

Oblivious PAKE: Efficient Handling of Password Trials

Franziskus Kiefer and Mark Manulis

Department of Computing, University of Surrey, United Kingdom
mail@franziskuskiefer.de, mark@manulis.eu

Abstract. In this work we introduce the notion of *Oblivious Password based Authenticated Key Exchange* (O-PAKE) and a compiler to transform a large class of PAKE into O-PAKE protocols. O-PAKE allows a client that shares *one* password with a server to use a *subset of passwords* within one PAKE session. It succeeds if and only if one of those input passwords matches the one stored on the server side. The term *oblivious* is used to emphasise that no information about any password, input by the client, is made available to the server. O-PAKE protocols can be used to improve the overall efficiency of login attempts using PAKE protocols in scenarios where users are not sure (e.g. no longer remember) which of their passwords has been used at a particular web server. Using special processing techniques, our O-PAKE compiler reaches nearly constant run time on the server side, independent of the size of the client’s password set; in contrast, a naive approach to run a new PAKE session for each login attempt would require linear run time for both parties. We prove security of the O-PAKE compiler under standard assumptions using the latest game-based PAKE model by Abdalla, Fouque and Pointcheval (PKC 2005), tailored to our needs. We identify the requirements that PAKE protocols must satisfy in order to suit the compiler and give two concrete O-PAKE instantiation.

1 Introduction

Authentication with passwords is the most common (and perhaps most critical) authentication mechanism on the modern Internet. The dominating approach today is when clients send passwords (or some function thereof) to the server over a secure channel (e.g. TLS [28]). This approach requires PKI and its security relies solely on the secure channel and the client’s ability to correctly verify the server’s certificate. Any impersonation of the certificate leads to password exposure. Even if no impersonation takes place, any password input on the client side is revealed to the server. This creates a different problem based on statistics, indicating that many users operate with a small set of passwords but often do not remember their correct mapping to the servers. If a user types in a password that is not shared with this server but with another one then its exposure may lead to subsequent impersonation attacks on the client. The studies in [30, 31] show that every user has 6.5 passwords on average, used on 25 different websites and that on average 2.4 password trials are required until the user types in the correct password. These numbers suggest that in case where a server limits a number of failed attempts to say 3, in the worst case roughly 2 passwords from the client’s set could potentially be revealed to the server within a single TLS session — a significant threat for the client. Note that the amount of work for processing failed login attempts on the server side is negligible since all trials are performed through the same secure channel.

The notion of *Password-based Authenticated Key Exchange (PAKE)*, introduced by Bellare and Merritt [13], initially formalised in [11, 20], and later explored in numerous further works (cf. Section 1), is considered as a more secure alternative to the above approach. The standard model of PAKE does not require any PKI and assumes that only a human-memorable password is shared between both parties. PAKE protocols solve the problem of potential password leakage, inherent

to the previously described approach. They aim to protect against offline dictionary attacks but require the same method of protection against online dictionary attacks as the aforementioned TLS-based approach, namely by restricting the number of failed password trials. While passwords can be retransmitted and checked by the server, using the same TLS channel, the only way for current PAKE protocols to deal with failed password trials is to repeat the entire protocol. This however implies that the computational costs on the server side, in particular for (costly) public key-operations that are inherent to all PAKE protocols, increase *linearly* with the number of attempts. This can be seen as a reason for the limited progress on the adoption of PAKE on the Internet (in addition to unrelated issues such as browser incompatibility, patent considerations, and the lack of adopted standards).

While handling multiple password trials with PAKE may seem like a pure implementation problem at first sight, the problem becomes non-trivial if we want to avoid linear increase of public key operations on the server side. This seems to be avoidable only if in a single PAKE execution the client can use several passwords, while the server would use only the one password, shared with the client. Yet this idea alone is not sufficient for breaking the linear bound on the server side: for instance, assume that one PAKE execution is built out of n independent (possibly parallelised) runs of some secure PAKE protocol, where the client uses a different password in each run but the server uses the same one in all of them. The amount of work for the server in this case would still remain $O(n)$. Therefore, something non-trivial must additionally happen in order to reduce the amount of work on the server side to $O(1)$.

However, we still need to fulfil basic PAKE requirements like addressing the persistent threat of online dictionary attacks by enforcing that the number of passwords that can be tested by the client in one session remains below some threshold, which is set by the server. For the server there is no difference whether a client is given the opportunity to perform at most c independent PAKE sessions (password trials) with one input password per session, or only one session but with at most c input passwords. Finally, we must be able to prevent a possibly malicious server from obtaining any password from the set of passwords input by the client.

Related Work The concept of PAKE was introduced by Bellare [13] and corresponding security models were initially developed by Bellare, Pointcheval, and Rogaway [11], Boyko, MacKenzie, and Patel [20], and Goldreich and Lindell [35]. The Find-then-Guess PAKE model from [11], strengthened and simplified in [7] towards the Real-or-Random PAKE model aims at Authenticated Key Exchange (AKE)-security, the main security property of PAKE protocols. The models in [7, 11] remain the most popular game-based PAKE models, adopted in the analysis of many protocols, including the random oracle-based protocols [2, 6] and protocols requiring a common reference string (CRS) [32, 33, 39]. The only (but fairly inefficient) protocol that is built from general secure multi-party computation techniques but does not require any setup assumptions nor random oracles is due to Goldreich and Lindell [35], which was subsequently simplified at the cost of weakened security in [44]. A stronger PAKE model in the framework of Universal Composability (UC) [23] is due to Canetti, Halevi, Katz, Lindell and MacKenzie [24]. UC-secure PAKE protocols require setup assumptions, with CRS being the most popular one [40], albeit ideal ciphers / random oracles [4] and stronger hardware-based assumptions [27] have also been used. In general, all PAKE models (see [45] for a recent overview) take into account unavoidable online dictionary attacks and aim to guarantee security against offline dictionary attacks. While many PAKE constructions

require a constant number of communication rounds [8, 32, 33, 39, 40]; recent frameworks by Katz and Vaikuntanathan [40] and Benhamouda et al. [16] offer optimal one-round PAKE.

There exist further PAKE flavours, including three-party protocols [3, 7] and group protocols [1, 9]. Furthermore, threshold-based PAKE protocols, e.g. [6, 10], where the client’s password is shared amongst two (or possibly more) servers that jointly authenticate the client. In addition to the aforementioned approaches that are tailored to the password-based setting there exist several more general authentication and key exchange frameworks such as [18, 22] that also lend themselves to the constructions of (somewhat less practical) PAKE protocols. Note that no protocols or frameworks exist that would allow to perform an authentication with a single password on the server and a set of passwords on the client side efficiently. We therefore observe that for many of these aforementioned password-based concepts efficient processing of failed password attempts (without repeating the execution) remains an open problem and our technique could possibly be used in their contexts.

1.1 Oblivious PAKE and Our Contributions

We solve the aforementioned problem of efficient handling of password trials on the server side by proposing a compiler that transforms suitable PAKE protocols in a black-box way into what we call an *Oblivious PAKE* (O-PAKE). To describe and analyse the proposed O-PAKE notion we introduce a new algorithmic way to model PAKE protocols that also allows for easy compilation as done with O-PAKE, and real-world implementation. The functionality of O-PAKE protocols resembles that of PAKE except that the client inputs a set \mathbf{pw} of $n \in [1, c]$ passwords from some dictionary while the input on the server-side is limited to a single password pw . Note that this does *not* increase the probability for online dictionary attacks because the server still limits the number of trials and the *maximum number* of passwords c input by the client. Note that the client is still allowed to input less than c passwords, e.g., only one if the client is confident about its validity. The protocol execution of O-PAKE succeeds if and only if the server’s password pw is part of the client’s password set \mathbf{pw} . We utilise the well-known (game-based) PAKE model by Bellare, Pointcheval, and Rogaway [11] in its (stronger) Real-or-Random flavour from [7] and update it to match the O-PAKE setting by considering \mathbf{pw} as input on the client side. In this model passwords are assumed to be distributed uniformly at random. In practice, a client that uses passwords with different strengths in the same O-PAKE session would obviously lower the overall security to the least secure password.

The crucial idea behind our transformation is to let each client execute n sessions of some given secure PAKE protocol in parallel and let the server execute only *one* PAKE session. The challenging part is to enable the server to actually identify the correct PAKE session in which the client used the correct password pw , while preserving security against offline dictionary attacks for all passwords in the client’s password set \mathbf{pw} . This is the trickiest part of the compiler. Intuitively, if the server can recover the messages of the correct PAKE session, it can answer them according to the specification of the PAKE protocol. By repeating this approach in each communication round of the given PAKE protocol both parties will be able to successfully accomplish the protocol. If identification of the correct PAKE session by the server requires only a constant amount of (costly) operations, then the total amount of server’s work in the resulting O-PAKE protocol will also remain constant. The amount on the client side remains linear in the size n of input passwords. This stems from the obvious fact that the client has to compute messages for all PAKE sessions without knowing the correct password. To exemplify the O-PAKE compiler we give two concrete

instantiations. The first instantiation is based on the random-oracle-based SPAKE protocol from [8] and the second one is based on the standard model protocol from [41].

2 Oblivious PAKE Model

In this section we recall the PAKE security model from [7], tailored to the needs of O-PAKE. The security model for O-PAKE protocols is in the multi-user setting and utilises the Real-or-Random approach for AKE-security from [7, 11]. Note that the AKE-security definition addresses the aforementioned security against malicious servers, trying to retrieve client passwords. A server learning information about the additional passwords in the client’s password set \mathbf{pw} can easily break AKE-security by using this password in another session with the same client.

Participants and Passwords An O-PAKE protocol is executed between two parties P and P' , chosen from the universe of participants $\Omega = \mathcal{S} \cup \mathcal{C}$, where \mathcal{S} denotes the universe of servers and \mathcal{C} the universe of clients, such that if $P \in \mathcal{C}$ then $P' \in \mathcal{S}$, and vice versa. We assume the scenario where every client in \mathcal{C} is registered with every servers from \mathcal{S} . For each such pair $(P, P') \in \mathcal{C} \times \mathcal{S}$, a password $\text{pw}_{P,P'}$ (shared between client P and server P') is drawn uniformly at random from the dictionary \mathcal{D} of size $|\mathcal{D}|$. Execution of an oblivious PAKE protocol between P and P' uses $\text{pw}_{P,P'}$ on the server and a password vector $\mathbf{pw}_P \subseteq \{\text{pw}_{P,P'_{x_1}}, \dots, \text{pw}_{P,P'_{x_n}}\}$ for $1 \leq n \leq c$ and client-server pairs (P, P'_{x_i}) for $i \in [2, n]$ on the client side. For the protocol to be successful it is necessary that $\text{pw}_{P,P'} \in \mathbf{pw}_P$. The value c is a global parameter with $c \leq |\mathcal{S}|$. We will sometimes write \mathbf{pw} and pw instead of \mathbf{pw}_P and $\text{pw}_{P,P'}$ when the association with the participants is clear or if it applies to every participant. We will further write PAKE for O-PAKE protocols with $n = 1$, i.e. standard class of PAKE protocols where the client uses $\mathbf{pw}_P = \text{pw}_{P,P'}$.

Protocol Instances For $i \in \mathbb{N}$, we denote by P_i the i -th instance of $P \in \Omega$. In order to model uniqueness of P_i within the model we use i as a counter. For each instance P_i we consider further a list of parameters:

- pid_P^i is the partner id of P_i , defined upon initialisation, subject to following restriction: if $P_i \in \mathcal{C}$ then $\text{pid}_P^i \in \mathcal{S}$, and if $P_i \in \mathcal{S}$ then $\text{pid}_P^i \in \mathcal{C}$.
- sid_P^i is the session id of P_i , modelled as ordered (partial) protocol transcript $[m_{\text{in}}^1, m_{\text{out}}^1, \dots, m_{\text{in}}^r, m_{\text{out}}^r]$ of incoming and outgoing messages of P_i in rounds 1 to r . sid_P^i is thus updated on each sent or received protocol message.
- k_P^i is the value of the session key of instance P_i , which is initialised to **null**.
- state_P^i is the internal state of instance P_i .
- used_P^i indicates whether P_i has already been used.
- role_P^i indicates whether P_i acts as a **client** or a **server**.

