

Mitigating Server Breaches in Password-Based Authentication: Secure and Efficient Solutions

Olivier Blazy¹, Céline Chevalier^{2(✉)}, and Damien Vergnaud³

¹ Université de Limoges, XLim, Limoges, France

² Université Panthéon-Assas, Paris, France

`celine.chevalier@ens.fr`

³ ENS, CNRS, INRIA and PSL Research University, Paris, France

Abstract. *Password-Authenticated Key Exchange* allows users to generate a strong cryptographic key based on a shared “human-memorable” password without requiring a public-key infrastructure. It is one of the most widely used and fundamental cryptographic primitives. Unfortunately, mass password theft from organizations is continually in the news and, even if passwords are salted and hashed, brute force breaking of password hashing is usually very successful in practice.

In this paper, we propose two efficient protocols where the password database is somehow shared among two servers (or more), and authentication requires a distributed computation involving the client and the servers. In this scenario, even if a server compromise is doable, the secret exposure is not valuable to the adversary since it reveals only a share of the password database and does not permit to brute force guess a password without further interactions with the parties for each guess. Our protocols rely on *smooth projective hash functions* and are proven secure under classical assumption in the standard model (*i.e.* do not require idealized assumption, such as random oracles).

Keywords: Password-authenticated key exchange · Distributed computation · Decision diffie-hellman · Smooth projective hashing

1 Introduction

Authenticated Key Exchange protocols enable two parties to establish a shared cryptographically strong key over an insecure network under the complete control of an adversary. This primitive is one of the most widely used and fundamental cryptographic primitives and it obviously requires the parties to have authentication means, *e.g.* (public or secret) cryptographic keys or short (*i.e.*, low-entropy) secret keys.

PAKE, for *Password-Authenticated Key Exchange*, allows users to generate a strong cryptographic key based on a shared “human-memorable” password without requiring a public-key infrastructure. In this setting, an adversary controlling all communication in the network should not be able to mount an *offline dictionary attack*. More precisely, an eavesdropper should not obtain enough

information to be able to brute force guess a password without further interactions with the parties for each guess. Note that *online dictionary attacks* in which an adversary simply attempts to log-in repeatedly, trying each possible low-entropy password can be dealt with using other computer security methods (such as limiting the number of attempts). In particular, strong security can be obtained even using passwords chosen from a small set of possible values (a four-digit pin, for example).

Incidents of sensitive customer information “hacking” (including leaking of passwords) in e-commerce systems are frequently revealed in the newspaper. In addition to major reputational damage, a company with a significant data breach may be sued by its clients for the breach and may be suspended or disqualified from future public sector or government work.

To alleviate the threat that stored passwords are revealed immediately in case of a server compromise, many servers adopt the approach for storing passwords in a hashed form with a random salt. When the database of hashed password is compromised, the offline dictionary attack requires a more important computational effort but remains usually possible. The notion of *Verifier-based PAKE*, where the client owns a password pw and the server knows a one-way transformation v of the password only were proposed by Bellare and Merritt [BM92]. The two players eventually agree on a common high entropy secret if and only if pw and v match together. It prevents massive password recovering in case of server corruption and it forces the attacker who breaks into the server and is willing to recover passwords to perform an additional costly offline dictionary attack.

We consider an alternative approach inspired by the multi-party computation paradigm (and first suggested by Ford and Kaliski [FK00]). The password database on the server side is somehow shared among two servers (or more, but we focus here on two for sake of simplicity), and authentication requires a distributed computation involving the client – who still does not need an additional cryptographic device capable of storing high-entropy secret keys – and the two servers who will use some additional shared secret information. The interaction is performed using a *gateway* that does not know any secret information and ends up in the gateway and the client sharing a common key. The lifetime of the protocol is divided into distinct periods (for simplicity, one may think of these time periods as being of equal length; e.g. one day) and at the beginning of each period, the two servers interact and update their sharing of the password database. Similarly to proactive schemes in multi-party computation, we allow the adversary multiple corruptions of each server, limiting only the corruptions to one server for each period. The user does not need to update his password nor to perform any kind of computations and its interaction with the two servers (performed using the gateway) remains the same for the lifetime of the protocol. In this scenario, even if a server compromise is doable, the secret exposure is not valuable to the adversary since it reveals only a share of the password database and does not permit to run an offline dictionary attack.

The goal of our paper is to present practical realizations based on classical cryptographic assumptions in the standard security model.

Related Work. EKE (Encrypted Key Exchange) is the most famous instantiation of *Password-Authenticated Key Exchange*. It has been proposed by Bellare and Merritt [BM92] and consists of a Diffie-Hellman key exchange [DH76], where the flows are symmetrically encrypted under the shared password.

A first formal security model was proposed by Bellare, Pointcheval and Rogaway [BPR00] (the BPR model), to deal with offline dictionary attacks. It essentially says that the best attack should be the online exhaustive search, consisting in trying all the passwords by successive executions of the protocol with the server. Several variants of EKE with BPR-security proofs have been proposed in the ideal-cipher model or the random-oracle model (see the survey [Poi12] for details). Katz, Ostrovsky and Yung [KOY01] proposed the first practical scheme, provably secure in the standard model under the Decision Diffie-Hellman assumption (DDH). It has been generalized by Gennaro and Lindell [GL03], making use of smooth projective hash functions.

As mentioned above, Ford and Kaliski [FK00] were the first to propose to distribute the capability to test passwords over multiple servers. Building on this approach, several such protocols were subsequently proposed in various settings (e.g. [Jab01, MSJ02, BJKS03, DG06, SK05, KMTG12, ACFP05, KM14]) and it is worth noting that the protocol from [BJKS03] is commercially available as EMC’s *RSA Distributed Credential Protection*. Recently, Camenisch, Enderlein and Neven [CEN15] revisited this approach and proposed a scheme in the universal composability framework [Can01] (which has obvious advantages for password-based protocols since users often use related passwords for many providers). Camenisch *et al.* gave interesting details about the steps that need to be taken when a compromise actually occurs. Unfortunately, due to the inherent difficulties of construction of the simulator in the universal composability framework, their scheme is inefficient since users and servers have to perform a few hundred exponentiations each.

Our Contributions. In order to achieve practical constructions in the standard security model, we consider the variant of the BPR model¹ in the distributed setting proposed by Katz, MacKenzie, Taban and Gligor in [KMTG12]. In this security model, we assume that the communication between the client and the authentication servers, is carried on a basically insecure network. Messages can be tapped and modified by an adversary and the communication between the clients and the servers is asynchronous. The adversary should not be able to brute force guess a password without further interactions with the client for each guess even if he corrupts and impersonates a server in an active way.

