

PAKE Combiners and Efficient Post-Quantum Instantiations

Julia Hesse¹^{*}
IBM Research Europe – Zurich
jhs@zurich.ibm.com

Michael Rosenberg²
research@mrosenberg.pub

Version 1.6, May 2, 2025[†]

Abstract

Much work has been done recently on developing password-authenticated key exchange (PAKE) mechanisms with post-quantum security. However, modern guidance recommends the use of *hybrid* schemes—schemes which rely on the combined hardness of a post-quantum assumption, e.g., Learning with Errors (LWE), and a more traditional assumption, e.g., decisional Diffie-Hellman. To date, there is no known hybrid PAKE construction, let alone a general method for achieving such.

In this paper, we present two efficient PAKE combiners—algorithms that take two PAKEs satisfying mild assumptions, and output a third PAKE with combined security properties—and prove these combiners secure in the Universal Composability (UC) model. Our sequential combiner, instantiated with efficient existing PAKEs such as CPace (built on Diffie-Hellman-type assumptions) and CAKE (built on lattice assumptions), yields the first known hybrid PAKE.

Keywords: key agreement, password-based cryptography, post-quantum cryptography

1 Introduction

Memorable passwords are the one of most commonly used forms of authentication today, and exist across a wide variety of applications. While some applications, such as website login, have users send their plaintext password, many large-scale applications utilize password-authenticated key exchange (PAKE) and its variants to limit the amount of information that can be intercepted. These include iCloud Keychain escrow, 1Password user authentication, Facebook Messenger

^{*}The author was supported by the Swiss National Science Foundation (SNSF) under the AMBIZIONE grant “Cryptographic Protocols for Human Authentication and the IoT”.

[†]Version 1.0 of this paper was researched concurrently with and independently of [LL24], whose combiners are identical to ours.

chat history sharing, WhatsApp backup encryption, and passport chip access control [fan24].

As a result of the widespread push from industry and government to move towards post-quantum cryptography, many efficient lattice- and isogeny-based PAKEs have been developed in recent years [DAL⁺17, BCP⁺23, ABJS24, AEK⁺22, IY23, LLH24]. However, these constructions are not necessarily the end goal of the post-quantum efforts. The largest rollout of a post-quantum protocol to date is that of TLS 1.3 with *hybrid* key exchange [ABBO24], i.e., key exchange which enjoys the protection of both newer post-quantum constructions and classical constructions. The purpose of the hybrid in this application is two-fold. Firstly, the newer, less tested post-quantum hardness assumptions may turn out to be easier than expected. Secondly, the newer implementations of the post-quantum algorithms may turn out to have exploitable bugs. Perhaps unintuitively, the second is considered by some as the more likely scenario [Wes24], and in fact already has occurred [GJN20, GHJ⁺22, HSC⁺23, BBB⁺24]. For these reasons, some governments recommend the use of hybrid constructions for secure protocols [Fre22, Ger24]. Unfortunately, no hybrid PAKE currently exists.

1.1 Non-solutions

The most common way to develop a primitive from a hybrid assumption is to simply *combine* two copies of the primitive, with each copy relying on a different assumption. However, the obvious combiners from other cryptographic primitives do not work for PAKEs.

Concatenation combiner. Hybrid KEMs and hybrid digital signature schemes can be constructed from the *concatenation combiner*: separately execute two copies of the primitive, and output the hash or concatenation of the outputs, respectively [BCD⁺24, BH23]. Let us attempt this for PAKE: let P be the PAKE given by running some PAKE P_1 and another PAKE P_2 in parallel using the same password pw and then outputting the hash $H(K_1, K_2)$, where K_i is the session key computed in P_i . In general, this straightforward combiner is only guaranteed to be as strong as the *weaker* of P_1, P_2 . Intuitively, the reason is that we made the situation worse by trusting *two* PAKEs with protecting the password. For a concrete example, consider P_1 to be KC-SPAKE2 [Sho20] and P_2 to be MLWE-PAKE [RG22]. In the design of KC-SPAKE2, the responder sends a MAC tag derived from their password-dependent Diffie-Hellman shared secret. An attacker with a discrete-logarithm oracle (e.g., from a quantum computer) can break P by merely observing one transcript of the KC-SPAKE2 subprotocol, and mount an offline attack where guesses are confirmed against the MAC tag. MLWE-PAKE, on the other hand, sends samples from the Module Learning with Errors (MLWE) distribution, blinded with the hashed password. An attacker with a decision-MLWE oracle who wishes to break P can observe one transcript of the MLWE-PAKE subprotocol to mount an offline attack, because the MLWE oracle will reject with overwhelming probability when the password guess is

wrong. Hence, *both* assumptions have to hold in order for the concatenation combiner to protect the password.

KEM-based PAKEs. Another potential method to develop a hybrid PAKE is to use an existing KEM-based PAKE construction, e.g., CAKE [BCP⁺23], and instantiate it with a hybrid KEM. Similar to the above, this does not necessarily have the combined security of the two underlying KEMs. This is because CAKE requires its KEM to have *anonymity* and *fuzziness* (essentially, public keys and ciphertexts must look uniform), and it is not clear how to construct a KEM with hybrid anonymity and fuzziness.

Consider the case where the hybrid KEM is a concatenated KEM, as many of the existing hybrids are. In the first communication round, CAKE encrypts the KEM public key, i.e., a concatenation of KEM_1 and KEM_2 public keys, under pw . Thus, if *either* of the underlying public keys has structure that makes it distinguishable from uniform, then trial decryption on the CAKE encrypted public key yields an offline attack. Concretely, take KEM_1 to be a DH KEM over the Ristretto prime order group [LHT16] with points represented using the Elligator2 encoding [BHKL13], and KEM_2 to be FrodoKEM [BCD⁺16]. While the DH public keys are statistically indistinguishable from a uniform 32-byte string, FrodoKEM public keys are only *computationally* indistinguishable from uniform under the LWE assumption.¹ If a passive attacker possesses an LWE oracle, they may trial-decrypt the encrypted concatenated public key, ignore the DH public key, and check if the decrypted value is an LWE sample. Thus, CAKE with this hybrid KEM is at most as secure as LWE. A similar argument shows that the same is true for other recent KEM-based PAKEs such as OCAKE, CHIC, and NoIC [BCP⁺23, ABJS24, ABJ25].

1.2 Contributions

In this paper we give a generalizable method for constructing hybrid PAKEs and demonstrate efficient instantiations. We present two PAKE combiners: **ParComb**, a parallel combiner, executing both PAKEs in parallel and hashing the results; and **SeqComb**, a sequential combiner, executing the first PAKE and feeding the resulting session key into the second PAKE. We prove the combiners' security in the universal composability (UC) model with random oracles and static corruptions.

ParComb. For our **ParComb** combiner, which follows the blueprint of the usual concatenation combiner, we need to circumvent the above issues. To this end, we formalize a property of a PAKE that is sufficient to prove **ParComb** secure. Intuitively, we demand that the password protection of both PAKEs used in **ParComb** not fatally break if any of the underlying computational assumptions break, i.e., their transcripts information-theoretically hide the password. This ensures that none of the PAKEs leak information about the input password. We identify a handful of PAKE protocols from the literature that satisfy this

¹This is not particular to FrodoKEM. It appears that all standards-track lattice- and code-based KEMs achieve, at best, computational public key uniformity.

unconditional hiding property. Unfortunately, none of these are post-quantum PAKEs. So **ParComb**, while having zero round complexity overhead over its building blocks, does not yet yield hybrid PAKE. We conclude by pointing to concurrent work which shows that **ParComb** realizes an incomparable notion of PAKE when instantiated with a specific isogeny-based PAKE.

SeqComb. These limitations motivate looking into sequential combiner designs. Our combiner **SeqComb** derives a shared key from a password through the first PAKE, and then uses this shared key as a password in the second PAKE to derive another shared key. The final key is, again, the hashed concatenation of both PAKE keys. As in **ParComb**, the first PAKE consumes the plain password and hence must enjoy the same unconditional hiding property as both PAKEs in **ParComb**. For the second PAKE, however, we can get away with a weaker hiding property. This is because we used the first PAKE as an entropy amplifier to derive unguessable preshared keys from the passwords, and hence we do not rely on offline guess protection from the second PAKE to protect these keys. We formalize a more mild unconditional hiding property, namely that the second PAKE’s transcript does not reveal whether the first PAKE run was successful, that is sufficient to prove the security of **SeqComb**. We identify KEM-based PAKEs with this property, namely CAKE and EKE-PRF [BCP⁺23, JRX24]. Thus, we get a hybrid PAKE with a round complexity that is the sum of the underlying PAKEs’, and computational overhead of only two hashes. To the authors’ knowledge, this is the first (concurrent with [LL24]) description of a hybrid PAKE.

Beyond PAKE. Finally, we remark that hybrid PAKE implies hybrid instances of other flavors of PAKE. [GMR06] and [LLH25] describe combiners which construct an augmented PAKE (aPAKE) from PAKE, authenticated encryption, and either a signature scheme or KEM, respectively. Similarly, the LATKE identity-binding PAKE (iPAKE) framework [KR24] takes a generic PAKE and identity-based key exchange (IBKE) protocol. Since hybrid signatures [BH23], KEMs [BCD⁺24], and IBKEs [KR24] are known, hybrid PAKE implies hybrid aPAKE and iPAKE.

