Round-Optimal Password-Based Authenticated Key Exchange

Jonathan Katz* Vinod Vaikuntanathan[†]

Abstract

We show a general framework for constructing password-based authenticated key exchange protocols with *optimal* round complexity — one message per party, sent simultaneously — in the standard model, assuming the existence of a common reference string. When our framework is instantiated using bilinear-map based cryptosystems, the resulting protocol is also (reasonably) efficient. Somewhat surprisingly, our framework can be adapted to give protocols in the standard model that are *universally composable* while still using only one (simultaneous) round.

1 Password-Based Authenticated Key Exchange

Protocols for *authenticated key exchange* enable two parties to generate a shared, cryptographically strong key while communicating over an insecure network under the complete control of an adversary. Such protocols are among the most widely used and fundamental cryptographic primitives; indeed, agreement on a shared key is necessary before "higher-level" tasks such as encryption and message authentication become possible.

Parties must share *some* information in order for authenticated key exchange to be possible. It is well known that shared cryptographic keys — either in the form of public keys or a long, uniformly random symmetric key — suffice, and several protocols in this model, building on the classic Diffie-Hellman protocol [18] (which protects only against an eavesdropping adversary and provides no authentication at all) are known; see, e.g., [7, 4].

Password-based protocols allow users to "bootstrap" even a weak (e.g., short) shared secret into a (much longer) cryptographic key. The canonical application here is authentication using passwords, though protocols developed in this context can be useful even when the shared secret has high min-entropy (but is not uniform) [11]. The security guaranteed by password-based protocols (roughly speaking) is that if the password is chosen uniformly¹ from a dictionary of size D then an adversary who initiates Q "on-line" attacks — i.e., who actively interferes in Q sessions — has "advantage" at most Q/D. (This is inherent, as an adversary can always carry out Q impersonation attempts and succeed with this probability.) In particular, "off-line" dictionary attacks where an

^{*}Dept. of Computer Science, University of Maryland. Work done while visiting IBM. Email: jkatz@cs.umd.edu. Research supported by NSF grant #0627306 and NSF CAREER award #0447075.

[†]Dept. of Computer Science, University of Toronto. Work done while at IBM and Microsoft Research. Email: vinodv@cs.toronto.edu.

¹Although the usual presentation of PAK assumes a uniform password, known protocols work with passwords chosen from any (efficiently sampleable) distribution.

adversary enumerates passwords from the dictionary of potential passwords, and tries to match observed protocol executions to each one, are of no use.

Early work [23, 27] considered a "hybrid" setting where users share public keys in addition to a password. In the setting where *only* a password is shared, Bellovin and Merritt [6] proposed the first protocols for password-based authenticated key exchange (PAK) with heuristic arguments for their security. Several years later, provably secure PAK protocols were constructed [3, 12, 33] in the random oracle/ideal cipher models, and many improvements and generalizations of these protocols are known. In contrast, only a handful of PAK protocols are known in the so-called "standard model" (i.e., without random oracles):

- General assumptions: Goldreich and Lindell [22] gave the first PAK protocol in the standard model (and without requiring any additional setup). Their work does not handle concurrent executions of the protocol by the same party, but Goyal, Jain, and Ostrovsky [24] recently showed a new protocol without this restriction. Work of Barak et al. [2] shows a general feasibility result for computation over unauthenticated networks which implies a solution for PAK as a special case, assuming a common reference string (CRS) is available to the parties.² All of these approaches are impractical in terms of communication, computation, and round complexity. (Nguyen and Vadhan [35] show some efficiency improvements, but achieve a weaker notion of security. Their protocol is still far from practical.)
- Efficient protocols: Katz, Ostrovsky, and Yung [31] demonstrated the first *efficient* PAK protocol with a proof of security based on standard assumptions; extensions and improvements of their protocol were given in [21, 15, 30, 20, 32]. Different constructions of efficient PAK protocols were given in [29, 25]. These works all require a CRS.

Round/message complexity of existing protocols. We distinguish between *rounds* and *messages*. Differing somewhat from the usual convention in the two-party setting (but matching the usual convention in the multi-party setting), we let a round consist of one message sent by each party simultaneously; note that in a one-round protocol each honest party's message cannot depend on the other party's message. We stress, however, that the adversary is always assumed to be *rushing*; i.e., the adversary may wait to receive an honest party's first-round message before sending its own.

Determining the optimal round complexity of key-exchange protocols is of both theoretical and practical interest, and has been studied in various settings. The original Diffie-Hellman protocol [18], which provides security against a passive eavesdropper, can be run in one round; oneround authenticated key exchange based on shared public/symmetric keys is also possible [28, 36]. One-round PAK protocols are also known (e.g., [3]) in the random oracle model. All prior PAK protocols based on standard assumptions, though, require three or more rounds. We remark that the protocols in [29, 25] achieve *explicit* authentication in three rounds (whereas the protocols of [31, 21, 20, 32] achieve only *implicit* authentication in three rounds, and require an additional round for explicit authentication), but the round complexity of these protocols cannot be further reduced even if only implicit authentication is desired.

 $^{^{2}}$ In practice, a CRS is not a serious drawback in the context of PAK where the CRS can be hard-coded into an implementation of the protocol. We note also that reliance on a CRS (or some other setup) is inherent for achieving *universally composable* PAK [15].

1.1 Our Results

We show a new framework for constructing *one-round* PAK protocols in the standard model (assuming a CRS), where each party may send their message *simultaneously*. (Once again, we stress that our security model allows for a "rushing" adversary who waits to see the message sent by a party before sending its response.) Our protocols achieve implicit authentication but can be extended to give explicit authentication using one additional round; it is not hard to see that explicit authentication is impossible in one round without stronger setup assumptions (e.g., a global clock).

Our framework relies on non-interactive zero-knowledge (NIZK) proofs and so, in general, is computationally inefficient. When instantiating our framework using bilinear maps, however, we obtain a reasonably efficient solution communicating a constant number of group elements.

Somewhat surprisingly, we can extend our framework to give a universally composable PAK protocol [14] without increasing the round complexity at all (and still without relying on random oracles). In contrast, the work of [15] shows a method for obtaining universal composability (used also in [25]) that requires additional messages/rounds. Abdalla et al. [1] show a universally composable PAK protocol, proven secure in the random oracle model, that requires three rounds. To the best of our knowledge, no prior universally composable PAK protocol (whether in the random oracle model or not) can be run in only one round.

1.2 Our Techniques

At a basic level we rely on smooth projective hash functions [16], as used for PAK in [21] (and implicitly in [31]); see Section 2.2 for a definition. The basic structure of previous protocols [31, 21], omitting many important details, is as follows:

First round: The client sends an encryption C of the password pw.

Second round: The server sends an encryption C' of pw, and a projected key $s' = \alpha(k', C, pw)$ corresponding to the hash key k'.

Third round: The client sends a projected key $s = \alpha(k, C', pw)$ corresponding to the hash key k.

The client computes the session key as $H_k(C', pw) \cdot H_{s'}(C, pw, r)$, and the server computes the session key as $H_s(C', pw, r') \cdot H_{k'}(C, pw)$. (Here, r, r' is the randomness used to compute C, C', respectively.) Properties of the smooth projective hash function ensure that these are equal.

Two difficulties must be overcome in order to collapse a protocol of the above form to one round:

• In the smooth projective hash functions used in prior work, the "projection function" α was *adaptive*, and depended on both the hash key k and the element being hashed (i.e., (C, pw) in the above example). This leads to protocols requiring three rounds just to ensure correctness.

Here we show a construction of CCA-secure encryption schemes with associated smooth projective hash functions whose projection function is *non-adaptive*, and depends only on the hash key k. This allows us to obtain the *functionality* of PAK in a single round, by having the client send $(\alpha(k), C)$ and the server send $(\alpha(k'), C')$ simultaneously.

• The above addresses correctness, but says nothing about security. The main technical difficulty is that an honestly generated client message msg = (s, C) might be forwarded by an adversary to *multiple* server instances (and similarly for server messages), and the session keys computed in all these instances should look random and independent to the adversary. (This issue does not arise in prior work because, roughly speaking, messages are *bound* to a single session by virtue of a signature verification key sent in the first round [31, 21] or a MAC derived from the shared session key [20]. Neither approach is viable if we want the entire protocol to take place in a single round.)

To address the above difficulty, we rely on a technical lemma (which may be of independent interest) regarding the re-use of both the hash keys and the inputs to the smooth projective hash function.

Additional ideas are needed to obtain a *universally composable* protocol without increasing the number of rounds. We refer the reader to Section 5.1 for an overview of the techniques used there.

1.3 Outline of the Paper

In Section 2 we present a standard definition of security for PAK due to Bellare et al. [3]. We also review there the notion of smooth projective hashing, and prove a technical lemma regarding its usage. In Section 3 we describe our basic framework for constructing one-round PAK protocols, and prove security of this approach according to the definition of [3]. We discuss in Section 4 two instantiations of our framework: one based on the decisional Diffie-Hellman assumption, and a second, more efficient instantiation based on bilinear maps. In Section 5 we describe an extension of our framework that yields one-round, universally composable PAK protocols.

2 Definitions and Background

Throughout, we denote the security parameter by n.

2.1 Password-Based Authenticated Key Exchange

Here we present a definition of security for PAK due to Bellare, Pointcheval, and Rogaway [3], based on prior work of [4, 5]. The text here is taken almost verbatim from [31].

Participants, passwords, and initialization. Prior to any execution of the protocol there is an initialization phase during which public parameters and a CRS are established. We assume a fixed set User of protocol participants (also called principals or users). For every distinct $U, U' \in$ User, users U and U' share a password $pw_{U,U'}$. We assume that each $pw_{U,U'}$ is chosen independently and uniformly from the set $[D] \stackrel{\text{def}}{=} \{1, \ldots, D\}$ for some integer D. (Our proof of security extends to more general cases, and we implicitly consider arbitrary password distributions in the setting of universal composability.)

Execution of the protocol. In the real world, a protocol determines how principals behave in response to input from their environment. In the formal model, these inputs are provided by the adversary. Each principal can execute the protocol multiple times (possibly concurrently) with different partners; this is modeled by allowing each principal to have an unlimited number of *instances* with which to execute the protocol. We denote instance i of user U as Π_U^i . Each instance may be used only once. The adversary is given oracle access to these different instances; furthermore, each instance maintains (local) state which is updated during the course of the experiment. In particular, each instance Π_U^i is associated with the following variables:

- $\operatorname{sid}_{U}^{i}$, $\operatorname{pid}_{U}^{i}$, and $\operatorname{sk}_{U}^{i}$ denote the session *id*, partner *id*, and session key for an instance, respectively. The session *id* is simply a way to keep track of different executions; we let $\operatorname{sid}_{U}^{i}$ be the (ordered) concatenation of all messages sent and received by Π_{U}^{i} . The partner *id* denotes the user with whom Π_{U}^{i} believes it is interacting. (Note that $\operatorname{pid}_{U}^{i}$ can never equal U.)
- $\operatorname{acc}_{U}^{i}$ and $\operatorname{term}_{U}^{i}$ are boolean variables denoting whether a given instance has accepted or terminated, respectively.

The adversary's interaction with the principals (more specifically, with the various instances) is modeled via access to *oracles* that we describe now:

• Send(U, i, msg) — This sends message msg to instance Π_U^i . This instance runs according to the protocol specification, updating state as appropriate. The message output by Π_U^i is given to the adversary.

The adversary can "prompt" instance Π^i_U to initiate the protocol with partner U' by querying Send(U, i, U'). In response to this query, instance Π^i_U outputs the first message of the protocol.

- Execute(U, i, U', j) If Π_U^i and $\Pi_{U'}^j$ have not yet been used, this oracle executes the protocol between these instances and gives the transcript of this execution to the adversary. This oracle call represents passive eavesdropping of a protocol execution.
- Reveal(U, i) This outputs the session key skⁱ_U, modeling leakage of session keys due to, e.g., improper erasure of session keys after use, compromise of a host computer, or cryptanalysis.
- Test(U, i) This oracle does not model any real-world capability of the adversary, but is instead used to define security. A random bit b is chosen; if b = 1 the adversary is given sk_U^i , and if b = 0 the adversary is given a session key chosen uniformly from the appropriate space.

Partnering. Let $U, U' \in \text{User}$. Instances Π_U^i and $\Pi_{U'}^j$ are partnered if: (1) $\text{sid}_U^i = \text{sid}_{U'}^j \neq \text{NULL}$; and (2) $\text{pid}_U^i = U'$ and $\text{pid}_{U'}^j = U$.

Correctness. To be viable, a key-exchange protocol must satisfy the following notion of correctness: if Π^i_U and $\Pi^j_{U'}$ are partnered then $\operatorname{acc}^i_U = \operatorname{acc}^j_{U'} = \operatorname{TRUE}$ and $\operatorname{sk}^i_U = \operatorname{sk}^j_{U'}$, i.e., they both accept and conclude with the same session key.

