⁴) by $\stackrel{c}{\equiv}$.

We say that a real-valued function $\epsilon(\cdot)$ defined over the integers is *negligible* if for every constant $c \geq 0$ there exists an integer n_c such that for all $n > n_c$ $\epsilon(n) < n^{-c}$.

2.1 Chosen-Ciphertext Secure Public-Key Encryption

A public key encryption scheme is a tuple of three algorithms $\text{PKE} = (\mathcal{K}, \mathcal{E}, \mathcal{D})$. The key generation algorithm $\mathcal{K}$ generates a pair $(PK, SK) \in_R \mathcal{K}(1^n)$, where PK is a public key and SK is a secret key.

We use *labeled* encryption, which means that the encryption algorithm $\mathcal{E}$ takes a public key PK a plaintext m , and a label ℓ and returns a ciphertext $c \in_R \mathcal{E}_{PK}(m, \ell)$. The decryption algorithm $\mathcal{D}$ takes a secret key SK , a ciphertext c and a label ℓ , and returns $\mathcal{D}_{SK}(c, \ell)$ which is either a message m or *reject*. If $c \in_R \mathcal{E}_{PK}(m, \ell)$ then $m = \mathcal{D}_{SK}(c, \ell)$.

The adaptive chosen ciphertext attack (IND-CCA) game is defined as follows. A key pair is generated by the key generation algorithm: $(PK, SK) \in_R \mathcal{K}(1^n)$. Then a PPT adversary $\mathcal{A}$, on input the public key PK , queries a pair of equal length messages m_0 and m_1 and a label ℓ^* to an encryption oracle. The encryption oracle chooses $b \in_R \{0, 1\}$ and computes a challenge ciphertext $c^* \in_R \mathcal{E}_{PK}(m_b, \ell^*)$, which is given to $\mathcal{A}$. In the course of the game the adversary $\mathcal{A}$ is given access to a decryption oracle, $\mathcal{D}_{SK}(\cdot, \cdot)$ which $\mathcal{A}$ can query on any ciphertext/label pair except the challenge ciphertext/label pair c^*, ℓ^* . The game ends with the adversary outputting a bit $\tilde{b}$.

We say that the encryption scheme is secure against (adaptive) chosen-ciphertext attack if for any adversary $\mathcal{A}$, the probability that $b = \tilde{b}$ is negligible (in the security parameter n).

Notice that the adversary is allowed to query the decryption oracle on *any* ciphertext/label pair which is not the target pair. In particular this definition guarantees that the adversary will not get any information from querying a ciphertext with a label different from the one used when the ciphertext was created.

⁴ All of our results also hold with respect to uniform adversaries.

Notice that in our notation the first argument of the encryption algorithm is always the message, the second argument is the label, and the random coins are implicit.

2.2 Message Authentication Codes

A message authentication code MAC is a function

$$MAC : \{0, 1\}^n \times \{0, 1\}^* \longrightarrow \{0, 1\}^n.$$

The first input is the key $k \in \{0, 1\}^n$, and the second input is the message $m \in \{0, 1\}^*$. The output is called a “tag” $t = MAC_k(m)$.

The chosen message attack (CMA) game is defined as follows. A key is selected uniformly at random $k \in_R \{0, 1\}^n$. The adversary $\mathcal{A}$ is given $t^* = MAC_k(m^*)$ for many adaptively adversarially chosen m^* , after which the adversary outputs a pair (m, t) . We say that (m, t) is a *forgery* if $m \neq m^*$ for all the queried m^* and $t = MAC_k(m)$.

We say that a MAC is secure if for every adversary $\mathcal{A}$ the probability of computing a forgery is negligible. We say that a MAC is *1-time secure* if the adversary in the above game is restricted to querying a single message. We note that 1-time secure MACs can be constructed unconditionally.

3 Smooth Projective Hash Functions

Following Gennaro and Lindell [19] we use a modified version of the notion of *smooth projective hashing* introduced by Cramer and Shoup [12]. We recall the definition from [19] which is needed here and refer the reader to [19] for a description of the differences between this definition and the original one from [12].