Our first construction uses a similar approach to the schemes from [Jab01, MSJ02, BJKS03, DG06, SK05, KMTG12, ACFP05, KM14]: the user generates information theoretic shares of his password and sends them to the servers. In the authentication phase, the parties run a dedicated protocol to verify that

¹ Our schemes can be adapted to achieve security in universal composability framework using techniques similar to those used in [CEN15]. The resulting schemes are slightly more efficient but are unfortunately still not practical.

the provided password equals the priorly shared one. Our solution then consists in some sort of three-party PAKE, in which (1) the user implicitly checks (using a smooth projective hash function) that its password is indeed the sum of the shares owned by the two servers, and (2) each server implicitly checks that its share is the difference of the password owned by the user and the share owned by the other server. Contrary to the popular approach initiated in [KOY01, GL03] for PAKE, we cannot use two smooth projective hash functions (one for the client and one for the server) so we propose a technique in order to combine in a secure way six smooth projective hash functions. This new method (which may be of independent interest) allows us to prove the security of this construction under classical cryptographic assumptions (namely the DDH assumption) in the standard security model from [KMTG12] (without any idealized assumptions).

The main weakness of this first solution is that at each time period, the servers have to refresh the information-theoretic sharing of the password of all users. This can be handled easily using well-known techniques from proactive multi-party computation but if the number of users is large, this can be really time-consuming (in particular if the time period is very short). Our second construction (which is the main contribution of the paper) is built on the ideas from the first one but passwords are now encrypted using a public-key encryption scheme where the corresponding secret key is shared among the servers. At the beginning of each time period, the servers only need to refresh the sharing of this secret key but the password database is not modified (and can actually be public). Password verification and the authenticated key exchange is then carried out without ever decrypting the database. A secure protocol is run to verify that the password sent by the user matches the encrypted password. It is similar to the protocol we design for the first construction except that the user encrypts its password and the parties implicitly check (using in this case five smooth projective hash functions) that the message encrypted in this ciphertext is the same as the message encrypted in the database (using the secret key shared upon the servers). Both constructions consist in only two flows (one from the client and one from the servers) and a (private) flow from the servers to the gateway.

2 Preliminaries

In this section we recall various classical definitions, tools used throughout this paper. We use classical notions and notations and the familiar reader may skip this section.

Public-Key Encryption Scheme. An encryption scheme \mathcal{E} is described by four algorithms (Setup, KeyGen, Encrypt, Decrypt):

- Setup($1^{\mathfrak{K}}$), where \mathfrak{K} is the security parameter, generates the global parameters param of the scheme;
- KeyGen(param) outputs a pair of keys, a (public) encryption key ek and a (private) decryption key dk;

- $\text{Encrypt}(\text{ek}, M; \rho)$ outputs a ciphertext \mathcal{C} , on the message M , under the encryption key ek , with randomness ρ ;
- $\text{Decrypt}(\text{dk}, \mathcal{C})$ outputs the plaintext M , encrypted in the ciphertext \mathcal{C} or \perp .

Such encryption scheme is required to have the classical properties, *Correctness* and *Indistinguishability under Chosen Plaintext Attack* IND-CPA [GM84]: One might want to increase the requirements on the security of an encryption, in this case the IND-CPA notion can be strengthened into Indistinguishability under Adaptive Chosen Ciphertext Attack IND-CCA (see the full version [BCV16] for formal definitions).

Smooth Projective Hash Functions. SPHF [CS02] were introduced by Cramer and Shoup. A projective hashing family is a family of hash functions that can be evaluated in two ways: using the (secret) hashing key, one can compute the function on every point in its domain, whereas using the (public) *projected* key one can only compute the function on a special subset of its domain. Such a family is deemed *smooth* if the value of the hash function on any point outside the special subset is independent of the projected key.

Smooth Projective Hashing System: A Smooth Projective Hash Function over a language $\mathcal{L} \subset X$, onto a set \mathbb{G} , is defined by five algorithms (Setup , HashKG , ProjKG , Hash , ProjHash):

- $\text{Setup}(1^{\mathfrak{R}})$ where \mathfrak{R} is the security parameter, generates the global parameters param of the scheme, and the description of an \mathcal{NP} language \mathcal{L} ;
- $\text{HashKG}(\mathcal{L}, \text{param})$, outputs a hashing key hk for the language \mathcal{L} ;
- $\text{ProjKG}(\text{hk}, (\mathcal{L}, \text{param}), W)$, derives the projection key hp , possibly depending on the word W [GL03, ACP09] thanks to the hashing key hk .
- $\text{Hash}(\text{hk}, (\mathcal{L}, \text{param}), W)$, outputs a hash value $v \in \mathbb{G}$, thanks to the hashing key hk , and W
- $\text{ProjHash}(\text{hp}, (\mathcal{L}, \text{param}), W, w)$, outputs the hash value $v' \in \mathbb{G}$, thanks to the projection key hp and the witness w that $W \in \mathcal{L}$.

In the following, we consider \mathcal{L} as a hard-partitioned subset of X , *i.e.* it is computationally hard to distinguish a random element in \mathcal{L} from a random element in $X \setminus \mathcal{L}$. A Smooth Projective Hash Function SPHF should satisfy the following properties:

- *Correctness:* Let $W \in \mathcal{L}$ and w a witness of this membership. For all hashing keys hk and associated projection keys hp we have $\text{Hash}(\text{hk}, (\mathcal{L}, \text{param}), W) = \text{ProjHash}(\text{hp}, (\mathcal{L}, \text{param}), W, w)$.
- *Smoothness:* For all $W \in X \setminus \mathcal{L}$ the following distributions are statistically indistinguishable:

$$\left\{ (\mathcal{L}, \text{param}, W, \text{hp}, v) \left| \begin{array}{l} \text{param} = \text{Setup}(1^{\mathfrak{R}}), \text{hk} = \text{HashKG}(\mathcal{L}, \text{param}), \\ \text{hp} = \text{ProjKG}(\text{hk}, (\mathcal{L}, \text{param}), W), \\ v = \text{Hash}(\text{hk}, (\mathcal{L}, \text{param}), W) \end{array} \right. \right\} \\ \simeq \left\{ (\mathcal{L}, \text{param}, W, \text{hp}, v) \left| \begin{array}{l} \text{param} = \text{Setup}(1^{\mathfrak{R}}), \text{hk} = \text{HashKG}(\mathcal{L}, \text{param}), \\ \text{hp} = \text{ProjKG}(\text{hk}, (\mathcal{L}, \text{param}), W), v \xleftarrow{\$} \mathbb{G} \end{array} \right. \right\}.$$

- *Pseudo-Randomness*: If $W \in \mathcal{L}$, then without a witness of membership the two previous distributions should remain computationally indistinguishable.