Advantage of the adversary. Informally, the adversary succeeds if it can guess the bit b used by the Test oracle. To formally define the adversary's success, we first define a notion of *freshness*. An instance Π_U^i is *fresh* unless one of the following is true at the conclusion of the experiment: (1) at some point, the adversary queried Reveal(U, i); or (2) at some point, the adversary queried Reveal(U', j), where Π_U^j and Π_U^i are partnered. We allow the adversary to succeed only if its Test query is made to a fresh instance; this is necessary for any reasonable definition of security.

An adversary \mathcal{A} succeeds if it makes a single query $\mathsf{Test}(U, i)$ to a fresh instance Π_U^i , and outputs a bit b' with b' = b (recall that b is the bit chosen by the **Test** oracle). We denote this event by Succ. The *advantage* of \mathcal{A} in attacking protocol Π is given by $\mathsf{Adv}_{\mathcal{A},\Pi}(k) \stackrel{\text{def}}{=} 2 \cdot \Pr[\mathsf{Succ}] - 1$, where the probability is taken over the random coins used by the adversary and the random coins used during the course of the experiment (including the initialization phase).

It remains to define a secure protocol. A probabilistic polynomial-time (PPT) adversary can always succeed with probability 1 by trying all passwords one-by-one; this is possible since the size of the password dictionary is small. Informally, a protocol is secure if this is the best an adversary can do. Formally, an instance Π^i_U represents an *on-line attack* if both the following are true at the time of the Test query: (1) at some point, the adversary queried Send(U, i, *); and (2) at some point, the adversary queried Reveal(U, i) or Test(U, i). The number of on-line attacks represents a bound on the number of passwords the adversary could have tested in an on-line fashion.

Definition 1. Protocol Π is a secure protocol for password-based authenticated key exchange *if*, for all dictionary sizes D and for all PPT adversaries \mathcal{A} making at most Q(n) on-line attacks, it holds that $\mathsf{Adv}_{\mathcal{A},\Pi}(n) \leq Q(n)/D + \mathsf{negl}(n)$.

2.2 Smooth Projective Hash Functions

We provide a self-contained definitional treatment of smooth projective hash functions. These were introduced by Cramer and Shoup [16], and our discussion here is based on that of Gennaro and Lindell [21]. Rather than aiming for utmost generality, we tailor the definitions to our application.

A hard subset-membership problem. Fix some integer D. Let (Gen, Enc, Dec) be a CCAsecure labeled encryption scheme (cf. Appendix A.1). For a given public key pk, we let C_{pk} denote the set of pairs of valid labels and ciphertexts with respect to pk, and require that this set be efficiently recognizable. For a given public key pk, define sets X and $\{L_{pw}\}_{pw \in [D]}$ as follows:

1. $X \stackrel{\text{def}}{=} \{ (\mathsf{label}, C, pw) \}$, where $(\mathsf{label}, C) \in C_{pk}$ and $pw \in \{1, \dots, D\}$.

2. $L_{pw} \stackrel{\text{def}}{=} \{ (\text{label}, \text{Enc}_{pk}(\text{label}, pw), pw) \}, \text{ where } \text{label} \in \{0, 1\}^*.$

That is, X consists of all tuples of labels, ciphertexts, and passwords, while L_{pw} consists of all tuples where the third component is pw and the second component is an encryption of pw. Let $L = \bigcup_{pw=1}^{D} L_{pw}$, and note that $L \subset X$. It follows from CCA security of (Gen, Enc, Dec) that the following is negligible for any polynomial-time \mathcal{A} :

$$\begin{vmatrix} (pk, sk) \leftarrow \mathsf{Gen}(1^n); \\ (\mathsf{label}, pw) \leftarrow \mathcal{A}^{\mathsf{Dec}_{sk}(\cdot, \cdot)}(pk); \\ C \leftarrow \mathsf{Enc}_{pk}(\mathsf{label}, pw) \end{vmatrix} : \mathcal{A}^{\mathsf{Dec}_{sk}(\cdot, \cdot)}(C) = 1 \\ \end{vmatrix} \\ - \Pr \begin{bmatrix} (pk, sk) \leftarrow \mathsf{Gen}(1^n); \\ (\mathsf{label}, pw) \leftarrow \mathcal{A}^{\mathsf{Dec}_{sk}(\cdot, \cdot)}(pk); \\ C \leftarrow \mathsf{Enc}_{pk}(\mathsf{label}, 0) \end{vmatrix} : \mathcal{A}^{\mathsf{Dec}_{sk}(\cdot, \cdot)}(C) = 1 \\ \end{vmatrix}$$

where \mathcal{A} is disallowed from querying (label, C) to its decryption oracle.

Smooth projective hash functions. Fix pk and sets X, L as above. A smooth projective hash function $\mathcal{H} = \{H_k\}_{k \in K}$ is a keyed function mapping elements in X to elements in some group G, along with a projection function $\alpha : K \to S$. Informally, if $x \in L$ then the value of $H_k(x)$ is uniquely determined by $s = \alpha(k)$ and x, whereas if $x \in X \setminus L$ then the value of $H_k(x)$ is statistically close to uniform given $\alpha(k)$ and x (assuming k was chosen uniformly in K). A smooth projective hash function is formally defined by a sampling algorithm that, given pk, outputs $(K, G, \mathcal{H} = \{H_k : X \to G\}_{k \in K}, S, \alpha : K \to S)$ such that:

1. There are efficient algorithms for (1) sampling a uniform $k \in K$, (2) computing $H_k(x)$ for any $k \in K$ and $x \in X$, and (3) computing $\alpha(k)$ for $k \in K$.

- 2. For all (label, C, pw) $\in L$, the value of $H_k(label, C, pw)$ is uniquely determined by $\alpha(k)$. Moreover, there is an efficient algorithm that takes as input $s = \alpha(k)$ and (label, C, pw, r) for which $C = \text{Enc}_{pk}(label, pw; r)$, and outputs $H_k(label, C, pw)$. (In other words, when $(label, C, pw) \in L$ then $H_k(label, C, pw)$ can be computed in two ways: either using k itself, or using $\alpha(k)$ and the randomness used to generate C.)
- 3. For any (even unbounded) function $f: S \to X \setminus L$, the following distributions have statistical difference negligible in n:

$$\left\{k \leftarrow K; s := \alpha(k) : \left(s, H_k(f(s))\right)\right\} \text{ and } \left\{k \leftarrow K; s := \alpha(k); g \leftarrow G : (s, g)\right\}.$$
(1)

Above we have modified the definition from [21] in two ways: first, α is *non-adaptive* and depends on k only (rather than both k and x); second, we require Equation (1) to hold even for *adaptive* choice of $f(s) \notin L$. The first modification allows us to achieve the functionality of password-based authenticated key exchange in one round, whereas the second is used for proving security of the resulting protocol.

A technical lemma. We now prove a technical lemma regarding smooth projective hash functions. Somewhat informally, Gennaro and Lindell [21] showed that, for randomly generated pk and any (label, pw), the distribution

$$\left\{k \leftarrow K; s := \alpha(k); C \leftarrow \mathsf{Enc}_{pk}(\mathsf{label}, pw) : (s, C, H_k(\mathsf{label}, C, pw))\right\}$$

is computationally indistinguishable from the distribution

$$\{k \leftarrow K; s := \alpha(k); C \leftarrow \mathsf{Enc}_{pk}(\mathsf{label}, pw); g \leftarrow G : (s, C, g)\}.$$

(Note this holds even though $H_k(\mathsf{label}, C, pw)$ is uniquely determined by s and C.) Here we show that this continues to hold even if hash keys and ciphertexts are *re-used* multiple times. That is, at a high level (ignoring labels and technical details), we show that the distribution

$$\left\{ \begin{array}{l} k_1, \dots, k_\ell \leftarrow K; \forall i: s_i := \alpha(k_i); \\ C_1, \dots, C_\ell \leftarrow \mathsf{Enc}_{pk}(pw) \end{array} : \left(\{s_i\}, \{C_i\}, \{H_{k_i}(C_j, pw)\}_{i,j=1}^\ell \right) \right\}$$

is computationally indistinguishable from the distribution

$$\left\{ \begin{array}{c} k_1, \dots, k_\ell \leftarrow K; \forall i: s_i := \alpha(k_i); \\ C_1, \dots, C_\ell \leftarrow \mathsf{Enc}_{pk}(pw); \forall i, j: g_{i,j} \leftarrow G \end{array} : \left(\{s_i\}, \{C_i\}, \{g_{i,j}\}_{i,j=1}^\ell \right) \right\}$$

Formally, fix a function $\ell = \ell(n)$, let \mathcal{A} be an adversary, and let $b \in \{0, 1\}$. Consider the following experiment Expt_b :

- 1. Compute $(pk, sk) \leftarrow \text{Gen}(1^n)$ and let $(K, G, \mathcal{H} = \{H_k : X \to G\}_{k \in K}, S, \alpha : K \to S)$ be a smooth projective hash function for pk. Give pk to \mathcal{A} .
- 2. Sample $k_1, \ldots, k_\ell \leftarrow K$, and let $s_i := \alpha(k_i)$ for all *i*. Give s_1, \ldots, s_ℓ to \mathcal{A} .
- 3. \mathcal{A} may adaptively query a (modified) encryption oracle that takes as input (label, pw) for $pw \in [D]$, and outputs a ciphertext $C \leftarrow \mathsf{Enc}_{pk}(\mathsf{label}, pw)$ along with:

- (a) If b = 0, the values $H_{k_i}(\mathsf{label}, C, pw)$ for i = 1 to ℓ .
- (b) If b = 1, random values $g_1, \ldots, g_\ell \leftarrow G$.
- 4. \mathcal{A} can also query a decryption oracle $\mathsf{Dec}_{sk}(\cdot, \cdot)$ at any point, except that it may not query any pair (label, C) where C was obtained from the encryption oracle on query (label, pw).
- 5. At the end of the experiment, \mathcal{A} outputs a bit b'. We say \mathcal{A} succeeds if b' = b.

Lemma 1. Let (Gen, Enc, Dec) be a CCA-secure labeled encryption scheme. For any polynomial ℓ and polynomial-time \mathcal{A} , we have $\Pr[\mathcal{A} \text{ succeeds}] \leq \frac{1}{2} + \operatorname{negl}(n)$.

Proof: Let ℓ' be a (polynomial) bound on the number of encryption queries asked by \mathcal{A} , and let $C_i \leftarrow \mathsf{Enc}_{pk}(\mathsf{label}_i, pw_i)$ be the ciphertext returned in response to the *i*th query of \mathcal{A} . When b = 0, the values given to \mathcal{A} include pk, s_1, \ldots, s_ℓ , the ciphertexts $C_1, \ldots, C_{\ell'}$, and the values

$$\left(\begin{array}{cccc} H_{k_1}(\mathsf{label}_1, C_1, pw_1) & \cdots & \cdots & H_{k_1}(\mathsf{label}_{\ell'}, C_{\ell'}, pw_{\ell'}) \\ \vdots & \ddots & \ddots & \vdots \\ H_{k_\ell}(\mathsf{label}_1, C_1, pw_1) & \cdots & \cdots & H_{k_\ell}(\mathsf{label}_{\ell'}, C_{\ell'}, pw_{\ell'}) \end{array}\right).$$

We show that this is computationally indistinguishable from the experiment where \mathcal{A} gets pk, s_1, \ldots, s_ℓ , ciphertexts $C_1, \ldots, C_{\ell'}$, and a matrix of $\ell \cdot \ell'$ uniform and independent elements of G.

To prove this we show that, for arbitrary i, j, the experiment in which \mathcal{A} is given

$$\begin{pmatrix} g_{1,1} & \vdots & \vdots \\ \vdots & g_{i-1,j} & \vdots \\ \vdots & \cdots & H_{k_i}(\mathsf{label}_j, C_j, pw_j) & \cdots & \vdots \\ \vdots & H_{k_{i+1}}(\mathsf{label}_j, C_j, pw_j) & \vdots \\ g_{\ell,1} & \vdots & H_{k_\ell}(\mathsf{label}_{\ell'}, C_{\ell'}, pw_{\ell'}) \end{pmatrix}$$

is computationally indistinguishable from the experiment in which \mathcal{A} is given

$$\begin{pmatrix} g_{1,1} & \vdots & \vdots \\ \vdots & g_{i-1,j} & \vdots \\ \vdots & \cdots & g_{i,j} & \cdots & \vdots \\ \vdots & H_{k_{i+1}}(\mathsf{label}_j, C_j, pw_j) & \vdots \\ g_{\ell,1} & \vdots & H_{k_{\ell}}(\mathsf{label}_{\ell'}, C_{\ell'}, pw_{\ell'}) \end{pmatrix}$$

where the $g_{a,b}$ denote uniform and independent elements of G. (In both cases, \mathcal{A} is also given pk, s_1, \ldots, s_ℓ , and $C_1, \ldots, C_{\ell'}$, and may access the decryption oracle as in the original experiment.) Once we show this, the lemma follows by a standard hybrid argument.