Subset membership problems. Intuitively, a hard subset membership problem is a problem for which “hard instances” can be efficiently sampled. More formally, a subset membership problem $\mathcal{I}$ specifies a collection $\{I_n\}_{n \in \mathbb{N}}$ such that for every n , I_n is a probability distribution over *problem instance descriptions* Λ . A problem instance description defines a set and a hard language for that set. Formally, each instance description Λ specifies the following:

1. Finite, non-empty sets $X_n, L_n \subseteq \{0, 1\}^{\text{poly}(n)}$ such that $L_n \subset X_n$, and distributions $D(L_n)$ over L_n and $D(X_n \setminus L_n)$ over $X_n \setminus L_n$.
2. A witness set $W_n \subseteq \{0, 1\}^{\text{poly}(n)}$ and an NP-relation $R_n \subseteq X_n \times W_n$. R_n and W_n must have the property that $x \in L_n$ if and only if there exists $w \in W_n$ such that $(x, w) \in R_n$.

We are interested in subset membership problems $\mathcal{I}$ which are efficiently samplable. That is, the following algorithms must exist:

1. *Problem instance samplability*: a probabilistic polynomial-time algorithm that upon input 1^n , samples an instance $\Lambda = (X_n, D(X_n \setminus L_n), L_n, D(L_n), W_n, R_n)$ from I_n .
2. *Instance member samplability*: a probabilistic polynomial-time algorithm that upon input 1^n and an instance $(X_n, D(X_n \setminus L_n), L_n, D(L_n), W_n, R_n)$, samples $x \in L_n$ according to distribution $D(L_n)$, together with a witness w for which $(x, w) \in R_n$.
3. *Instance non-member samplability*: a probabilistic polynomial-time algorithm that upon input 1^n and an instance $(X_n, D(X_n \setminus L_n), L_n, D(L_n), W_n, R_n)$, samples $x \in X_n \setminus L_n$ according to distribution $D(X_n \setminus L_n)$.

We are now ready to define hard subset membership problems:

Definition 1. (hard subset membership problems): Let $V(L_n)$ be the following random variable: Choose a problem instance Λ according to I_n , a value $x \in L_n$ according to $D(L_n)$ (as specified in Λ), and then output (Λ, x) . Similarly, define $V(X_n \setminus L_n)$ as follows: Choose a problem instance Λ according to I_n , a value $x \in X_n \setminus L_n$ according to $D(X_n \setminus L_n)$ (as specified in Λ) and then output (Λ, x) . Then, we say that a subset membership problem $\mathcal{I}$ is **hard** if

$$\{V(L_n)\}_{n \in \mathbb{N}} \stackrel{c}{=} \{V(X_n \setminus L_n)\}_{n \in \mathbb{N}}.$$

In other words, $\mathcal{I}$ is hard if random members of L_n cannot be distinguished from random non-members. In order to simplify notation, from here on we drop the subscript of n from all sets. However, all mention of sets X and L etc., should be understood as having been sampled according to the security parameter n .

Smooth projective hash functions. Loosely speaking a smooth projective hash function is a function with two keys. The first key maps the entire set X to some set G . The second key (called the projection key) is such that it can be used to correctly compute the mapping of L to G . However, it gives no information about the mapping of $X \setminus L$ to G . In fact, given the projection key, the distribution over the mapping of $X \setminus L$ to G is statistically close to uniform (or “smooth”). We now present the formal definition.