2 Preliminaries

2.1 Notation

We write probabilistic algorithms as $\text{Alg}(x; r)$ where x denotes the input and r denotes the random coins. We write $y \leftarrow \text{Alg}(x)$ to denote sampling uniform r and setting $y := \text{Alg}(x; r)$, and $x \leftarrow_{\$} S$ to denote sampling a uniform value from a set S . For our security proofs we will consider probabilistic polynomial time (PPT) adversaries, and denote them with calligraphic letters \mathcal{A} . We use λ to denote the security parameter.

In our pseudocode for ideal functionalities, **assert** checks the given condition, and early-returns with “assertion failed” on failure. **retrieve** denotes retrieval of a record with a specific marking; if no such record is present, this early-

returns with “no record.” We mark new variables introduced into the scope by a **retrieve** using square brackets. We write $\boxed{\mathsf{P}^{\text{sid}}}$ to denote execution of a PAKE protocol P with session identifier sid .

2.2 Universal composability

The Universal Composability (UC) model [Can01] provides an alternative to game-based security definitions. The model frames cryptographic protocols as idealized *functionalities*, which can be thought of as black boxes with a tightly constrained interface to the outside world. In the security game showing Π *UC-realizes* the functionality \mathcal{F} , the goal of an interactive Turing machine called the *environment* \mathcal{Z} is to distinguish between the *ideal world* and the *real world*, which are defined as follows.

Real / ideal. In the *real world* all the parties participate in Π , \mathcal{Z} may view parties’ outputs, and is permitted to give arbitrary instructions to a separate interactive Turing machine, called the *adversary* \mathcal{A} . The environment can arbitrarily ask the adversary to view/modify/delay/drop messages between parties, corrupt parties, and interact with any ideal functionalities \mathcal{F}'_i used to instantiate Π . In the *ideal world* all the parties are *dummies*, speaking directly to \mathcal{F} . In addition, the adversary may also only interact with \mathcal{F} , though it is still permitted to corrupt parties.

Π *UC-realizes* \mathcal{F} if for any \mathcal{A} , there exists a simulator \mathcal{S} of \mathcal{A} such that any \mathcal{Z} has at most negligible advantage in distinguishing between \mathcal{A} and \mathcal{S} . That is,

$$\text{IDEAL}_{\mathcal{S}, \mathcal{Z}}^{\mathcal{F}} \stackrel{c}{\approx} \text{EXEC}_{\mathcal{A}, \mathcal{Z}}$$

where the LHS refers to the probability ensemble consisting of the view of the environment in the ideal world, and the RHS in the real world.

By [Can01, Theorem 11], it suffices to imagine that the adversary \mathcal{A} is the *dummy adversary* \mathcal{A}_D , which simply delivers backdoor messages generated by the environment to the specified recipients, and delivers to the environment all backdoor messages generated by the protocol parties, as well as the sender machine’s identity.

Corruption. There are two models in which the environment can corrupt parties. In the *static corruption model*, the adversary may not corrupt any parties during protocol execution. In the *adaptive corruption model*, the adversary may corrupt parties at any time. In this work, we will consider adversaries who are allowed static corruptions. We note that, since PAKE sessions are so short lived, and static corruptions permit corruptions in between executions, this security model is still quite strong.

Sufficiency of single-session functionality. The ideal functionalities we consider in this work are *single-session*, i.e., they permit at most one key exchange for a single pair of parties, after which they cease to function. It is known that a single-session protocol can be lifted to a multi-session (i.e., multiple pairs of parties, multiple key exchanges per pair) protocol with joint state as long as

the multi-session extensions of the helper (formally, *hybrid*) functionalities exist [CR03, Hes20]. Since our reductions only use the random oracle functionality, which has a known multi-session extension, we get these theorems for free.

Session identifiers. To represent different instances of the same scheme, the UC model uses *session identifiers*. Each party is given a session identifier sid on activation and will only interact with other parties with the same sid . For clarity of presentation, we will assume in our protocol definitions that session identifier establishment has already occurred. Further, we will assume that every pair of parties, and hence every execution in the single-session model, has a unique sid (this can be done, e.g., by exchanging random nonces or appending unique personal identifiers to the existing sid).

We note that single-session functionalities do not usually make use of the session identifier sid , since it does not serve to disambiguate sessions. Regardless, we include sid as a random oracle input and ideal cipher domain separator in our pseudocode. We do this so that an implementor who does not read this section will be more likely to implement the protocols securely, i.e., in a way that yields more than single-session security.

2.3 PAKE

A (*balanced*) *password authenticated key-exchange protocol* (PAKE) is a two-party key-exchange protocol where parties use mutual knowledge of a low-entropy password pw to establish a high-entropy *session key* K .

2.3.1 PAKE protocols

We define a rudimentary notion of a PAKE protocol by specifying its input-output behavior and correctness. Looking ahead, the following definition captures the minimal guarantees that we can expect from a single PAKE building block of our combiners, even when its security breaks.

Definition 1 (PAKE protocol). *Let P be an interactive protocol between two parties, \mathcal{P} and \mathcal{P}' , producing keys of bit length λ . We say P is a PAKE protocol if it satisfies the following properties.*

1. P takes input pw from \mathcal{P} and pw' from \mathcal{P}'
2. P produces output K for \mathcal{P} and K' for \mathcal{P}'
3. If $\text{pw} = \text{pw}'$ then $K = K'$ except with probability negligible in λ

For any interactive protocol, a *protocol round* is the set of all messages that can be sent in parallel from any point in the protocol [Gon93]. Thus, in a one-round protocol, no message depends on another. We require a minimal syntax $\mathsf{P} = (\text{Start}, \text{Finish})$ for one-round PAKEs, where minimal means that the PAKE may consist of more algorithms (e.g., for computing a setup), and the algorithms may take more inputs. The minimal requirements are as follows:

- **Start** takes as input a security parameter 1^λ , the session identifier sid , and a password pw and outputs a message x and a state st ; and
- **Finish** takes as input st and a message y and outputs a key.

2.3.2 Ideal PAKE functionality

The security goals for a PAKE are (1) to establish a high-entropy shared session key when both parties are honest, (2) to prevent passive adversaries from learning anything about the password, and (3) to limit active adversaries to one (or another small constant) password guess(es) per protocol instance, even if given access to session keys. These guarantees can be captured by the ideal functionality $\mathcal{F}_{\text{PAKE}}$ [CHK⁺05]. Some widely used protocols use a slight weakening of this functionality, called *lazy extraction PAKE* ($\mathcal{F}_{\text{lePAKE}}$) [ABB⁺20], which permits adversaries complete active attacks even after a session is already finished. In this work we also make use of a novel functionality, a *relaxed* version of lazy extraction PAKE ($\mathcal{F}_{\text{relPAKE}}$), where the adversary additionally learns whether the active attack was successful. We state all functionalities in Figure 1, and include the small modifications from [AHH21] regarding adversary-controlled keys.

Isolated parties. We introduce another small modification to our PAKE ideal functionality. In the original definition, it is possible for a party to output a session key before the counterparty has even started its session. In our definition, we prevent these *isolated parties* from getting their session key. This modification is crucial to using $\mathcal{F}_{\text{PAKE}}$ in a combiner, since we rely on timely extraction from a particular sub-PAKE, and cannot deal with output keys that were produced without any interaction.

This new functionality is slightly stronger than the original, but this is not a problem: any PAKE which UC-realizes the original $\mathcal{F}_{\text{PAKE}}$ UC-realizes the new one as well. Consider a UC simulator for a such a PAKE. We will show that **NewKey** is never called on an isolated party. First, note if the simulator calls **NewKey** on an isolated party, it is because the environment \mathcal{Z} has triggered the end to the protocol (this is because **NewKey** gives \mathcal{Z} the session key, so \mathcal{Z} knows precisely when **NewKey** is called). Further, if **NewSession** has not been called for the counterparty, it is because \mathcal{Z} has not activated that party. So the only case in which **NewKey** can be called on an isolated party is if only that party was activated and the protocol has finished, i.e., the protocol is at most a one-message PAKE. Since such PAKEs are impossible in the UC model,² this scenario never occurs.

²More specifically, PAKEs with negligible correctness error and at most one message are impossible. An efficient offline attack is as follows. Let \mathcal{P} be the initiator in the execution of a one-message PAKE, sending message m to \mathcal{P}' . If \mathcal{Z} permits \mathcal{P} to send m , then \mathcal{P}' immediately outputs its session key K . This K is the output of a function $f(\text{pw}, \text{sid}, m)$ (w.l.o.g., this is deterministic, otherwise we have noticeable correctness error). Since \mathcal{Z} knows sid and m , it can test various pw until it receives an output that matches K .