Partnered Instances Two instances P_i and P'_j are *partnered* if all of the following holds: (i) $(P, P') \in \mathcal{C} \times \mathcal{S}$, (ii) $\text{pid}_P^i = P'$ and $\text{pid}_{P'}^j = P$, and (iii) $\text{match}(\text{sid}_P^i, \text{sid}_{P'}^j) = 1$, where Boolean algorithm match is defined according to the matching conversations from [12], i.e. outputs 1 if and only if round messages (in temporal order) in sid_P^i equal to the corresponding round messages in $\text{sid}_{P'}^j$ except for the final round, in which the incoming message of one instance may differ from the outgoing message of another instance.

Oblivious PAKE We define O-PAKE using an initialisation algorithm `init` and a stateful interactive algorithm `next`, which handles protocol messages and eventually outputs the session key.

Definition 1 (Oblivious PAKE). An O-PAKE protocol $\mathsf{O-PAKE} = (\mathsf{init}, \mathsf{next})$ over a message space $\mathcal{M} = (\bigcup_r \mathcal{M}_C^r) \cup (\bigcup_r \mathcal{M}_S^r)$, where \mathcal{M}_C^r resp. \mathcal{M}_S^r denotes the space of outgoing server's resp. client's messages in the r -th invocation of `next`, a dictionary \mathcal{D} , and a key space \mathcal{K} consists of two polynomial-time algorithms:

$P_i \leftarrow \mathsf{init}(\mathbf{pw}, \mathbf{role}, P', \mathbf{par})$: On input \mathbf{pw} , $\mathbf{role} \in \{\mathbf{client}, \mathbf{server}\}$, $P' \in \Omega$ and the public parameters \mathbf{par} , the algorithm initialises a new instance P_i with the internal O-PAKE state information `state`, defines the intended partner id as $\mathbf{pid}_P^i = P'$ and session key $\mathbf{k}_P^i = \mathbf{null}$, and stores protocol parameters \mathbf{par} . The `role` indicates whether the participant acts as `client` or `server`.

$(m_{\text{out}}, \mathbf{k}_P^i) \leftarrow \mathsf{next}(m_{\text{in}})$: On input $m_{\text{in}} \in \mathcal{M}_{[S,C]}^r \cup \emptyset$ with implicit access to internal `state`, the algorithm outputs the next protocol message $m_{\text{out}} \in \mathcal{M}_{[S,C]}^{r+1} \cup \emptyset$ and updates \mathbf{k}_P^i with $\mathbf{k}_P^i \in \mathcal{K} \cup \mathbf{null} \cup \perp$. As long as the instance has not terminated the key \mathbf{k}_P^i is `null`. If m_{in} leads to acceptance then \mathbf{k}_P^i is from \mathcal{K} , otherwise $\mathbf{k}_P^i = \perp$. We also assume that `next` implicitly updates the internal `state` prior to each output and sets `used` to `true`.

Note that $\mathcal{M} = (\bigcup_r \mathcal{M}_S^r) \cup (\bigcup_r \mathcal{M}_C^r)$ is the union of outgoing client's message spaces \mathcal{M}_C^r and server's message spaces \mathcal{M}_S^r over all protocol rounds r . We may further view each round's message space \mathcal{M}_C^r as a Cartesian product $\mathcal{M}_C^{r,1} \times \dots \times \mathcal{M}_C^{r,l}$ for up to l different classes of message components, e.g. to model labels, identities, group elements, etc. When clear from the context, we will write \mathcal{M}_C^r instead of $\mathcal{M}_C^{r,1} \times \dots \times \mathcal{M}_C^{r,l}$.

Correctness Let P_i be an instance initialised through $\mathsf{init}(\mathbf{pw}_P, \mathbf{client}, P', \mathbf{par})$ and P'_j be an instance initialised through $\mathsf{init}(\mathbf{pw}_{P,P'}, \mathbf{server}, P, \mathbf{par})$ where $P \in \mathcal{C}$, $P' \in \mathcal{S}$, and $\mathbf{pw}_{P,P'} \in \mathbf{pw}_P$. Assume that all outgoing messages, generated by `next` are faithfully transmitted between P_i and P'_j so that the instances become partnered. An $\mathsf{O-PAKE} = (\mathsf{init}, \mathsf{next})$ is said to be *correct* if for all partnered P_i and P'_j it holds that $\mathbf{k}_P^i \in \mathcal{K}$ and $\mathbf{k}_P^i = \mathbf{k}_{P'}^j$.

Adversary Model The adversary \mathcal{A} is modelled as a probabilistic-polynomial time (PPT) algorithm, with access to the following oracles:

$m_{\text{out}} \leftarrow \mathsf{Send}(P, i, m_{\text{in}})$: the oracle processes the incoming message $m_{\text{in}} \in \mathcal{M}_{[C,S]}^r$ for the instance P_i and returns its outgoing message $m_{\text{out}} \in \mathcal{M}_{[C,S]}^{r+1} \cup \emptyset$. If P_i does not exist, a new session is created with P' as partner, where P' is given in m_{in} .

$\mathbf{trans} \leftarrow \mathsf{Execute}(P, P')$: if $(P, P') \in \mathcal{C} \times \mathcal{S}$ the oracle creates two new instances P_i and P'_j via appropriate calls to `init` and returns the transcript `trans` of their protocol execution, obtained through invocations of corresponding `next` algorithms and faithful transmission of generated messages amongst the two instances.

$\mathbf{pw} \leftarrow \mathsf{Corrupt}(P, P')$: if $P \in \mathcal{C}$ and $P' \in \mathcal{S}$ then return $\mathbf{pw}_{P,P'}$ and mark (P, P') as a corrupted pair.

AKE-Security The following definition of AKE-security follows the *Real-Or-Random* (ROR) approach from [7], which provides the adversary multiple access to the Test oracle for which the randomly chosen bit $b \in_R \{0, 1\}$ is fixed in the beginning of the experiment:

$\mathbf{k}_A \leftarrow \text{Test}_b(P, i)$, depending on the values of bit b and \mathbf{k}_P^i , this oracle responds with key \mathbf{k}_A defined as follows:

- If, while $\mathbf{k}_P^i = \text{null}$, either (P, P') or (pid_P^i, P) were queried to the **Corrupt** oracle for, w.l.o.g., any client-server pair (P, P') with $\text{pw}_{P, P'} \in \mathbf{pw}_P$, then abort. Note that this prevents \mathcal{A} from obtaining any $\text{pw}_{P, P'} \in \mathbf{pw}$ and then testing new instances of P and P' , or instances that were still in the process of establishing session keys when corruption took place.
- If some previous query $\text{Test}(P', j)$ was asked for an instance P'_j , which is partnered with P_i , then return the same response as to that query. Note that this guarantees consistency of oracle responses.
- If $\mathbf{k}_P^i \in \mathcal{K}$ then if $b = 1$, return \mathbf{k}_P^i , else if $b = 0$, return a randomly chosen element from \mathcal{K} and store it for later use.
- Else return \mathbf{k}_P^i . Note that in this case \mathbf{k}_P^i is either \perp or **null**.

According to [7] a session is an online session when \mathcal{A} queried the **Send** oracle on one of the participants.

Definition 2 (AKE-Security). *An O-PAKE protocol Π with up to c passwords on client side is AKE-secure if for all dictionaries \mathcal{D} with corresponding universe of participants Ω and for all PPT adversaries \mathcal{A} using at most t online sessions there exists a negligible function $\varepsilon(\cdot)$ such that:*

$$\text{Adv}_{\Pi, \mathcal{A}}^{\text{AKE}}(\lambda) = \left| \Pr[\text{Exp}_{\Pi, \mathcal{A}}^{\text{AKE}}(\lambda) = 1] - \frac{1}{2} \right| \leq \frac{c \cdot \mathcal{O}(t)}{|\mathcal{D}|} + \varepsilon(\lambda).$$

$\text{Exp}_{\Pi, \mathcal{A}}^{\text{AKE}}(\lambda) : c \in \mathbb{N}; b \in_R \{0, 1\}; \forall (P, P') \in \mathcal{C} \times \mathcal{S}$ choose $\text{pw}_{P, P'} \in_R \mathcal{D}; b' \leftarrow \mathcal{A}^{\text{Send, Execute, Corrupt, Test}_b}(\lambda, c)$; return $b = b'$.

The above definition (without **Corrupt**) reverts to RoR AKE-security from [7] for $c = 1$. We have to factor in the maximal size of $|\mathbf{pw}| = n \leq c$ into the original adversarial advantage bound $\mathcal{O}(t)/|\mathcal{D}|$ to account for the adversarial possibility of testing up to c passwords per session in the role of the client.

PAKE vs O-PAKE The actual relation between common PAKE and O-PAKE security may not be immediately evident. For clarification, we discuss the relation between O-PAKE and the simple repetition of a PAKE protocol c times, and the implication of user's password choice.

The advantage of an adversary that is allowed to query up to c passwords in one session is not greater than the advantage of an adversary that runs c online sessions using one password in each of them. The typical advantage of a PAKE adversary \mathcal{A} in an AKE-security experiment, e.g. [7, 11], is bounded by $\mathcal{O}(t)/|\mathcal{D}| + \varepsilon(\lambda)$. In contrast, we limit the advantage of an O-PAKE adversary to $c \cdot \mathcal{O}(t)/|\mathcal{D}| + \varepsilon'(\lambda)$. We give the following lemma to formalise the relation between the two notions.

Lemma 1. $\text{Adv}_{\Pi_c, \mathcal{A}}^{\text{AKE}} \leq c \cdot \text{Adv}_{\Pi, \mathcal{A}}^{\text{AKE}}$ for O-PAKE protocol Π_c allowing up to c passwords in one session, built from PAKE protocol Π .

Proof. The lemma follows directly from the following observations. O-PAKE can be realised in the naïve way by running c separate PAKE sessions. That results in an advantage of at most $c \cdot \text{Adv}_{\Pi, \mathcal{A}}^{\text{AKE}} = c \cdot \mathcal{O}(t)/|\mathcal{D}| + \varepsilon'(\lambda)$. Information gathered from `Send` and `Execute` oracle invocations are the same for the O-PAKE and PAKE adversary. `Corrupt` and `Testb` queries of the O-PAKE adversary return one password, respectively key, independent from c , while the PAKE adversary gets c passwords, respectively keys. Thus, the resulting advantage of the O-PAKE adversary is at most $c \cdot \mathcal{O}(t)/|\mathcal{D}| + \varepsilon'(\lambda)$, but depending on the implementation most probably lower. \square

Assuming malicious servers one may also be concerned about the client’s password choice considering a client entering passwords with different levels of entropy. Similar to the standard PAKE case the weakest password from \mathbf{pw} would determine the security of O-PAKE. However, the used model considers uniformly at random chosen passwords from one dictionary such that the case of varying password probabilities can not be adequately addressed in this model (as is also the case for the models in [7, 11]).

3 Transforming PAKE Protocols into O-PAKE

Recall that one may realise O-PAKE in a naïve way by running the input PAKE protocol n times, which is not efficient on the server side due to the linearly increasing round complexity. The idea of the O-PAKE compiler is to mix the n PAKE messages on client side such that the server can extract the “right” message using the shared password and reply only to that. This, however, is a non-trivial problem because PAKE messages do not provide information that would allow the server to check locally whether a given password was used in their computation; as this would offer the possibility of offline dictionary attacks. Note that we assume throughout this section that $n \geq 2$ and $\text{pw}_{P, P'} \in \mathbf{pw}_P$. Our solution for the identification of the “right” PAKE session is a careful composition of two encoding techniques that were introduced in a different context yet allow us to generically construct AKE-secure O-PAKE protocols from (suitable) AKE-secure PAKE protocols, preserving constant round complexity and offering nearly constant server load.