Classical Instantiations. For our needs, we consider discrete-logarithm based encryption schemes and related smooth projective hash functions. The underlying setting is a group \mathbb{G} (denoted multiplicatively) of prime order p and we denote g a random generator of $\mathbb{G} = \langle g \rangle$. The security of our constructions will rely on the standard Decisional Diffie Hellman problems in \mathbb{G} :

Decisional Diffie Hellman (DDH) [Bon98]: The Decisional Diffie-Hellman hypothesis states that in a group (p, \mathbb{G}, g) (written in multiplicative notation), given (g^μ, g^ν, g^ψ) for unknown $\mu, \nu \xleftarrow{\$} \mathbb{Z}_p$, it is hard to decide whether $\psi = \mu\nu$.

ElGamal encryption [ElG84] is defined by the following four algorithms:

- **Setup**($1^{\mathbb{R}}$): The scheme needs a multiplicative group (p, \mathbb{G}, g) . The global parameters **param** consist of these elements (p, \mathbb{G}, g) .
- **KeyGen**(**param**): Chooses one random scalar $\alpha \xleftarrow{\$} \mathbb{Z}_p$, which define the secret key $\text{dk} = \alpha$, and the public key $\text{pk} = h = g^\alpha$.
- **Encrypt**($\text{pk} = h, M; r$): For a message $M \in \mathbb{G}$ and a random scalar $r \xleftarrow{\$} \mathbb{Z}_p$, computes the ciphertext as $C = (c_1 = h^r M, c_2 = g^r)$.
- **Decrypt**($\text{dk} = \alpha, C = (c_1, c_2)$): One computes $M = c_1 / (c_2^\alpha)$.

As shown by Boneh [Bon98], this scheme is IND-CPA under the hardness of DDH.

Cramer-Shoup encryption scheme [CS98] is an IND-CCA version of the ElGamal Encryption.

- **Setup**($1^{\mathbb{R}}$) generates a group \mathbb{G} of order p , with a generator g
- **KeyGen**(**param**) generates $(g_1, g_2) \xleftarrow{\$} \mathbb{G}^2$, $\text{dk} = (x_1, x_2, y_1, y_2, z) \xleftarrow{\$} \mathbb{Z}_p^5$, and sets, $c = g_1^{x_1} g_2^{x_2}$, $d = g_1^{y_1} g_2^{y_2}$, and $h = g^z$. It also chooses a Collision-Resistant hash function \mathfrak{H}_K in a hash family \mathcal{H} (or simply a Universal One-Way Hash Function). The encryption key is $\text{ek} = (g_1, g_2, c, d, h, \mathfrak{H}_K)$.
- **Encrypt**($\text{ek}, M; r$), for a message $M \in \mathbb{G}$ and a random scalar $r \in \mathbb{Z}_p$, the ciphertext is $C = (\mathbf{u} = (g_1^r, g_2^r), e = M \cdot h^r, v = (cd^\xi)^r)$, where v is computed afterwards with $\xi = \mathfrak{H}_K(\mathbf{u}, e)$.
- **Decrypt**(ℓ, dk, C): one computes $\xi = \mathfrak{H}_K(\mathbf{u}, e)$ and checks whether $u_1^{x_1 + \xi y_1} \cdot u_2^{x_2 + \xi y_2} \stackrel{?}{=} v$. If the equality holds, one computes $M = e / (u_1^z)$ and outputs M . Otherwise, one outputs \perp .

The security of the scheme is proven under the DDH assumption and the fact the hash function used is a Universal One-Way Hash Function (see [CS98]).

3 Security Model

Distributed PAKE. In a distributed PAKE system, we consider as usual a client (owning a password) willing to interact with a gateway, such as a website.

The difference compared to a non-distributed system is that the gateway itself interacts with two servers, and none of the three owns enough information to be able to recover the passwords of the clients on its own². Such a scheme is correct if the interaction between a client with a correct password and the gateway succeeds. An honest execution of a distributed PAKE protocol should result in the client holding a session key K_U and the gateway holding a session key $K_G = K_U$.

We propose in this paper two settings that describe well this situation. In a first setting, we consider that the passwords of the clients are shared information-theoretically between the servers, such as $\pi = \pi_1 + \pi_2$ (if the password π belongs to an appropriate group) or with the help of any secret sharing protocol. At the beginning of each time period, the shares are updated, in a probabilistic way, using a public function **Refresh**, depending on the sharing protocol used.

In a second setting, we consider that the gateway owns a database of encrypted passwords (which can be considered public), and the servers each own a share of the corresponding private keys (obtained by a secret sharing protocol). At the beginning of each time period, the shares are updated, in a probabilistic way, using a public function **Refresh**, depending on the sharing protocol used.

Since the security of our schemes is not analyzed in the universal composability framework (contrary to the recent paper [CEN15]), the **Refresh** procedure can be handled easily using classical techniques from computational proactive secret sharing (see [OY91, HJKY95] for instance).

Security Model. We consider the classical model [BPR00] for authenticated key-exchange, adapted to the two-server setting by [ACFP05, KMTG12]. In the latter model, the authors assume that every client in the system shares its password with exactly two servers. We loosen this requirement here, depending on the setting considered, as described above. We refer the interested reader to these articles for the details and we give the high-level ideas in [BCV16].

4 Our Simple Protocol

In this first setting, we consider a client U owning a password π and willing to interact with a gateway G . The gateway interacts with two servers S_1 (owning π_1) and S_2 (owning π_2), such that $\pi = \pi_1 + \pi_2$. It should be noted that only the client’s password is assumed to be small and human-memorable. The two “passwords” owned by the servers can be arbitrarily big. The aim of the protocol is to establish a shared session key between the client and the gateway.

A simple solution to this problem consists in considering some sort of three-party PAKE, in which the client implicitly checks (using an SPHF) whether its password is the sum of the shares owned by the two servers, and the servers implicitly check (also using an SPHF) whether their share is the difference of the

² Note that the gateway can be merged with one server.

password owned by the client and the share owned by the other server. For sake of simplicity, we denote the client U as S_0 and its password π as π_0 .

4.1 Building Blocks

Cramer-Shoup Encryption and SPHF. We consider Cramer-Shoup encryption as described in Sect. 2. The public key is denoted by $\text{ek} = (g_1, g_2, c, d, h, \mathfrak{H}_K)$ and the private key by $\text{dk} = (x_1, x_2, y_1, y_2, z) \xleftarrow{\$} \mathbb{Z}_p^5$. The public parameters $(\mathbb{G}, p, g, \text{ek})$ are given as a common reference string.

We denote the ciphertext of a message $M \in \mathbb{G}$ with the scalar $r \in \mathbb{Z}_p$ by $\mathcal{C} = \text{CS}_{\text{ek}}(M; r) = (u_1, u_2, e, v)$, with $v = cd^\xi$ and $\xi = \mathfrak{H}_K(u_1, u_2, e)$.