<p>Session management</p> <p>On (NewSession, sid, \mathcal{P}_j, pw_i) from \mathcal{P}_i</p> <p>send (NewSession, sid, \mathcal{P}_i, \mathcal{P}_j) to \mathcal{A}</p> <p>if $\nexists(\text{Session, sid, } \mathcal{P}_i, \cdot) \wedge \neg \text{NoMoreSessions}$:</p> <p style="padding-left: 2em;">// After two NewSessions stop accepting more</p> <p>if $\exists(\text{Session, } \cdot \cdot \cdot)$: record NoMoreSessions</p> <p>record (Session, sid, \mathcal{P}_i, \mathcal{P}_j, pw_i)</p> <p>and mark it fresh</p> <p>Key generation and authentication</p> <p>On (NewKey, sid, \mathcal{P}_i, K') from \mathcal{A}</p> <p>retrieve (Session, sid, \mathcal{P}_i, [\mathcal{P}_j], [pw_i])</p> <p>with mark $m \neq \text{completed}$</p> <p>if $m = \text{fresh}$: assert $\exists(\text{Session, sid, } \mathcal{P}_j, \mathcal{P}_i, \cdot)$</p> <p>if $m = \text{compromised}$: $K_i := K'$</p> <p>elif $m = \text{fresh}$ and</p> <p style="padding-left: 2em;">$\exists(\text{Key, sid, } [\mathcal{P}_j], [\text{pw}_j], [K_j], \text{fresh})$ s.t. pw_i = pw_j :</p> <p style="padding-left: 4em;">$K_i := K_j$</p> <p>else : $K_i \leftarrow_{\\$} \{0, 1\}^\lambda$</p> <p>record (Key, sid, \mathcal{P}_i, pw_i, K_i, m)</p> <p>Mark the session completed</p> <p>send (sid, K_i) to \mathcal{P}_i</p>	<p>Active session attack</p> <p>On (TestPwd, sid, \mathcal{P}_i, pw') from \mathcal{A}</p> <p>retrieve (Session, sid, \mathcal{P}_i, [\mathcal{P}_j], [pw_i]) marked fresh</p> <p>if pw_i = pw' :</p> <p style="padding-left: 2em;">Mark session compromised</p> <p style="padding-left: 2em;">send “correct” to \mathcal{A}</p> <p>else</p> <p style="padding-left: 2em;">Mark session interrupted</p> <p style="padding-left: 2em;">send “wrong” to \mathcal{A}</p> <p>Completed session attack</p> <p>On (RegisterTest, sid, \mathcal{P}) from \mathcal{A}</p> <p>retrieve (Session, sid, \mathcal{P}, \cdot, \cdot) marked fresh</p> <p>Mark the record interrupted and flag it tested</p> <p>On (LateTestPwd, sid, \mathcal{P}, pw') from \mathcal{A}</p> <p>retrieve (Session, sid, \mathcal{P}, \cdot, \cdot) flagged tested</p> <p>Remove the flag</p> <p>retrieve (Key, sid, \mathcal{P}, [pw], [K], \cdot)</p> <p>if pw = pw' : set $K' := K$</p> <p>else :</p> <p style="padding-left: 2em;">Set $K' \leftarrow_{\\$} \{0, 1\}^\lambda$</p> <p style="padding-left: 2em;">Set $K' := \perp$</p> <p>send (sid, K') to \mathcal{A}</p>
---	---

Figure 1: The $\mathcal{F}_{\text{PAKE}}$, $\mathcal{F}_{\text{lePAKE}}$, and $\mathcal{F}_{\text{rePAKE}}$ ideal functionalities. The original $\mathcal{F}_{\text{PAKE}}$ functionality excludes all gray and dashed parts. Adding the [dashed] interfaces without the line in dark gray yields the lazy password extraction functionality $\mathcal{F}_{\text{lePAKE}}$. Adding the line in dark gray yields the relaxed lazy extraction functionality $\mathcal{F}_{\text{rePAKE}}$. In this paper, we include the line in light gray in *all* functionalities, to prevent isolated parties from outputting a key.

2.3.3 Hiding properties of PAKEs

For our combiner constructions, we will require two specific security properties from the underlying PAKEs.

Statistical password hiding. The first property we require is that of *statistical password hiding*—that PAKE execution transcripts unconditionally hide the password (as opposed to relying on a computational indistinguishability assumption). Crucially, we demand this property from PAKE protocols *even in situations where its security breaks*, e.g., through the discovery of an algorithm that solves the underlying computational assumption in polynomial time. Hence, we define a standalone property rather than modifying the ideal PAKE functionality. We restrict the definition to one-round PAKEs for simplicity and because the popular PAKEs satisfying this property are all one-round anyway.³

We formally define statistical password hiding through the existence of a

³In fact, we do not know of a PAKE with this property that is not one round. The failure modes were discussed briefly in the introduction, using KC-SPAKE2 and CAKE as examples.

simulator that produces nearly identically distributed protocol messages without access to the password, but is still able to compute keys from simulated messages and passwords, using a simulation trapdoor. We note that the existence of a simulator for producing *computationally* indistinguishable transcripts follows from the fact that a PAKE protocol UC-realizes $\mathcal{F}_{\text{PAKE}}$ or $\mathcal{F}_{\text{lePAKE}}$. Essentially, statistical password hiding demands that this part of the PAKE never breaks.

Definition 2 (Statistically password-hiding one-round PAKE). *Let $P = (\text{Start}, \text{Finish})$ be a one-round PAKE protocol. We call P statistically password-hiding iff there exist PPT algorithms $\text{Sim} = (\text{Start}, \text{ComputeKey})$ such that:*

- *Sim.Start is a probabilistic algorithm that takes as input a security parameter 1^λ and session identifier sid , and outputs a tuple (x, τ) .*
- *Sim.ComputeKey is an algorithm that takes as input $(\text{sid}, x, \tau_x, y, \text{pw})$ and outputs a session key.*
- *Password hiding. The simulator's messages are indistinguishable from those of the real protocol with any password. Formally, for any sid and pw ,*

$$P.\text{Start}(1^\lambda, \text{sid}, \text{pw})|_1 \approx \text{Sim.Start}(1^\lambda, \text{sid})|_1,$$

where $|_1$ denotes the first component of the output. As a consequence, a message created by the real protocol is equally explained by any password.

- *Trapdoor correctness. The simulator's session key matches the real protocol's session key when the transcript and passwords are the same. Formally, pick any sid , pw , $(x, \tau_x) \leftarrow \text{Sim.Start}(1^\lambda, \text{sid})$, and (possibly malicious) message y . Then for all $(x', \text{st}') \leftarrow P.\text{Start}(1^\lambda, \text{sid}, \text{pw})$, it is the case that*

$$x = x' \implies \text{Sim.ComputeKey}(\text{sid}, x, \tau_x, y, \text{pw}) = P.\text{Finish}(\text{st}', y)$$

When the statistical difference in the password hiding relation the above relations is 0, we say the PAKE is perfectly password-hiding. For brevity, we will avoid writing sid in the inputs of the above procedures when it is otherwise clear.

EKE, CPace, and SPAKE2 (Figure 2) are perfectly password-hiding one-round PAKEs.⁴ This property follows immediately from the fact they have UC proofs and their messages (Diffie-Hellman key shares) are uniformly distributed for any fixed password, i.e., any password equally explains any message. To demonstrate, we provide the proof for CPace below.

CPace is perfectly password-hiding. We will assume the group has prime order for simplicity, though this also holds for groups with nontrivial cofactor.

⁴We note that a SPAKE2 initiator uses a different basepoint from the responder, but our definition of password hiding requires both initiator and responder to behave identically. For simplicity we keep our definition symmetric, since its extension is clear and the difference has no effect on our analysis.

<u>Alice(sid, pw)</u>	<u>Bob(sid, pw)</u>
$G := H_0(\text{sid}, \text{pw})$	$G := H_0(\text{sid}, \text{pw})$
$r \leftarrow_{\$} \mathbb{F}$	$s \leftarrow_{\$} \mathbb{F}$
$R := rG$	$S := sG$
$Z := rS$	$Z := sR$
$K := H_1(\text{sid}, R, S, Z)$	$K := H_1(\text{sid}, R, S, Z)$
return K	return K

(a) The CPace PAKE [HL19]. The random oracle H_0 outputs values in a prime-order group.

<u>Alice(sid, pw)</u>	<u>Bob(sid, pw)</u>
$r \leftarrow_{\$} \mathbb{F}$	$s \leftarrow_{\$} \mathbb{F}$
$R := E_{\text{pw}}^{\text{sid}}(rG)$	$S := E_{\text{pw}}^{\text{sid}}(sG)$
$Z := r \cdot D_{\text{pw}}^{\text{sid}}(S)$	$Z := s \cdot D_{\text{pw}}^{\text{sid}}(R)$
$K := H_1(\text{sid}, R, S, Z)$	$K := H_1(\text{sid}, R, S, Z)$
return K	return K

(b) The EKE PAKE [BM92, JRX24]. G is a group generator, and (E, D) are the encryption and decryption algorithms of an ideal cipher.

<u>Alice(sid, pw)</u>	<u>Bob(sid, pw)</u>
$r \leftarrow_{\$} \mathbb{F}$	$s \leftarrow_{\$} \mathbb{F}$
$R := rG + \text{pw} \cdot M$	$S := sG + \text{pw} \cdot N$
$Z := r \cdot (S - \text{pw} \cdot N)$	$Z := s \cdot (R - \text{pw} \cdot M)$
$K := H_1(\text{sid}, \text{pw}, R, S, Z)$	$K := H_1(\text{sid}, \text{pw}, R, S, Z)$
return K	return K

(c) The SPAKE2 PAKE [AP05]. $G, M,$ and N are unrelated generators of a prime-order group. pw is interpreted as an integer (perhaps via hashing).

Figure 2: Three perfectly password-hiding PAKEs.

We define Sim as follows. First, we let the simulator respond to H_0 queries by selecting a scalar u , recording u , and returning uB where B is a fixed generator. Sim.Start picks a uniform scalar r and computes $R := rB$. The output is $(x = R, \tau = r)$.

Sim.ComputeKey takes r, R, S, pw . Let \tilde{u} be the recorded discrete log of $H_0(\text{sid}, \text{pw})$. The procedure computes $\tilde{r} := r/\tilde{u}$, and $Z := \tilde{r}S$. The output is $H_1(\text{sid}, R, S, Z)$.