Let X and G be finite, non-empty sets and let $\mathcal{H} = \{H_k\}_{k \in K}$ be a collection of hash functions from X to G . We call K the key space of the family. Now, let L be a non-empty, proper subset of X (i.e., L is a language). Then, we define a *key projection* function $\alpha : K \times X \rightarrow S$, where S is the space of key projections. Informally, the above system defines a projective hash system if for $x \in L$, the projection key $s_x = \alpha(k, x)$ uniquely determines $H_k(x)$. (Ignoring issues of efficiency, this means that $H_k(x)$ can be computed given only s_x and $x \in L$.) We stress that the projection key $s_x = \alpha(k, x)$ is only guaranteed to determine $H_k(x)$ for $x \in L$, and nothing is guaranteed for $x' \neq x$. Formally,

Definition 2. (projective hash functions): The family $(\mathcal{H}, K, X, L, G, S, \alpha)$ is a projective hash family if for all $k \in K$ and $x \in L$, it holds that the value of $H_k(x)$ is uniquely determined by $\alpha(k, x)$ and x .

Of course, projective hash functions can always be defined by taking $\alpha(\cdot, \cdot)$ to be the identity function. However, we will be interested in *smooth* projective hash functions which have the property that for every $x \notin L$, the projection key $s_x = \alpha(k, x)$ reveals (almost) nothing about $H_k(x)$. More exactly, for every $x \notin L$, the distribution of $H_k(x)$ given $\alpha(k, x)$ should be statistically close to uniform. Formally,

Definition 3. (smooth projective hash functions [12]): Let $(\mathcal{H}, K, X, L, G, S, \alpha)$ be a projective hash family. Then, for every $x \in X \setminus L$ define the random variable $V(x, \alpha(k), H_k(x))$ by choosing $k \in_R K$ and output $(x, \alpha(k, x), H_k(x))$. Similarly, define $V(x, \alpha(k, x), g)$ as follows: choose $k \in_R K$, $g \in_R G$ and output $(x, \alpha(k, x), g)$. Then, the projective hash family $(\mathcal{H}, K, X, L, G, S, \alpha)$ is smooth if for every $x \in X \setminus L$:

$$\left\{ V(x, \alpha(k), H_k(x)) \right\}_{n \in \mathbb{N}} \stackrel{s}{=} \left\{ V(x, \alpha(k), g) \right\}_{n \in \mathbb{N}}.$$

Efficient smooth projective hash functions. We say that a smooth projective hash family is efficient if the following algorithms exist:

1. *Key sampling*: a probabilistic polynomial-time algorithm that upon input 1^n samples $k \in K$ uniformly at random.
2. *Projection computation*: a deterministic polynomial-time algorithm that upon input 1^n , $k \in K$ and $x \in X$ outputs $s = \alpha(k, x)$.
3. *Efficient hashing from key*: a deterministic polynomial-time algorithm that upon input 1^n , $k \in K$ and $x \in X$, outputs $H_k(x)$.
4. *Efficient hashing from projection key and witness*: a deterministic polynomial-time algorithm that upon input 1^n , $x \in L$ with a witness w such that $(x, w) \in R$, and $\alpha(k, x)$ (for some $k \in K$), computes $H_k(x)$.

We note an interesting and important property of such hash functions. For $x \in L$, it is possible to compute $H_k(x)$ in two ways: either by knowing the key k (as in item 3 above) or by knowing the projection s_x of the key, and a witness for x (as in item 4 above). This property plays a central role in our password-based protocol.

Another interesting property formalized by Gennaro and Lindell in [19] is that these are the only ways to compute $H_k(x)$. Specifically, for $x \in_R D(L)$ (where an appropriate witness w is not known), the value $H_k(x)$ is *computationally* indistinguishable from random, given the projection s_x .

Since we use smooth projective hashing in our password protocol, it is necessary to prove the above statement even when the adversary sees many tuples $(x, s_x, H_k(x))$ with $x \in_R D(L)$. Let M be a (non-uniform) polynomial-time oracle machine. Define the following two experiments.