We use the SPHF described in [BBC+13] for the language of the valid ciphertexts of M under the public key ek . Its main advantage is that it can be computed without using the associated ciphertext, and in particular before having seen it. This allows all the participants to send their ciphertext and their projected keys in only one flow. The classical use of this SPHF is as follows: user U (owning a message M) and V (owning a message M') are supposed to share a common message, so that $M = M'$. User U wants to implicitly check this equality. To this aim, user V sends an encryption \mathcal{C} of M' under randomness r . In order for U to implicitly check that \mathcal{C} is a valid encryption of M , it chooses a hash key hk and computes and sends a projection key hp to V . If $M = M'$ and if the encryption was computed correctly, then the hash value H computed by U using the private value hk is the same as the projected hash value H' computed by V using the public value hp and its private witness r . The SPHF is described by the following algorithms.

$$\begin{aligned} \text{Setup}(1^\kappa): \text{param} &= (\text{ek}, M) \\ \mathcal{L} &= \{\mathcal{C} = (u_1, u_2, e, v) \in \mathbb{G}^4 \mid \exists r \in \mathbb{Z}_p \text{ s. t.} \\ &\quad \mathcal{C} = \text{CS}_{\text{ek}}(M; r)\} \end{aligned}$$

$$\text{HashKG}(\mathcal{L}, (\text{ek}, M)): \text{hk} = (\eta, \gamma, \theta, \lambda, \kappa) \xleftarrow{\$} \mathbb{Z}_p^5$$

$$\text{ProjKG}(\text{hk}, (\mathcal{L}, (\text{ek}, M))): \text{hp} = (\text{hp}_1 = g_1^\eta g_2^\theta h^\lambda c^\kappa, \text{hp}_2 = g_1^\gamma d^\kappa) \in \mathbb{G}^2$$

$$\text{Hash}(\text{hk}, (\mathcal{L}, (\text{ek}, M)), \mathcal{C}): H = \text{Hash}(\text{hk}, (\text{ek}, M), \mathcal{C}) = u_1^{(\eta+\xi\gamma)} u_2^\theta (e/M)^\lambda v^\kappa$$

$$\text{ProjHash}(\text{hp}, (\mathcal{L}, (\text{ek}, M')), \mathcal{C}, r): H' = (\text{hp}_1 \text{hp}_2^\xi)^r$$

It has been known to be correct, smooth and pseudo-random since [BBC+13].

Main Idea of the Construction. In our setting, we denote by $\text{pw}_b = g^{\pi_b}$. The main idea of the protocol is depicted on Fig. 1. For sake of readability, the participants which have a real role in the computations are directly linked by arrows in the picture, but one should keep in mind that all the participants (U , S_1 and S_2) only communicate with G , which then broadcasts all the messages.

In a classical SPHF-based two-party key-exchange between U and G , the client and the gateway would compute a Cramer-Shoup encryption of their password: $\mathcal{C}_0 = \text{CS}_{\text{ek}}(\text{pw}_0; r_0)$ and $\mathcal{C}_G = \text{CS}_{\text{ek}}(\text{pw}_G; r_G)$. The gateway would then send a projection key $\text{hp}_{G,0}$ in order to implicitly check via an SPHF whether \mathcal{C}_0 is a valid Cramer-Shoup encryption of pw_G , and the client would send a projection key $\text{hp}_{0,G}$ in order to implicitly check via an SPHF whether \mathcal{C}_G is a valid Cramer-Shoup encryption of pw_0 .

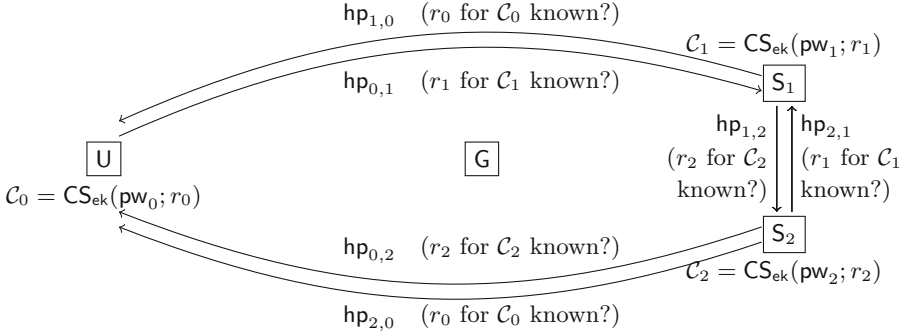


Fig. 1. Main idea of the construction

Here, since S_0 owns $\text{pw}_0 = \text{pw}_1 \cdot \text{pw}_2$, so that the players do not share the same password, we consider an SPHF between each pair of players (S_i, S_j) , in which player S_i computes the ciphertext $C_i = \text{CS}_{\text{ek}}(\text{pw}_i; r_i)$, the keys $\text{hk}_{i,j}$ and $\text{hp}_{i,j}$ and sends $(C_i, \text{hp}_{i,j})$ to S_j . It also computes the hash value $H_{i,j} = \text{Hash}(\text{hk}_{i,j}, (\text{ek}, \text{pw}_i), C_j)$ and the projected hash value $H'_{j,i} = \text{ProjHash}(\text{hp}_{j,i}, (\text{ek}, M_i), C_i, r_i)$. Formally, for each pair of users (S_i, S_j) , the language checked on S_j by S_i is defined as follows: $C_j \in \mathcal{L}_{i,j} = \{C = (u_1, u_2, e, v) \in \mathbb{G}^4 \mid \exists r \in \mathbb{Z}_p \text{ such that } C = \text{CS}_{\text{ek}}(\text{pw}_j; r)\}$ but it cannot be checked directly by a unique SPHF since the passwords are different (and thus S_i does not know pw_j). Rather, we combine in the protocol the six SPHF described to globally ensure the correctness (each one remaining smooth and pseudo-random), as described in the next part. The correctness of the SPHF for the pair (S_i, S_j) implies that if everything was computed honestly, then one gets the equalities $H_{i,j}(\text{pw}_i/\text{pw}_j)^{\lambda_i} = H'_{i,j}$ and $H_{j,i}(\text{pw}_j/\text{pw}_i)^{\lambda_j} = H'_{j,i}$.