This satisfies the trapdoor correctness requirement. Suppose $(R', r') \leftarrow \text{P.Start}(1^\lambda, \text{sid}, \text{pw})$ such that $R = R' = r'H_0(\text{sid}, \text{pw})$. This implies $rB = r'\tilde{u}B$, and thus $r' = \tilde{r}$, since basepoint scalar multiplication is injective. Thus,

$$\text{Sim.ComputeKey}(r, R, S, \text{pw}) = Z = \tilde{r}S = r'S = \text{P.Finish}(r', S)$$

This is also perfectly hiding, since both Sim.Start and P.Start output a uniformly random group element.

PSK equality hiding. The second property we require is that of *statistical PSK equality hiding*—that a PAKE, when executed by two parties using high-entropy passwords (a.k.a., a preshared key, or PSK), does not leak whether the passwords match. As before, we require this property even when the PAKE

<pre> Game $G_{b,P}^{\text{PEH}}(\mathcal{A})$ ----- $\text{psk}_0 \leftarrow_{\\$} \{0, 1\}^\lambda$ if $b = 0$: $\text{psk}_1 \leftarrow_{\\$} \{0, 1\}^\lambda$ else: $\text{psk}_1 := \text{psk}_0$ $i := 0$ //wlog the initiator is participant 0 $(\text{st}_0, \text{st}_1) := (\text{P}_0.\text{Start}(\text{psk}_0), \text{P}_1.\text{Start}(\text{psk}_1))$ $\text{st}_{\mathcal{A}} := \emptyset$ $m := \emptyset$ </pre>	<pre> while $(\text{done}_0, \text{done}_1) \neq (\text{true}, \text{true})$: $(\text{st}'_i, \text{done}'_i, m') := \text{P}_i.\text{Next}(\text{st}_i, m)$ $(\text{st}'_{\mathcal{A}}, m) := \mathcal{A}(\text{st}_{\mathcal{A}}, m')$ $\text{st}_i := \text{st}'_i$ $\text{st}_{\mathcal{A}} := \text{st}'_{\mathcal{A}}$ $\text{done}_i := \text{done}'_i$ $i := i \oplus 1$ $b' := \mathcal{A}(\text{st}_{\mathcal{A}}, \text{final})$ return $b = b'$ </pre>
--	---

Figure 3: The PSK equality hiding (PEH) security game for the PAKE P . \mathcal{A} may read and modify any messages, and has access to any ideal functionality, e.g., a random oracle, that P has access to.

is otherwise insecure. Note statistical password hiding implies statistical PSK equality hiding, but not vice-versa. It is thus a strictly weaker property. We provide a proper definition below. Unlike with password hiding, we define this property to include multi-round PAKEs.

Definition 3 (PSK Equality Hiding). *Let P be a PAKE protocol. We say that P is PSK equality hiding iff an active adversary observing two honest parties interacting cannot distinguish between the case where the parties share a high-entropy password and the case where they have distinct high-entropy passwords.*

Formally, we define the advantage of an adversary \mathcal{A} as

$$\text{Adv}_{\text{P}}^{\text{PEH}}(\mathcal{A}) = |2 \cdot \Pr[G_{b,P}^{\text{PEH}}(\mathcal{A}) = b \mid b \leftarrow_{\$} \{0, 1\}] - 1|,$$

where G^{PEH} is the PSK equality hiding game defined in Figure 3. Note we permit the adversary to read and modify messages, but not see the final protocol output.

We say that P is statistically PSK equality hiding if $\text{Adv}_{\text{P}}^{\text{PEH}}(\mathcal{A}) = \text{negl}(\lambda)$ for any unbounded \mathcal{A} .

As with password hiding, the computational version of PSK equality hiding is guaranteed by the fact that the PAKE UC-realizes one of the PAKE functionalities, but the statistical version is not.⁵

Fortunately, some recent efficient universally composable PAKEs such as CAKE [BCP⁺23] and EKE-PRF [JRX24] (Figure 4) achieve this property. When instantiated with an appropriate post-quantum KEM and post-quantum key agreement protocol, respectively, these constructions yield post-quantum PAKEs. At a high level, this property holds because both protocols are two rounds, and in

⁵In fact, this is a strict separation. Take any secure PAKE and add $(\text{pk}, \text{Enc}_{\text{pk}}(\text{pw}))$ to both the initiator's and responder's first messages, where pk is a public key in some public key encryption scheme. Then it is clear that this new PAKE is secure, but at best has PSK equality hiding that depends on the computational hardness of decrypting a ciphertext in the asymmetric encryption scheme. This sort of construction occurs in existing PAKEs. The one-round PAKE from smooth projective hash functions presented in [KV11] is neither statistically password hiding nor statistically PSK equality hiding for precisely this reason.

<u>Alice(sid, pw)</u>		<u>Bob(sid, pw)</u>
$(pk, sk) := \text{Keygen}(1^\lambda)$		
$epk := E_{pw}^{\text{sid}\parallel 1}(pk)$	$\xrightarrow{\text{epk}}$	$pk := D_{pw}^{\text{sid}\parallel 1}(epk)$
$c := D_{pw}^{\text{sid}\parallel 2}(ec)$	\xleftarrow{ec}	$(c, Z) := \text{Encap}(pk)$
$Z := \text{Decap}_{sk}(c)$		$ec := E_{pw}^{\text{sid}\parallel 2}(c)$
$K := H(\text{sid}, epk, ec, Z)$		$K := H(\text{sid}, epk, ec, Z)$
return K		return K

(a) The CAKE KEM-based PAKE [BCP⁺23]. (E, D) are the encryption and decryption algorithms of an ideal cipher.

<u>Blindⁱ(sid, pw, v)</u>		<u>Unblindⁱ(sid, pw, (s, t))</u>
$r \leftarrow_{\mathcal{S}} \{0, 1\}^{3\lambda}$		$h' := H_1(\text{sid}, i, pw, t); r := s \oplus h'$
$h := H_0^{\mathcal{M}_i}(\text{sid}, i, pw, r); t := v \odot h^{-1}$		$h := H_0^{\mathcal{M}_i}(\text{sid}, i, pw, r); v := h \odot t$
$h' := H_1(\text{sid}, i, pw, t); s := r \oplus h'$		return v
return (s, t)		
<u>Alice(sid, pw)</u>		<u>Bob(sid, pw)</u>
$a \leftarrow_{\mathcal{S}} \mathcal{R}$		$b \leftarrow_{\mathcal{S}} \mathcal{R}$
$A := \text{KA.msg}_1(a)$		$A := \text{Unblind}^1(\text{sid}, pw, \phi)$
$\phi \leftarrow \text{Blind}^1(\text{sid}, pw, A)$	$\xrightarrow{\phi}$	$B := \text{KA.msg}_2(b, A)$
$B := \text{Unblind}^2(\text{sid}, pw, \phi')$	$\xleftarrow{\phi'}$	$\phi' \leftarrow \text{Blind}^2(\text{sid}, pw, B)$
$K := \text{KA.key}_1(a, B)$		$K := \text{KA.key}_2(b, A)$
return $\text{PRF}_K(\phi')$		return $\text{PRF}_K(\phi')$

(b) The EKE-PRF PAKE [JRX24]. KA is a two-round key agreement protocol with randomness space \mathcal{R} , message spaces $\mathcal{M}_1, \mathcal{M}_2$ (first and second messages, respectively), both of which are abelian groups (e.g., the module lattice used in ML-KEM, with addition) with group operation denoted \odot . The random oracle $H_0^{\mathcal{M}_i}$ takes on values in \mathcal{M}_i and H_1 takes on values in $\{0, 1\}^{3\lambda}$. PRF is a pseudorandom function.

Figure 4: Two post-quantum statistically PSK-equality-hiding PAKEs.

each round, both messages are encrypted under the password using independent ideal ciphers or random oracles. In other words, the parties in these protocols are communicating over a confidential channel.

We provide the formal statistical PSK equality hiding argument for CAKE below, as well as a sketch for EKE-PRF. We remark that proving statistical PSK equality hiding for OEKE-style PAKEs such as OEKE-PRF, OCAKE, CHIC, and NoIC [JRX24, BCP⁺23, ABJS24, ABJ25], is more subtle and would require a careful analysis of whether the unencrypted second message reveals statistical differences in uniform versus honestly generated KEM public keys.⁶ We thus leave this to future work.

CAKE is statistically PSK-equality-hiding. Let game 0 be the PEH game with $b = 0$, i.e., the passwords are different. In game 1, rather than sending $ec := E_{pw_A}^{sid||1}(pk)$, Bob sends a uniformly chosen ciphertext. Since the ideal cipher is a family of uniformly chosen permutations, and both ideal cipher instances are independent, the advantage of this hop is bounded by the ability of the adversary to guess pw_B , which is $2^{-\lambda}$ times the number of ideal cipher queries.

In game 2, replace Bob’s copy of pk with a uniformly chosen public key. Since the ciphertext c is discarded anyway, this hop is perfect.

In game 3, replace Alice’s copy of epk with a uniformly random public key. This does not affect Bob’s computations, since epk is discarded. Further, since the ideal cipher is a family of uniformly chosen permutations and the instances are independent the advantage of this hop is bounded by the ability of the adversary to guess pw_A , which is $2^{-\lambda}$ times the number of ideal cipher queries.

In game 4, we set $b = 1$, i.e., we make Alice and Bob’s passwords equal. This does not affect the transcript or any computations.

In games 5, 6, and 7, we undo games 3, 2, and 1, respectively, and we are done.

EKE-PRF is statistically PSK-equality-hiding. This argument operates nearly identically. Since EKE-PRF includes many more intermediate variables, we present a sketch.