Expt-Hash(M): An instance $\Lambda = (X, D(X \setminus L), L, D(L), W, R)$ of a hard subset membership problem is chosen from I_n . Then, the machine M is given access to two oracles: Ω_L and $\text{Hash}(\cdot)$. The Ω_L oracle receives an empty input and

returns $x \in L$ chosen according to the distribution $D(L)$. The Hash oracle receives an input x . It first checks that x was previously output by the Ω_L oracle. If no, then it returns nothing. Otherwise, it chooses a key $k \in_R K$ and returns the pair $(\alpha(k, x), H_k(x))$. We stress that the Hash oracle *only* answers for inputs x that were generated by Ω_L . The output of the experiment is whatever machine M outputs.

Expt-Unif(M): This experiment is defined exactly as above except that the Hash oracle is replaced by the following Unif oracle. On input x , Unif first checks that x was previously output by the Ω_L oracle. If no, it returns nothing. Otherwise, it chooses a key $k \in_R K$ and a random element $g \in_R G$, and returns the pair $(\alpha(k, x), g)$. As above, the output of the experiment is whatever M outputs.

In [19] it is proven that no efficient M can distinguish between the experiments. In other words, when $x \in_R D(L)$, the value $H_k(x)$ is pseudorandom in G , even given $\alpha(k, x)$. This lemma is used a number of times in the proof of our password protocol.

Lemma 1. *Assume that $\mathcal{I}$ is a hard subset membership problem. Then, for every (non-uniform) polynomial-time oracle machine M it holds that,*

$$|\Pr[\text{Expt-Hash}(M) = 1] - \Pr[\text{Expt-Unif}(M) = 1]| < \text{negl}(n).$$

Hard partitioned subset membership problems. We now consider a variant of hard subset membership problems, where the set X can be *partitioned* into disjoint subsets of hard problems. That is, assume that the set X contains pairs of the form (i, x) , where $i \in \{1, \dots, \ell\}$ is an index. We denote by $X(i)$ the subset of pairs in X of the form (i, x) . Furthermore, we denote by $L(i)$ the subset of pairs in the language L of the form (i, x) . (We also associate sampling distributions $D(L(i))$ and $D(X(i) \setminus L(i))$ to each partition.) Then, such a problem constitutes a hard **partitioned** subset membership problem if for *every* i , it is hard to distinguish $x \in_R D(L(i))$ from $x \in_R D(X(i) \setminus L(i))$. (In the notation of Definition 1, we require that for every i , the ensembles $\{V(L(i))\}$ and $\{V(X(i) \setminus L(i))\}$ are computationally indistinguishable.) We stress that the definition of smooth projective hashing is unchanged when considered in the context of hard partitioned subset problems. That is, the smoothness is required to hold with respect to the entire sets X and L , and not with respect to individual partitions.

Lemma 1 also holds for hard partitioned subset membership problems (see [19]). Specifically, the definitions of the oracles in the experiments are modified as follows. The Ω_L oracle is modified so that instead of receiving the empty input, it is queried with an index i , and returns $x \in_R D(L(i))$. Likewise, $\Omega_{X \setminus L}$ receives an index i and returns an element $x \in_R D(X(i) \setminus L(i))$. Notice that in this scenario, the distinguishing machine M is given some control over the choice of x . Specifically, M can choose the index i that determines from which partition an element x is sampled.

Corollary 1. *Assume that $\mathcal{I}$ is a family of hard partitioned subset membership problem. Then, for every (non-uniform) polynomial-time oracle machine M*

$$|\Pr[\text{Expt-Hash}(M) = 1] - \Pr[\text{Expt-Unif}(M) = 1]| < \text{negl}(n).$$

4 The Protocol

Our protocol uses a chosen-ciphertext secure public-key labeled encryption scheme $\mathcal{E}$. The common reference string for the protocol is a public key PK for $\mathcal{E}$.

We then use a family of smooth projective functions $\mathcal{H} = \{H_k\}$ such that for every k in the key space K , $H_k : C_{PK} \times M \rightarrow \{0, 1\}^{2n}$, where M is the message space, C_{PK} is an efficiently recognizable superset of the ciphertext space. Notice that we are assuming that the projective hash function outputs $2n$ -bit strings⁵. If sk is a $2n$ -bit string we denote with $sk^{(1)}$ and $sk^{(2)}$ the first and second half of it respectively.