4.2 Login Procedure

- Each participant S_b picks r_b at random and computes a Cramer-Shoup encryption of its password $C_b = \text{CS}_{\text{ek}}(\text{pw}_b; r_b)$, with $v_b = cd^{\xi_b}$. It also chooses, for $i \in \{0, 1, 2\} \setminus \{b\}$, a random hash key $\text{hk}_{b,i} = (\eta_{b,i}, \gamma_{b,i}, \theta_{b,i}, \lambda_b, \kappa_{b,i})$ and sets $\text{hp}_{b,i} = (\text{hp}_{b,i;1}, \text{hp}_{b,i;2}) = (g_1^{\eta_{b,i}} g_2^{\theta_{b,i}} h^{\lambda_b} c^{\kappa_{b,i}}, g_1^{\gamma_{b,i}} d^{\kappa_{b,i}})$ as the projection key intended to the participant S_i . It sends $(C_b, (\text{hp}_{b,i})_{i \in \{0,1,2\} \setminus \{b\}})$ to the gateway G , which broadcasts these values to the other participants.
- After receiving the first flow from the servers, the client computes, for $i \in \{1, 2\}$, $H'_{i,0} = \text{hp}_{i,0;1}^{r_0} \text{hp}_{i,0;2}^{\xi_0 r_0}$. It also computes $H_0 = H_{0,1} \cdot H_{0,2} \cdot \text{pw}_0^{\lambda_0}$, and sets its session key K_U as $K_U = K_0 = H'_{1,0} \cdot H'_{2,0} \cdot H_0$. After receiving the first flow from the other server and the client, the server S_b computes, for $i \in \{0, 3-b\}$, $H'_{i,b} = \text{hp}_{i,b;1}^{r_b} \text{hp}_{i,b;2}^{\xi_b r_b}$. It also computes $H_b = H_{b,0} / (H_{b,3-b} \cdot \text{pw}_b^{\lambda_b})$, and sets its partial key K_b as $K_b = H'_{0,b} \cdot H'_{3-b,b} \cdot H_b$. It privately sends this value K_b to the gateway G .

– The gateway finally sets $K_G = K_1 \cdot K_2$.

Correctness. Recall that $H'_{i,j} = H_{i,j}(\text{pw}_i/\text{pw}_j)^{\lambda_i}$ for all pairs $(i,j) \in \{0,1,2\}^2$, so that the session key of the gateway is equal to $K_G = K_1 \cdot K_2$, *i.e.*

$$K_G = \frac{H'_{0,1}H'_{2,1}H_{1,0}}{H_{1,2}\text{pw}_1^{\lambda_1}} \frac{H'_{0,2}H'_{1,2}H_{2,0}}{H_{2,1}\text{pw}_2^{\lambda_2}} = H_{0,1}H_{1,0}H_{0,2}H_{2,0} \left(\frac{1}{\text{pw}_2}\right)^{\lambda_1} \left(\frac{1}{\text{pw}_1}\right)^{\lambda_2} \left(\frac{(\text{pw}_0)^2}{\text{pw}_1\text{pw}_2}\right)^{\lambda_0}$$

while the session key of the client is $K_U = H'_{1,0} \cdot H'_{2,0} \cdot H_0$, *i.e.*

$$K_U = H'_{1,0} \cdot H'_{2,0} \cdot H_{0,1} \cdot H_{0,2} \cdot \text{pw}_0^{\lambda_0} = H_{0,1}H_{1,0}H_{0,2}H_{2,0} \left(\frac{\text{pw}_1}{\text{pw}_0}\right)^{\lambda_1} \left(\frac{\text{pw}_2}{\text{pw}_0}\right)^{\lambda_2} (\text{pw}_0)^{\lambda_0}$$

and these two values are equal as soon as $\text{pw}_0 = \text{pw}_1 \cdot \text{pw}_2$.

Complexity. The following table sums up the number of group elements needed for each participant. It should be noted that in case one of the servers is the gateway, its communication costs are reduced.

	Ciphertext (Broadcast)	Projection Keys (Overall)	To the Gateway
Client	4	4	0
Server (each)	4	4	1

Proof of Security. The proof follows the spirit of the proofs given in [KV11, BBC+13]. For lack of space, a sketch is given in [BCV16].

5 Our Efficient Protocol

In this second setting, we consider again a client U owning a password π and willing to interact with a gateway G . The gateway owns a public database of encrypted passwords, and it interacts with two servers S_1 and S_2 , each owning a share of the secret key of the encryption scheme. The aim of the protocol is to establish a shared session key between the client and the gateway.

The idea is similar to the protocol described in the former section, except that only the client needs to compute a ciphertext, the other ciphertext being publicly available from the database. The participants implicitly check (using several SPHF) that the message encrypted in the ciphertext of the client is the same as the message encrypted in the database (using the secret key shared upon the servers).

5.1 Building Blocks

Cramer-Shoup Encryption and SPHF. We consider Cramer-Shoup encryption as the previous construction. We use here the simpler SPHF described in [GL03] for the language of the valid ciphertexts of M under the public key ek , in which the participants need to see the ciphertext before being able to compute their projection keys. The SPHF is described by the following algorithms.

Setup($1^{\mathbb{R}}$): param = (ek, M)

$$\mathcal{L} = \{\mathcal{C} = (u_1, u_2, e, v) \in \mathbb{G}^4 \mid \exists r \in \mathbb{Z}_p \text{ s. t.} \\ \mathcal{C} = \text{CS}_{\text{ek}}(M; r)\}$$

HashKG(\mathcal{L} , (ek, M)): hk = ($\eta, \theta, \lambda, \kappa$) $\xleftarrow{\$} \mathbb{Z}_p^4$

ProjKG(hk, (\mathcal{L} , (ek, M)), \mathcal{C}): hp = $g_1^\eta g_2^\theta h^\lambda (cd^\xi)^\kappa \in \mathbb{G}$

Hash(hk, (\mathcal{L} , (ek, M)), \mathcal{C}): $H = u_1^\eta u_2^\theta (e/M)^\lambda v^\kappa$

ProjHash(hp, (\mathcal{L} , (ek, M')), \mathcal{C} , r): $H' = \text{hp}^r$

It has been known to be correct, smooth and pseudo-random since [GL03].

El Gamal Encryption and SPHF. We consider El Gamal encryption as described in Sect. 2. The public key is denoted by $\text{pk} = h = g^\alpha$ and the private key by $\text{sk} = \alpha \xleftarrow{\$} \mathbb{Z}_p$. The public parameters $(\mathbb{G}, p, g, \text{pk})$ are given as a common reference string. We denote the ciphertext of a message $M \in \mathbb{G}$ with the scalar $r \in \mathbb{Z}_p$ by $\mathcal{C} = \text{EG}_{\text{pk}}(M; r) = (e, u) = (h^r M, g^r)$.

The regular SPHF used throughout the literature [CS02] on the language $\{\mathcal{C} = (e, u) \in \mathbb{G}^2 \mid \exists r \in \mathbb{Z}_p \text{ such that } \mathcal{C} = \text{EG}_{\text{pk}}(M; r)\}$ of the valid encryptions of M , with r being the witness of the word \mathcal{C} , usually allows a server to check whether the ciphertext was honestly computed on the value M by a client, by generating a pair of keys (hk, hp) and sending hp to the client.