We may replace ϕ' with randomness, since both components are perfectly blinded with independent outputs of a random oracle (namely, h and h'), whose input includes pw_B , which is high-entropy. We may then replace Bob’s copy of A with randomness, since no function of Bob’s A appears in the transcript anymore. We then replace ϕ with randomness, since it does not affect Bob’s computation, and both components are perfectly blinded with independent outputs of a random oracle whose input includes pw_A , which is high-entropy.

⁶There are KEMs for which these constructions are provably *not* statistically PSK-equality-hiding. Consider Regev’s public key encryption scheme [Reg05] as an IND-CPA-secure KEM. In all the aforementioned protocols, an encapsulated key is sent in the clear. Consider the case where the (high-entropy) passwords are equal, i.e., encapsulation is to an honestly generated public key. In Regev’s scheme, every encrypted 0 bit is an LWE sample, and every encrypted 1 bit is not. So an adversary with an LWE oracle will observe 50% LWE samples. However, in the case where the passwords are unequal, encapsulation is to a random element in the space of public keys. With overwhelming probability *none* of the encrypted bits are LWE samples. Thus, an adversary can win the PEH game if given an LWE oracle.

We then set the passwords equal, which does not affect anything, and then we undo all the above steps in reverse order.

3 Constructions

In this section we present our two PAKE combiners, **ParComb** (*parallel combiner*) and **SeqComb** (*sequential combiner*). Unlike the failed attempt at a combiner in our introduction, these combiners will require specific hiding properties of their underlying PAKEs. With these assumptions, we show that both combiners achieve security in the UC model.

3.1 The ParComb combiner

We return to the insecure parallel combiner of our introduction. Recall the reason it fails is, if the underlying security assumption were broken, the transcripts of the underlying PAKEs would contain enough auxiliary information to check password guesses. For KC-SPAKE2, this auxiliary information is a MAC whose key is derived from the Diffie-Hellman shared secret, and for CAKE, this is the fact that KEM ciphertexts are LWE samples.

However, not every PAKE has this auxiliary information. In fact, some widely deployed, well-standardized PAKEs such as SPAKE2 or CPace [Lad23, AHH24] do not. The lack of such auxiliary information is what we formalize in Definition 2 (statistical password hiding). As it turns out, this property is necessary and sufficient in order to make our parallel combiner work. At a high level, **ParComb** executes two statistically password-hiding one-round PAKEs P_1 and P_2 in parallel, and hashes the resulting keys. This combiner is highly efficient, requiring only one extra hash, and has round complexity bounded by the maximum round complexity of the underlying PAKEs. We give the full construction in Figure 5.

3.1.1 Security

Before presenting the security theorem, we provide intuition for each assumption and design decision.

Statistical password hiding. We know from our broken combiner example why statistical password hiding is necessary for P_1 and P_2 . We provide intuition now for why it is sufficient. Consider what happens when P_1 is broken and P_2 remains secure. An active adversary can engage in P_1 with Bob and be able to efficiently compute the guessing function $\text{pw} \mapsto K'_1$ (in EKE, for example, this requires solving computational Diffie-Hellman). However, since the messages are statistically password-hiding, the only way to check a guess is to compare it to the output key $K := H(\text{sid}, \text{tr}, K_1, K_2)$. So the adversary is in the scenario where it can make guesses of pw and see (a hash of) K_2 . This is precisely the UC PAKE game for P_2 , which is assumed to be secure.

The same argument shows that P is secure when P_2 is broken and P_1 remains secure.

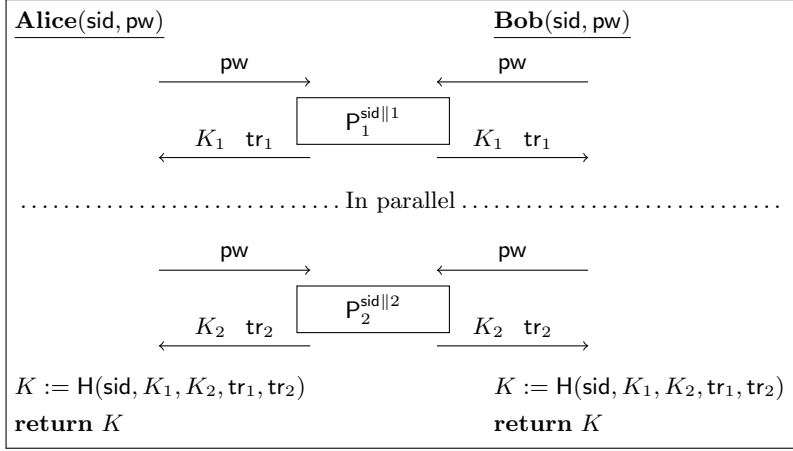


Figure 5: The ParComb combiner for statistically password-hiding one-round PAKEs P_1 and P_2 . tr_i is the protocol transcript of the execution of P_i . The random oracle H outputs values in $\{0, 1\}^\lambda$. sid is unique to every execution.

Hashing the transcript. We include the full protocol transcript in the final key computation so that the session key can be simulated when messages are mauled. Suppose P_1 is broken, and the adversary has modified one of the P_1 messages. If a simulator does not know either party’s password, then it cannot decide whether P_1 should succeed or fail. And since P_1 is broken, the simulator cannot rely on the PAKE UC-realizing $\mathcal{F}_{\text{PAKE}}$. Hashing the transcript resolves the ambiguity, since any mauled messages will cause the protocol to fail (i.e., return different keys) with overwhelming probability.

We now state the security theorem for ParComb. The arguments of the proof are sketched above, and the full proof can be found in Appendix A.1.

Theorem 1. *Let P_1 and P_2 be statistically password-hiding one-round PAKE protocols, and let $P := \text{ParComb}[P_1, P_2]$. Then the following hold in the static-corruptions setting:*

1. *If P_1 UC-realizes $\mathcal{F}_{\text{PAKE}}$, then so does P (in the \mathcal{F}_{RO} -hybrid model)*
2. *If P_2 UC-realizes $\mathcal{F}_{\text{PAKE}}$, then so does P (in the \mathcal{F}_{RO} -hybrid model).*

In other words, $\text{ParComb}[P_1, P_2]$ is at least as strong as the strongest of P_1, P_2 , and is hence a good combiner.

3.1.2 Instantiations

We now discuss the PAKEs that can be used to instantiate ParComb with hybrid security.

(Not) Achieving $\mathcal{F}_{\text{PAKE}}$. It is not currently possible to instantiate ParComb to achieve hybrid security and realize $\mathcal{F}_{\text{PAKE}}$. One of the underlying PAKEs can be

straightforwardly instantiated with EKE [BM92], since it is one-round, perfectly password hiding, and UC-realizes $\mathcal{F}_{\text{PAKE}}$ [JRX24]. The other underlying PAKE remains a question, though, as there is no known post-quantum, one-round, statistically password-hiding PAKE protocol that UC-realizes $\mathcal{F}_{\text{PAKE}}$.

One line of work includes the EKE-NIKE construction from [BGHJ24], using the lattice-based NIKE, Swoosh [GdKQ⁺24]. However, this does not have statistical password hiding, since messages are of the form $E_{\text{pw}}(\mathbf{t})$, where \mathbf{t} is an LWE sample and E is an ideal cipher. This has the same problem as the CAKE construction in the introduction.⁷ Thus, we do not yet have a one-round hybrid construction UC-realizing $\mathcal{F}_{\text{PAKE}}$.

Achieving $\mathcal{F}_{\text{blePAKE}}$. A concurrent work [LL24] shows that `ParComb` can be a *bare lazy-extraction PAKE* (blePAKE), i.e., a lazy-extraction PAKE that does not need a unique `sid`. Specifically, the authors show that, when `ParComb` is instantiated with an lePAKE and a blePAKE (both with statistical password hiding), then it UC-realizes $\mathcal{F}_{\text{blePAKE}}$. They also show that X-GA-PAKE (isogeny-based), SPAKE2 (ECDH-based), and TBPEKE (ECDH-based) [AEK⁺22, AP05, PW17], are blePAKEs with statistical password hiding. Thus, `ParComb` can be instantiated as a one-round hybrid construction UC-realizing $\mathcal{F}_{\text{blePAKE}}$.

3.2 The SeqComb combiner

To arrive at the `SeqComb` combiner, we again consider the minimum possible amount of security necessary for a combined PAKE. Beginning our protocol with a one-round statistically password-hiding PAKE cannot hurt, since even when it is broken, it is difficult to exploit for a password guess. To handle the case that this PAKE P_1 is broken, we may pass `pw` along with the P_1 's session key to a second PAKE P_2 , receive a second session key, and output the hash of everything. If P_2 is secure, then including its session key in the final hash will be sufficient to ensure that the overall protocol is secure. If the P_2 is broken then including P_1 's session key in the final hash will be sufficient, so long as P_2 is not so broken as to reveal `pw` outright.

This is almost exactly the definition of `SeqComb`. That is, `SeqComb` executes a statistically password-hiding PAKE P_1 , feeds its output into a statistically PSK equality hiding PAKE P_2 , and hashes everything at the end. Like `ParComb`, this combiner is highly efficient, requiring only two extra hashes. Unlike `ParComb`, its round complexity is the sum of the number of rounds of the underlying PAKEs, rather than the maximum. We give the full construction in Figure 6.

3.2.1 Security

Again, before presenting the security theorem, we provide intuition for why each component of the protocol is necessary.

P_1 statistical password hiding. We require P_1 to be statistically password-hiding for the same reason as in `ParComb`: it directly deals with `pw`, and so any

⁷While it appears no post-quantum PAKE has this property, we do not believe it is inherent. See Section 4 for potential new directions.