Our first building block is *Index-Hiding Message Encoding (IHME)* [42, 43]. An IHME scheme assigns a different index to each given message and encodes the resulting index-message pairs into a single structure from which messages can be recovered on the receiver side using the corresponding indices. The IHME structure hides indices that were used for encoding and therefore all encoded messages must contain enough entropy to prevent dictionary attacks over the index space. An IHME scheme consists of two algorithms `iEncode` and `iDecode`. The `iEncode` algorithm takes as input a set of index-message pairs $(ix_1, m_1), \dots, (ix_n, m_n)$ and outputs a structure S whereas the `iDecode` algorithm can extract $m_j, j \in [1, n]$ from S using the corresponding index ix_j . For formal definitions surrounding IHME we refer to the original work or Appendix B and only mention that the original IHME construction in [42] assumes $(ix_j, m_j) \in \mathbb{F}$ for a prime-order finite field \mathbb{F} and defines the IHME structure S through coefficients of the interpolated polynomial by treating index-message pairs as its points. There exists a more efficient IHME version from [43] for longer messages, which uses $(ix_j, m_j) \in \mathbb{F} \times \mathbb{F}^\nu$ and thus splits m_j into ν components each being an element of \mathbb{F} . The corresponding index-hiding property demands that no information about indices ix_j is leaked to the adversary that doesn’t know the corresponding messages m_j and is defined for messages that are chosen uniformly from the IHME message space. For the aforementioned IHME schemes the message space is given by \mathbb{F} (or \mathbb{F}^ν) and their index-hiding property is perfect (in the information-

theoretic sense). Note that this approach still allows the server to learn which of the n PAKE sessions is the correct one without revealing any password to the server.

In order to enable encoding of PAKE messages using IHME with passwords as indices we apply our second building block, namely *admissible encoding* [19, 21, 29]. Briefly, a function $F : S \rightarrow R$ is an ϵ -admissible encoding for (S, R) with $|S| > |R|$ when for all uniformly distributed $r \in R$, the distribution of the inverse transformation $\mathcal{I}_F(r)$ is ϵ -statistically indistinguishable from the uniform distribution over S . We refer to [21, 29] or Appendix A for more details. \mathcal{I}_F enables us to map PAKE messages into the IHME message space where necessary. In Section 3.5 we will discuss suitable PAKE message spaces and their admissible encodings offering compatibility with the message space \mathbb{F} of the IHME schemes from [42, 43].

In the following we describe our compiler that transforms suitable AKE-secure PAKE protocols into AKE-secure O-PAKE protocols. The intuition behind the compiler is to let the client run n PAKE sessions, one session for each of the n input passwords \mathbf{pw} , and apply an index-hiding message encoding on each message-password pair. The server can apply the shared password \mathbf{pw} as index to IHME to extract the “right” PAKE message. For this message the server executes the algorithm `next` of the given PAKE protocol and returns the resulting PAKE message to the client. As soon as the algorithm `next` terminates, the server generates a confirmation message, which is then used by the client to derive the final session key.

3.1 Requirements on PAKE

Our O-PAKE can be used to convert any AKE-secure R -round PAKE protocol Π where in each round $r \in [1, \dots, R]$ the client sends messages from \mathcal{M}_C^r that can be processed using a compatible admissible encoding $F^r : \mathcal{M}^{\text{IHME}, r} \rightarrow \mathcal{M}_C^r$. In order to guarantee that client messages from \mathcal{M}_C^r , when mapped into $\mathcal{M}^{\text{IHME}, r}$ using the inverse transformation \mathcal{I}_F , are uniformly distributed over $\mathcal{M}^{\text{IHME}, r}$, the underlying Π itself must output client messages whose joint distribution over all R rounds remains indistinguishable from a distribution where for each round r the output client message is chosen uniformly at random from \mathcal{M}_C^r . For this purpose Π must satisfy a stronger notion of AKE security that in addition to the indistinguishability of session keys requires indistinguishability of client messages. This requirement is formalised in Definition 3 that extends the AKE-security experiment for PAKE from Definition 2, using `Execute`, `Send`, `Corrupt` and `Testb` definitions from there. We assume that $c = 1$ and define two oracles `Sendb` and `Executeb` that are parameterised with the bit b as used in the `Testb` oracle. Any query `Sendb(P, i, min)` for a client $P \in \mathcal{C}$ made by the adversary \mathcal{A} first triggers the invocation of $m_{\text{out}} \leftarrow \text{Send}(P, i, m_{\text{in}})$. If \mathcal{A} queried `Corrupt(P, pidPi)` or `Corrupt(pidPi, P)` while $\mathbf{k}_P^i = \text{null}$ or if $b = 1$ then m_{out} is returned to \mathcal{A} without any modification. The additional condition on the `Corrupt` queries prevents \mathcal{A} from trivially distinguishing the client messages by corrupting passwords and then communicating with client instances that were still in the process of establishing the session keys. If $b = 0$ then m_{out} is set to a random message from \mathcal{M}_C^r and returned to \mathcal{A} . Any `Executeb(P, P')` query first triggers the invocation of $\mathbf{trans} \leftarrow \text{Execute}(P, P')$. If $b = 0$ then for each round r the corresponding client’s message in \mathbf{trans} is replaced with an independently at random chosen message from \mathcal{M}_C^r , else if $b = 1$ then \mathbf{trans} is forwarded without any modification. Note that if \mathcal{A} mounts an online attack with a correct password then it can easily distinguish so that the lower bound of $\frac{\mathcal{O}(t)}{|\mathcal{D}|}$ that accounts for online dictionary attacks still applies in the definition.

Definition 3 (AKE-Security with Indistinguishable Client Messages). *A PAKE protocol Π is AKE-secure with indistinguishable client messages if for all dictionaries \mathcal{D} with corresponding*

universe of participants Ω and for all PPT adversaries \mathcal{A} using at most t online sessions there exists a negligible function $\varepsilon(\cdot)$ such that:

$$\text{Adv}_{\Pi, \mathcal{A}}^{\text{AKE-ICM}}(\lambda) = \left| \Pr[\text{Exp}_{\Pi, \mathcal{A}}^{\text{AKE-ICM}}(\lambda) = 1] - \frac{1}{2} \right| \leq \frac{\mathcal{O}(t)}{|\mathcal{D}|} + \varepsilon(\lambda).$$

$\text{Exp}_{\Pi, \mathcal{A}}^{\text{AKE}}(\lambda) : c = 1; b \in_R \{0, 1\}; \forall (P, P') \in \mathcal{C} \times \mathcal{S}$ choose $\text{pw}_{P, P'}$;
 $b' \leftarrow \mathcal{A}^{\text{Send}_b, \text{Execute}_b, \text{Corrupt}, \text{Test}_b}(\lambda, c)$; return $b = b'$.

The above requirement is stronger than AKE-security. In particular, it cannot be satisfied by PAKE protocols where client messages depend on those of the server or where client messages sent in later rounds depend on client messages that were sent in previous rounds. Nonetheless, there exist efficient AKE-secure PAKE protocols with indistinguishable client messages as discussed in Section 3.5. In particular, for AKE-secure *one-round* PAKE protocols, where the client can send its message independently of the server's message the indistinguishability property can be argued based on the uniformity of the client's message in the message space.

Note that we do not consider verifier-based PAKE protocols, such as [14, 17, 34, 37], where only some verification information derived from the password is stored on the server.

3.2 The O-PAKE Compiler

Our compiler takes as input a PAKE protocol Π and outputs its O-PAKE version, denoted C_Π . The compiled protocol C_Π follows Definition 1 and consists of the two algorithms $C_\Pi.\text{init}$ and $C_\Pi.\text{next}$. For the passwords in \mathbf{pw} used as input to $C_\Pi.\text{init}$ we assume that each $\mathbf{pw}[i] = (\text{ix}, \pi) \in \mathbb{F} \times \mathcal{D}_\Pi$, where ix denotes an *index* and π the corresponding *password* for the underlying PAKE protocol Π , whereby the distributions of ix and π are independent and no two pairs $(\text{ix}_1, \pi_1), (\text{ix}_2, \pi_2) \in \mathbf{pw}$ have $\text{ix}_1 = \text{ix}_2$. For each PAKE round r the compiler uses a corresponding instance IHME^r with message space $\mathcal{M}^{\text{IHME}^r}$ and a compatible admissible encoding $F^r : \mathcal{M}^{\text{IHME}^r} \rightarrow \mathcal{M}_C^r$ where \mathcal{M}_C^r is the space of clients messages of Π in that round. In the following we assume that the underlying $\Pi.\text{next}$ algorithm outputs messages that can be seen as one element and thus can be processed using one instance (F^r, IHME^r) in each round. Note that this allows for a more comprehensible description and is not a restriction of the O-PAKE compiler. We discuss the case of multi-set messages $\mathcal{M}_C^r = \mathcal{M}_C^{r,1} \times \dots \times \mathcal{M}_C^{r,l}$ that will require composition of up to l instances of encoding schemes per round in Section 3.6.

The $C_\Pi.\text{next}$ algorithm on the client side computes corresponding PAKE round messages for all passwords in \mathbf{pw} using the original $\Pi.\text{next}$ algorithm and encodes them with \mathcal{I}_{F^r} and $\text{IHME}^r.\text{iEncode}$ prior to transmission to the server. On the server side $C_\Pi.\text{next}$ decodes the incoming PAKE message using F^r and $\text{IHME}^r.\text{iDecode}$ (using its input $\mathbf{pw}[i].\text{ix}$ as index) and replies with the message output by $\Pi.\text{next}$. Note that the server only decodes messages but never encodes them. If $pw \in \mathbf{pw}$ then at the end of its n PAKE sessions the client will hold n intermediate PAKE keys, whereas the server holds only one such key. The additional key confirmation and key derivation steps allow the client to determine which of its n PAKE session keys matches the one held by the server, in which case both participants will derive the same session key. In the following we describe the two algorithms $C_\Pi.\text{init}$ and $C_\Pi.\text{next}$ more in detail.

a $C_{\Pi}.\text{next}(m_{\text{in}})$ — Client	b $C_{\Pi}.\text{next}(m_{\text{in}})$ — Server
Input: m_{in} Output: $(m_{\text{out}}, \mathbf{k})$ $E = \emptyset; m_{\text{out}} = \emptyset$ for $i = 1 \dots n$ do if $\Pi[i]$ has not finished then $(m'_{\text{out}}, \Pi[i].\mathbf{k}) \leftarrow \Pi[i].\text{next}(m_{\text{in}})$ if $m'_{\text{out}} \neq \emptyset$ then $E = E \cup \{(\mathbf{pw}[i].\text{ix}, \mathcal{I}_{F^r}(m'_{\text{out}}))\}$ else if $\Pi[i].\mathbf{k} \in \mathcal{K}_{\Pi}$ and $m_{\text{in}} = \text{PRF}_{\Pi[i].\mathbf{k}}(\text{sid}_{P'}^i P_i \text{pid}_{P'}^i 0)$ then $\mathbf{k} = \text{PRF}_{\Pi[i].\mathbf{k}}(\text{sid}_{P'}^i P_i \text{pid}_{P'}^i 1)$ else $\mathbf{k} = \perp$ if $E \neq \emptyset$ then $m_{\text{out}} = \text{IHME}^r.\text{iEncode}(E)$ return $(m_{\text{out}}, \mathbf{k})$	Input: m_{in} Output: $(m_{\text{out}}, \mathbf{k})$ $m_{\text{out}} = \emptyset$ if Π has not finished then $m \leftarrow \text{IHME}^r.\text{iDecode}(\text{pw}.\text{ix}, m_{\text{in}})$ $m' = F^r(m)$ $(m_{\text{out}}, \Pi.\mathbf{k}) \leftarrow \Pi.\text{next}(m')$ if $\Pi.\mathbf{k} \in \mathcal{K}_{\Pi}$ then $m_{\text{out}} = \text{PRF}_{\Pi.\mathbf{k}}(\text{sid}_{P'}^j \text{pid}_{P'}^j P_j 0)$ $\mathbf{k} = \text{PRF}_{\Pi.\mathbf{k}}(\text{sid}_{P'}^j \text{pid}_{P'}^j P_j 1)$ else $\mathbf{k} = \perp$ return $(m_{\text{out}}, \mathbf{k})$

Fig. 1: $C_{\Pi}.\text{next}$ Algorithms

Algorithm $C_{\Pi}.\text{init}$ The algorithm makes n calls to $\Pi.\text{init}$, one for each password $\mathbf{pw}[i].\pi$, to generate corresponding **states** for each of the n PAKE sessions that are stored in $\mathbf{state}_{P'}^i$. An i th session of Π run by the client using the corresponding password $\mathbf{pw}[i].\pi$ is denoted by $\Pi[i]$. The partner id $\text{pid}_{P'}^i$ is set to P' and the instance P_i with the given **role** and a vector of n local states in $\mathbf{state}_{P'}^i$ is established. We require that no two passwords in \mathbf{pw} are identical, which is necessary to ensure the correctness of the IHME step. Note that if **role** = **server** then $n = 1$, i.e. servers run only one PAKE session.