Finally, we assume that there is a mechanism that enables the parties to differentiate between different concurrent executions and to identify who they are interacting with. This can easily be implemented by having P_i choose a sufficiently long random string r and send the pair (i, r) to P_j along with its first message. P_i and P_j will then include r in any future messages of the protocol. We stress that the security of the protocol does not rest on the fact that these values are not modified by the adversary. Rather, this just ensures correct communication for protocols that are not under attack. The protocol appears in Figure 2.

Intuitive Security Proof. First notice that both A and B can compute the session key as instructed. Specifically, A can compute $H_k(c, w, A \circ B)$ because it has the projection key s and the witness (coins) for c . Furthermore, it can compute $H_{k'}(c', w, c \circ s)$ because it has the key k' (and therefore does not need the witness for c'). Likewise, B can also correctly compute both the hash values (and thus the session key). Second, when both parties A and B see the same messages (c, s, c', s', t) the session keys that they compute are the same. This is because the same hash value is obtained when using the hash keys (k and k') and when using the projection keys (s and s'). This implies that the correctness property holds for the protocol.

We now proceed to motivate why the adversary cannot distinguish a session key from a random key with probability greater than $Q_{\text{send}}/|\mathcal{D}|$, where Q_{send} equals the number of **Send** oracle calls made by the adversary to different protocol instances and $\mathcal{D}$ is the password dictionary. In order to see this, notice that if A , for example, receives c' that is not an encryption to w with label $c \circ s$ under PK , then A 's component of the session key sk_A will be statistically close to uniform.

⁵ We note that the constructions in [19] output values in a large algebraic group G . It is a standard application of randomness extraction (e.g. via universal hashing) to map such elements into $2n$ -bit strings, assuming the group G is large enough.

RG-PaKE

- **Common reference string:** The public key PK for a chosen-ciphertext secure encryption scheme $\mathcal{E}$. A description of a smooth projective hashing family H_k over the set X of ciphertext/password pairs (c, w) . The NP language L is composed of the tuples (c, m, ℓ) where $c = \mathcal{E}_{PK}(m, \ell)$ i.e. c is an encryption of m with label ℓ under PK . A message authentication code MAC .
- **Common input:** a shared (low-entropy) password w .
- **The protocol:**
 1. Party A computes an encryption $c = \mathcal{E}_{PK}(w, A \circ B)$ and sends it to party B .
 2. Party B chooses a key k for the smooth projective hash function (for the language L described above), and computes its projection $s = \alpha(k, c)$. Also B computes the projective hash over $(c, w, A \circ B)$, i.e. $sk_B = H_k(c, w, A \circ B)$. Finally B computes another encryption of the password with label $c \circ s$ i.e., $c' = \mathcal{E}_{PK}(w, c \circ s)$. B sends s, c' to party A .
 3. Party A chooses another key k' for the smooth projective hash function (for the language L described above), and computes its projection $s' = \alpha(k', c')$. Also A computes the projective hash over $(c', w, c \circ s)$, i.e. $sk_A = H_{k'}(c', w, c \circ s)$. A also computes sk_B using the projection s and its knowledge of a witness for the fact that c is an encryption to the string w with label $A \circ B$ (it knows a witness because it generated c). Set $t = MAC_{sk_B^{(1)}}(c, s, c', s') \oplus sk_A^{(1)}$. A sends s' to party B .
- **Session Key Definition:**
 1. Party B computes sk_A using the projection s' and its knowledge of a witness for the fact that c' is an encryption to the password w with label $c \circ s$ (it knows a witness because it generated c'). It tests if $t = MAC_{sk_B^{(1)}}(c, s, c', s') \oplus sk_A^{(1)}$. If the test fails it outputs an error message, otherwise it outputs $sk = sk_A^{(2)} \oplus sk_B^{(2)}$.
 2. Party A outputs $sk = sk_A^{(2)} \oplus sk_B^{(2)}$.