We denote the first experiment, in which the (i, j)th entry of the matrix is computed as $H_{k_i}(\mathsf{label}_j, C_j, pw_j)$, by $\mathsf{real}_{i,j}$; we denote the second experiment, in which the (i, j)th entry of the matrix is a random $g_{i,j}$, by $\mathsf{rand}_{i,j}$. To prove that these two experiments are indistinguishable, we introduce two additional experiments. Experiment $\mathsf{real}'_{i,j}$ (resp., $\mathsf{rand}'_{i,j}$) is identical to $\mathsf{real}_{i,j}$ (resp.,

 $\operatorname{rand}_{i,j}$) except that now the *j*th ciphertext C_j returned by the encryption oracle is computed as an encryption of 0 (i.e., $C_j \leftarrow \operatorname{Enc}_{pk}(\operatorname{label}_j, 0)$). It follows from the CCA-security of the encryption scheme that $\operatorname{real}_{i,j}$ and $\operatorname{real}'_{i,j}$ (resp., $\operatorname{rand}_{i,j}$) and $\operatorname{rand}'_{i,j}$) are computationally indistinguishable.

To complete the proof that $\operatorname{real}_{i,j}$ and $\operatorname{rand}_{i,j}$ are computationally indistinguishable, we show that $\operatorname{real}'_{i,j}$ and $\operatorname{rand}'_{i,j}$ are statistically close. To see this, consider the following experiment involving an algorithm B who is given $s_i = \alpha(k_i)$ for unknown (random) k_i , outputs (label, C, pw) $\in X \setminus L$, and is given in response an element $h_{i,j} \in G$:

- 1. Choose random $k_i \leftarrow K$ and give $s_i = \alpha(k_i)$ to B.
- 2. B computes $(pk, sk) \leftarrow \text{Gen}(1^n)$ and internally runs \mathcal{A} on initial input pk. Then B samples $k_1, \ldots, k_{i-1}, k_{i+1}, \ldots, k_{\ell} \leftarrow K$, sets $s_m = \alpha(k_m)$ for all $m \neq i$, and gives s_1, \ldots, s_{ℓ} to \mathcal{A} .
- 3. When \mathcal{A} queries its encryption oracle with (label, pw), then B does:
 - (a) For the kth such query, with k < j, we have B compute $C_k \leftarrow \mathsf{Enc}_{pk}(\mathsf{label}, pw)$ and give to \mathcal{A} the ciphertext C_k along with random $g_{1,k}, \ldots, g_{\ell,k} \leftarrow G$.
 - (b) For the *j*th such query (so (label_j, pw_j) = (label, pw)), B computes C_j ← Enc_{pk}(label_j, 0), outputs (label_j, C_j, pw_j), and receives h_{i,j}. It then chooses g_{1,j},..., g_{i-1,j} ← G, and gives to A the values g_{1,j},..., g_{i-1,j}, h_{i,j}, H_{ki+1}(label_j, C_j, pw_j),..., H_{kℓ}(label_j, C_j, pw_j). Note that B can compute these latter values since it knows k_{i+1},..., k_ℓ.
 - (c) For the kth such query, with k > j, we have B compute $C_k \leftarrow \mathsf{Enc}_{pk}(\mathsf{label}, pw)$ and give to \mathcal{A} the ciphertext C_k along with $H_{k_1}(\mathsf{label}, C_k, pw), \ldots, H_{k_\ell}(\mathsf{label}, C_k, pw)$. Note that B can compute $H_{k_i}(\mathsf{label}, C_k, pw)$, even though it does not know k_i , because of the fact that $(\mathsf{label}, C_k, pw) \in L_{pw}$ and B knows the randomness used to compute C_k .

B answers any decryption queries of \mathcal{A} using sk.

Note that $pw_j \neq 0$ and so $(\mathsf{label}_j, C_j, pw_j) \in X \setminus L$. Observe that the view of \mathcal{A} is distributed according to $\mathsf{real}'_{i,j}$ if $h_{i,j} = H_{k_i}(\mathsf{label}_j, C_j, pw_j)$, and is distributed according to $\mathsf{rand}'_{i,j}$ if $h_{i,j}$ is chosen uniformly from G. It follows from Equation (1) that $\mathsf{real}'_{i,j}$ and $\mathsf{rand}'_{i,j}$ are statistically close.

3 A Framework for One-Round PAK Protocols


Our protocol uses a CCA-secure labeled public-key encryption scheme (Gen, Enc, Dec), and a smooth projective hash function as described in Section 2.2.

Public parameters. The public parameters consist of a public key pk generated by $\text{Gen}(1^n)$. No one need know or store the associated secret key. (For the specific instantiations given in Section 4, a public key can be derived from a common random string.) Let $(K, G, \mathcal{H} = \{H_k : X \to G\}_{k \in K}, S, \alpha : K \to S)$ be a smooth projective hash function for pk.

Protocol execution. Consider an execution of the protocol between users U and $U' \neq U$ holding a shared password pw. Our protocol is symmetric, and so we describe the execution from the point of view of U; see also Figure 1.

First, U chooses random hash key $k \leftarrow K$ and computes $s := \alpha(k)$. It then sets $\mathsf{label} := (U, U', s)$ and computes the ciphertext $C \leftarrow \mathsf{Enc}_{pk}(\mathsf{label}, pw)$. It sends the message (s, C).

	Public parameters: pk	
$\underline{\text{User }U}$		<u>User U'</u>
$k \leftarrow K; s := \alpha(k)$		$k' \leftarrow K; s' := \alpha(k')$
label := (U, U', s)		label' := (U', U, s')
$C \leftarrow Enc_{pk}(label, pw)$	s, C	$C' \leftarrow Enc_{pk}(label', pw)$
	s', C'	_
label' := (U', U, s')		label := (U, U', s)
$SK_U := H_k(label', C', pw)$		$SK_{U'} := H_k(label', C', pw)$
$\cdot H_{k'}(label, C, pw)$)	$\cdot H_{k'}(label, C, pw)$



Upon receiving the message (s', C'), user U does the following. If C' is not a valid ciphertext or $s' \notin S$, then U simply rejects. Otherwise, U sets |abe|' := (U', U, s') and computes

$$\mathsf{sk}_U := H_k(\mathsf{label}', C', pw) \cdot H_{k'}(\mathsf{label}, C, pw)$$

Note that U computes $H_k(\mathsf{label}', C', pw)$ using k, and can compute $H_{k'}(\mathsf{label}, C, pw)$ using $s' = \alpha(k')$ and the randomness it used to generate C. Correctness follows immediately from the definition of smooth projective hashing.

Theorem 1. If (Gen, Enc, Dec) is a CCA-secure labeled encryption scheme and $(K, G, \mathcal{H} = \{H_k : X \to G\}_{k \in K}, S, \alpha : K \to S)$ is a smooth projective hash function, then the protocol in Figure 1 is a secure protocol for password-based authenticated key exchange.

Proof: Let Π denote the protocol in Figure 1, and fix a polynomial-time adversary \mathcal{A} attacking Π . We construct a sequence of experiments $\mathsf{Expt}_0, \ldots, \mathsf{Expt}_5$, with the original experiment corresponding to Expt_0 . Let $\mathsf{Adv}_{\mathcal{A},i}(n)$ denote the advantage of \mathcal{A} in experiment Expt_i . To prove the desired bound on $\mathsf{Adv}_{\mathcal{A},\Pi}(n) = \mathsf{Adv}_{\mathcal{A},0}(n)$, we bound the effect of each change in the experiment on the advantage of \mathcal{A} , and then show that $\mathsf{Adv}_{\mathcal{A},5}(n) \leq Q(n)/D$ (where, recall, Q(n) denotes the number of on-line attacks made by \mathcal{A} , and D denotes the dictionary size).

Experiment Expt₁: Here we change the way Execute queries are answered. Specifically, the ciphertexts C, C' sent by the two parties U, U' are computed as encryptions of 0 instead of being computed as encryptions of the correct password $pw_{U,U'}$. (Recall that the space of legal passwords is $\{1, \ldots, D\}$, and so 0 is never a valid password.) The (common) session key is computed as

$$\mathsf{sk}_U := \mathsf{sk}_{U'} := H_k(\mathsf{label}', C', pw) \cdot H_{k'}(\mathsf{label}, C, pw),$$

where both values are computed using the (known) keys k, k'. A proof of the following is immediate:

Claim 1. If (Gen, Enc, Dec) is semantically secure, $|\mathsf{Adv}_{\mathcal{A},0}(n) - \mathsf{Adv}_{\mathcal{A},1}(n)|$ is negligible.

Experiment Expt_2 : Here, we again change the way $\mathsf{Execute}$ queries are answered. Now, the (common) session key $\mathsf{sk}_U = \mathsf{sk}_{U'}$ is chosen uniformly from G.

Claim 2. $|\operatorname{Adv}_{\mathcal{A},1}(n) - \operatorname{Adv}_{\mathcal{A},2}(n)|$ is negligible.

Proof: The claim follows from the properties of the smooth projective hash function. Consider a single call to the Execute oracle (in either Expt_1 or Expt_2), where the transcript given to the adversary is (s, C, s', C') with $C \leftarrow \mathsf{Enc}_{pk}(\mathsf{label}, 0)$ and $C' \leftarrow \mathsf{Enc}_{pk}(\mathsf{label}', 0)$. In Expt_1 the session keys are computed as

$$\mathsf{sk}_U := \mathsf{sk}_{U'} := H_k(\mathsf{label}', C', pw) \cdot H_{k'}(\mathsf{label}, C, pw),$$

where $pw = pw_{U,U'}$ is the password shared by U and U'. Since $(\mathsf{label}', C', pw)$ is not in L, it follows (cf. Equation (1)) that $(s, H_k(\mathsf{label}', C', pw))$ is statistically close to (s, g), where g is uniform in G. This means that, in Expt_1 , $\mathsf{sk}_U = \mathsf{sk}_{U'}$ is statistically close to uniform in G, even conditioned on the given transcript. Since this is how $\mathsf{sk}_U, \mathsf{sk}_{U'}$ are chosen in Expt_2 , the claim follows.

Before continuing, we distinguish between two possible types of Send oracle queries. We let $Send_0(U, i, U')$ denote a "prompt" query that causes instance Π^i_U of user U to initiate the protocol with user U'. In response to a Send₀ query, the adversary is given the message sent by U to U'. This query also has the effect of setting $pid^i_U = U'$.

The second type of Send query, $\text{Send}_1(U, i, \text{msg})$, denotes the event where \mathcal{A} sends the message msg to instance Π^i_U . In response, a session key sk^i_U is computed. (Nothing is output in response to this query, but the value of the computed session key can be learned via a subsequent Reveal or Test query.) For a query $\text{Send}_1(U, i, \text{msg})$ with $\text{pid}^i_U = U'$, we say a valid msg is previously used if it was output by a previous oracle query $\text{Send}_0(U', \star, U)$. In any other case, we say a valid msg is adversarially generated. (An invalid message is always ignored by the instance that receives it.)

Experiment Expt₃: We first modify the experiment so that, when the public parameters pk are generated, the simulator stores the associated secret key sk. (This is just a syntactic change.) We then modify the way queries to the Send₁ oracle are handled. Specifically, in response to the query Send₁(U, i, msg) where msg = (s', C'), we distinguish the following three cases (in all the following, let $pid_{U}^{i} = U'$, let label' = (U', U, s'), and let $pw = pw_{U,U'}$):

- 1. If msg is adversarially generated, then compute $pw' := \mathsf{Dec}_{sk}(\mathsf{label}', C')$. Then:
 - (a) If pw' = pw, the simulator declares that \mathcal{A} succeeds and terminates the experiment.
 - (b) If $pw' \neq pw$, the simulator chooses sk_U^i uniformly from G.
- 2. If msg is previously used, then in particular the simulator knows a value k' such that $s' = \alpha(k')$. The simulator computes $\mathsf{sk}_U^i := H_k(\mathsf{label}', C', pw) \cdot H_{k'}(\mathsf{label}, C, pw)$, but using k' to compute $H_{k'}(\mathsf{label}, C, pw)$ (rather than using the randomness used to generate C, as done in Expt₂).

Invalid messages are treated as before, and no session key is computed.

Claim 3. $\operatorname{Adv}_{\mathcal{A},2}(n) \leq \operatorname{Adv}_{\mathcal{A},3}(n) + \operatorname{negl}(n)$.