Here, on the contrary, we now want a client to implicitly check that a server knows the decryption key of the encryption scheme. This means that the client computes both the ciphertext (or the ciphertext is publicly available in a database, as here) and the pair of keys (hk, hp) and sends the ciphertext as well as hp to the server, which now uses the decryption key as a witness. This implies the following modifications to the algorithms of the SPHF:

Setup($1^{\mathbb{R}}$): param = (pk, M)

$$\mathcal{L} = \{\mathcal{C} = (e, u) \in \mathbb{G}^2 \mid \exists \alpha \in \mathbb{Z}_p \text{ such that} \\ h = g^\alpha \text{ and } e/u^\alpha = M\}$$

HashKG(\mathcal{L} , (pk, M)): hk = (λ, μ) $\xleftarrow{\$} \mathbb{Z}_p^2$

ProjKG(hk, (\mathcal{L} , (pk, M)), \mathcal{C}): hp = $u^\lambda g^\mu \in \mathbb{G}$

Hash(hk, (\mathcal{L} , (pk, M)), \mathcal{C}): $H = h^\mu (e/M)^\lambda$

ProjHash(hp, (\mathcal{L} , (pk, M')), \mathcal{C} , α): $H' = \text{hp}^\alpha$

and we show that this SPHF satisfies the usual properties:

- *Correctness*: if $\mathcal{C} \in \mathcal{L}$ with witness α , one directly gets $H = H'$;
- *Smoothness*: assume $\mathcal{C} \notin \mathcal{L}$. It can then be parsed as $\mathcal{C} = (e, u)$ with $u = g^r$ and $e = h^r M' = g^{\alpha r} M g^{\delta_M}$ (with $M' \neq M$ and thus $\delta_M \neq 0$), so that $\text{hp} = u^\lambda g^\mu = g^{r\lambda + \mu}$ and $H = h^\mu (e/M)^\lambda = g^{\alpha\mu} g^{(\alpha r + \delta_M)\lambda}$.

The smoothness is then easy to see by considering the following equality of matrices:

$$\begin{pmatrix} \log_g(\text{hp}) \\ \log_g(H) \end{pmatrix} = \begin{pmatrix} r & 1 \\ \alpha r + \delta_M & \alpha \end{pmatrix} \cdot \begin{pmatrix} \lambda \\ \mu \end{pmatrix}$$

which shows that as soon as the word is not a valid encryption of M , and so δ_M is not equal to 0, the Hash value H is not in the span of the projection key hp so that the two distributions described in Sect. 2 are indistinguishable.

- *Pseudo-Randomness.* Since El Gamal encryption is IND-CPA, it is impossible to distinguish between a real encryption and a random value without knowing the decryption key. Since the decryption key is the witness of the SPHF, without knowing the witness, the two distributions remain indistinguishable.

Main Idea of the Construction. Again, we denote the client U as S_0 and its password π as π_0 . In our setting, we denote by $\text{pw}_k = g^{\pi k}$ for all k . The database contains El Gamal encryptions of each ciphertext pw_{U_i} , under randomness s_{U_i} : $\mathcal{C}_{U_i}^{db} = \text{EG}_{\text{pk}}(\text{pw}_{U_i}; s_{U_i}) = (h^{s_{U_i}} \text{pw}_{U_i}, g^{s_{U_i}})$, so that here, $\mathcal{C}_U^{db} = \text{EG}_{\text{pk}}(\text{pw}_U; s_U) = (h^{s_U} \text{pw}_U, g^{s_U})$. The client computes a Cramer-Shoup encryption of its password: $\mathcal{C}_0 = \text{CS}_{\text{ek}}(\text{pw}_0; r_0) = (u_1, u_2, e, v)$ with $v = cd^{\xi}$. The execution of the protocol should succeed if these encryptions are correct and $\text{pw}_0 = \text{pw}_U$. Recall that the server \mathcal{S}_i knows α_i such that $\alpha = \alpha_1 + \alpha_2$ is the decryption key of the El Gamal encryption.

The main idea is depicted on Fig. 2. For sake of readability, the participants which have a real role in the computations are directly linked by arrows in the picture, but one should keep in mind that all the participants (U , S_1 and S_2) only communicate with G , which then broadcasts all the messages.

In a classical SPHF-based two-party key-exchange between U and G , the gateway would check whether \mathcal{C}_0 is a valid Cramer-Shoup encryption of pw_U . Since here the password pw_U is unknown to the servers S_1 and S_2 , this is done in our setting by two SPHF, using hp_1^{CS} (sent by S_1) and hp_2^{CS} (sent by S_2), where the servers use the first term of the public encryption $\mathcal{C}_U^{\text{DB}}(h^{s_U} \text{pw}_U)$ in order to cancel the unknown pw_U .

In a classical SPHF-based two-party key-exchange between U and G , the client would also check whether $\mathcal{C}_U^{\text{DB}}$ is a valid El Gamal encryption of its password pw_0 , *i.e.* whether the gateway knows a witness for its ciphertext $\mathcal{C}_U^{\text{DB}}$ (s_U in the usual constructions, α here). Since α is unknown to the gateway, this is done in our setting by the combination of three SPHF, using hp_0^{EG} (sent by the client), hp_1^{EG} (sent by S_1) and hp_2^{EG} (sent by S_2). These three SPHF allow the client and the servers to implicitly check that the servers know α_1 and α_2 such that $\mathcal{C}_U^{\text{DB}}$ can be decrypted (using the decryption key $\alpha = \alpha_1 + \alpha_2$) to the same password pw_0 than the one encrypted in \mathcal{C}_0 sent by the client. Formally, the languages checked are as follows:

- by the client:
 $\mathcal{C}_U^{\text{DB}} \in \mathcal{L}_0 = \{\mathcal{C} = (e, u) \in \mathbb{G}^2 \mid \exists \alpha \in \mathbb{Z}_p \text{ such that } h = g^\alpha \text{ and } e/u^\alpha = \text{pw}_0\}$
- by server \mathcal{S}_i (with respect to the client \mathcal{S}_0 and server \mathcal{S}_j):
 $\mathcal{C}_0 \in \mathcal{L}_{i,0} = \{\mathcal{C} = (u_1, u_2, e, v) \in \mathbb{G}^4 \mid \exists r \in \mathbb{Z}_p \text{ such that } \mathcal{C} = \text{CS}_{\text{ek}}(\text{pw}_U; r)$
and $\mathcal{C}_U^{\text{DB}} \in \mathcal{L}_{i,j} = \{\mathcal{C} = (e, u) \in \mathbb{G}^2 \mid \exists \alpha_j \in \mathbb{Z}_p \text{ such that } h = g^{\alpha_i + \alpha_j}\}$ and $e/u^{\alpha_i + \alpha_j} = \text{pw}_U\}$

but they cannot be checked directly by a unique SPHF since the value pw_U appearing in the languages is unknown to the verifier \mathcal{S}_i . Rather, the server \mathcal{S}_i will use the first term of the public encryption $\mathcal{C}_U^{\text{DB}}(h^{s_U} \text{pw}_U)$ in order to cancel this unknown pw_U . To achieve this goal, we combine the five SPHF described