Session-Identifier Definition: Both parties take the series of messages (c, s, c', s') to be their session identifiers.

Fig. 2. Improved Password-Based Session-Key Exchange

This is because A computes $H_{k'}(c', w, c \circ s)$ for $c' \notin \mathcal{E}_{PK}(w, c \circ s)$ i.e. on an input outside the language. Therefore, by the definition of smooth projective hashing, $\{c', w, \alpha(k, c'), H_k(c', w, c \circ s)\}$ is statistically close to $\{c', w, \alpha(k, c'), r\}$, where r is a random $2n$ -bit string.

The same argument holds if B receives c that is not a encryption of w with label $A \circ B$. It therefore follows that if the adversary is to distinguish the session key from a random element, it must hand the parties encryptions of the valid messages (and in particular containing the correct passwords). One way for the adversary to do this is to copy (valid) commitments that are sent by the honest parties in the protocol executions. However, in this case, the adversary does not know the random coins used in generating the commitment, and once again the result of the projective hash function is a pseudorandom $2n$ -bit string (see Lemma 1). This means that the only option left to the adversary is to come up with valid commitments that were not previously sent by honest parties. However, by the non-malleability of the encryption scheme, the adversary cannot succeed in doing this with probability non-negligibly greater than just a priori guessing the password. Thus, its success probability is limited to $Q_{\text{send}}/|\mathcal{D}| + \text{negl}(n)$.

This intuitive explanation of the security of the protocol is not complete. Indeed it does not address the use of message authentication codes in the protocol. The MAC is needed to prevent further malleability attacks. Indeed while the ciphertexts containing the password are not malleable (because of the strong security of the encryption scheme used), the computation of the projective hash function could be malleable, and by recycling messages from previous executions the adversary could gain some knowledge about a session key. For example, it is possible that for some smooth projective hash family it holds that for every k , $H_{2k}(x) = 2H_k(x)$, and that by seeing $s = \alpha(k, c)$, the value $\hat{s} = \alpha(2k, c)$ is efficiently computable.

If this were the case, in the basic protocol described in Figure 1 an adversary could cause two instances to accept with *different* session identifiers and *related* session-keys. For example, given the message $s = \alpha(k, c), c'$ by party B in Round 2, the adversary could forward to A the message $\hat{s} = \alpha(2k, c), c'$. By requesting a Reveal for one of the instances, it could then distinguish the other instance's session-key from random, in contradiction to the security requirements.

Notice that the projection s is protected by malleability attacks because is incorporated as a label in the ciphertext c' . The surprising thing is that at this point the value sk_B is already pseudo-random to the eyes of the adversary, and known to both parties A and B . The most intuitive thing would be to use it to MAC the other projection s' .

But if A were to send $t = \text{MAC}_{sk_B}(c, s, c', s')$ the adversary could perform the following off-line attack. The adversary would start a session with A pretending to be B and obtain the commitment c . Next, the adversary $\mathcal{A}$ chooses k and returns $s = \alpha(k, c)$ to A together with an incorrect encryption c' . The response from A is s' and t computed as above with $sk_B = H_k(c, w, A \circ B)$. Now the adversary can traverse the entire dictionary $\mathcal{D}$ and for all possible w 's compute $sk_B = H_k(c, w, A \circ B)$ (it can do this because it knows k and so can compute $H_k(c, w)$ without a witness for c). The right password is the one for which sk_B verifies the above MAC t .

The final solution is then to “mask” the MAC, using sk_A which is not known to the adversary, and cannot even be computed off-line by traversing the dictionary (because the adversary does not know the coins used to produce c'). If the

adversary modifies s' and makes B accept then it must produce a MAC forgery with key sk_B . We note that a specific MAC key is used to MAC a single message, so it is sufficient to assume 1-time security for the MAC algorithm.

Of course this is just an intuition and the proof presented in [17] works out all the details⁶.