Algorithm $C_{\Pi}.\text{next}$ We distinguish between $C_{\Pi}.\text{next}$ specifications for clients (Algorithm 1a) and servers (Algorithm 1b) as they are significantly different. We write $\Pi[i].\text{next}$ for the invocation of $\Pi.\text{next}$ for the i th session of Π run by the client using $\mathbf{pw}[i].\pi$. On the client side $C_{\Pi}.\text{next}$ computes messages m'_{out} for all running PAKE sessions and encodes them. The server decodes the incoming IHME structure and computes its response using $\Pi.\text{next}$. If any PAKE session $\Pi[i]$ at the client has finished with $\Pi[i].\mathbf{k} \in \mathcal{K}_{\Pi}$ then the client expects a valid confirmation message from the server prior to derivation of the resulting session key \mathbf{k} with PRF using $\Pi[i].\mathbf{k}$. An invalid confirmation message implies that \mathbf{k} is set to \perp . This confirmation message is generated on the server side using PRF only if and immediately after $\Pi.\text{next}$ outputs $\Pi[i].\mathbf{k} \in \mathcal{K}_{\Pi}$; in which case a valid resulting session key \mathbf{k} is also derived. If, however, Π finishes with $\Pi[i].\mathbf{k} = \perp$ then \mathbf{k} will also be set to \perp .

3.3 Relation to LAKE

A Language Authenticated Key Exchange (LAKE) protocol, proposed by Benhamouda et al. in [15], authenticates two parties, client C and server S holding each a word in an algebraic languages. In particular, let $R : \{0, 1\}^* \times P \times W \rightarrow \{0, 1\}$ denote a relation and $L_R(\text{pub}, \text{priv}) \subseteq W$ a language with $\text{pub} \in \{0, 1\}^*$ and $\text{priv} \in P$. A word $w \in W$ is in the language L_R iff $R(\text{pub}, \text{priv}, w) = 1$. The client holds a word w_c for relation R_C and the server holds a word w_s for relation R_S . They agree on public parameters pub , exchange ephemeral public keys, and *think* of a value priv'_C , resp. priv'_S , they

expect to be used by the other party. To instantiate the LAKE framework it is necessary to specify client and server languages and according commitments with associated smooth projective hash functions (SPHF) [26]. We briefly recall how to instantiate LAKE with passwords from [15, Section 6.2], i.e. how to build PAKE protocols in the LAKE framework. The languages are defined as $L_C = \{w_c\}$ for the client and $L_S = \{w_s\}$ for the server, such that $\text{priv}'_C = \text{priv}'_S = w_c = w_s$ is the password and the relations are $R_C = R_S = (\emptyset, \text{priv}, w) = 1 \iff \text{priv} = w$, i.e. equality test for the password.

To instantiate O-PAKE in LAKE we define client relation $R_C(\emptyset, \text{pw}', \mathbf{pw}) = 1 \iff \text{pw}' \in \mathbf{pw}$ and server relation $R_S(\emptyset, \text{pw}', \text{pw}) = 1 \iff \text{pw}' = \text{pw}$. While the server relation stays the same as in PAKE, the client language $L_{R_C}(\emptyset, \text{pw}') \subseteq \{\text{pw}_1, \dots, \text{pw}_n\} = \mathbf{pw}$ uses a relation that takes a set of passwords \mathbf{pw} and an *expected* password pw' as input, and is fulfilled iff $\text{pw}' \in \mathbf{pw}$. Following [15, Figure 4] we realise O-PAKE in the LAKE framework as follows: First, the client (initiator) generates a multiDLCSCom' commitment (C_C, C'_C) on word w_c , i.e. a multi-commitment to all passwords $\text{pw} \in (\text{pw}_1, \dots, \text{pw}_n)$, as well as a Pedersen commitment C''_C on C'_C , and sends (C_C, C''_C) to the server S . The server replies with $(C_S, \varepsilon, \mathbf{k}_{p_S}, \sigma_S)$, computed as follows: C_S is a multi-LCS commitment on $w_s = \text{pw}_S$; ε is a challenge vector on C_C of length n ; \mathbf{k}_{p_S} is a projection key for a suitable SPHF for C_C ; and σ_S is a signature on all flows. In the final round, the client checks σ_S before returning $(C'_C, t, \mathbf{k}_{p_C}, \sigma_C)$ to the server, which is computed as follows: (C'_C, t) is the decommitment to C''_C , where t is the used randomness; \mathbf{k}_{p_C} is a projection key for a suitable SPHF for C_S ; and σ_C is a signature on all flows. After checking all signatures and commitments, session keys are computed as multiplication of projection and hash function on $\text{Com}_C = C_C \cdot C'_C \varepsilon$ and $\text{Com}_S = C_S$.

So while it seems possible to instantiate O-PAKE in the LAKE framework (after specifying necessary primitives), the construction is rather inefficient. In particular, an instantiation of O-PAKE in LAKE needs four rounds, our O-PAKE compiler adds only one round to the round-complexity of the underlying PAKE, i.e. can be instantiated with three rounds. Further, server-complexity is linear in the number of client-passwords n . This stems from the observation that the projection key \mathbf{k}_{p_S} , as well as the computation of the hash function, requires a linear number of public key operations, e.g., exponentiations, in n . Performance of O-PAKE instantiated in the LAKE framework is therefore not more efficient than the naïve construction, and in particular does not fulfil our requirement of nearly constant server performance.

3.4 Security Analysis

AKE-security of the protocol generated with the O-PAKE compiler is established in Theorem 1.

Theorem 1. *If Π is an R -round AKE-secure PAKE protocol with pseudorandom client messages in \mathcal{M}_C^r for $r \in [1, \dots, R]$, $F^r : \mathcal{M}^{\text{IHME}^r} \rightarrow \mathcal{M}_C^r$ is an ϵ -admissible encoding, and IHME^r is an index-hiding message encoding, then C_Π is an $R + 1$ -round AKE-secure O-PAKE protocol.*

Proof (sketch). The proof uses a sequence of experiments Exp_i , $i = 1, \dots, 4$, where each Exp_i is based on a small modification of Exp_{i-1} . At a high level we first replace real client messages and session keys of underlying PAKE sessions $\Pi[i]$ with random messages and keys while ensuring the consistency against an adversary that corrupted passwords and then mounts online attacks. This modification remains unnoticeable to \mathcal{A} if the underlying PAKE protocol Π is AKE-secure with indistinguishable client messages. Then, we replace the outputs of the inverse transformations \mathcal{I}_{F^r}

applied in each round r by choosing random elements from the corresponding round’s IHME ^{r} message space and show that this remains unnoticeable assuming that F^r is an ϵ_{F^r} -admissible encoding F^r . Then, we replace each real password index in the IHME ^{r} encoding process with a random password and show that this remains unnoticeable due to the index-hiding property of each IHME ^{r} scheme. Finally, we modify the computation of the server’s confirmation message and of the session keys that are returned in Test _{b} queries by using random elements from the corresponding spaces. This remains unnoticeable due to the pseudorandomness of the PRF function that is used to derive their values. See Appendix D for the full proof.

3.5 Oblivious PAKE Instantiation

An AKE-secure PAKE protocol Π is suitable for our O-PAKE transformation if it is also AKE-ICM-secure and there exist admissible encodings to map those messages into the message space of the IHME scheme. In the following we list four sets R with suitable admissible encodings. Thus, any AKE-secure PAKE protocol whose client messages contain components from these four sets can be transformed into an O-PAKE protocol using our compiler.

Definition 4 (Admissible Encodings for Client Messages). *An admissible encoding $F : \{0, 1\}^{\ell(\lambda)} \rightarrow R$ with polynomial $\ell(\lambda)$ exists for any of the following four sets:*

- (1) Set $R = \{0, \dots, N - 1\} = \mathbb{Z}_N$ of natural numbers, for arbitrary $N \in \mathbb{N}$. (cf. [29, Lemma 12])
- (2) The set of quadratic residues modulo safe primes p , i.e. $R = QR(p) \subseteq \mathbb{Z}_p^\times$. (cf. [29, Lemma 12])
- (3) Arbitrary subgroups $G \subseteq \mathbb{Z}_p^\times$ of prime order q . (cf. [29, Lemma 12])
- (4) The set $R = E(\mathbb{F})$ of rational points on (certain) elliptic curves, defined over a finite field (cf. [21]).

Computing Indices We require that password pw used in O-PAKE consists of two independent components ix and π . For instance, it is sufficient for the user to choose $\pi \in_R \mathcal{D}$ and compute the index pw.ix = $f(\rho, \text{pw}.\pi)$ using some fresh randomness ρ and a function f with output independent from π , i.e. the probability that π was used as input to f to produce ix must remain $1/|\mathcal{D}|$. Note that this approach requires a pre-flow to the protocol to exchange randomness ρ , which can however be easily integrated into the overall login process. Furthermore, it is crucial that randomness ρ is fresh for every execution of the protocol as any reuse of ρ would offers an attacker the possibility to distinguish between real and simulated O-PAKE messages.

3.6 Processing Multi-Component Messages

In the following we describe how the compiler can handle PAKE protocol messages consisting of multiple elements, possibly from different sets. We observe that any such PAKE message can be seen as an element of a combined message space that is formed through a Cartesian product of those sets and distinguish between two types of message components, namely components that represent constants and components that depend on passwords, including integer values and group elements. Since constants are password-independent they do not need to be processed by the compiler and can be communicated directly. All other message components have to be encoded according to the compiler specification. In order to encode those components we use ν -fold IHME introduced in [43], which allows to encode a list of ν message components from the same finite field (cf. Appendix B). The compiler splits message components from different finite fields into corresponding classes

and applies appropriate IHME encoding to each class separately in order to compute the corresponding IHME structure. The IHME structures for all message components are then concatenated and treated as a single compiler message. This processing of multi-component messages requires existence of admissible encodings and index-hiding message encodings for each component class m_j of m . In order to process the components, a loop over m_1, \dots, m_l adds $(\mathbf{pw}[i], \mathcal{I}_{F^{r,j}}(m_j))$ to the input set of ν -fold-IHME $_j^r$.iEncode according to their classes (e.g. finite fields). Likewise, the output message m_{out} of the `next` algorithm is the concatenation of the encoded component classes. Upon receiving a client message m_{in} , the server has to decompose it to retrieve the IHME encoded messages. After decoding the message parts with $m_j \leftarrow \nu$ -fold-IHME $_j^r$.iDecode($\mathbf{pw}_{P,P'}, m_{\text{in}}^j$) the original PAKE message of Π is reassembled by decoding messages $F^{r,j}(m_j)$.