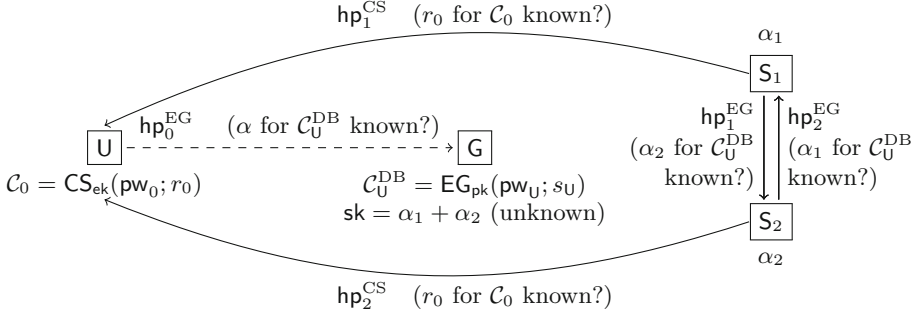


Fig. 2. Main idea of the construction

to globally ensure the correctness (each one remaining smooth and pseudo-randomness), as described in the next part.

5.2 Login Procedure

- The client $U = S_0$ chooses r_0 at random and computes a Cramer-Shoup encryption of its password $C_0 = CS_{ek}(pw_0; r_0) = (u_1, u_2, e, v)$ with $v = cd^\xi$. It also chooses a random (El Gamal) hash key $hk_0^{EG} = (\lambda_0, \mu_0)$ and computes the corresponding projected key on the ciphertext $C_U^{DB} = EG_{pk}(pw_U; s_U) = (h^{s_U} pw_U, g^{s_U})$ contained in the database: $hp_0^{EG} = g^{s_U \lambda_0} g^{\mu_0}$. It sends (C_0, hp_0^{EG}) to the gateway, which broadcasts these values to the servers.
- After receiving this flow from the client, the servers S_1 and S_2 also choose a random (El Gamal) hash key $hk_b^{EG} = (\lambda_b, \mu_b)$ and compute the corresponding projected key on the ciphertext C_U^{DB} contained in the database: $hp_b^{EG} = g^{s_U \lambda_b} g^{\mu_b}$. The servers S_1 and S_2 also choose a random (Cramer-Shoup) hash key $hk_b^{CS} = (\eta_b, \theta_b, \lambda_b, \kappa_b)$ (with the same value λ_b) and compute the corresponding projected key on the ciphertext C_0 sent by the client: $hp_b^{CS} = (g_1^{\eta_b} g_2^{\theta_b} h^{\lambda_b} (cd^\xi)^{\kappa_b})$. Each server sends (hp_b^{EG}, hp_b^{CS}) to the gateway G , which broadcasts these values to the other participants.
- After receiving the projected keys from the servers, the client computes, for $i \in \{1, 2\}$, $H'_{i,0} = (hp_i^{CS})^{r_0}$. It also computes $H_0 = (h^{s_U} pw_U / pw_0)^{\lambda_0} h^{\mu_0}$ by dividing the first term of the ciphertext C_U^{DB} contained in the database by its password pw_0 . It finally sets its session key K_U as $K_U = K_0 = H'_{1,0} \cdot H'_{2,0} \cdot H_0$. After receiving the first flow from the other server and the client, the server S_b computes, for $i \in \{0, 3-b\}$, $H'_{i,b} = (hp_i^{EG})^{\alpha_b}$. It also computes $H_{b,0} = (u_1)^{\eta_b} (u_2)^{\theta_b} [e / (h^{s_U} pw_U)]^{\lambda_b} v^{\kappa_b}$ and $H_b = H_{b,0} \cdot (hp_b^{EG})^{\alpha_b} / h^{\mu_b}$, and sets its partial key K_b as $K_b = H'_{0,b} \cdot H'_{3-b,b} \cdot H_b$. It privately sends this value K_b to the gateway G .
- The gateway finally sets $K_G = K_1 \cdot K_2$.

Correctness. Due to the correctness of the Cramer-Shoup SPHF, we have the equalities $H'_{b,0} = H_{b,0} \left(\frac{h^{s_U \text{pw}_U}}{\text{pw}_0} \right)^{\lambda_b}$ for $b \in \{1, 2\}$. The session key of the gateway is thus equal to $K_G = K_1 \cdot K_2 = (H'_{0,1} \cdot H'_{2,1} \cdot H_1) \cdot (H'_{0,2} \cdot H'_{1,2} \cdot H_2)$ where, after computation, $K_b = h^{(\lambda_0 + \lambda_1 + \lambda_2)\alpha_b} g^{(\mu_0 + \mu_1 + \mu_2)\alpha_b} H'_{b,0} (\text{pw}_0 / \text{pw}_U)^{\lambda_b} h^{-s_U \lambda_b} h^{-\mu_b}$.

If $\text{pw}_0 = \text{pw}_U$, $h = g^\alpha$ and $\alpha = \alpha_1 + \alpha_2$, $K_G = h^{s_U \lambda_0} h^{\mu_0} H'_{1,0} H'_{2,0}$, which is equal to the session key of the client $K_U = K_0 = H'_{1,0} \cdot H'_{2,0} \cdot H_0 = H'_{1,0} \cdot H'_{2,0} \cdot h^{s_U \lambda_0} h^{\mu_0}$.

Complexity. The following table sums up the number of group elements needed for each participant. It should be noted that in case one of the servers is the gateway, its communication costs are reduced.

	Ciphertext (Broadcast)	Projection Keys (Overall)	To the Gateway
Client	4	2	0
Server (each)	0	2	1

Compared to [KMTG12], the communication complexity of our protocol is decreased by more than 50% (9 group elements instead of 20 group elements). For efficiency, as in [KMTG12], we count exponentiations only, and assume a multi-exponentiation with up to 5 bases can be computed at the cost of at most 1.5 exponentiations. The client performs the equivalent of 8 full exponentiations, while each server performs 7 exponentiations (instead of 15 and 13 respectively in [KMTG12]).

Proof of Security. The proof follows the idea of the former one. For lack of space, a sketch is given in [BCV16].

6 Conclusion

We presented two constructions of distributed Password-Authenticated Key Exchange between a user and several servers. We focused on presenting them in a classical group setting (with only two servers). Very efficient implementations of our protocols can be readily obtained using standard cryptographic libraries and do not require pairings.