Theorem 1. *Assume that $\mathcal{E}$ is a public-key encryption secure against adaptive chosen ciphertext attack, MAC is a 1-time secure message authentication code and $\mathcal{H}$ is a family of smooth projective hash functions. Then, Protocol RG-PaKE in Figure 2 is a secure password-based session-key generation protocol.*

5 Extension and Conclusions

The reader is referred to [19] to see examples of efficient chosen-ciphertext secure encryption schemes that admit the type of projective hash functions needed in this protocol. They are based on the encryption schemes proposed by Cramer and Shoup in [11,12], and can be based on the DDH, Quadratic Residuosity and N -Residuosity Assumptions.

For simplicity we have presented the protocol using chosen-ciphertext secure encryption. It is possible using techniques used in [19] to prove the protocol assuming that $\mathcal{E}$ is non-malleable commitment scheme, which admits a projective hash function. The protocol needs to be modified and the proof is more complicated. However in practice it does not make much of a difference, as the only known efficient implementations of such commitment schemes are the ones mentioned above (i.e. based on chosen-ciphertext secure encryption and described by [19]).

Canetti et al. extend the KOY and GL protocol to the Universal Composability framework in [8]. Their protocol also uses one-time signatures to prevent malleability attacks and our modification is applicable to their protocol as well.

Conclusions. We have shown an improvement of the KOY and GL protocols, which does not require one-time signatures. Our protocol works in the common reference string and its proof does not require idealized assumptions such as the random oracle. For typical security parameters our protocol saves about 12 Kbytes of bandwidth, thus bringing provable security in the realm of password-authenticated key exchange one step closer to practical.

References

1. Bellare, M., Pointcheval, D., Rogaway, P.: Authenticated Key Exchange Secure Against Dictionary Attacks. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 139–155. Springer, Heidelberg (2000)
2. Bellare, M., Rogaway, P.: Entity Authentication and Key Distribution. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 232–249. Springer, Heidelberg (1994)

⁶ One final technicality: in order to protect the semantic security of the final session key we use the first half of sk_A and sk_B to perform the MAC test during the protocol and the second half to compute the actual session key.

3. Bellare, S.M., Merritt, M.: Encrypted Key Exchange: Password Based Protocols Secure Against Dictionary Attacks. In: Proceedings 1992 IEEE Symposium on Research in Security and Privacy, pp. 72–84. IEEE Computer Society, Los Alamitos (1992)
4. Bellare, S.M., Merritt, M.: Augmented Encrypted Key Exchange: A Password-Based Protocol Secure Against Dictionary Attacks and Password File Compromise. In: Proceedings of the 1st ACM Conference on Computer and Communication Security, pp. 244–250 (1993)
5. Boneh, D., Boyen, X.: Short Signatures Without Random Oracles. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 56–73. Springer, Heidelberg (2004)
6. Boyko, V., MacKenzie, P., Patel, S.: Provably Secure Password-Authenticated Key Exchange Using Diffie-Hellman. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 156–171. Springer, Heidelberg (2000)
7. Canetti, R., Goldreich, O., Halevi, S.: The Random Oracle Methodology, Revisited. *Journal of the ACM* 51(4), 557–594 (2004)
8. Canetti, R., Halevi, S., Katz, J., Lindell, Y., MacKenzie, P.: Universally Composable Password-Based Key Exchange. In: Cramer, R.J.F. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 404–421. Springer, Heidelberg (2005)
9. Canetti, R., Krawczyk, H.: Analysis of Key-Exchange Protocols and Their Use for Building Secure Channels. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 453–474. Springer, Heidelberg (2001)
10. Catalano, D., Di Raimondo, M., Fiore, D., Gennaro, R.: Some Theoretical and Experimental Results about Off-Line/On-Line Signatures. In : PKC 2008 (to appear, 2008)
11. Cramer, R., Shoup, V.: A Practical Public-Key Cryptosystem Secure Against Adaptive Chosen Ciphertexts Attacks. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 13–25. Springer, Heidelberg (1998)
12. Cramer, R., Shoup, V.: Universal Hash Proofs and a Paradigm for Adaptive Chosen Ciphertext Secure Public-Key Encryption. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 45–64. Springer, Heidelberg (2002)
13. Cramer, R., Shoup, V.: Design and Analysis of Practical Public-Key Encryption Schemes Secure Against Adaptive Chosen Ciphertext Attack. *SIAM Journal of Computing* 33, 167–226 (2003)
14. Cramer, R., Shoup, V.: Signature schemes based on the strong RSA assumption. *ACM Trans. Inf. Syst. Secur.* 3(3), 161–185 (2000)
15. Diffie, W., Hellman, M.E.: New Directions in Cryptography. *IEEE Trans. on Inf. Theory* IT-22, 644–654 (1976)
16. Dolev, D., Dwork, C., Naor, M.: Non-Malleable Cryptography. *SIAM Journal of Computing* 30(2), 391–437 (2000)
17. Gennaro, R.: Faster and Shorter Password-Authenticated Key Exchange, <http://eprint.iacr.org/2007/325>
18. Gennaro, R., Halevi, S., Rabin, T.: Secure Hash-and-Sign Signatures Without the Random Oracle. In: Stern, J. (ed.) EUROCRYPT 1999. LNCS, vol. 1592, pp. 123–139. Springer, Heidelberg (1999)
19. Gennaro, R., Lindell, Y.: A framework for password-based authenticated key exchange. *ACM Transactions on Information and System Security (TISSEC)* 9(2), 181–234 (2006)
20. Goldreich, O., Lindell, Y.: Session Key Generation using Human Passwords Only. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 408–432. Springer, Heidelberg (2001)