Adopting this approach for multi-component messages, the AKE-security remains preserved. This is due to the following observation about the proof of Theorem 1: in the game-hopping sequence the adversary will be provided with l IHME encoded messages (one for each message element class that requires encoding). The corresponding index-hiding advantage will therefore be multiplied by l . The remaining parts of the proof remain as is.

4 Concrete Instantiation Examples

In this section we give two concrete instantiations of the O-PAKE compiler, using the random-oracle based SPAKE protocol from Abdalla and Pointcheval [8] and the common-reference-string-model protocol from Katz and Vaikuntanathan [41].

4.1 Oblivious SPAKE

We demonstrate how the compiler can be applied to PAKE protocols using the AKE-secure, random-oracle-based SPAKE protocol from [8]. The resulting O-SPAKE is specified in Figure 2 and involves steps of the original SPAKE protocol from [8, Sec. 5], which is a secure variant of [13], whose security has been proven in the random oracle model.¹ SPAKE uses a prime-order cyclic group G for which the Computational Diffie-Hellman (CDH) problem is assumed to be hard. The shared SPAKE password \mathbf{pw} is chosen from \mathbb{Z}_q . Let $M, N \in G$ denote two public group elements. The protocol proceeds in one round, where the client sends $X^* \leftarrow g^x \cdot M^{\mathbf{pw} \cdot \pi}$, $x \in_R \mathbb{Z}_q$ and the server responds with $Y^* \leftarrow g^y \cdot N^{\mathbf{pw} \cdot \pi}$, $y \in_R \mathbb{Z}_q$. The actual order of these messages does not matter since they are independent. The algorithm `next` computes an intermediate value s and derives the session key as $\Pi.k \leftarrow H(P, P', X^*, Y^*, \mathbf{pw}, s)$. We refer to the original work [8, Sec. 5] for more details on SPAKE. The SPAKE protocol is a suitable input PAKE protocol for our O-PAKE compiler since it can be instantiated using subgroups $G \subseteq \mathbb{Z}_p^\times$ of prime order q in which the CDH problem is believed to be hard. We can apply the admissible encodings (3) from Definition 4 due to the fact that client's SPAKE message $X^* = g^x \cdot M^{\mathbf{pw} \cdot \pi}$ is uniformly distributed in G , given the uniformity of $x \in \mathbb{Z}_q$. We formalise this by showing that SPAKE fulfils our definition of AKE-ICM, before defining suitable admissible encodings, which concludes the instantiation of O-SPAKE.

Lemma 2 (SPAKE is AKE-ICM secure). *The SPAKE protocol from [8, Sec. 5] is AKE-ICM secure.*

¹ Note that the very similar SOKE protocol from [2] can also be used in the O-PAKE compiler following the here given description of O-SPAKE.




Fig. 2: Oblivious SPAKE (O-SPAKE)

Public Input: $G, g, p, q, M, N, H, \text{IHME}, F$

Proof. The initial experiment in the proof for SPAKE security in [8] corresponds to the AKE-ICM experiment with $b = 1$. In the following we show that the proof in [8, Appendix C] can be modified without changing the adversary's advantage such that the final experiment is equal to the AKE-ICM experiment with $b = 0$, which concludes the proof. We first change experiment one by additionally simulating the **Corrupt** oracle and using a global bit b in simulating the **Test** oracle. This does not change the adversary's success probability. The second experiment, aborting on hash collisions, stays unchanged. In the following two experiments we have to make sure that the adversary does not win trivially by returning the correct key to **Test** queries on corrupted sessions and only modify oracle replies to uncorrupted sessions. While the original proof only changes the calculation of the session key in passive sessions to a random element in experiment three, we also change client messages produced in **Execute** queries to random elements. Note that this is implicitly already done in the original proof. However, we formalise it here again and change experiment three as follows: Invocations of the **Execute** oracle on uncorrupted parties are answered with uniformly at random chosen messages, i.e. $X^* = Ag^x$ and $Y^* = Bg^y$ with $x, y \in_R \mathbb{Z}_p$, for some DH instance (A, B) . Experiment three corresponds now to the AKE-ICM experiment with Send_1 , Test_b and Execute_0 . The lemma follows by noting that after our modifications of experiment three the last experiment of the AKE-security proof of SPAKE in [8, Appendix C] is equivalent to the AKE-ICM experiment with Execute_0 , Send_0 and Test_0 , i.e. the adversary only wins by guessing the correct password. \square

Admissible Encodings for SPAKE We use admissible encodings (1) and (3) from Definition 4 to encode SPAKE client messages. To implement the inverse encoding of (1) $:= \mathcal{I}_{F(1)} : \mathbb{Z}_N \rightarrow \{0, 1\}^{\ell(\lambda)}$ we use the inverse of encoding (3) $:= \mathcal{I}_{F(1)} : G \rightarrow \mathbb{Z}_p^\times$. This results in a combined inverse encoding of $\mathcal{I}_{F(3,1)} : G \rightarrow \mathbb{Z}_p^\times \rightarrow \{0, 1\}^{\ell(\lambda)}$ with $\ell(\lambda) > 2|N|$ and $p = N$. Implementation of $F^{(3,1)} : \mathbb{Z}_{q'} \rightarrow G$ and $\mathcal{I}_{F(3,1)}$ follows the specification from [29, Lemma 12] with prime $|q'| = \ell(\lambda) > 2|N|$ to meet IHME requirements.

4.2 Oblivious KV-PAKE

We give a second instantiation of the O-PAKE compiler using a round-optimal common-reference-string-model PAKE protocol from Katz and Vaikuntanathan [41]. Let Π denote the KV protocol, where client and server each encrypt their password \mathbf{pw} with a labelled CCA-secure encryption scheme, generate hash and projection keys for an according SPHF, and exchange projection keys




Fig. 3: Oblivious KV-PAKE (O-KV-PAKE) using CS Encryption

Public Input: $G, g_1, g_2, c, d, h, p, H, \text{IHME}, F$

and ciphertexts. The session key is then computed by multiplying hash and projection values of the two ciphertexts. To use KV-PAKE in our compiler we have to ensure that ciphertexts and projection keys are indistinguishable from random elements, and appropriate admissible encodings exist. Note that KV-PAKE messages consist of multiple elements such that every element is encoded separately (cf. Appendix 3.6). In the following we use a DDH-based KV-PAKE instantiation from [16] on a group G of prime order q , with generator g , using labelled Cramer-Shoup encryption [25]. Note that the DDH-based construction from [41, Section 4.1] is *not* usable in our compiler (cf. Section 4.2). The following lemma allows us to construct oblivious KV-PAKE (cf. Figure 3) with the same admissible encodings as for the aforementioned SPAKE instantiation.

Lemma 3. *The KV-PAKE from [41, Figure 1] instantiated with labelled Cramer-Shoup encryption [25] and smooth projective hashing from [16, Section 2.4] is AKE-ICM secure.*

Proof. This proof follows the proof of [41, Theorem 1] for the generic PAKE protocol. First note that we have to make sure that the adversary does not win trivially by returning the correct key to Test queries on corrupted sessions and only modify oracle replies to uncorrupted sessions. We change experiment one such that we do not encrypt 0 when simulating Execute oracles, but use random elements for ciphertexts, i.e. $C = (u_1, u_2, e, v)$ with $u_1 = g_1^\alpha, u_2 = g_2^\beta, e = h^\alpha g^\beta, v = (cd^\gamma)^\alpha$ for $\alpha, \beta, \gamma \in_R \mathbb{Z}_p$. It is not necessary to replace the projection keys k_p as they are already uniformly distributed in G . The claim follows, as in the original proof, directly from the semantic security of the CS encryption scheme. Experiment one corresponds to the AKE-ICM experiment with Execute₀, Send₁ and Test_b. The second experiment stays unchanged as here replies to the Test_b oracles on passive sessions are replaced with random keys. The third and fourth experiments can be adopted unmodified as well as they return random session keys in active sessions where the adversary does not now the correct password. Finally, we change the last experiment by replacing ciphertexts C in response to Send queries on uninitialised parties (this is called Send₀ in the original proof),

i.e. the `Send` oracle returns $C = (u_1, u_2, e, v)$ with $u_1 = g_1^\alpha, u_2 = g_2^\alpha, e = h^\alpha g^\beta, v = (cd^\gamma)^\alpha$ for $\alpha, \beta, \gamma \in_R \mathbb{Z}_p$ and a correct projection key (recall that projection keys are uniformly at random in G already). As in the original proof, this step is negligible, which follows from the CCA security of the CS encryption scheme and therefore the DDH assumption. In the last experiment all messages and keys are chosen at random from their respective spaces such that the adversary only wins by guessing the correct password. By observing that this experiment is equivalent to the AKE-ICM experiment with $b = 0$ the lemma follows. \square

4.3 Generalisations and Limitations

The design of the KV-PAKE protocol originates from the KOY protocol [38] and is used in many PAKE construction, e.g., [5, 16, 32, 33, 36]. It is therefore safe to assume that some of these protocols, instantiated with an encryption scheme that yields ciphertexts with pseudorandom elements and SPHF's with pseudorandom projection keys, can be used with our compiler. However, not all KOY-based constructions are usable in our compiler. The DDH-based instantiation of the KV-PAKE protocol in [41, Section 4.1] is not usable in our compiler as the ciphertext allows any party to verify its correctness, such that we can not replace it with random elements, i.e. this instantiation is not AKE-ICM secure. Similar observations result from analysing frameworks from [33, 38] that incorporate one-time signatures. The protocols from [33, 38] cannot be made oblivious because the use of one-time signatures introduces the possibility to publicly check the consistency of PAKE messages by performing signature verification. Note that existence of such public checks rules out any index-hiding encoding of those messages. That is, PAKE protocols from [33, 39, 41], if processed with our compiler, would become susceptible to offline dictionary attacks. In contrast, the MAC-based approach taken in RG-PAKE [32] and instantiation of KV-PAKE [41] with CS encryption do not admit public consistency checks of the resulting PAKE messages.

5 Conclusion

In this paper we addressed the problem of handling multiple password trials efficiently within the execution of PAKE protocols; in particular, aiming to optimise the amount of work on the server side. The proposed O-PAKE compiler results in almost constant computational complexity for the server without significantly increasing the computation costs on the client side, yet preserving all security guarantees offered by standard PAKE protocols. It can be used with PAKE protocols that fulfil our new definition of AKE-ICM security and whose client messages can be encoded through a suitable admissible encoding scheme. The security of the compiler has been proven under standard assumptions in an extension of the widely used PAKE model from [7] and exemplified on the PAKE protocols from [8] and [41].

References

1. M. Abdalla, J.-M. Bohli, M. I. G. Vasco, and R. Steinwandt. (Password) Authenticated Key Establishment: From 2-Party to Group. In *TCC'07*, volume 4392 of *LNCS*, pages 499–514. Springer, 2007.
2. M. Abdalla, E. Bresson, O. Chevassut, B. Möller, and D. Pointcheval. Provably secure password-based authentication in TLS. In *ASIACCS'06*, pages 35–45. ACM, 2006.
3. M. Abdalla, E. Bresson, O. Chevassut, B. Moller, and D. Pointcheval. Strong password-based authentication in TLS using the three-party group Diffie Hellman protocol. *Int. J. Secur. Netw.*, 2(3/4):284–296, Apr. 2007.