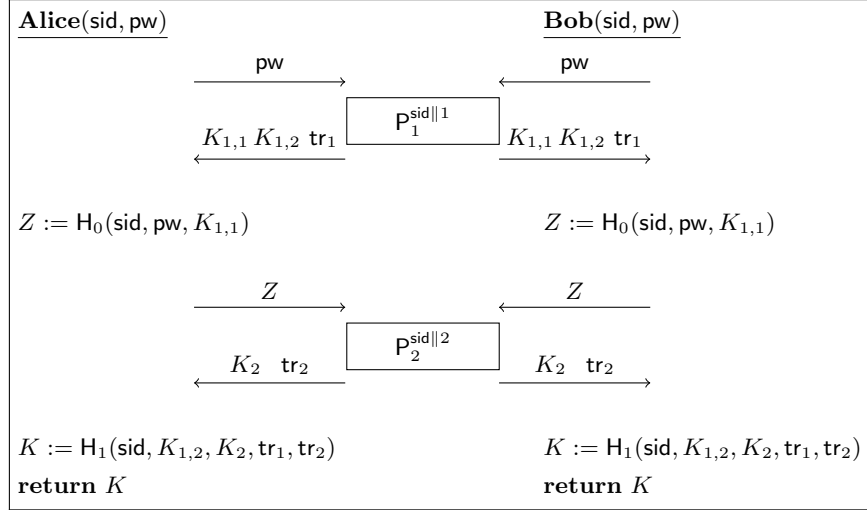


Figure 6: The SeqComb combiner for PAKEs P_1, P_2 , where P_1 is one round and statistically password-hiding, and P_2 is statistically PSK equality-hiding. tr_i is the protocol transcript of the execution of P_i . Both random oracles H_0, H_1 output values in $\{0, 1\}^\lambda$. sid is unique to every execution. We assume P_1 outputs two keys, each of bitlength λ (if not, it suffices to apply a PRG to stretch its output).

leakage here yields a leakage in the combined protocol. More rigorously, if P_1 is broken, an active adversary can efficiently compute the mapping $\text{pw} \mapsto Z'$, but not receive any extra information on whether its guess was correct. So the adversary is in the scenario where it has guesses for Z and can see (a hash of) K_2 . This is precisely the UC PAKE game for P_2 .

P_2 statistical PSK equality hiding. We require P_2 to be statistically PSK equality-hiding in order to permit its transcript to be simulated when P_2 is broken. In the case where both SeqComb parties are honest, and have completed P_1 , a simulator must simulate P_2 without knowing whether P_1 succeeded, i.e., without knowing whether the parties agree on the input to P_2 . Since this input is high-entropy by assumption (P_1 is a secure PAKE), then the PSK equality hiding property is sufficient to permit the simulator to simulate one case and have it be indistinguishable from the other.

Hashing $K_{1,1}$. We include $K_{1,1}$ in the hash Z because pw alone is insufficient in the case that P_2 is broken. If P_2 were so broken as to leak its password Z in its entirety (technically prevented by PSK equality hiding, but we may assume), then pw would be trivially revealed. We therefore hash $K_{1,1}$, which we assume is high entropy, into Z . Thus, even if Z leaks in its entirety, pw is unaffected.

Hashing pw . We include pw in the computation of Z because $K_{1,1}$ is not sufficient in the case that P_1 is broken. If P_1 is broken, then we must pessimistically

assume that both $K_{1,1}$ and $K_{1,2}$ are fully controlled by the adversary. Thus, simply feeding one of these values into P_2 yields no security. Instead, we hash pw with $K_{1,1}$ and feed that into P_2 .

Hashing $K_{1,2}$. We include $K_{1,2}$ in the final session key computation because K_2 is not sufficient in the case that P_2 is broken. It suffices to include some yet-unused entropy from the session key of P_1 .

We now state the security theorem for `SeqComb`. The arguments of the proof are sketched above, and the full proof can be found in Appendix A.2.

Theorem 2. *Let P_1 be a one-round statistically password-hiding PAKE protocol, let P_2 be a statistically PSK equality hiding PAKE protocol, and let $P := \text{SeqComb}[P_1, P_2]$. Then the following hold in the static-corruptions setting:*

1. *If P_1 UC-realizes $\mathcal{F}_{\text{PAKE}}$, then so does P (in the \mathcal{F}_{RO} -hybrid model).*
2. *If P_2 UC-realizes $\mathcal{F}_{\text{PAKE}}$, then so does P (in the \mathcal{F}_{RO} -hybrid model).*

In other words, $\text{SeqComb}[P_1, P_2]$ is at least as strong as the strongest of P_1, P_2 , and is hence a good combiner.

3.2.2 Using weaker P_1

Some of the most widely deployed PAKEs such as the Diffie-Hellman-based CPace and SPAKE2 do not enjoy full UC PAKE security but are so-called *lazy extraction* PAKEs (lePAKEs). These PAKEs allow adversaries to successfully complete an active attack even after the attacked party finished the protocol. Such prolonged attack completions do not seem to pose any real-world threat, and lazy extraction PAKEs are in widespread use (e.g., for Facebook Messenger chat history transfer [Fac23]). Hence, to make `SeqComb` more applicable, we would like to make it work when P_1 UC-realizes $\mathcal{F}_{\text{lePAKE}}$. This functionality is depicted in Figure 1 and is strictly weaker than $\mathcal{F}_{\text{PAKE}}$ due to the additional password test interface on completed sessions. This interface allows the “lazy” attacker to provide a password guess even after an actively attacked party output a session key, and it provides the attacker with that key or a random one depending on whether the guess was correct.

Perhaps surprisingly, `SeqComb` is not as tolerant of a lazy-extraction P_1 as one would hope for: $\text{SeqComb}[\mathcal{F}_{\text{lePAKE}}, P_2]$ does not UC-realize $\mathcal{F}_{\text{lePAKE}}$. This can be made more intuitive by looking at `SeqComb`’s design. In the case where P_2 cannot be trusted to protect K_2 , security is based solely on the output keys $K_{1,1}, K_{1,2}$ of P_1 . When these values are revealed to the environment through a correct late password guess on $\mathcal{F}_{\text{lePAKE}}$, the simulator has to explain how they lead to the actual Z and output key K of the honest party. On the other hand, if the guess was wrong, no connection should exist between $K_{1,1}, K_{1,2}$ and Z, K . Since $\mathcal{F}_{\text{lePAKE}}$ does not leak the information of whether the guess was correct or not, the simulator is trapped. We formalize this intuition in Appendix B with a formal distinguisher between `SeqComb` with a lazy extraction P_1 and $\mathcal{F}_{\text{lePAKE}}$.

This intuition implies a fix. It would suffice to let the simulator know whether a late password guess succeeded. A PAKE that allows one late password

guess with responses of the form “correct” or “wrong” is called a *relaxed PAKE* [ABB⁺20]. Combining this property with the lazy extraction property yields what we call *relaxed lazy extraction PAKE*, or *rllePAKE*. We depict the functionality in Figure 1. It has the same guarantees as $\mathcal{F}_{\text{lePAKE}}$, but additionally leaks the information whether or not the late password guess was correct—a weakening that does not seem to rule out any applications of PAKE protocols. We prove the following result in Appendix A.3.

Lemma 1. *Let P_1 be a statistically password-hiding one-round PAKE protocol, let P_2 be a statistically PSK equality-hiding PAKE protocol, and let $P := \text{SeqComb}[P_1, P_2]$. If P_1 UC-realizes $\mathcal{F}_{\text{lePAKE}}$, then P UC-realizes $\mathcal{F}_{\text{rllePAKE}}$ (in the \mathcal{F}_{RO} -hybrid model) in the static-corruptions setting.*

If P_2 UC-realizes $\mathcal{F}_{\text{PAKE}}$, which is stronger than $\mathcal{F}_{\text{rllePAKE}}$, we can use Theorem 2 to conclude that this is a good combiner which achieves at least $\mathcal{F}_{\text{rllePAKE}}$. This implies that SeqComb is safe to use when P_1 is instantiated with a perfectly password-hiding lazy extraction PAKE such as CPace or SPAKE2.

3.2.3 Concrete instantiations yielding hybrid PAKE

SeqComb can be instantiated with existing efficient protocols to produce a hybrid (rlle)PAKE. For P_1 , it is possible to use the Diffie-Hellman-based CPace, SPAKE2, or EKE. For P_2 , any statistically PSK equality-hiding PAKE will do, e.g., any plausibly post-quantum EKE-style protocols such as CAKE and EKE-PRF. We discuss the open problem of proving statistical PSK equality hiding for more PAKEs in Section 4.

3.2.4 Achieving security in the QROM

We note that our combiners require random oracles to function, and our instantiations include post-quantum PAKEs that rely on an ideal-cipher or random oracle. Since our goal is (hybrid) post-quantum security, it is important to consider models which more accurately depict the capabilities of a quantum adversary.

SeqComb was analyzed in the Quantum Random Oracle Model (QROM) in concurrent work [LL24]. The authors show that when P_1 (a statistically password-hiding PAKE) and P_2 UC-realize $\mathcal{F}_{\text{PAKE}}$, SeqComb nearly UC-realizes $\mathcal{F}_{\text{PAKE}}$ in the \mathcal{F}_{QRO} -hybrid model.⁸ We say nearly to indicate that the simulation proof is made with respect to a *restricted set of environments*, namely those that never start the protocol with mismatching passwords. While this does not yield full UC emulation and hence does not allow one to apply, e.g., the UC composition theorem to their construction, it does relax the requirements on the underlying PAKE. For this theorem to hold, P_2 can be an *arbitrary* PAKE, not just one with statistical PSK equality hiding. Thus, with this notion of

⁸ More precisely, a QROM reduction is only provided in the case that P_1 is broken and P_2 is secure. This is a relatively small distinction, though, since, if the opposite is true and P_1 is a classical PAKE, then the combined PAKE is already insecure against quantum adversaries.