Proof: Consider the three possible cases described above. Imagine first that only the change in Case 1(b) is made. This change introduces a negligible statistical difference between Expt_3 and Expt_2 ; the analysis is as in Claim 2, except that we now specifically use the fact that Equation (1) holds even under adaptive choice of (label', C', pw) $\notin L$.

The change in Case 2 does not affect the computed value sk_U^i since $(\mathsf{label}, C, pw) \in L$. The change in Case 1(a) can only increase the advantage of \mathcal{A} .

Experiment Expt₄: Once again we change how Send₁ queries are handled. In response to query Send₁(U, i, msg) where msg = (s', C') is previously used, let $pid_U^i = U'$ and proceed as follows:

- If there exists an instance $\Pi_{U'}^j$ partnered with Π_U^i (i.e., such that $\mathsf{sid}_{U'}^j$, the transcript of the protocol for instance $\Pi_{U'}^j$, is equal to sid_U^i), then set $\mathsf{sk}_U^i := \mathsf{sk}_{U'}^j$.
- Otherwise, choose sk_U^i uniformly from G.

Claim 4. If (Gen, Enc, Dec) is CCA-secure, $|\mathsf{Adv}_{\mathcal{A},3}(n) - \mathsf{Adv}_{\mathcal{A},4}(n)|$ is negligible.

Proof: The proof relies on Lemma 1 (cf. Section 2.2). Let ℓ be a polynomial upper bound on the number of Send queries issued by \mathcal{A} , and consider the following adversary \mathcal{S} interacting in the experiment defined in Lemma 1:

- 1. S is given pk and s_1, \ldots, s_ℓ .
- 2. S chooses random passwords $pw_{U,U'}$ for all pairs of parties U, U', and runs \mathcal{A} on input pk.
- 3. S responds to Execute queries as in Expt_2 by generating a transcript where C, C' are encryptions of 0, and where the (matching) session keys are chosen uniformly at random.
- 4. S responds to the *i*th Send₀ query Send₀(U, \star, U') as follows: Set label := (U, U', s_i). Submit (label, $pw_{U,U'}$) to the encryption oracle, and receive in return a ciphertext C_i along with values $h_{1,i}, \ldots, h_{\ell,i}$. Give to \mathcal{A} the message (s_i, C_i).
- 5. S responds to a query $\mathsf{Send}_1(U, j, \mathsf{msg})$, where $\mathsf{msg} = (s', C')$, as follows: If there exists an instance $\Pi_{U'}^k$ partnered with Π_U^j , then set $\mathsf{sk}_U^j := \mathsf{sk}_{U'}^k$. Otherwise, let $\mathsf{pid}_U^j = U'$ and $\mathsf{label}' = (U', U, s')$, and say the query $\mathsf{Send}_0(U, j, U')$ (i.e., the Send query that initiated instance Π_U^j) was the *i*th Send_0 query made by \mathcal{A} , and resulted in the response (s_i, C_i) . We now distinguish several cases based on $\mathsf{msg} = (s', C')$:
 - (a) If msg is invalid (i.e., C' is an invalid ciphertext, or $s' \notin S$), no session key is computed.
 - (b) If msg is previously used, then (by definition) it was output by some previous query $Send_0(U', \star, U)$. Say this was the *r*th $Send_0$ query made by \mathcal{A} , and so $msg = (s_r, C_r)$. Then \mathcal{S} computes $sk_U^j := h_{i,r} \cdot h_{r,i}$.
 - (c) If msg is adversarially generated, then S submits (label', C') to its decryption oracle and receives in return a value pw. If $pw \neq pw_{U,U'}$ then sk_U^j is chosen uniformly from G. If $pw = pw_{U,U'}$ then S declares that \mathcal{A} succeeds and terminates the experiment.
- 6. At the end of the experiment, \mathcal{S} outputs 1 if and only if \mathcal{A} succeeds.

Referring to the game defined in Section 2.2, note that if b = 0 then the view of \mathcal{A} in the above execution with \mathcal{S} is identical to the view of \mathcal{A} in Expt_3 . This is true since when b = 0 it holds in step 5(b), above, that $h_{i,r} = H_{k_i}(\text{label}', C_r, pw_{U,U'})$ and $h_{r,i} = H_{k_r}(\text{label}, C_i, pw_{U,U'})$, where $s_i = \alpha(k_i), s_r = \alpha(k_r)$, and C_i, C_r are encryptions of $pw_{U,U'}$.

On the other hand, when b = 1 the view of \mathcal{A} in the above execution with \mathcal{S} is identical to the view of \mathcal{A} in Expt_4 . To see this, recall that when b = 1 all the values $\{h_{i,j}\}$ received by \mathcal{S} are chosen uniformly and independently from G. We need to show that this yields a uniform and independent distribution on all the session keys computed in step 5(b). Consider a particular session key sk_U^j computed as in step 5(b). The only other time the value $h_{i,r}$ could be used in the experiment is if \mathcal{A} queries $\mathsf{Send}_1(U', \star, (s_i, C_i))$ to the instance $\Pi_{U'}^{\star}$ which sent (s_r, C_r) . But then $\Pi_{U'}^{\star}$ and Π_U^j are partnered, and so the session key $\mathsf{sk}_{U'}^{\star}$ will be set equal to sk_U^j (as in Expt_4). Since $h_{i,r}$ is random and used only once to compute a session key in step 5(b), we conclude that (when b = 1) any session keys computed in that step are independently uniform in G.

The claim follows from Lemma 1.

Experiment Expt₅: Here, we change how Send₀ queries are handled. Now, in response to a query Send₀(U, i, U'), we compute s as usual but let C be an encryption of 0. The following claim is immediate from CCA-security of (Gen, Enc, Dec).

Claim 5. If (Gen, Enc, Dec) is CCA-secure, $|\mathsf{Adv}_{\mathcal{A},4}(n) - \mathsf{Adv}_{\mathcal{A},5}(n)|$ is negligible.

In Expt_5 , the view of \mathcal{A} is independent of any of the user's passwords until it sends an adversarially generated message that corresponds to an encryption of the correct password (at which point \mathcal{A} succeeds). It therefore holds that $\mathsf{Adv}_{\mathcal{A},5}(n) \leq Q(n)/D$. Claims 1–5 thus imply that $\mathsf{Adv}_{\mathcal{A},0}(n) \leq Q(n)/D + \mathsf{negl}(n)$, completing the proof of the theorem.

4 Instantiating the Building Blocks

We now discuss two possible instantiations of the building blocks required by the protocol of Section 3.

- Our first instantiation is based on the decisional Diffie-Hellman (DDH) assumption and (generic) simulation-sound NIZK proofs [37]; see Appendix A.2 for a definition. (This instantiation could also be based on the quadratic residuosity assumption or the Paillier assumption, as in [21]. We omit further details.)
- Our second, more efficient instantiation is based on the decisional linear assumption [10] in groups with a bilinear map, and a specific simulation-sound NIZK proof system.

4.1 A Construction Based on the DDH Assumption

We first describe an encryption scheme and then the associated smooth projective hash function.

A CCA-secure encryption scheme. We construct a CCA-secure encryption scheme by applying the Naor-Yung/Sahai paradigm [34, 37] to the El Gamal encryption scheme. Briefly, the public key defines a group \mathbb{G} of prime order p along with generators $g_1, h_1, g_2, h_2 \in \mathbb{G}$. The public key also contains a common random string crs for a (one-time) simulation-sound NIZK proof system [37].

Fixing G, let $\mathsf{ElGamal}_{g,h}(m)$ denote an El Gamal encryption of $m \in \mathbb{G}$ with respect to (g, h); namely, $\mathsf{ElGamal}_{g,h}(m) \to (g^r, h^r \cdot m)$, where $r \in \mathbb{Z}_p$ is chosen uniformly at random. (We assume passwords can be represented as elements of G.) To encrypt a message $m \in \mathbb{G}$ in our CCA-secure scheme, the sender outputs the ciphertext

(ElGamal_{q1,h1}(m), ElGamal_{q2,h2}(m), π)

where π is a simulation-sound NIZK proof that the same *m* is encrypted in both cases. Labels can be incorporated by including the label in the proof π ; we omit the standard details.

Decryption of the ciphertext $(c_1, d_1, c_2, d_2, \pi)$ rejects if $c_1, d_1, c_2, d_2 \notin \mathbb{G}$ or if the proof π is invalid. (Note that the space of valid label/ciphertext pairs is efficiently recognizable without the secret key.) If the ciphertext is valid, then one of the two component ciphertexts is decrypted and the resulting message is output. The results of [37] show that this yields a CCA-secure (labeled) encryption scheme based on the DDH assumption and simulation-sound NIZK.

A smooth projective hash function. Fix a group \mathbb{G} and a public key $pk = (g_1, h_1, g_2, h_2, \operatorname{crs})$ as above, and define sets X and $\{L_{pw}\}$ as in Section 2.2. Define a smooth projective hash function as

follows. The set of keys K consists of all four-tuples of elements in \mathbb{Z}_p . Given a valid label/ciphertext pair (label, $C = (c_1, d_1, c_2, d_2, \pi)$) and key $k = (x_1, y_1, x_2, y_2)$, the hash function is defined as:

$$H_{(x_1,y_1,x_2,y_2)}\left(\mathsf{label},(c_1,d_1,c_2,d_2,\pi),pw\right) = c_1^{x_1} \cdot (d_1/pw)^{y_1} \cdot c_2^{x_2} \cdot (d_2/pw)^{y_2}.$$

(Thus, the range of H is the group \mathbb{G} .) The projection function α is defined as:

$$\alpha(x_1, y_1, x_2, y_2) = (g_1^{x_1} \cdot h_1^{y_1}, g_2^{x_2} \cdot h_2^{y_2}).$$

Lemma 2. $(K, \mathbb{G}, \mathcal{H} = \{H_k\}_{k \in K}, S, \alpha)$ as defined above is a smooth projective hash function for the hard subset-membership problem $(X, \{L_{pw}\})$.

Proof: Sampling a uniform $k \in K$, computing $H_k(x)$ given $k \in K$ and $x \in X$, and computing $\alpha(k)$ for $k \in K$ are all easy.

We now show that if $(\mathsf{label}, C, pw) \in L$, then $H_k(\mathsf{label}, C, pw)$ can be computed efficiently given $\alpha(k)$ and the randomness used to generate C. Since $(\mathsf{label}, C, pw) \in L$, we have that

$$C = (c_1, d_1, c_2, d_2, \pi) = (g_1^{r_1}, h_1^{r_1} \cdot pw, g_2^{r_2}, h_2^{r_2} \cdot pw, \pi)$$

for some $r_1, r_2 \in \mathbb{Z}_p$. For $k = (x_1, y_1, x_2, y_2)$ we have

$$\begin{split} H_k(\mathsf{label}, C, pw) &= c_1^{x_1} \cdot (d_1/pw)^{y_1} \cdot c_2^{x_2} \cdot (d_2/pw)^{y_2} \\ &= g_1^{r_1x_1} \cdot (h_1)^{r_1y_1} \cdot g_2^{r_2x_2} \cdot (h_2)^{r_2y_2} \\ &= (g_1^{x_1}h_1^{y_1})^{r_1} \cdot (g_2^{x_2}h_2^{y_2})^{r_2} \,. \end{split}$$

This can be computed easily given r_1, r_2 and $\alpha(k) = (g_1^{x_1} h_1^{y_1}, g_2^{x_2} h_2^{y_2}).$

Next, we show that if $(\mathsf{label}, C, pw) \in X \setminus L$, then the value of $H_k(\mathsf{label}, C, pw)$ is uniform conditioned on $\alpha(k)$. (This holds even if (label, C, pw) are chosen adaptively depending on $\alpha(k)$.) Fix any $\alpha(k) = (s_1, s_2)$. This constraints $k = (x_1, y_1, x_2, y_2)$ to satisfy

$$x_1 + (\log_{g_1} h_1) \cdot y_1 = \log_{g_1} s_1 \tag{2}$$

$$x_2 + (\log_{g_2} h_2) \cdot y_2 = \log_{g_2} s_2, \qquad (3)$$

where the above are modulo the group order p. For any $(label, C, pw) \in X \setminus L$, we can write $C = (g_1^{r_1}, h_1^{r_1} \cdot pw', g_2^{r_2}, h_2^{r_2} \cdot pw', \pi)$ for some $pw' \neq pw$. (We assume for simplicity that the same pw' is encrypted twice; since π is valid, this is the case with all but negligible probability.) Then:

$$\begin{aligned} H_k(\mathsf{label}, C, pw) &= g_1^{r_1 x_1} \cdot \left(h_1^{r_1} \cdot pw'/pw\right)^{y_1} \cdot g_2^{r_2 x_2} \cdot \left(h_2^{r_2} \cdot pw'/pw\right)^{y_2} \\ &= g_1^{r_1 x_1} \cdot h_1^{r_1' y_1} \cdot g_2^{r_2 x_2} \cdot h_2^{r_2' y_2}, \end{aligned}$$

for some $r'_1 \neq r_1$ and $r'_2 \neq r_2$. So, for any $g \in \mathbb{G}$, we have $H_k(\mathsf{label}, C, pw) = g$ iff

$$r_1 \cdot x_1 + (r_1' \cdot \log_{g_1} h_1) \cdot y_1 + (r_2 \cdot \log_{g_1} g_2) \cdot x_2 + (r_2' \cdot \log_{g_1} h_2) \cdot y_2 = \log_{g_1} g_2$$

Since this equation in the unknowns x_1, y_1, x_2, y_2 is linearly independent of Equations (2) and (3), we see that the probability that $H_k(\mathsf{label}, C, pw) = g$ is exactly $1/|\mathbb{G}|$, and so the distribution of $H_k(\mathsf{label}, C, pw)$ is uniform in \mathbb{G} .