4. M. Abdalla, D. Catalano, C. Chevalier, and D. Pointcheval. Efficient two-party password-based key exchange protocols in the UC framework. In *CT-RSA'08*, volume 4964 of *LNCS*, pages 335–351. Springer, 2008.
5. M. Abdalla, C. Chevalier, and D. Pointcheval. Smooth Projective Hashing for Conditionally Extractable Commitments. In *CRYPTO*, volume 5677 of *LNCS*, pages 671–689. Springer, 2009.
6. M. Abdalla, O. Chevassut, P.-A. Fouque, and D. Pointcheval. A simple threshold authenticated key exchange from short secrets. In *ASIACRYPT'05*, volume 3788 of *LNCS*, pages 566–584. Springer, 2005.
7. M. Abdalla, P.-A. Fouque, and D. Pointcheval. Password-based authenticated key exchange in the three-party setting. In *PKC'05*, volume 3386 of *LNCS*, pages 65–84. Springer, 2005.
8. M. Abdalla and D. Pointcheval. Simple password-based encrypted key exchange protocols. In *CT-RSA'05*, volume 3376 of *LNCS*, pages 191–208. Springer, 2005.
9. M. Abdalla and D. Pointcheval. A Scalable Password-Based Group Key Exchange Protocol in the Standard Model. In *ASIACRYPT'06*, pages 332–347. Springer, 2006.
10. A. Bagherzandi, S. Jarecki, N. Saxena, and Y. Lu. Password-protected secret sharing. In *CCS'11*, pages 433–444. ACM, 2011.
11. M. Bellare, D. Pointcheval, and P. Rogaway. Authenticated key exchange secure against dictionary attacks. In *EUROCRYPT'00*, volume 1807 of *LNCS*, pages 139–155. Springer, 2000.
12. M. Bellare and P. Rogaway. Entity Authentication and Key Distribution. In *CRYPTO'93*, volume 773 of *LNCS*, pages 232–249. Springer, 1994.
13. S. M. Bellovin and M. Merritt. Encrypted Key Exchange: Password-Based Protocols Secure Against Dictionary Attacks. In *IEEE Symposium on Research in Security and Privacy*, pages 72–84, 1992.
14. S. M. Bellovin and M. Merritt. Augmented encrypted key exchange: A password-based protocol secure against dictionary attacks and password file compromise. In *CCS'93*, pages 244–250. ACM, 1993.
15. F. Benhamouda, O. Blazy, C. Chevalier, D. Pointcheval, and D. Vergnaud. Efficient UC-Secure Authenticated Key-Exchange for Algebraic Languages. In *PKC'13*, volume 7778 of *LNCS*, pages 272–291. Springer, 2013.
16. F. Benhamouda, O. Blazy, C. Chevalier, D. Pointcheval, and D. Vergnaud. New Techniques for SPHF's and Efficient One-Round PAKE Protocols. In *CRYPTO'13*, volume 8042 of *LNCS*, pages 449–475. Springer, 2013.
17. F. Benhamouda and D. Pointcheval. Verifier-Based Password-Authenticated Key Exchange: New Models and Constructions. Cryptology ePrint Archive, Report 2013/833, 2013. <http://eprint.iacr.org/>.
18. O. Blazy, D. Pointcheval, and D. Vergnaud. Round-Optimal privacy-preserving protocols with smooth projective hash functions. In *TCC'12*, volume 7194 of *LNCS*, pages 94–111. Springer, 2012.
19. D. Boneh and M. K. Franklin. Identity-Based Encryption from the Weil Pairing. In *CRYPTO'01*, volume 2139 of *LNCS*, pages 213–229. Springer, 2001.
20. V. Boyko, P. MacKenzie, and S. Patel. Provably Secure Password-Authenticated Key Exchange using Diffie-Hellman. In *EUROCRYPT'00*, volume 1807 of *LNCS*, pages 156–171. Springer, 2000.
21. E. Brier, J.-S. Coron, T. Icart, D. Madore, H. Randriam, and M. Tibouchi. Efficient indifferentiable hashing into ordinary elliptic curves. In *CRYPTO'10*, volume 6223 of *LNCS*, pages 237–254. Springer, 2010.
22. J. Camenisch, N. Casati, T. Gross, and V. Shoup. Credential authenticated identification and key exchange. In *CRYPTO'10*, volume 6223 of *LNCS*, pages 255–276. Springer, 2010.
23. R. Canetti. Universally Composable Security: A New Paradigm for Cryptographic Protocols. In *FOCS'01*, page 136. IEEE Computer Society, 2001.
24. R. Canetti, S. Halevi, J. Katz, Y. Lindell, and P. MacKenzie. Universally Composable Password-Based Key Exchange. In *EUROCRYPT'05*, volume 3494 of *LNCS*, pages 404–421. Springer, 2005.
25. R. Cramer and V. Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In *CRYPTO'98*, volume 1462 of *LNCS*, pages 13–25. Springer, 1998.
26. R. Cramer and V. Shoup. Universal Hash Proofs and a Paradigm for Adaptive Chosen Ciphertext Secure Public-Key Encryption. In *EUROCRYPT'02*, volume 2332 of *LNCS*, pages 45–64. Springer, 2002.
27. Ö. Dagdelen and M. Fischlin. Intercepting Tokens: The Empire Strikes Back in the Clone Wars. *IACR Cryptology ePrint Archive*, 2012, 2012. <http://eprint.iacr.org/>.
28. T. Dierks and E. Rescorla. RFC 5246 - The Transport Layer Security (TLS) Protocol Version 1.2, aug 2008. Updated by RFCs 5746, 5878, 6176.
29. N. Fleischhacker, F. Günther, F. Kiefer, M. Manulis, and B. Poettering. Pseudorandom Signatures. In *ASIA CCS'13*, pages 107–118. ACM, 2013.
30. D. Florencio and C. Herley. A Large-Scale Study of Web Password Habits. In *16th international conference on World Wide Web*, WWW'07, pages 657–666. ACM, 2007.
31. S. Gaw and E. W. Felten. Password Management Strategies for Online Accounts. In *Symposium on Usable privacy and security*, SOUPS'06, pages 44–55. ACM, 2006.

32. R. Gennaro. Faster and Shorter Password-Authenticated Key Exchange. In *TCC'08*, volume 4948 of *LNCS*, pages 589–606. Springer, 2008.
33. R. Gennaro and Y. Lindell. A Framework for Password-Based Authenticated Key Exchange. *ACM Trans. Inf. Syst. Secur.*, 9(2):181–234, may 2006.
34. C. Gentry, P. D. MacKenzie, and Z. Ramzan. A Method for Making Password-Based Key Exchange Resilient to Server Compromise. In *CRYPTO'06*, volume 4117 of *LNCS*, pages 142–159. Springer, 2006.
35. O. Goldreich and Y. Lindell. Session-Key Generation Using Human Passwords Only. In *CRYPTO'01*, volume 2139 of *LNCS*, pages 408–432. Springer, 2001.
36. A. Groce and J. Katz. A new framework for efficient password-based authenticated key exchange. In *CCS'10*, pages 516–525. ACM, 2010.
37. D. P. Jablon. Extended Password Key Exchange Protocols Immune to Dictionary Attacks. In *WETICE*, pages 248–255. IEEE Computer Society, 1997.
38. J. Katz, R. Ostrovsky, and M. Yung. Efficient Password-Authenticated Key Exchange Using Human-Memorable Passwords. In *EUROCRYPT'01*, volume 2045 of *LNCS*, pages 475–494. Springer, 2001.
39. J. Katz, R. Ostrovsky, and M. Yung. Efficient and Secure Authenticated Key Exchange Using Weak Passwords. *J. ACM*, 57(1):3:1–3:39, nov 2009.
40. J. Katz and V. Vaikuntanathan. Round-optimal password-based authenticated key exchange. In *TCC'11*, volume 6597 of *LNCS*, pages 293–310. Springer, 2011.
41. J. Katz and V. Vaikuntanathan. Round-Optimal Password-Based Authenticated Key Exchange. *J. Cryptology*, 26(4):714–743, 2013.
42. M. Manulis, B. Pinkas, and B. Poettering. Privacy-preserving group discovery with linear complexity. In *ACNS'10*, volume 6123 of *LNCS*, pages 420–437. Springer, 2010.
43. M. Manulis and B. Poettering. Practical affiliation-hiding authentication from improved polynomial interpolation. In *ASIACCS'11*, pages 286–295. ACM, 2011.
44. M.-H. Nguyen and S. P. Vadhan. Simpler Session-Key Generation from Short Random Passwords. In *TCC'04*, volume 2951 of *LNCS*, pages 428–445. Springer, 2004.
45. D. Pointcheval. Password-based authenticated key exchange. In *PKC'12*, volume 7293 of *LNCS*, pages 390–397. Springer, 2012.

A Admissible Encoding

Definition 5 (Admissible Encoding [21]). *Let S and R denote two finite sets such that $|S| > |R|$. A function $F : S \rightarrow R$ is called an ϵ -admissible encoding for (S, R) if it satisfies the following properties:*

1. **EFFICIENT:** F is computable in deterministic polynomial time.
2. **INVERTIBLE :** There exists a polynomial time algorithm \mathcal{I}_F such that $\mathcal{I}_F(r) \in F^{-1}(r) \cup \{\perp\}$ for all $r \in R$
3. **UNIFORM:** For all r uniformly distributed in R the distribution of $\mathcal{I}_F(r)$ is ϵ -statistically indistinguishable from the uniform distribution over S .

If ϵ is a negligible function of the security parameter then F is called an admissible encoding.

B Index-Hiding Message Encoding (IHME)

The concept of an index-based message encoding scheme with the index-hiding property (IHME) was introduced in [42, 43]. Here we recall their definitions.

Definition 6 (Index-Based Message Encoding [42]). *An index-based message encoding scheme (iEncode , iDecode) over an index space \mathcal{I} and a message space \mathcal{M} consists of two efficient algorithms:*

$\mathcal{S} \leftarrow \text{iEncode}(\mathcal{P})$: On input consisting of a tuple of index-message pairs $\mathcal{P} = \{(i_1, m_1), \dots, (i_n, m_n)\} \subseteq \mathcal{I} \times \mathcal{M}$ with distinct indices i_1, \dots, i_n , this algorithm outputs an encoding \mathcal{S} .
 $m \leftarrow \text{iDecode}(\mathcal{S}, i)$: On input of an encoding \mathcal{S} and an index $i \in \mathcal{I}$, this algorithm outputs a message $m \in \mathcal{M}$.

An index-based message encoding scheme is correct if $\text{iDecode}(\text{iEncode}(\mathcal{P}), i_j) = m_j$ for all $j \in \{1, \dots, n\}$ and all tuples $\mathcal{P} = \{(i_1, m_1), \dots, (i_n, m_n)\} \subseteq \mathcal{I} \times \mathcal{M}$ with distinct indices i_j .

The index-hiding property of an index-based message encoding ensures that all indices are hidden from the receiver of the message, as long as he does not know which message to expect. Thus, a receiver of an index-hiding encoded message can only recover those messages, encoded with an index he knows while all other indices stay hidden. The index-hiding property of an index-based message encoding is given by the following definition.

Definition 7 (Index-Hiding Message Encoding (IHME) [42]). Let $\text{IHME} = (\text{iEncode}, \text{iDecode})$ denote a correct index-based message encoding scheme over index space \mathcal{I} and message space $\mathcal{M}^{\text{IHME}, r}$. Let $b \in_R \{0, 1\}$ be a randomly chosen bit and let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a PPT adversary. We say that IHME provides index-hiding if there exists a negligible function $\varepsilon(\cdot)$ such that:

$$\text{Adv}_{\text{IHME}, \mathcal{A}}^{\text{ihide}}(\lambda) := |\Pr[\text{Exp}_{\text{IHME}, \mathcal{A}}^{\text{ihide}, 0}(\lambda) = 1] - \Pr[\text{Exp}_{\text{IHME}, \mathcal{A}}^{\text{ihide}, 1}(\lambda) = 1]| \leq \varepsilon(\lambda).$$

Moreover, if $\text{Adv}_{\text{IHME}, \mathcal{A}}^{\text{ihide}}(\lambda) = 0$ for all λ , the IHME-scheme is called perfect.