Combiner	P_1	P_2	Combined	Model
ParComb	$\mathcal{F}_{\text{PAKE}+\text{PH}}$	$\mathcal{F}_{\text{PAKE}+\text{PH}}$	$\mathcal{F}_{\text{PAKE}}$	ROM Theorem 1
	$\mathcal{F}_{\text{blePAKE}+\text{PH}}$	$\mathcal{F}_{\text{lePAKE}+\text{PH}}$	$\mathcal{F}_{\text{blePAKE}}$	ROM [LL24, Thm. 2]
SeqComb	$\mathcal{F}_{\text{PAKE}+\text{PH}}$	$\mathcal{F}_{\text{PAKE}+\text{PEH}}$	$\mathcal{F}_{\text{PAKE}}$	ROM Theorem 2
	$\mathcal{F}_{\text{lePAKE}+\text{PH}}$	$\mathcal{F}_{\text{PAKE}+\text{PEH}}$	$\mathcal{F}_{\text{rlePAKE}}$	ROM Lemma 1
	$\mathcal{F}_{\text{PAKE}+\text{PH}}$	$\mathcal{F}_{\text{PAKE}}$	$(\mathcal{F}_{\text{PAKE}})$	QROM ⁸ [LL24, Thm. 4]

Table 1: A summary of combiner results from this work and concurrent work [LL24]. The +PH marker indicates the PAKE must be statistically password-hiding. The +PEH marker indicates the PAKE must be statistically PSK equality hiding. (\mathcal{F}) indicates the analysis was done w.r.t. a restricted set of UC environments.

security, SeqComb can be instantiated with a classical PAKE like EKE, and a post-quantum PAKE proven in the QROM, such as $\text{PAKE}_{\text{lwe}}^{\text{QRO}}$ [LLH24].⁹ In that sense, our work shows that the gap between standalone and universally composable security of SeqComb is upper bounded by the PSK equality hiding of the second PAKE.

For ease of reference, we provide a breakdown of all our results along with those of concurrent work in Table 1.

4 Future work

We identify some interesting problems which we do not address in this work.

Adaptive security. So far, we have only defined hybrid PAKE in the static corruptions setting. This may be inherent to the combiner construction. No one-round PAKE can achieve security in the UC model with adaptive corruptions. A simple attack is as follows. The environment \mathcal{Z} initiates a session between an honest Alice and Bob. \mathcal{Z} forwards Alice’s message to Bob, allowing Bob to finish the protocol and output K , and withholds Bob’s message to Alice. Now, \mathcal{Z} corrupts Alice. There is no way for a simulator to respond to this corruption. This is because the simulator must have used $\mathcal{F}_{\text{PAKE}}.\text{NewKey}$ to key Bob, since both parties were honest at the time. Thus, the simulator does not know K , and so cannot choose an appropriate response to \mathcal{Z} ’s corruption query. We conclude that any combiner which yields one-round PAKEs is limited to static corruptions.

We thus ask whether there exists a method for constructing adaptively secure PAKEs and, if so, whether it can be done generically.

⁹This PAKE is not statistically PSK equality hiding. If the underlying extractable lossy public key encryption (eLPKE) scheme is broken, then the responder’s message can be decrypted to reveal pw. This can be checked for equality against the initiator’s password by adding pw to the initiator’s public key and using an LWE oracle to check if it is an LWE sample.

PSK equality hiding for more PAKEs. We show statistical PSK equality hiding for two *EKE-style PAKEs*, i.e., two-round PAKEs wherein both messages are encrypted. We do not know, however, how to show this property for *OEKE-style PAKEs*, i.e., two-round PAKEs with only one encrypted message [MRR20, PZ23, BCJ⁺19, SGJ23, ABJS24, JRX24, ABJ25]. Since this is a large class of protocols, it would be useful to establish the properties necessary of the underlying cryptographic primitive (e.g., the structure of an MLWE-based KEM) to achieve statistical hiding.

One-round hybrid PAKE. In our assessment of `ParComb`, we concluded that the combiner cannot achieve $\mathcal{F}_{\text{PAKE}}$ using a post-quantum PAKE, since none has statistical password hiding, to the authors’ knowledge. This does not seem inherent, though. We conjecture that it is possible to modify EKE-NIKE[Swoosh] [BGHJ24, GdKQ⁺24] to statistically hide the password. This would involve sending two pairs of public keys, replacing public keys of the form $\mathbf{As} + \mathbf{e}$ with ones of the form \mathbf{As} for $\mathbf{A}, \mathbf{s}, \mathbf{e}$ of appropriate dimension. Effectively, this undoes the size improvements of [LP11] in exchange for statistically uniform public keys. Further, since the uniformity result relies on the leftover hash lemma [HILL99], which is not known to apply in the cyclotomic rings of typical module lattice constructions, the NIKE must use unstructured LWE. We estimate public key sizes in the resulting scheme to be 1GB at the 128-bit security level.

We leave as an open question whether such a scheme could be instantiated and improved on to be more practical.

Analysis against quantum adversaries. While our combiners can, in principle, use PAKEs that are secure against quantum adversaries, our reductions are classical and exist in the ROM. As mentioned earlier, concurrent work [LL24] analyzes `SeqComb` in the QROM and proves a theorem with respect to a restricted set of environments. We believe that this difference is small enough that the analysis may transfer immediately to our unrestricted environment setting, yielding a proof of (the second case of) Theorem 2 in the \mathcal{F}_{QRO} -hybrid model. We leave this for future work.

Combiners for *strong* PAKE variants. As mentioned earlier, hybrid PAKE implies hybrid aPAKE and iPAKE via the Ω -method [GMR06, LLH25] and LATKE [KR24]. However, it is not clear how to achieve the *strong* variants of these constructions (called *saPAKE* and *siPAKE*, respectively). A strong aPAKE is an aPAKE which is resistant to precomputation attacks—an adversary must compromise a server before they are able to begin brute-forcing the password. Existing constructions such as OPAQUE [JKX18] use an oblivious pseudorandom function (OPRF) to achieve strongness. However, no hybrid OPRF construction is currently known, and so it is not clear how to hybridize OPAQUE. More generally, it would be interesting to understand how much strongness relies on an OPRF-like construction, and how much can be achieved using generic combiners. For (s)aPAKE and (s)iPAKE in general, such black-box combiners would further be desirable in order to give developers the option to plug post-quantum protocols into their existing deployments.

Acknowledgements

We thank Phillip Gajland for the suggestion of the information-theoretic variant of Swoosh. We thank You Lyu and Shengli Liu for their helpful feedback after we found out we all wrote more or less the same paper. We also thank Jelle Vos and Stanislaw Jarecki (both supported by Apple) for catching an error in our proofs and pointing us to subtleties in proving PSK equality hiding for OEKE-style post-quantum PAKEs. Julia wants to say: thank you, Michael, for sharing your inspiring ideas with me—this has been an awesome project!

References

- [ABB⁺20] Michel Abdalla, Manuel Barbosa, Tatiana Bradley, Stanislaw Jarecki, Jonathan Katz, and Jiayu Xu. Universally composable relaxed password authenticated key exchange. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part I*, volume 12170 of *LNCS*, pages 278–307. Springer, Heidelberg, August 2020. doi:10.1007/978-3-030-56784-2_10.
- [ABBO24] David Adrian, Bob Beck, David Benjamin, and Devon O’Brien. Advancing Our Amazing Bet on Asymmetric Cryptography, 2024. URL: <https://blog.chromium.org/2024/05/advancing-our-amazing-bet-on-asymmetric.html>.
- [ABJ25] Afonso Arriaga, Manuel Barbosa, and Stanislaw Jarecki. NoIC: PAKE from KEM without ideal ciphers. Cryptology ePrint Archive, Paper 2025/231, 2025. URL: <https://eprint.iacr.org/2025/231>.
- [ABJS24] Afonso Arriaga, Manuel Barbosa, Stanislaw Jarecki, and Marjan Skrobot. C’est très chic: A compact password-authenticated key exchange from lattice-based kem. Cryptology ePrint Archive, Paper 2024/308, 2024. <https://eprint.iacr.org/2024/308>. URL: <https://eprint.iacr.org/2024/308>.
- [AEK⁺22] Michel Abdalla, Thorsten Eisenhofer, Eike Kiltz, Sabrina Kunzweiler, and Doreen Riepel. Password-authenticated key exchange from group actions. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 699–728. Springer, Heidelberg, August 2022. doi:10.1007/978-3-031-15979-4_24.
- [AHH21] Michel Abdalla, Björn Haase, and Julia Hesse. Security analysis of CPace. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part IV*, volume 13093 of *LNCS*, pages 711–741. Springer, Heidelberg, December 2021. doi:10.1007/978-3-030-92068-5_24.
- [AHH24] Michel Abdalla, Björn Haase, and Julia Hesse. CPace, a balanced composable PAKE. Internet Draft draft-irtf-cfrg-cpace-13, Internet Engineering Task Force, October 2024. Num Pages: 96. URL: <https://datatracker.ietf.org/doc/draft-irtf-cfrg-cpace-13>.
- [AP05] Michel Abdalla and David Pointcheval. Simple password-based encrypted key exchange protocols. In Alfred Menezes, editor, *CT-RSA 2005*, volume 3376 of *LNCS*, pages 191–208. Springer, Heidelberg, February 2005. doi:10.1007/978-3-540-30574-3_14.