4.2 A Construction Based on the Decisional Linear Assumption

We now present a more efficient construction based on bilinear maps. The efficiency advantage is obtained by using a *specific* simulation-sound NIZK proof system, constructed using techniques adapted from [26, 13]. Our construction here relies on the *decisional linear assumption* as introduced by Boneh et al. [10]; we refer the reader there for a precise statement of the assumption.

A CPA-secure encryption scheme. We start by describing a semantically secure encryption scheme, due to Boneh et al. [10], based on the decisional linear assumption; we then convert this into a CCA-secure encryption scheme via the same paradigm as above, but using an efficient simulation-sound NIZK proof system. The bilinear map is used only in the construction of the simulation-sound NIZK.

Fix groups \mathbb{G}, \mathbb{G}_T of prime order p, and a bilinear map $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$. The public key is $pk = (f, g, h) \in \mathbb{G}^3$, and the secret key is (α, β) such that $f = h^{1/\alpha}$ and $g = h^{1/\beta}$. A message $m \in \mathbb{G}$ is encrypted by choosing random $r, s \in \mathbb{Z}_p$ and computing the ciphertext $(f^r, g^s, h^{r+s} \cdot m)$. Given a ciphertext (c_1, c_2, c_3) , we can recover m as $c_3/c_1^{\alpha}c_2^{\beta}$.

A simulation-sound NIZK proof of plaintext equality. We can construct a (one-time) simulation-sound NIZK proof of plaintext equality for the encryption scheme described above using the techniques of [26, 13]. Details of the construction (which, while not entirely straightforward, are not the focus of this work) are given in Appendix B.

A CCA-secure encryption scheme. We obtain a CCA-secure encryption scheme by using the Naor-Yung/Sahai paradigm, as described previously. (The following discussion relies on the results of Appendix B.) The public key consists of group elements (f_1, g_1, f_2, g_2, h) used for encryption, in addition to any group elements needed for the CRS of the simulation-sound NIZK proof. Encryption of m, as described in Appendix B, is done by choosing $r_1, s_1, r_2, s_2 \in \mathbb{Z}_p$ and computing the ciphertext

$$(f_1^{r_1}, g_1^{s_1}, h^{r_1+s_1} \cdot m, f_2^{r_2}, g_2^{s_2}, h^{r_2+s_2} \cdot m, \pi),$$

where π denotes a simulation-sound NIZK proof that the same *m* was encrypted both times. (Once again, the space of valid label/ciphertext pairs is efficiently recognizable without the secret key.) It follows from [34, 37] that this yields a CCA-secure scheme under the decisional linear assumption. Ciphertexts consist of 66 group elements altogether (see Appendix B).

A smooth projective hash function. Fix \mathbb{G} , \mathbb{G}_T , and a public-key $pk = (f_1, g_1, f_2, g_2, h)$ as above, and define sets X and $\{L_{pw}\}$ as in Section 2.2. We define a smooth projective hash function as follows. The set of keys K is the set of six-tuples of elements in \mathbb{Z}_p . Given a valid label/ciphertext pair (label, $C = (c_1, d_1, e_1, c_2, d_2, e_2, \pi)$) and a key $k = (x_1, y_1, z_1, x_2, y_2, z_2) \in \mathbb{Z}_p^6$, the hash function is defined as

$$H_{(x_1,y_1,z_1,x_2,y_2,z_2)}(\mathsf{label}, C, pw) = c_1^{x_1} \cdot d_1^{y_1} \cdot (e_1/pw)^{z_1} \cdot c_2^{x_2} \cdot d_2^{y_2} \cdot (e_2/pw)^{z_2} \cdot d_2^{y_2} \cdot d_2^{y_2} \cdot (e_2/pw)^{z_2} \cdot d_2^{y_2} \cdot d$$

(The range of H is \mathbb{G} itself.) The projection function $\alpha: K \to \mathbb{G}^4$ is defined as:

$$\alpha(x_1, y_1, z_1, x_2, y_2, z_2) = (f_1^{x_1} h^{z_1}, g_1^{y_1} h^{z_1}, f_2^{x_2} h^{z_2}, g_2^{y_2} h^{z_2}).$$

Lemma 3. $(K, \mathbb{G}, \mathcal{H} = \{H_k\}_{k \in K}, S, \alpha)$ as defined above is a smooth projective hash function for the hard subset-membership problem $(X, \{L_{pw}\})$.

Proof: Sampling a uniform $k \in K$, computing $H_k(x)$ given $k \in K$ and $x \in X$, and computing $\alpha(k)$ for $k \in K$ are all easy.

We show that if $(\mathsf{label}, C, pw) \in L$, then $H_k(\mathsf{label}, C, pw)$ can be computed efficiently given $\alpha(k)$ and the randomness used to generate C. Since $(\mathsf{label}, C, pw) \in L$, we have

$$C = (c_1, d_1, e_1, c_2, d_2, e_2, \pi) = \left(f_1^{r_1}, g_1^{s_1}, h^{r_1 + s_1} \cdot pw, f_2^{r_2}, g_2^{s_2}, h^{r_2 + s_2} \cdot pw, \pi\right)$$

for some $r_1, s_1, r_2, s_2 \in \mathbb{Z}_p$. For a hash key $k = (x_1, y_1, z_1, x_2, y_2, z_2)$ we have

$$\begin{aligned} H_k(\mathsf{label}, C, pw) &= c_1^{x_1} d_1^{y_1} \cdot (e_1/pw)^{z_1} \cdot c_2^{x_2} \cdot d_2^{y_2} \cdot (e_2/pw)^{z_2} \\ &= (f_1^{x_1} h^{z_1})^{r_1} \cdot (g_1^{y_1} h^{z_1})^{s_1} \cdot (f_2^{x_2} h^{z_2})^{r_2} \cdot (g_2^{y_2} h^{z_2})^{s_2} \end{aligned}$$

This can be computed easily given r_1, s_1, r_2, s_2 and $\alpha(k) = (f_1^{x_1} h^{z_1}, g_1^{y_1} h^{z_1}, f_2^{x_2} h^{z_2}, g_2^{y_2} h^{z_2})$.

Next, we show that if $(\mathsf{label}, C, pw) \in X \setminus L$, then the value of $H_k(\mathsf{label}, C, pw)$ is uniform conditioned on $\alpha(k)$. (This holds even if (label, C, pw) are chosen adaptively depending on $\alpha(k)$.) Fix any $\alpha(k) = (S_1, S_2, S_3, S_4)$. Letting $\alpha_i = \log_h f_i$ and $\beta_i = \log_h g_i$, this value of $\alpha(k)$ constrains $k = (x_1, y_1, z_1, x_2, y_2, z_2)$ to satisfy

$$\begin{pmatrix} \alpha_1 & 0 & 1 & 0 & 0 & 0 \\ 0 & \beta_1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & \alpha_2 & 0 & 1 \\ 0 & 0 & 0 & 0 & \beta_2 & 1 \end{pmatrix} \begin{pmatrix} x_1 \\ y_1 \\ z_1 \\ x_2 \\ y_2 \\ z_2 \end{pmatrix} = \begin{pmatrix} \gamma_1 \\ \gamma_2 \\ \gamma_3 \\ \gamma_4 \end{pmatrix},$$
(4)

where $\gamma_i = \log_h S_i$. For any $(\mathsf{label}, C, pw) \in X \setminus L$, we can write

$$C = (f_1^{r_1}, g_1^{s_1}, h^{r_1+s_1} \cdot pw', f_2^{r_2}, g_2^{s_2}, h^{r_2+s_2} \cdot pw', \pi)$$

for some $pw' \neq pw$. (We assume for simplicity that the same pw' is encrypted twice; since π is valid, this is the case with all but negligible probability. The proof holds even when this is not the case.) We then have

$$H_{k}(\mathsf{label}, C, pw) = f_{1}^{r_{1}x_{1}} \cdot g_{1}^{s_{1}y_{1}} \cdot h^{(r_{1}+s_{1})z_{1}} \cdot (pw'/pw)^{z_{1}} \cdot f_{2}^{r_{2}x_{2}} \cdot g_{2}^{s_{2}y_{2}} \cdot h^{(r_{2}+s_{2})z_{2}} \cdot (pw'/pw)^{z_{2}} \\ = S_{1}^{r_{1}}S_{2}^{s_{1}}S_{3}^{r_{2}}S_{4}^{s_{2}} \cdot (pw'/pw)^{z_{1}+z_{2}},$$
(5)

where $pw'/pw \neq 1$. For any $g \in \mathbb{G}$, we have $(pw'/pw)^{z_1+z_2} = g$ iff

$$\log_h(pw'/pw) \cdot z_1 + \log_h(pw'/pw) \cdot z_2 = \log_h g. \tag{6}$$

Since Equation (6) is linearly independent of the system of equations given by (4), the probability that $(pw'/pw)^{z_1+z_2} = g$ is exactly $1/|\mathbb{G}|$ even conditioned on the value $\alpha(k)$. Looking at Equation (5), and noting that $S_1^{r_1}S_2^{s_1}S_3^{r_2}S_4^{s_2}$ is entirely determined by $\alpha(k)$ and C, we conclude that the distribution of $H_k(\mathsf{label}, C, pw)$ is uniform in \mathbb{G} .

5 A One-Round, Universally Composable PAK Protocol

Canetti et al. [15] gave a definition of security for password-based authenticated key exchange in the universal composability (UC) framework [14]. Their definition guarantees a strong, simulationbased notion of security that, in particular, guarantees that security is maintained even when the protocol is run concurrently with arbitrary other protocols. For the specific case of password-based key exchange, the definition also has the advantage of *automatically* handling arbitrary (efficiently sampleable) distributions on passwords, and even correlations between passwords of different users. We refer to [15] for a more complete discussion.

A brief review of the UC framework and the password-based key-exchange functionality \mathcal{F}_{pwKE} are given in Appendix C. We let $\hat{\mathcal{F}}_{pwKE}$ denote the multi-session extension of \mathcal{F}_{pwKE} .

5.1 Overview of the Construction

We do not know how to prove that the protocol from Section 3 is universally composable. The main difficulty is that the definition of PAK in the UC framework requires simulation even if the adversary guesses the correct password. (In contrast, in the proof of security in Section 3 we simply "give up" in case this ever occurs.) To see the problem more clearly, consider what happens in the UC setting when the simulator sends the first message of the protocol to the adversary, before the simulator knows the correct password. The simulator must send some ciphertext C as part of the first message, and this "commits" the simulator to some password pw. When the adversary sends the reply, the simulator can extract the adversary's "password guess" pw' and submit this guess to the ideal functionality. If this turns out to be the correct password, however, the simulator is stuck: it needs to compute a session key that matches the session key the adversary would compute, but the simulator is (information-theoretically!) unable to do so because it sent an incorrect ciphertext in the first message.

In prior work [15], the issue above was resolved by having one party send a "pre-commitment" to the password, and then running a regular PAK protocol along with a proof that the password being used in the protocol is the same as the password to which it "pre-committed". (The proof is set up in such a way that the simulator can equivocate this proof, but the adversary cannot.) This requires at least one additional round.

We take a different approach that does not affect the round complexity at all. Roughly, we modify the protocol from Figure 1 by having each party include as part of its message an encryption C_1 of its hash key k, along with a proof that C_1 encrypts a value k for which $\alpha(k) = s$. Now, even if the simulator is wrong in its guess of the password it will still be able to compute a session key by extracting this hash key from the adversary's message. A full description of the protocol is given in the following section.

While we do not describe in detail any instantiation of the components, we remark that it should be possible to use the same techniques as in Appendix B to construct (reasonably) efficient realizations of the necessary components using bilinear maps. We leave this for future work.

5.2 Description of the Protocol

In addition to the building blocks used in Section 3, here we also rely on an unbounded simulationsound [17] NIZK proof system (CRSGen, \mathcal{P}, \mathcal{V}) for a language L^* defined later below; refer to Appendix A.2 for definitions.