$\text{Exp}_{\text{IHME}, \mathcal{A}}^{\text{ihide}, b}(\lambda)$: Let $(I_0, I_1, M', St) \leftarrow \mathcal{A}_1(1^\lambda)$ with $I_0, I_1 \subseteq \mathcal{I}$, $|I_0| = |I_1| = n$, $M' = (m'_1, \dots, m'_{|I_0 \cap I_1|})$, $m'_j \in \mathcal{M}^{\text{IHME}^r}$ and $\{i_1, \dots, i_r\} = I_b \setminus I_{1-b}$. Choose m_1, \dots, m_r uniformly at random from $\mathcal{M}^{\text{IHME}^r}$, execute $\mathcal{S} \leftarrow \text{iEncode}(\{(i_j, m'_j) | i_j \in I_0 \cap I_1\} \cup \{(i_j, m_j) | i_j \in I_b \setminus I_{1-b}\})$ and $b' \leftarrow \mathcal{A}_2(St, \mathcal{S})$, and return $b' = b$.

ν -fold IHME There exists an optimised version of IHME for longer messages, denoted ν -fold IHME [43] with two advantages. First, the original instantiation of IHME requires that the index and message space correspond, i.e. $\mathcal{I} = \mathcal{M}^{\text{IHME}^r} = \mathbb{F}$. Secondly, ν -fold IHME is significantly faster than the original IHME implementation, as shown in [43]. Furthermore, it provides us with an easy way to encode multiple elements of the same protocol message (cf. Appendix 3.6).

Definition 8 (ν -fold IHME [43]). For an arbitrary finite field \mathbb{F} and $\nu \in \mathbb{N}$, after setting $\mathcal{I} = \mathbb{F}$ and $\mathcal{M}^{\text{IHME}^r} = \mathbb{F}^\nu$, an index-hiding message encoding scheme $\text{IHME} = (\text{iEncode}, \text{iDecode})$ with index space \mathcal{I} and message space $\mathcal{M}^{\text{IHME}^r}$ is constructed from standard $\text{IHME}' = (\text{iEncode}', \text{iDecode}')$ over \mathbb{F} as follows:

$\mathcal{S} \leftarrow \text{iEncode}(\mathcal{P})$: On input of $\mathcal{P} = \{(i_1, (m_{1,1}, \dots, m_{1,\nu})), \dots, (i_n, (m_{n,1}, \dots, m_{n,\nu}))\} \subseteq \mathcal{I} \times \mathcal{M}^{\text{IHME}^r} = \mathbb{F} \times \mathbb{F}^\nu$, the encoding is defined as the list $\mathcal{S} = (\mathcal{S}_1, \dots, \mathcal{S}_\nu)$ of IHME' -encodings $\mathcal{S}_k = \text{iEncode}'(\{i_j, m_{j,k}\}_{1 \leq j \leq n})$, for $1 \leq k \leq \nu$.
 $m \leftarrow \text{iDecode}(\mathcal{S}, i)$: On input of an encoding $\mathcal{S} = (\mathcal{S}_1, \dots, \mathcal{S}_\nu)$ and an index $i \in \mathcal{I}$, this algorithm outputs a message $(m_1, \dots, m_\nu) = m \in \mathcal{M}^{\text{IHME}^r}$ where $m_k = \text{iDecode}'(\mathcal{S}_k, i)$, for $1 \leq k \leq \nu$.

C Pseudorandom Functions

Definition 9 (Pseudorandom Function). A function $\text{PRF} : \{0, 1\}^k \times \{0, 1\}^m \rightarrow \{0, 1\}^m$ is a secure pseudorandom function if for all distinguishers D and truly random function $R : \{0, 1\}^m \rightarrow \{0, 1\}^m$ there exists a negligible function $\varepsilon(\cdot)$ such that

$$\text{Adv}_{\text{PRF}} = \left| \Pr[D^{\text{PRF}_k(\cdot)}() = 1] - \Pr[D^{R(\cdot)}() = 1] \right| \leq \varepsilon(\lambda).$$

D Proof of Theorem 1

Let Exp_0 denote the original AKE-security experiment for O-PAKE protocols from Definition 2 and \mathcal{A} be the corresponding AKE-security adversary. The proof uses a sequence of experiments Exp_i , $i = 1, \dots, 4$, where each Exp_i is based on a small modification of Exp_{i-1} . By $\text{Succ}_{i,\mathcal{A}} = \Pr[\text{Exp}_i = 1]$ we denote the success probability of \mathcal{A} in Exp_i . At a high level we first replace real client messages and session keys of underlying PAKE sessions $\Pi[i]$ with random messages and keys while ensuring the consistency against an adversary that corrupted passwords and then mounts online attacks. This modification remains unnoticeable to \mathcal{A} if the underlying PAKE protocol Π is AKE-secure with indistinguishable client messages. Then, we replace the outputs of the inverse transformations \mathcal{I}_{F^r} applied in each round r by choosing random elements from the corresponding round's IHME^r message space and show that this remains unnoticeable assuming that F^r is an ϵ_{F^r} -admissible encoding F^r . Then, we replace each real password index in IHME^r encoding process with a random password and show that this remains unnoticeable due to the index-hiding property of each IHME^r scheme. Finally, we modify the computation of the server's confirmation message and of the session keys that are returned in Test_b queries by using random elements from the corresponding spaces. This remains unnoticeable due to the pseudorandomness of the PRF function that is used to derive their values. Note that we only change oracle responses for *fresh* queries, i.e. neither $\text{Corrupt}(\text{pid}_P^i, P)$, nor $\text{Corrupt}(P, P')$ for any P' with $\text{pw}_{P,P'} \in \mathbf{pw}$ were queried before. In the following we describe each experiment in more details.

Exp₁: The experiment modifies the computation of $(m'_{\text{out}}, \Pi[i].\mathbf{k}) \leftarrow \Pi[i].\text{next}(m_{\text{in}})$, $i \in [1, \dots, n]$ in the client's $C_{\Pi}.\text{next}(m_{\text{in}})$ algorithm and the computation of $(m'_{\text{out}}, \Pi.\mathbf{k}) \leftarrow \Pi.\text{next}(m_{\text{in}})$ in the corresponding server's algorithm in response to *fresh* $\text{Execute}(P, P')$ and $\text{Send}(P, j, m_{\text{in}})$ queries of \mathcal{A} for any client $P \in \mathcal{C}$ and server $P' \in \mathcal{S}$ depending on round r . Otherwise, m'_{out} in client's algorithm is replaced by a randomly chosen message from $\mathcal{M}_{\mathcal{C}}^r$ and additionally, if $\Pi[i].\mathbf{k}' \notin \{\mathbf{null}, \perp\}$, then $\Pi[i].\mathbf{k}$ is replaced by a randomly chosen key from \mathcal{K}_{Π} (the session key space of Π). In the server's algorithm if $\Pi.\mathbf{k} \notin \{\mathbf{null}, \perp\}$, then $\Pi.\mathbf{k}$ get's replaced by a randomly chosen key from \mathcal{K}_{Π} . In order to ensure consistency the same random key is used for both $\Pi[i].\mathbf{k}$ and $\Pi.\mathbf{k}$ if the corresponding client and server have matching session ids. Note that server's output messages m'_{out} are not modified.

Claim (1). $|\text{Succ}_{0,\mathcal{A}} - \text{Succ}_{1,\mathcal{A}}| \leq c \cdot \text{Adv}_{\Pi,\mathcal{A}'}^{\text{AKE-ICM}}(\lambda) \leq \frac{c \cdot \mathcal{O}(t)}{|\mathcal{D}|} + \varepsilon'(\lambda)$

Proof. We start analysing modifications of client's $(m'_{\text{out}}, \Pi[i].\mathbf{k})$ and consider server's $\Pi.\mathbf{k}$ towards the end of the proof. For modifications on the client side, the only difference between the views of \mathcal{A} in Exp_0 and Exp_1 is that in Exp_1 client messages and session keys in the underlying PAKE sessions $\Pi[i]$ run by the client are replaced with random values chosen from the respective round

message spaces \mathcal{M}_C^r and key space \mathcal{K}_Π whereas in Exp_0 they are computed honestly according to the specification of Π . The condition on the **Corrupt** queries prevents \mathcal{A} from trivially distinguishing between the two views in an online attack with a corrupted password $\mathbf{pw}[i]$. Note that this condition ensures consistency with conditions on **Test_b** queries of \mathcal{A} . We show that any \mathcal{A} that can distinguish between the two views can be used to break the AKE-ICM notion of security for any of the underlying PAKE protocol sessions $\Pi[i]$. We construct an AKE-ICM adversary \mathcal{A}' against Π as follows. \mathcal{A}' chooses $\mathbf{pw}[i].\text{ix}$ for all $i \in [1, n]$ and stores them for later use, i.e. the values are the same for all executions of one client P . **Send**(P, j, m_{in}) queries from \mathcal{A} to P with $P.\text{role} = \text{client}$ and alleged sender $P' \in \mathcal{S}$ given in m_{in} are answered by \mathcal{A}' as follows: \mathcal{A}' queries its **Send_b**(Q, j, m_{in}) oracle with $n - 1$ random clients Q to get PAKE messages and on P , and constructs the outgoing message m_{out} of client P according to the compiler specification, i.e. applies admissible encodings and IHME using $\mathbf{pw}[i].\text{ix}$, and computes the key if necessary (note that this may include queries of \mathcal{A}' to appropriate **Test_b** oracles). **Execute**(P, P') queries from \mathcal{A} are handled similar. Note also that \mathcal{A}' is given access to **Corrupt** and **Test_b** oracles in the AKE-ICM experiment that are used to answer the corresponding queries of \mathcal{A} . Let b denote the bit used in the AKE-ICM experiment. If $b = 1$, then $(m'_{\text{out}}, \Pi[i].\mathbf{k})$ received by \mathcal{A}' for **Send**(Q, j, m_{in}) and **Send**(P, i, m_{in}), as well as **Execute** queries, are real values and their distribution is identical to their distribution in Exp_0 ; whereas if $b = 0$, they are random and have the distribution as in Exp_1 . Therefore, the output bit b' of \mathcal{A} at the end of the experiment can be used by \mathcal{A}' as a guess for b . If \mathcal{A} is successful in breaking the AKE security of C_Π then \mathcal{A}' is successful in breaking the AKE-ICM security of Π .

As for the modification of server's session keys $\Pi.\mathbf{k}$ we can apply similar argument with respect to the AKE security property of the underlying PAKE protocol Π . The corresponding AKE adversary \mathcal{A}'' would simulate the experiment for \mathcal{A} using access to its oracles **Execute** and **Send** and uses the output bit b' of \mathcal{A} as its own guess for b . Observe that any AKE-ICM secure Π is also AKE secure, i.e., $\text{Adv}_{\Pi, \mathcal{A}''}^{\text{AKE}}(\lambda) \leq \text{Adv}_{\Pi, \mathcal{A}'}^{\text{AKE-ICM}}(\lambda)$ for an AKE adversary \mathcal{A}'' . Therefore, it is sufficient to use $\text{Adv}_{\Pi, \mathcal{A}'}^{\text{AKE-ICM}}(\lambda)$ as a bound in the estimation of the difference on success probabilities between Exp_0 and Exp_1 . Since in Exp_1 client messages and keys are modified for all PAKE instances $\Pi[i]$, $i \in [1, n]$ run by the client and since $n \leq c$ (maximal size of \mathbf{pw}) we also need to factor in c in the computation of $|\text{Succ}_{0, \mathcal{A}} - \text{Succ}_{1, \mathcal{A}}| \leq c \cdot \text{Adv}_{\Pi, \mathcal{A}'}^{\text{AKE-ICM}}(\lambda)$. \square

As a consequence intermediate messages m'_{out} and keys $\Pi[i].\mathbf{k}$ used on the client side and keys $\Pi.\mathbf{k}$ on the server side in the $C_\Pi.\text{next}(m_{\text{in}})$ algorithm are uniformly distributed in respective message and key spaces and thus also independent from the PAKE sessions $\Pi[i]$ run by the client and Π by the server. In the following we can thus focus on the encoding of m'_{out} and computation of C_Π session keys \mathbf{k} .