21. Goldwasser, S., Micali, S., Rivest, R.L.: A Digital Signature Scheme Secure Against Adaptive Chosen-Message Attacks. *SIAM Journal on Computing* 17(2), 281–308 (1988)
22. Halevi, S., Krawczyk, H.: Public-Key Cryptography and Password Protocols. *ACM Transactions on Information and System Security (TISSEC)* 2(3), 230–268 (1999)
23. Jablon, D.P.: Strong Password-Only Authenticated Key Exchange. *SIGCOMM Computer Communication Review* 26(5), 5–26 (1996)
24. Katz, J., Ostrovsky, R., Yung, M.: Practical Password-Authenticated Key Exchange Provably Secure under Standard Assumptions. In: Pfitzmann, B. (ed.) *EUROCRYPT 2001*. LNCS, vol. 2045, pp. 475–494. Springer, Heidelberg (2001)
25. Kurosawa, K., Schmidt-Samoa, K.: New Online/Offline Signature Schemes Without Random Oracles. In: Yung, M., Dodis, Y., Kiayias, A., Malkin, T.G. (eds.) *PKC 2006*. LNCS, vol. 3958, pp. 330–346. Springer, Heidelberg (2006)
26. Lamport, L.: Constructing digital signatures from a one-way function. Technical Report CSL-98, SRI International (October 1979)
27. Lucks, S.: Open Key Exchange: How to Defeat Dictionary Attacks Without Encrypting Public Keys. In: Christianson, B., Lomas, M. (eds.) *Security Protocols 1997*. LNCS, vol. 1361, pp. 79–90. Springer, Heidelberg (1998)
28. Patel, S.: Number Theoretic Attacks on Secure Password Schemes. In: *Proceedings of the 1997 IEEE Symposium on Security and Privacy*, pp. 236–247 (1997)
29. Steiner, M., Tsudik, G., Waidner, M.: Refinement and Extension of Encrypted Key Exchange. *ACM SIGOPS Oper. Syst. Rev.* 29(3), 22–30 (1995)
30. Wu, T.: The Secure Remote Password Protocol. In: *1998 Internet Society Symposium on Network and Distributed System Security*, pp. 97–111 (1998)