Public Parameters: (pk_1, pk_2, crs)			
$\underline{\text{User }U}$		User U'	
$\begin{aligned} k \leftarrow K; s := \alpha(k) \\ C_1 \leftarrow Enc_{pk_1}(k) \\ \pi := \mathcal{P}_{crs}((s, C_1) \in L^*) \\ label := (sid, U, U', s, C_1, \pi) \\ C_2 \leftarrow Enc_{pk_2}(label, pw) \end{aligned}$	s, C_1, π, C_2	$\begin{split} k' &\leftarrow K; s' := \alpha(k') \\ C'_1 &\leftarrow \operatorname{Enc}_{pk_1}(k') \\ \pi' &:= \mathcal{P}_{\operatorname{crs}}((s', C'_1) \in L^*) \\ \operatorname{label}' &:= (\operatorname{sid}, U', U, s', C'_1, \pi') \\ C'_2 &\leftarrow \operatorname{Enc}_{pk_2}(\operatorname{label}', pw) \end{split}$	
	s',C_1',π',C_2'		
$\begin{split} label' &:= (sid, U', U, s', C_1', \pi') \\ sk_U &:= H_k(label', C_2', pw) \\ &\cdot H_{k'}(label, C_2, pw) \end{split}$		$\begin{aligned} label &:= (sid, U, U', s, C_1, \pi) \\ sk_{U'} &:= H_k(label', C_2', pw) \\ &\cdot H_{k'}(label, C_2, pw) \end{aligned}$	

Figure 2: A universally composable protocol for password-based authenticated key exchange.

Public parameters. The public parameters consist of two public keys pk_1, pk_2 generated by $Gen(1^n)$ and a common random string crs for the simulation-sound NIZK proof system. Let $(K, G, \mathcal{H} = \{H_k : X \to G\}_{k \in K}, S, \alpha : K \to S)$ be a smooth projective hash function for pk_2 .

Protocol execution. Consider an execution of the protocol between users U and $U' \neq U$ holding a shared password pw and a common session identifier sid. (The sid is an artifact of the UC framework, and it is guaranteed that (1) parties communicating with each other begin holding matching sids, and (2) each sid is used only once. Existence of these sids is not essential to our proof of security, though it does make the proof somewhat simpler.) Our protocol is symmetric, and so we describe the execution from the point of view of U; see Figure 2.

First, U chooses a random hash key $k \leftarrow K$ and computes $s := \alpha(k)$. It then computes an encryption of k, namely $C_1 \leftarrow \mathsf{Enc}_{pk_1}(k)$. Define a language L^* as follows.

$$L^* \stackrel{\text{def}}{=} \{(s, C_1) : \exists k \in K \text{ and } \omega \text{ s.t } s = \alpha(k) \text{ and } C_1 = \mathsf{Enc}_{pk_1}(k; \omega)\}$$

U computes an NIZK proof π that $(C_1, s) \in L^*$, using crs. It then sets label := $(sid, U, U', s, C_1, \pi)$ and computes the ciphertext $C_2 \leftarrow \mathsf{Enc}_{pk_2}(\mathsf{label}, pw)$. The message it sends is (s, C_1, π, C_2) .

Upon receiving the message (s', C'_1, π', C'_2) , user U does the following. If the message is invalid (i.e., if verification of π' fails, or C'_2 is not a valid ciphertext, or $s' \notin S$), then U simply rejects. Otherwise, U sets label' := $(sid, U', U, s', C'_1, \pi')$ and computes $sk_U := H_k(label', C'_2, pw) \cdot H_{k'}(label, C_2, pw)$. Note U can compute $H_k(label', C'_2, pw)$ since it knows k, and can compute $H_{k'}(label, C_2, pw)$ using $s' = \alpha(k')$ and the randomness used to generate C_2 . Correctness follows from the definition of smooth projective hashing.

Theorem 2. If (Gen, Enc, Dec) is CCA-secure, (CRSGen, \mathcal{P}, \mathcal{V}) is an unbounded simulation-sound NIZK proof system, and $(K, G, \mathcal{H} = \{H_k : X \to G\}_{k \in K}, S, \alpha : K \to S)$ is a smooth projective hash function, then the protocol in Figure 2 securely realizes $\hat{\mathcal{F}}_{\mathsf{pwKE}}$ in the $\mathcal{F}_{\mathsf{crs}}$ -hybrid model.

Proof: Let \mathcal{A} be an adversary that interacts with the parties running the protocol. We construct an ideal-world adversary (i.e., simulator) \mathcal{S} interacting with the ideal functionality $\hat{\mathcal{F}}_{pwKE}$, such that

no PPT environment \mathcal{Z} can distinguish an interaction with \mathcal{A} in the real world from an interaction with \mathcal{S} in the ideal world.

 \mathcal{S} starts by invoking a copy of \mathcal{A} and running a simulated interaction of \mathcal{A} with \mathcal{Z} and the parties in the network. \mathcal{S} forwards all messages to/from \mathcal{A} and \mathcal{Z} in the usual way.

Generating the public parameters. S generates pk_1 and pk_2 along with their corresponding secret keys sk_1 and sk_2 . It also runs $(crs, \tau) \leftarrow S_1(1^k)$, where S_1 is the initial simulator for the simulation-sound NIZK proof system. The public parameters (pk_1, pk_2, crs) are given to \mathcal{A} , and then \mathcal{S} responds to the messages of \mathcal{A} as described below.

Receiving a (NewSession, sid, P_i, P_j) message from $\hat{\mathcal{F}}_{pwKE}$. Upon receiving such a message (indicating that P_i should initiate the protocol with P_j), \mathcal{S} proceeds as follows. Choose a random hash key $k \leftarrow K$ and compute $s := \alpha(k)$. Compute the ciphertext $C_1 \leftarrow \mathsf{Enc}_{pk_1}(0)$ and a simulated NIZK proof π for the statement $(C_1, s) \in L^*$. Set label := $(\mathsf{sid}, P_i, P_j, s, C_1, \pi)$ and compute $C_2 \leftarrow \mathsf{Enc}_{pk_2}(\mathsf{label}, 0)$. Give the message (s, C_1, π, C_2) to \mathcal{A} .

Receiving a message msg' = (s', C'_1, π', C'_2) from \mathcal{A} . Let P_i denote the (uncorrupted) user to whom \mathcal{A} sends this message, and let sid denote the session ID with which this message is associated. Let P_i denote the partner of P_i for this session.

If msg' is invalid then S does nothing. Otherwise, we say msg' is *previously used* if it was sent by S (on behalf of P_j) upon receiving the message (NewSession, sid, P_j, P_i) from $\hat{\mathcal{F}}_{pwKE}$. In any other case we say msg' is *adversarially generated*. To respond to this message, S does:

- 1. If msg' is previously used, then S sends (NewKey, sid, P_i, \perp) to the functionality $\hat{\mathcal{F}}_{pwKE}$. (This has the effect of choosing a random session key for this instance of P_i if it terminates before the partnered session at P_j , or if the passwords used by these instances of P_i and P_j differ, and otherwise setting the session key for this instance of P_i equal to the session key already computed for the partnered instance of P_i ; cf. Appendix C.1.1.)
- 2. If msg' is adversarially generated, then S decrypts the ciphertext C'_2 using the secret key sk_2 to obtain a password pw'. Then S queries the functionality $\hat{\mathcal{F}}_{pwKE}$ on input (TestPwd, sid, P_i, pw'), which replies with either "correct guess" or "wrong guess".
 - (a) If the reply is "correct guess", then \mathcal{S} decrypts C'_1 using sk_1 to obtain k'. It then computes $\mathsf{sk} := H_k(\mathsf{label}', C'_2, pw') \cdot H_{k'}(\mathsf{label}, C_2, pw')$, where $H_{k'}(\mathsf{label}, C_2, pw')$ is computed using k'. Finally, \mathcal{S} sends (NewKey, sid, P_i , sk) to $\hat{\mathcal{F}}_{\mathsf{pwKE}}$.
 - (b) If the reply is "wrong guess", then \mathcal{S} sends (NewKey, sid, P_i, \perp) to $\hat{\mathcal{F}}_{\mathsf{pwKE}}$.

Let $\mathrm{IDEAL}_{\hat{\mathcal{F}}_{\mathsf{pwKE}},\mathcal{S},\mathcal{Z}}$ denote the view of the environment \mathcal{Z} in the ideal world when interacting with \mathcal{S} , and let $\mathrm{EXEC}_{\Pi,\mathcal{A},\mathcal{Z}}$ denote the view of \mathcal{Z} in the real world when the protocol from Figure 2 is being run. Our aim is to show that these distributions are computationally indistinguishable. We do this by considering a sequence of experiments $\mathrm{Expt}_1, \ldots, \mathrm{Expt}_5$ (where Expt_0 corresponds to the real-world execution). Let $\mathrm{EXEC}_{i,\mathcal{A},\mathcal{Z}}$ denote the view of \mathcal{Z} in Expt_i . We show that $\mathrm{EXEC}_{i,\mathcal{A},\mathcal{Z}} \approx_c$ $\mathrm{EXEC}_{i+1,\mathcal{A},\mathcal{Z}}$ for all i (where this denotes computational indistinguishability), and then argue that $\mathrm{EXEC}_{5,\mathcal{A},\mathcal{Z}}$ is identical to $\mathrm{IDEAL}_{\hat{\mathcal{F}}_{\mathsf{pwkE}},\mathcal{S},\mathcal{Z}}$. This completes the proof.

Experiment Expt_0 : Recall, this experiment involves the environment \mathcal{Z} interacting with the adversary \mathcal{A} , who in turn interacts with parties running the real protocol as specified in Figure 2. The

view of \mathcal{Z} consists of the public parameters and all the protocol messages (forwarded to it by \mathcal{A}) as well as all the session keys produced by parties during the course of the experiment.

Experiment Expt₁: We change the distribution of the public parameters and the messages generated by the parties in the protocol. Specifically, the common random string **crs** is replaced with a simulated one, and the proof π in every outgoing message is replaced with a simulated proof. It follows from the zero-knowledge properties of the proof system that $\text{EXEC}_{1,\mathcal{A},\mathcal{Z}} \approx_c \text{EXEC}_{0,\mathcal{A},\mathcal{Z}}$.

Experiment Expt₂: Here, we again change the distribution of the outgoing messages by always computing the ciphertext C_1 as an encryption of 0 (rather than an encryption of the hash key k). It follows immediately from semantic security of (Gen, Enc, Dec) that $\text{EXEC}_{2,\mathcal{A},\mathcal{Z}} \approx_c \text{EXEC}_{1,\mathcal{A},\mathcal{Z}}$.

Before continuing, we define the notion of a *previously used* message. Note that the definition here is slightly different from the definition used in the proof of Theorem 1.

Consider a valid $msg' = (s', C'_1, \pi', C'_2)$ sent to an instance of (uncorrupted user) P_i who is using session ID sid and is partnered with P_j . We say msg' is *previously used* if it was sent by an instance of the (uncorrupted user) P_j that is using the same session ID sid, and is partnered with P_i . We say a previously used msg' is *associated with pw* if pw is the password that was used by the instance of P_j that sent msg'. If msg' is not previously used, we say it is *adversarially generated*.

Experiment Expt₃: We change the way session keys are computed. Specifically, consider an instance of user P_i , with session ID sid and partner P_j , who receives an incoming message $msg' = (s', C'_1, \pi', C'_2)$. If msg' is invalid, then it is handled as before (and no session key is computed). Otherwise, set label' = (sid, $P_j, P_i, s', C'_1, \pi')$ and let pw be the password being used by the current instance of P_i . There are several sub-cases:

- 1. If msg' is adversarially generated, compute $pw' := \mathsf{Dec}_{sk_2}(\mathsf{label}', C'_2)$. Then:
 - (a) If pw' = pw, compute $k' := \text{Dec}_{sk_1}(C_1)$. Then compute $sk := H_k(\text{label}', C'_2, pw) \cdot H_{k'}(\text{label}, C_2, pw)$, using k' to compute the second hash value.
 - (b) If $pw' \neq pw$, choose sk uniformly from G.
- 2. If msg' is previously used, then the simulator knows the password pw' associated with msg', and also knows a value k' such that $s' = \alpha(k')$. Then:
 - (a) If pw' = pw, compute $\mathsf{sk} := H_k(\mathsf{label}', C'_2, pw) \cdot H_{k'}(\mathsf{label}, C_2, pw)$, using k' to compute the second hash value (rather than using the randomness used to generate C_2 , as in Expt_2).
 - (b) If $pw' \neq pw$, choose sk uniformly from G.

Claim 6. If (CRSGen, \mathcal{P}, \mathcal{V}) is simulation-sound, $\text{EXEC}_{3,\mathcal{A},\mathcal{Z}} \approx_c \text{EXEC}_{2,\mathcal{A},\mathcal{Z}}$.