Exp₂: The experiment modifies the encoding value $\mathcal{I}_{F^r}(m'_{\text{out}})$ for each PAKE session $\Pi[i]$ in the client's $C_\Pi.\text{next}(m_{\text{in}})$ algorithm in response to *fresh* **Execute**(P, P') and **Send**(P, j, m_{in}) queries of \mathcal{A} for any client $P \in \mathcal{C}$, depending on round r , i.e. $\mathcal{I}_{F^r}(m'_{\text{out}})$ is replaced by a random element M from $\mathcal{M}^{\text{IHME}^r}$.

Claim (2). $|\text{Succ}_{1, \mathcal{A}} - \text{Succ}_{2, \mathcal{A}}| \leq c \sum_{r=1}^R \epsilon_{F^r}$

Proof. This claim follows directly from the properties of the assumed admissible encoding $F^r : \mathcal{M}^{\text{IHME}^r} \rightarrow \mathcal{M}_C^r$ used in the r th round. More precisely, from the assumption that the distribution of the inverse transformation $\mathcal{I}_{F^r}(m'_{\text{out}})$ is ϵ_{F^r} -statistically indistinguishable from the uniform distribution over $\mathcal{M}^{\text{IHME}, r}$ if the input m'_{out} is uniformly distributed in \mathcal{M}_C^r . The uniform distribution

of all r -round messages m'_{out} over \mathcal{M}_C^r for any PAKE instance $\Pi[i]$ is already guaranteed by Exp_1 . Hence, the distribution of $\mathcal{I}_{F^r}(m'_{\text{out}})$ used in the r th round of $\Pi[i]$ in Exp_1 can be distinguished from the corresponding distribution in Exp_1 with probability no greater than ϵ_{F^r} . Considering that Exp_2 modifies $\mathcal{I}_{F^r}(m'_{\text{out}})$ for all uncorrupted instances $\Pi[i]$, $i \in [1, n]$ and rounds $r \in [1, R]$ the overall probability for distinguishing the two distributions is no greater than $c \sum_{r=1}^R \epsilon_{F^r}$. \square

As a consequence intermediate values $\mathcal{I}_{F^r}(m'_{\text{out}})$ used in the $C_{\Pi}.\text{next}(m_{\text{in}})$ algorithm on the client side are uniformly distributed in message spaces $\mathcal{M}^{\text{IHME}^r}$.

Exp₃: This experiment is equivalent to Exp_2 but instead of honestly simulating the IHME computations of the compiler, the simulator plays an adversary \mathcal{B} against the index-hiding experiment from Definition 7. The experiment modifies all pairs $(\mathbf{pw}[i], M)$ used in the client's $C_{\Pi}.\text{next}(m_{\text{in}})$ algorithm in response to *fresh* $\text{Execute}(P, P')$ and $\text{Send}(P, j, m_{\text{in}})$ queries of \mathcal{A} , depending on round r , i.e. random password indices $\text{ix}_i^* \in_R \mathbb{F}$ are chosen and consistently used as an index for IHME^r encoding instead of $\text{pw}_{P, P'}.\text{ix}$.

Claim (3). $|\text{Succ}_{2, \mathcal{A}} - \text{Succ}_{3, \mathcal{A}}| \leq c \sum_{r=1}^R \text{Adv}_{\text{IHME}^r, \mathcal{B}}^{\text{hide}}$

Proof. Let Exp'_j for $j \in [1, n]$ denote the experiment where the first j IHME indices are replaced by random indices ix_i , i.e. all ix_i for $i \leq j$ are chosen at random, different from $\mathbf{pw}[i].\text{ix}$, and all ix_i for $i > j$ are simulated honestly. In the following we bound the adversary's advantage of distinguishing between Exp'_j and Exp'_{j-1} by $\sum_{r=1}^R \text{Adv}_{\text{IHME}^r, \mathcal{B}}^{\text{hide}}$, such that the claim follows from the observation that $\text{Exp}'_0 = \text{Exp}_2$ and $\text{Exp}'_n = \text{Exp}_3$ by a hybrid argument.

We therefore replace $\mathbf{pw}[j].\text{ix}$ with a random ix_j^* such that the corresponding IHME index-message pair (ix_j^*, M) is part of the set E that is used to compute the outgoing client's message $m_{\text{out}} \leftarrow \text{IHME}^r.\text{iEncode}(E)$ in round r . The only difference between Exp'_{j-1} and Exp'_j is that Exp'_{j-1} uses $\mathbf{pw}[j].\text{ix}$ as an IHME index whereas Exp'_j uses ix_j^* . We show that any \mathcal{A} that can distinguish between the views from Exp'_{j-1} and Exp'_j can be used to break the index-hiding property of any IHME^r used in C_{Π} . We construct an index-hiding adversary \mathcal{B} against IHME^{r^*} for some $r^* \in [1, R]$ as follows. \mathcal{B} in its first stage picks two randomly chosen password indices $\text{pw}_{P, P'}.\text{ix}, \text{pw}_{P, P'}^*.\text{ix}$ for each client-server pair (P, P') and treats $\text{pw}_{P, P'}$ as a real password in the simulation of C_{Π} . In particular, $\text{pw}_{P, P'}$ will be returned to \mathcal{A} if the latter decides to corrupt the password shared by P and P' . \mathcal{B} simulates the execution of $C_{\Pi}.\text{next}(m_{\text{in}})$. In response to $\text{Execute}(P, P')$ and $\text{Send}(P, j, m_{\text{in}})$ in any round $r \in [1, R]$, $r \neq r^*$ it computes $m_{\text{out}} \leftarrow \text{IHME}^r.\text{iEncode}(E)$ using E that contains the real password index $\text{pw}_{P, P'}.\text{ix}$ and honestly generated and encoded message. In round r^* it defines $I_0 = \{\text{ix}_1, \dots, \text{ix}_{j-1}, \mathbf{pw}[j].\text{ix}, \dots, \mathbf{pw}[n].\text{ix}\}$ and $I_1 = \{\text{ix}_1, \dots, \text{ix}_j, \mathbf{pw}[j+1].\text{ix}, \dots, \mathbf{pw}[n].\text{ix}\}$ and outputs (I_0, I_1) as its challenge in the index-hiding experiment against IHME^{r^*} . The index-hiding experiment picks random elements from $\mathcal{M}^{\text{IHME}^{r^*}}$ as messages and returns an IHME encoding that \mathcal{B} uses as m_{out} in round r^* . Let b denote the bit used in index-hiding experiment. If $b = 0$ then m_{out} was computed with the real password index $\mathbf{pw}[j].\text{ix}$ and so the distribution of the set E in round r^* is identical to the one from Exp_{j-1} . If $b = 1$, then m_{out} was computed using a random password index ix_j and so the distribution of E in round r^* is the same as in Exp_j . Therefore, the output bit b' of \mathcal{A} at the end of the experiment can be used by \mathcal{B} as a guess for b . If \mathcal{A} is successful in breaking the AKE security of C_{Π} then \mathcal{B} is successful in breaking the index-hiding property of IHME^{r^*} . Since in Exp_j pairs $(\mathbf{pw}[j], M)$ are modified for all rounds $r \in [1, R]$ the overall probability for distinguishing the two experiments is no greater than $\sum_{r=1}^R \text{Adv}_{\text{IHME}^r, \mathcal{B}}^{\text{IHME}}$. \square

In the following experiment we focus on the computation of the confirmation message c and the final session key \mathbf{k} that are computed using a pseudorandom function PRF. Let Adv_{PRF} denote the probability with which the outputs of this PRF can be distinguished from random.

Exp₄: This experiment modifies the computation of the confirmation message $m_{\text{out}} = \text{PRF}_{\Pi.\mathbf{k}}(\text{sid}_{P'}^j || \text{pid}_{P'}^j || P'_j || 0)$ and of the session key $\mathbf{k} = \text{PRF}_{\Pi.\mathbf{k}}(\text{sid}_{P'}^j || \text{pid}_{P'}^j || P'_j || 1)$ on the server side and of the session key $\mathbf{k} = \text{PRF}_{\Pi.\mathbf{k}}(\text{sid}_P^i || P_i || \text{pid}_P^i || 1)$ on the client side of the algorithm $C_{\Pi}.\text{next}(m_{\text{in}})$ in response to corresponding *fresh* $\text{Execute}(P, P')$ and $\text{Send}(P', j, m_{\text{in}})$ resp. Test_b queries of \mathcal{A} for any server $P' \in \mathcal{S}$ and client $P \in \mathcal{C}$. In particular, m_{out} and \mathbf{k} are set to be random elements from the output space of PRF if $\Pi.\mathbf{k} \in \mathcal{K}_{\Pi}$. In any other case m_{out} and \mathbf{k} are not changed, i.e. in the case of $m_{\text{out}} = \emptyset$ and $\mathbf{k} = \perp$. In order to ensure consistency the same random key \mathbf{k} is used for partnered client-server pairs that have matching session ids.

Claim (4). $|\text{Succ}_{3,\mathcal{A}} - \text{Succ}_{4,\mathcal{A}}| \leq 2\text{Adv}_{\text{PRF}}$

Proof. First recall that the intermediate keys $\Pi[i].\mathbf{k}$ on client side and $\Pi.\mathbf{k}$ on server side are random for uncorrupted sessions according to Exp_1 . We show how to build distinguishers D_1 and D_2 that are able to distinguish between real confirmation messages and random ones, and real and random keys respectively. Let D_1 denote a distinguisher against $\text{PRF}_{\Pi.\mathbf{k}}$ of the confirmation message m_{out} . Instead of honestly generating the confirmation message, D_1 uses its oracle to get random elements for m_{out} on input of $(\text{sid}_{P'}^j || \text{pid}_{P'}^j || P'_j || 0)$, and outputs the guess b of \mathcal{A} . Distinguisher D_2 is build accordingly for \mathbf{k} . The probability to distinguish m_{out} or \mathbf{k} from truly random elements in Exp_4 is therefore bounded by Adv_{PRF} each. \square

In Exp_4 , the key \mathbf{k} is chosen uniformly at random such that the success probability of the adversary to determine the bit b is the same as guessing b . Thus, the advantage of \mathcal{A} is given by $\text{Adv}_{C_{\Pi},\mathcal{A}}^{\text{AKE}}(\lambda) \leq c \cdot (\text{Adv}_{\Pi}^{\text{AKE-ICM}}(\lambda) + \sum_{r=1}^R \epsilon_F + \sum_{r=1}^R \text{Adv}_{\text{IHME}}^{\text{ihide}}) + 2\text{Adv}_{\text{PRF}}$, and since $\text{Adv}_{\text{IHME}}^{\text{ihide}}$, Adv_{PRF} and ϵ_F are negligible this results in

$$\text{Adv}_{C_{\Pi},\mathcal{A}}^{\text{AKE}}(\lambda) \leq \frac{c \cdot \mathcal{O}(t)}{|\mathcal{D}|} + \varepsilon(\lambda).$$

\square