Our methods can be generalized to the setting where n servers share the (encryption of the) password. SPHF can further handle polynomials of variables and the use of secret sharing techniques *à la* Shamir, allows to share polynomials evaluation between n servers and to provide a threshold distributed PAKE such that security is ensured as long as less than a certain arbitrary threshold $t \in \{1, \dots, n\}$ of servers are compromised (contrary to the protocol from [DG06] which requires an honest majority of the servers).

Smooth projective hashing is mostly used in a classical discrete-logarithm-based setting (or pairing-based setting) but constructions were also proposed for Paillier encryption [CS02] and for LWE encryption [KV09]. These SPHF would allow a readily adaptation of our techniques to other classical settings of cryptography.

Acknowledgements. This work was supported in part by the French ANR Project ANR-14-CE28-0003 EnBiD.

References

- [ACFP05] Abdalla, M., Chevassut, O., Fouque, P.-A., Pointcheval, D.: A simple threshold authenticated key exchange from short secrets. In: Roy, B. (ed.) ASIACRYPT 2005. LNCS, vol. 3788, pp. 566–584. Springer, Heidelberg (2005)
- [ACP09] Abdalla, M., Chevalier, C., Pointcheval, D.: Smooth projective hashing for conditionally extractable commitments. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 671–689. Springer, Heidelberg (2009)
- [BBC+13] Benhamouda, F., Blazy, O., Chevalier, C., Pointcheval, D., Vergnaud, D.: New techniques for SPHF and efficient one-round PAKE protocols. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 449–475. Springer, Heidelberg (2013)
- [BCV16] Blazy, O., Chevalier, C., Vergnaud, D.: Mitigating server breaches in password-based authentication: secure and efficient solutions. In: Sako, K. (eds.) Topics in Cryptology, CT-RSA 2016, pp. 3–18. Springer, Heidelberg (2016)
- [BJKS03] Brainard, J.G., Juels, A., Kaliski, B., Szydlo, M.: A new two-server approach for authentication with short secrets. In: Proceedings of the 12th USENIX Security Symposium, Washington, D.C., USA, 4–8 August 2003 (2003)
- [BM92] Bellare, S.M., Merritt, M.: Encrypted key exchange: password-based protocols secure against dictionary attacks. In: 1992 IEEE Symposium on Security and Privacy, pp. 72–84. IEEE Computer Society Press, May 1992
- [Bon98] Boneh, D.: The decision Diffie-Hellman problem. In: Buhler, J.P. (ed.) ANTS 1998. LNCS, vol. 1423, pp. 48–63. Springer, Heidelberg (1998)
- [BPR00] 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)
- [Can01] Canetti, R.: Universally composable security: a new paradigm for cryptographic protocols. In: 42nd FOCS, pp. 136–145. IEEE Computer Society Press, October 2001
- [CEN15] Camenisch, J., Enderlein, R.R., Neven, G.: Two-server password-authenticated secret sharing UC-secure against transient corruptions. In: Katz, J. (ed.) PKC 2015. LNCS, vol. 9020, pp. 283–307. Springer, Heidelberg (2015)
- [CS98] Cramer, R., Shoup, V.: A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 13–25. Springer, Heidelberg (1998)
- [CS02] 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)
- [DG06] Di Raimondo, M., Gennaro, R.: Provably secure threshold password-authenticated key exchange. *J. Comput. Syst. Sci.* **72**(6), 978–1001 (2006)
- [DH76] Diffie, W., Hellman, M.E.: New directions in cryptography. *IEEE Trans. Inf. Theo.* **22**(6), 644–654 (1976)

- [ElG84] El Gamal, T.: A public key cryptosystem and a signature scheme based on discrete logarithms. In: Blakely, G.R., Chaum, D. (eds.) CRYPTO 1984. LNCS, vol. 196, pp. 10–18. Springer, Heidelberg (1985)
- [FK00] Ford, W., Kaliski Jr., B.S.: Server-assisted generation of a strong secret from a password. In: 9th IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE 2000), Gaithersburg, MD, USA, 4–16 June 2000, pp. 176–180 (2000)
- [GL03] Gennaro, R., Lindell, Y.: A framework for password-based authenticated key exchange. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 524–543. Springer, Heidelberg (2003)
- [GM84] Goldwasser, S., Micali, S.: Probabilistic encryption. *J. Comput. Syst. Sci.* **28**(2), 270–299 (1984)
- [HJKY95] Herzberg, A., Jarecki, S., Krawczyk, H., Yung, M.: Proactive secret sharing or: how to cope with perpetual leakage. In: Coppersmith, D. (ed.) CRYPTO 1995. LNCS, vol. 963, pp. 339–352. Springer, Heidelberg (1995)
- [Jab01] Jablon, D.P.: Password authentication using multiple servers. In: Naccache, D. (ed.) CT-RSA 2001. LNCS, vol. 2020, pp. 344–360. Springer, Heidelberg (2001)
- [KM14] Kiefer, F., Manulis, M.: Distributed smooth projective hashing and its application to two-server password authenticated key exchange. In: Boureanu, I., Owesarski, P., Vaudenay, S. (eds.) ACNS 2014. LNCS, vol. 8479, pp. 199–216. Springer, Heidelberg (2014)
- [KMTG12] Katz, J., MacKenzie, P.D., Taban, G., Virgil, D.: Two-server password-only authenticated key exchange. *J. Comput. Syst. Sci.* **78**(2), 651–669 (2012)
- [KOY01] Katz, J., Ostrovsky, R., Yung, M.: Efficient password-authenticated key exchange using human-memorable passwords. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 475–494. Springer, Heidelberg (2001)
- [KV09] Katz, J., Vaikuntanathan, V.: Smooth projective hashing and password-based authenticated key exchange from lattices. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 636–652. Springer, Heidelberg (2009)
- [KV11] Katz, J., Vaikuntanathan, V.: Round-optimal password-based authenticated key exchange. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 293–310. Springer, Heidelberg (2011)
- [MSJ02] MacKenzie, P.D., Shrimpton, T., Jakobsson, M.: Threshold password-authenticated key exchange. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 385–400. Springer, Heidelberg (2002)
- [OY91] Ostrovsky, R., Yung, M.: How to withstand mobile virus attacks (extended abstract). In: 10th ACM PODC, pp. 51–59. ACM, August 1991
- [Poi12] Pointcheval, D.: Password-based authenticated key exchange. In: Fischlin, M., Buchmann, J., Manulis, M. (eds.) PKC 2012. LNCS, vol. 7293, pp. 390–397. Springer, Heidelberg (2012)
- [SK05] Szydło, M., Kaliski, B.: Proofs for two-server password authentication. In: Menezes, A. (ed.) CT-RSA 2005. LNCS, vol. 3376, pp. 227–244. Springer, Heidelberg (2005)