Proof: The proof is similar to that of Claim 3, though there are some differences due to the particularities of the UC setting. Imagine first that only the changes in Case 1(b) and 2(b) are made. These changes introduce a negligible statistical difference between Expt_3 and Expt_2 , exactly as in Claim 3 (though note that Case 2(b) cannot occur in the context of the previous claim since parties are assumed to always use matching passwords there).

Consider next the change in Case 1(a). Since the proof system is simulation-sound, with all but negligible probability S extracts a value k' such that $\alpha(k') = s'$; assuming this occurs, the session key computed in Expt₃ is identical to the session key that would have been computed in Expt₂. The same holds (always) in Case 2(a).

Experiment Expt₄: Here, in every outgoing message the ciphertext C_2 is generated as an encryption of 0, rather than as an encryption of pw. It follows readily from the CCA-security of the encryption scheme used that $\text{EXEC}_{4,\mathcal{A},\mathcal{Z}} \approx_c \text{EXEC}_{3,\mathcal{A},\mathcal{Z}}$.

Experiment Expt₅: We once again change the way session keys are computed. Consider an instance of user P_i , with session ID sid, partner P_j , and using password pw, who receives an incoming message $msg' = (s', C'_1, \pi', C'_2)$ that is previously used and is associated with pw. Then:

- If the corresponding partnered instance of P_j has already computed a session key sk, then set the session key of the current instance of P_i equal to sk.
- Otherwise, choose the session key for the current instance of P_i uniformly from G.

Claim 7. $\text{EXEC}_{5,\mathcal{A},\mathcal{Z}} \approx_c \text{EXEC}_{4,\mathcal{A},\mathcal{Z}}$.

Proof: Since msg' is previously used, the ciphertext C'_2 is an encryption of 0. Thus $(\mathsf{label}', C'_2, pw)$ is not in L, and it follows (cf. Equation (1)) that $(s, H_k(\mathsf{label}', C'_2, pw))$ is statistically close to (s, g), where $s = \alpha(k)$ and g is a uniform element in G. This means that the secret key is statistically close to a uniform element in G, even conditioned on the given transcript. The claim follows.

The distribution $\text{EXEC}_{5,\mathcal{A},\mathcal{Z}}$ is identical to the distribution produced by the simulator, namely $\text{IDEAL}_{\hat{\mathcal{F}}_{\text{pwKE}},\mathcal{S},\mathcal{Z}}$. This follows by inspection. Indeed, the only difference is a syntactic one: in experiment Expt_5 the session keys are computed and stored locally by the simulator, whereas in the ideal world the functionality computes the session keys and sends them to the (dummy) parties.

The above shows that $\text{EXEC}_{\Pi,\mathcal{A},\mathcal{Z}} \equiv \text{EXEC}_{0,\mathcal{A},\mathcal{Z}} \approx_c \text{EXEC}_{5,\mathcal{A},\mathcal{Z}} \equiv \text{IDEAL}_{\hat{\mathcal{F}}_{\mathsf{pwKE}},\mathcal{S},\mathcal{Z}}$, completing the proof of the theorem.

References

- M. Abdalla, D. Catalano, C. Chevalier, and D. Pointcheval. Efficient two-party passwordbased key exchange protocols in the UC framework. In *Cryptographers' Track — RSA 2008*, volume 4964 of *LNCS*, pages 335–351. Springer, 2008.
- [2] B. Barak, R. Canetti, Y. Lindell, R. Pass, and T. Rabin. Secure computation without authentication. Journal of Cryptology, 24(4):720–760, 2011.
- [3] M. Bellare, D. Pointcheval, and P. Rogaway. Authenticated key exchange secure against dictionary attacks. In Advances in Cryptology — Eurocrypt 2000, volume 1807 of LNCS, pages 139–155. Springer, 2000.
- [4] M. Bellare and P. Rogaway. Entity authentication and key distribution. In Advances in Cryptology — Crypto '93, volume 773 of LNCS, pages 232–249. Springer, 1994.
- [5] M. Bellare and P. Rogaway. Provably secure session key distribution: The three party case. In 27th Annual ACM Symposium on Theory of Computing (STOC), pages 57–66. ACM Press, 1995.
- [6] S. M. Bellovin and M. Merritt. Encrypted key exchange: Password-based protocols secure against dictionary attacks. In *IEEE Symposium on Security & Privacy*, pages 72–84. IEEE, 1992.

- [7] R. Bird, I. Gopal, A. Herzberg, P. Janson, S. Kutten, R. Molva, and M. Yung. Systematic design of two-party authentication protocols. *IEEE J. on Selected Areas in Communications*, 11(5):679–693, 1993.
- [8] M. Blum, A. De Santis, S. Micali, and G. Persiano. Noninteractive zero-knowledge. SIAM Journal on Computing, 20(6):1084–1118, 1991.
- [9] M. Blum, P. Feldman, and S. Micali. Proving security against chosen cyphertext attacks. In Advances in Cryptology — Crypto '88, volume 403 of LNCS, pages 256–268. Springer, 1990.
- [10] D. Boneh, X. Boyen, and H. Shacham. Short group signatures. In Advances in Cryptology Crypto 2004, volume 3152 of LNCS, pages 41–55. Springer, 2004.
- [11] X. Boyen, Y. Dodis, J. Katz, R. Ostrovsky, and A. Smith. Secure remote authentication using biometric data. In Advances in Cryptology — Eurocrypt 2005, volume 3494 of LNCS, pages 147–163. Springer, 2005.
- [12] V. Boyko, P. D. MacKenzie, and S. Patel. Provably secure password-authenticated key exchange using Diffie-Hellman. In Advances in Cryptology — Eurocrypt 2000, volume 1807 of LNCS, pages 156–171. Springer, 2000.
- [13] J. Camenisch, N. Chandran, and V. Shoup. A public key encryption scheme secure against key dependent chosen plaintext and adaptive chosen ciphertext attacks. In A. Joux, editor, Advances in Cryptology — Eurocrypt 2009, volume 5479 of LNCS, pages 351–368. Springer, 2009.
- [14] R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In 42nd Annual Symposium on Foundations of Computer Science (FOCS), pages 136–145. IEEE, 2001. Full version at http://eprint.iacr.org/2000/067/.
- [15] R. Canetti, S. Halevi, J. Katz, Y. Lindell, and P. D. MacKenzie. Universally composable password-based key exchange. In Advances in Cryptology — Eurocrypt 2005, volume 3494 of LNCS, pages 404–421. Springer, 2005.
- [16] R. Cramer and V. Shoup. Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In Advances in Cryptology — Eurocrypt 2002, volume 2332 of LNCS, pages 45–64. Springer, 2002.
- [17] A. De Santis, G. Di Crescenzo, R. Ostrovsky, G. Persiano, and A. Sahai. Robust non-interactive zero knowledge. In Advances in Cryptology — Crypto 2001, volume 2139 of LNCS, pages 566– 598. Springer, 2001.
- [18] W. Diffie and M. E. Hellman. New directions in cryptography. IEEE Trans. Information Theory, 22(6):644–654, 1976.
- [19] U. Feige, D. Lapidot, and A. Shamir. Multiple non-interactive zero knowledge proofs under general assumptions. SIAM Journal on Computing, 29(1):1–28, 1999.
- [20] R. Gennaro. Faster and shorter password-authenticated key exchange. In 5th Theory of Cryptography Conference — TCC 2008, volume 4948 of LNCS, pages 589–606. Springer, 2008.

- [21] R. Gennaro and Y. Lindell. A framework for password-based authenticated key exchange. ACM Trans. Information and System Security, 9(2):181–234, 2006.
- [22] O. Goldreich and Y. Lindell. Session-key generation using human passwords only. Journal of Cryptology, 19(3):241–340, 2006.
- [23] L. Gong, T. M. A. Lomas, R. M. Needham, and J. H. Saltzer. Protecting poorly chosen secrets from guessing attacks. *IEEE J. Selected Areas in Communications*, 11(5):648–656, 1993.
- [24] V. Goyal, A. Jain, and R. Ostrovsky. Password-authenticated session-key generation on the internet in the plain model. In Advances in Cryptology — Crypto 2010, volume 6223 of LNCS, pages 277–294. Springer, 2010.
- [25] A. Groce and J. Katz. A new framework for efficient password-based authenticated key exchange. In 17th ACM Conf. on Computer and Communications Security (CCS), pages 516–525. ACM Press, 2010.
- [26] J. Groth and A. Sahai. Efficient non-interactive proof systems for bilinear groups. In Advances in Cryptology — Eurocrypt 2008, volume 4965 of LNCS, pages 415–432. Springer, 2008.
- [27] S. Halevi and H. Krawczyk. Public-key cryptography and password protocols. ACM Trans. Information and System Security, 2(3):230–268, 1999.
- [28] I. R. Jeong, J. Katz, and D. H. Lee. One-round protocols for two-party authenticated key exchange. In 2nd Intl. Conference on Applied Cryptography and Network Security (ACNS), volume 3089 of LNCS, pages 220–232. Springer, 2004.
- [29] S. Jiang and G. Gong. Password based key exchange with mutual authentication. In 11th Annual International Workshop on Selected Areas in Cryptography (SAC), volume 3357 of LNCS, pages 267–279. Springer, 2004.
- [30] J. Katz, P. D. MacKenzie, G. Taban, and V. D. Gligor. Two-server password-only authenticated key exchange. In 3rd Intl. Conference on Applied Cryptography and Network Security (ACNS), volume 3531 of LNCS, pages 1–16. Springer, 2005.
- [31] J. Katz, R. Ostrovsky, and M. Yung. Efficient and secure authenticated key exchange using weak passwords. J. ACM, 57(1):78–116, 2009.
- [32] J. Katz and V. Vaikuntanathan. Smooth projective hashing and password-based authenticated key exchange from lattices. In Advances in Cryptology — Asiacrypt 2009, volume 5912 of LNCS, pages 636–652. Springer, 2009.
- [33] P. D. MacKenzie, S. Patel, and R. Swaminathan. Password-authenticated key exchange based on RSA. In Advances in Cryptology — Asiacrypt 2000, volume 1976 of LNCS, pages 599–613. Springer, 2000.
- [34] M. Naor and M. Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In 22nd Annual ACM Symposium on Theory of Computing (STOC), pages 427–437. ACM Press, 1990.

- [35] M.-H. Nguyen and S. Vadhan. Simpler session-key generation from short random passwords. *Journal of Cryptology*, 21(1):52–96, 2008.
- [36] T. Okamoto. Authenticated key exchange and key encapsulation in the standard model. In Advances in Cryptology — Asiacrypt 2007, volume 4833 of LNCS, pages 474–484. Springer, 2007.
- [37] A. Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In 40th Annual Symposium on Foundations of Computer Science (FOCS), pages 543–553. IEEE, 1999.

A Additional Definitions

A.1 Labeled CCA-Secure Encryption

We use the standard notion of chosen-ciphertext security for public-key encryption, though adapted to support the inclusion of *labels* when generating ciphertexts.

Definition 2. A public-key encryption scheme *supporting labels is a tuple of* PPT *algorithms* (Gen, Enc, Dec) *such that:*

- The key-generation algorithm Gen takes as input a security parameter 1ⁿ and returns a public key pk and a secret key sk.
- The encryption algorithm Enc takes as input a public key pk, a label label, and a message m. It returns a ciphertext C ← E_{pk}(label, m).
- The decryption algorithm Dec takes as input a secret key sk, a label label, and a ciphertext C. It returns a message m or a distinguished symbol ⊥. We write this as m = D_{sk}(label, C).

We require that for all pk, sk output by $gen(1^n)$, any $label \in \{0, 1\}^*$, all m in the (implicit) message space, and any C output by $Enc_{pk}(label, m)$ we have $Dec_{sk}(label, C) = m$.

Our definition of security against chosen-ciphertext attacks is standard except for our inclusion of labels. In the following, we define a left-or-right encryption oracle $\mathbb{E}_{pk,b}(\cdot, \cdot, \cdot)$ (where $b \in \{0, 1\}$) as follows:

$$\operatorname{Enc}_{pk,b}(\operatorname{label}, m_0, m_1) \stackrel{\text{der}}{=} \operatorname{Enc}_{pk}(\operatorname{label}, m_b).$$

Definition 3. A public-key encryption scheme (Gen, Enc, Dec) is secure against adaptive chosenciphertext attacks (CCA-secure) if the following is negligible for all PPT algorithms A:

$$2 \cdot \Pr[(pk, sk) \leftarrow \mathsf{Gen}(1^n); b \leftarrow \{0, 1\} : \mathcal{A}^{\mathsf{Enc}_{pk, b}(\cdot, \cdot, \cdot), \mathsf{Dec}_{sk}(\cdot, \cdot)}(1^n, pk) = b] - 1$$

where \mathcal{A} 's queries are restricted as follows: if \mathcal{A} makes a query $\operatorname{Enc}_{pk,b}(\operatorname{label}, m_0, m_1)$ then $|m_0| = |m_1|$; furthermore, if \mathcal{A} receives ciphertext C in response to this query, then \mathcal{A} cannot later query $\operatorname{Dec}_{sk}(\operatorname{label}, C)$ (but it is allowed to query $(\operatorname{Dec}_{sk}(\operatorname{label}', C)$ with $\operatorname{label}' \neq \operatorname{label}$).

Simulation-Sound Non-Interactive Zero Knowledge (NIZK) A.2

Simulation-sound NIZK was introduced in [37, 17]. Intuitively, a simulation-sound NIZK proof system is a NIZK proof system with the extra property that a polynomially bounded cheating prover is incapable of convincing the verifier of a false statement, even after seeing any number of simulated proofs of her choosing. We first recall the notion of (adaptive) NIZK:

Definition 4 ([19, 9, 8]). A tuple of PPT algorithms $\Pi = (\mathsf{CRSGen}, \mathcal{P}, \mathcal{V}, \mathcal{S}_1, \mathcal{S}_2)$ is an efficient NIZK proof system for a language $L \in \mathcal{NP}$ with witness relation R if the following hold:

- Completeness: For all n, all $x \in L \cap \{0,1\}^n$, all w such that R(x,w) = 1, and all strings $\sigma \leftarrow \mathsf{CRSGen}(1^n), \text{ it holds that } \mathcal{V}_{\sigma}(x, \mathcal{P}_{\sigma}(x, w)) = 1.$
- Adaptive Soundness: For all adversaries \mathcal{A} , the following is negligible (in n):

 $\Pr[\sigma \leftarrow \mathsf{CRSGen}(1^n); (x, \pi) \leftarrow \mathcal{A}(\sigma) : V_{\sigma}(x, \pi) = 1 \land \pi \notin L].$

• Adaptive Zero Knowledge: For all PPT adversaries A, the following is negligible

$$\left| \Pr[\mathsf{Expt}_{\mathcal{A}}(n) = 1] - \Pr[\mathsf{Expt}_{\mathcal{A}}^{Sim}(n) = 1] \right|$$

where experiment $\mathsf{Expt}_{A}(n)$ is defined as:

$$\sigma \leftarrow \mathsf{CRSGen}(1^n)$$

Return $\mathcal{A}^{\mathcal{P}_{\sigma}(\cdot,\cdot)}(1^n,\sigma)$

and experiment $\mathsf{Expt}^{Sim}_{\mathcal{A}}(n)$ is defined as:

 \bot

$$(\sigma, \tau) \leftarrow \mathcal{S}_{1}(1^{n})$$

$$Return \ \mathcal{A}^{\mathcal{S}'_{\sigma,\tau}(\cdot, \cdot)}(1^{n}, \sigma),$$
where $S'_{\sigma,\tau}(x, w) = \begin{cases} \mathcal{S}_{2}(x, \sigma, \tau) & R(x, w) = 1 \land x \in \{0, 1\}^{n} \\ \bot & otherwise \end{cases}$

otherwise

We next define the notion of simulation-sound NIZK. (Note that although one-time simulation soundness would suffice for our applications in Section 4, we only define the stronger notion of unbounded simulation-sound NIZK. See [37] for a definition of the former.)

Definition 5. Let $\Pi = (CRSGen, \mathcal{P}, \mathcal{V}, \mathcal{S}_1, \mathcal{S}_2)$ be an efficient NIZK proof system for a language $L \in \mathcal{NP}$. We say Π is simulation sound if for all PPT adversaries \mathcal{A} it holds that $\Pr[\mathsf{Expt}_{\mathcal{A},\Pi}(n) = 1]$ is negligible, where $\mathsf{Expt}_{\mathcal{A},\Pi}(n)$ denotes the following experiment:

$$\begin{aligned} (\sigma,\tau) &\leftarrow \mathcal{S}_1(1^n) \\ (x,\pi) &\leftarrow \mathcal{A}^{\mathcal{S}_2(\cdot,\sigma,\tau)}(1^n,\sigma) \\ Let \ Q \ be \ the \ list \ of \ proofs \ returned \ by \ \mathcal{S}_2, \ above \\ Return \ 1 \ iff \ (\pi \notin Q \ and \ x \notin L \ and \ \mathcal{V}_{\sigma}(x,\pi) = 1) . \end{aligned}$$

Assuming the existence of enhanced trapdoor permutations, every language in \mathcal{NP} has a simulation-sound NIZK proof system [17].

B A Simulation-Sound NIZK Proof of Plaintext Equality

Fix groups \mathbb{G} , \mathbb{G}_T of prime order p, and a bilinear map $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$ as in Section 4.2. Fix also two public keys $pk_1 = (f_1, g_1, h)$ and $pk_2 = (f_2, g_2, h)$. We encrypt a message m with respect to pk_1 by choosing random r, s and computing the ciphertext $(f_1^r, g_1^s, h^{r+s} \cdot m)$. We encrypt a message mwith respect to pk_2 by choosing random $r, s \in \mathbb{Z}_p$ and computing the ciphertext $(f_2^r, g_2^s, h^{r+s} \cdot m)$. We stress that the public keys use the same value h.

We first describe a (potentially malleable) NIZK proof of plaintext equality. That is, given two ciphertexts (F_1, G_1, H_1) and (F_2, G_2, H_2) encrypted with respect to pk_1, pk_2 , respectively, we describe a proof that these ciphertexts encrypt the same message. The observation is that plaintext equality is equivalent to the existence of $r_1, s_1, r_2, s_2 \in \mathbb{Z}_p$ such that:

$$F_1 = f_1^{r_1} (7)$$

$$G_1 = g_1^{s_1} (8)$$

$$F_2 = f_2^{r_2} (9)$$

$$G_2 = g_2^{s_2} \tag{10}$$

$$H_1/H_2 = h^{r_1 + s_1 - r_2 - s_2}. (11)$$

As shown in [26] (see also [13, Section 4.4] for a self-contained description), NIZK proofs of satisfiability (with a CRS) can be constructed for a system of equations as above; since, in our case, we have 5 linear equations in 4 variables, proofs would contain 22 group elements.³

Camenisch et al. [13] show a construction of an *unbounded* simulation-sound NIZK proof system. For our purposes in Section 3, a simpler construction that is *one-time* simulation sound [37] suffices. Let (Gen, Sign, Vrfy) be a one-time signature scheme, where for simplicity we assume verification keys are elements of \mathbb{G} (this can always be achieved using an extra step of hashing). To make the above (one-time) simulation-sound, we add to the CRS group elements (f, g, h, F, G, H). Roughly, proofs of plaintext equality now contain:

- 1. A fresh signature verification key vk.
- 2. A proof that either there exists a satisfying assignment to Equations (7)–(11), or that the given tuple (f, g, h, F, G, H) is an encryption of vk, i.e., that there exist r, s such that:

$$F = f^r, \qquad G = g^s, \qquad H/vk = h^{r+s}.$$
(12)

3. A signature σ (with respect to vk) on the proof from the previous step.

Noting that Equation (12) describes a system of 3 linear equations in 2 variables, and using the techniques from [13, Appendix A.2], an NIZK proof as required in step 2 can be obtained using⁴ 58 group elements, for a total of 60 group elements for the entire simulation-sound NIZK proof (assuming signatures are one group element for simplicity). See also footnote 3.

³Our calculations here are based on the decisional linear assumption (the 2-linear assumption in the terminology of [13]). If we are willing to use the 1-linear assumption, the efficiency of our proofs can be improved.

⁴In detail: Applying the OR-transformation from [13, Appendix A.2.2] to the two systems of equations given by Equations (7)–(11) and Equation (12), respectively, we obtain one system of 8 linear equations and 1 quadratic equation in 7 variables. Applying the transformation from [13, Appendix A.2.1], we get a system of 11 linear equations in 11 variables, plus 3 extra group elements that need to be sent. From [13, Section 4.4] we see that an NIZK proof for the former can be done using 55 group elements.

C Universally Composable Password-Based Key Exchange

C.1 The Universal Composability Framework

We provide a brief review of the universally composable security framework [14]. The framework allows for defining the security properties of cryptographic tasks so that security is maintained under general composition with an unbounded number of instances of arbitrary protocols running concurrently. In the UC framework, the security requirements of a given task are captured by specifying an ideal functionality run by a "trusted party" that obtains the inputs of the participants and provides them with the desired outputs. Informally, then, a protocol securely carries out a given task if running the protocol in the presence of a real-world adversary amounts to "emulating" the desired ideal functionality.

The notion of emulation in the UC framework is considerably stronger than that considered in previous models. As usual, the real-world model includes the parties running the protocol and an adversary \mathcal{A} who controls their communication and potentially corrupts parties, while the idealworld includes a simulator \mathcal{S} who interacts with an ideal functionality \mathcal{F} and also with dummy players who simply send input to/receive output from \mathcal{F} ; "emulating an ideal process" requires that for any adversary \mathcal{A} there should exist a simulator \mathcal{S} that causes the outputs of the parties in the ideal process to have a "similar" (i.e., computationally-indistinguishable) distribution to the outputs of the parties in a real-world execution of the protocol. In the UC framework, the requirement on \mathcal{S} is more stringent than this. Specifically, there is also an additional entity called the environment \mathcal{Z} . This environment generates the inputs to all parties, observes all their outputs, and interacts with the adversary in an arbitrary way throughout the computation. A protocol π is said to securely realize an ideal functionality \mathcal{F} if for any real-world adversary \mathcal{A} that interacts with \mathcal{Z} and real players running π , there exists an ideal-world simulator \mathcal{S} that interacts with \mathcal{Z} , the ideal functionality \mathcal{F} , and the "dummy" players communicating with \mathcal{F} , such that no poly-time environment \mathcal{Z} can distinguish whether it is interacting with \mathcal{A} (in the real world) or \mathcal{S} (in the ideal world). \mathcal{Z} thus serves as an "interactive distinguisher" between a real-world execution of the protocol π and an ideal execution of functionality \mathcal{F} . A key point is that \mathcal{Z} cannot be re-wound by \mathcal{S} ; in other words, \mathcal{S} must provide a so-called "straight-line" simulation.

The following universal composition theorem is proven in [14]. Consider a protocol π that operates in the \mathcal{F} -hybrid model, where parties can communicate as usual and in addition have ideal access to an unbounded number of *copies* of the functionality \mathcal{F} . Let ρ be a protocol that securely realizes \mathcal{F} as sketched above, and let π^{ρ} be identical to π with the exception that the interaction with *each copy* of \mathcal{F} is replaced with an interaction with a *separate instance* of ρ . Then, π and π^{ρ} have essentially the same input/output behavior. In particular, if π securely realizes some functionality \mathcal{G} in the \mathcal{F} -hybrid model then π^{ρ} securely realizes \mathcal{G} in the standard model (i.e., without access to any functionality).

C.1.1 Universally-Composable Password-Based Key Exchange

In Figure 3 we present the functionality \mathcal{F}_{pwKE} for password-based key exchange, taken directly from [15]. We refer the reader to their work for extensive discussion regarding the particular choices made regarding this formulation of this functionality.

Functionality \mathcal{F}_{pwKE}

The functionality $\mathcal{F}_{\mathsf{pwKE}}$ is parameterized by a security parameter *n*. It interacts with an adversary \mathcal{S} and a set of parties via the following queries:

Upon receiving a query (NewSession, sid, P_i , P_j , pw) from party P_i :

Send (NewSession, sid, P_i , P_j) to S. In addition, if this is the first NewSession query, or if this is the second NewSession query and there is a record (P_j, P_i, pw') , then record (P_i, P_j, pw) and mark this record fresh.

Upon receiving a query (TestPwd, sid, P_i, pw') from the adversary S:

If there is a record of the form (P_i, P_j, pw) which is fresh, then do: If pw' = pw, mark the record compromised and reply to S with "correct guess". If $pw \neq pw'$, mark the record interrupted and reply with "wrong guess".

Upon receiving a query (NewKey, sid, P_i , sk) from S, where |sk| = n:

- If there is a record of the form (P_i, P_j, pw) , and this is the first NewKey query for P_i , then:
- If this record is compromised, or either P_i or P_j is corrupted, then output (sid, sk) to player P_i .
- If this record is fresh, and there is a record (P_j, P_i, pw') with pw' = pw, and a key sk' was sent to P_j , and (P_j, P_i, pw) was fresh at the time, then output (sid, sk') to P_i .
- In any other case, pick a new random key $\mathsf{sk}' \leftarrow \{0,1\}^n$ and send $(\mathsf{sid},\mathsf{sk}')$ to P_i .
- In all cases, mark the record (P_i, P_j, pw) as completed.

Figure 3: The password-based key-exchange functionality \mathcal{F}_{pwKE} .