- [BBB⁺24] Daniel J. Bernstein, Karthikeyan Bhargavan, Shivam Bhasin, Anupam Chattopadhyay, Tee Kiah Chia, Matthias J. Kannwischer, Franziskus Kiefer, Thales Paiva, Prasanna Ravi, and Goutam Tamvada. KyberSlash: Exploiting secret-dependent division timings in kyber implementations. *Cryptology ePrint Archive*, Paper 2024/1049, 2024. URL: <https://eprint.iacr.org/2024/1049>.
- [BCD⁺16] Joppe W. Bos, Craig Costello, Léo Ducas, Ilya Mironov, Michael Naehrig, Valeria Nikolaenko, Ananth Raghunathan, and Douglas Stebila. Frodo: Take off the ring! Practical, quantum-secure key exchange from LWE. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 1006–1018. ACM Press, October 2016. doi: 10.1145/2976749.2978425.
- [BCD⁺24] Manuel Barbosa, Deirdre Connolly, João Diogo Duarte, Aaron Kaiser, Peter Schwabe, Karolin Varner, and Bas Westerbaan. X-wing. *IACR Communications in Cryptology*, 1(1), 2024. doi: 10.62056/a3qj89n4e.
- [BCJ⁺19] Tatiana Bradley, Jan Camenisch, Stanislaw Jarecki, Anja Lehmann, Gregory Neven, and Jiayu Xu. Password-authenticated public-key encryption. In Robert H. Deng, Valérie Gauthier-Umaña, Martín Ochoa, and Moti Yung, editors, *ACNS 19*, volume 11464 of *LNCS*, pages 442–462. Springer, Heidelberg, June 2019. doi: 10.1007/978-3-030-21568-2_22.
- [BCP⁺23] Hugo Beguinet, Céline Chevalier, David Pointcheval, Thomas Ricosset, and Mélissa Rossi. GeT a CAKE: Generic transformations from key encapsulation mechanisms to password authenticated key exchanges. In Mehdi Tibouchi and Xiaofeng Wang, editors, *ACNS 23, Part II*, volume 13906 of *LNCS*, pages 516–538. Springer, Heidelberg, June 2023. doi: 10.1007/978-3-031-33491-7_19.
- [BGHJ24] Manuel Barbosa, Kai Gellert, Julia Hesse, and Stanislaw Jarecki. Bare pake: Universally composable key exchange from just passwords. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology – CRYPTO 2024*, pages 183–217, Cham, 2024. Springer Nature Switzerland.
- [BH23] Nina Bindel and Britta Hale. A note on hybrid signature schemes. *Cryptology ePrint Archive*, Paper 2023/423, 2023. <https://eprint.iacr.org/2023/423>. URL: <https://eprint.iacr.org/2023/423>.
- [BHKL13] Daniel J. Bernstein, Mike Hamburg, Anna Krasnova, and Tanja Lange. Elligator: elliptic-curve points indistinguishable from uniform random strings. In Ahmad-Reza Sadeghi, Virgil D. Gligor,

- and Moti Yung, editors, *ACM CCS 2013*, pages 967–980. ACM Press, November 2013. doi:10.1145/2508859.2516734.
- [BM92] S.M. Bellare and M. Merritt. Encrypted key exchange: password-based protocols secure against dictionary attacks. In *Proceedings 1992 IEEE Computer Society Symposium on Research in Security and Privacy*, pages 72–84, 1992. doi:10.1109/RISP.1992.213269.
- [Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145. IEEE Computer Society Press, October 2001. doi:10.1109/SFCS.2001.959888.
- [CHK⁺05] Ran Canetti, Shai Halevi, Jonathan Katz, Yehuda Lindell, and Philip D. MacKenzie. Universally composable password-based key exchange. In Ronald Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 404–421. Springer, Heidelberg, May 2005. doi:10.1007/11426639_24.
- [CR03] Ran Canetti and Tal Rabin. Universal composition with joint state. In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 265–281. Springer, Heidelberg, August 2003. doi:10.1007/978-3-540-45146-4_16.
- [DAL⁺17] Jintai Ding, Saed Alsayigh, Jean Lancrenon, Saraswathy RV, and Michael Snook. Provably secure password authenticated key exchange based on RLWE for the post-quantum world. In Helena Handschuh, editor, *CT-RSA 2017*, volume 10159 of *LNCS*, pages 183–204. Springer, Heidelberg, February 2017. doi:10.1007/978-3-319-52153-4_11.
- [Fac23] Facebook. The Labyrinth Encrypted Message Storage Protocol. December 2023. URL: https://engineering.fb.com/wp-content/uploads/2023/12/TheLabyrinthEncryptedMessageStorageProtocol_12-6-2023.pdf.
- [fan24] Deployments of fancy cryptography, September 2024. URL: <https://github.com/fancy-cryptography/fancy-cryptography>.
- [Fre22] French Cybersecurity Agency. ANSSI views on the post-quantum cryptography transition, March 2022. https://cyber.gouv.fr/sites/default/files/document/EN_Position.pdf.
- [GdKQ⁺24] Phillip Gajland, Bor de Kock, Miguel Quaresma, Giulio Malavolta, and Peter Schwabe. SWOOSH: Efficient Lattice-Based Non-Interactive key exchange. In *33rd USENIX Security Symposium (USENIX Security 24)*, pages 487–504, Philadelphia, PA, August 2024. USENIX Association. URL: <https://www.usenix.org/conference/usenixsecurity24/presentation/gajland>.

- [Ger24] German Federal Office for Information Security. BSI TR-02102-1, cryptographic mechanisms: Recommendations and key lengths, February 2024. Version 2024-01, <https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Publications/TechGuidelines/TG02102/BSI-TR-02102-1.pdf>.
- [GHJ⁺22] Qian Guo, Clemens Hlauschek, Thomas Johansson, Norman Lahr, Alexander Nilsson, and Robin Leander Schröder. Don't reject this: Key-recovery timing attacks due to rejection-sampling in HQC and BIKE. *IACR TCHES*, 2022(3):223–263, 2022. doi:10.46586/tches.v2022.i3.223-263.
- [GJN20] Qian Guo, Thomas Johansson, and Alexander Nilsson. A key-recovery timing attack on post-quantum primitives using the Fujisaki-Okamoto transformation and its application on FrodoKEM. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 359–386. Springer, Heidelberg, August 2020. doi:10.1007/978-3-030-56880-1_13.
- [GMR06] Craig Gentry, Philip MacKenzie, and Zulfikar Ramzan. A method for making password-based key exchange resilient to server compromise. In Cynthia Dwork, editor, *CRYPTO 2006*, volume 4117 of *LNCS*, pages 142–159. Springer, Heidelberg, August 2006. doi:10.1007/11818175_9.
- [Gon93] Li Gong. Lower bounds on messages and rounds for network authentication protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *ACM CCS 93*, pages 26–37. ACM Press, November 1993. doi:10.1145/168588.168592.
- [Hes20] Julia Hesse. Separating symmetric and asymmetric password-authenticated key exchange. In Clemente Galdi and Vladimir Kolesnikov, editors, *SCN 20*, volume 12238 of *LNCS*, pages 579–599. Springer, Heidelberg, September 2020. doi:10.1007/978-3-030-57990-6_29.
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.
- [HL19] Björn Haase and Benoît Labrique. AuCPace: Efficient verifier-based PAKE protocol tailored for the IIoT. *IACR TCHES*, 2019(2):1–48, 2019. <https://tches.iacr.org/index.php/TCHES/article/view/7384>. doi:10.13154/tches.v2019.i2.1-48.
- [HSC⁺23] Senyang Huang, Rui Qi Sim, Chitchanok Chuengsatiansup, Qian Guo, and Thomas Johansson. Cache-timing attack against HQC.

- IACR TCHES*, 2023(3):136–163, 2023. doi:10.46586/tches.v2023.i3.136–163.
- [IY23] Ren Ishibashi and Kazuki Yoneyama. Compact password authenticated key exchange from group actions. In Leonie Simpson and Mir Ali Rezazadeh Bae, editors, *ACISP 23*, volume 13915 of *LNCS*, pages 220–247. Springer, Heidelberg, July 2023. doi:10.1007/978-3-031-35486-1_11.
- [JKX18] Stanislaw Jarecki, Hugo Krawczyk, and Jiayu Xu. OPAQUE: An asymmetric PAKE protocol secure against pre-computation attacks. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 456–486. Springer, Heidelberg, April / May 2018. doi:10.1007/978-3-319-78372-7_15.
- [JRX24] Jake Januzelli, Lawrence Roy, and Jiayu Xu. Under what conditions is encrypted key exchange actually secure? *Cryptology ePrint Archive*, Paper 2024/324, 2024. URL: <https://eprint.iacr.org/2024/324>.
- [KR24] Jonathan Katz and Michael Rosenberg. Latke: A framework for constructing identity-binding pakes. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology – CRYPTO 2024*, pages 218–250, Cham, 2024. Springer Nature Switzerland.
- [KV11] Jonathan Katz and Vinod Vaikuntanathan. Round-optimal password-based authenticated key exchange. In Yuval Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 293–310. Springer, Heidelberg, March 2011. doi:10.1007/978-3-642-19571-6_18.
- [Lad23] Watson Ladd. SPAKE2, a Password-Authenticated Key Exchange. Request for Comments RFC 9382, Internet Engineering Task Force, September 2023. Num Pages: 17. URL: <https://datatracker.ietf.org/doc/rfc9382>, doi:10.17487/RFC9382.
- [LHT16] A. Langley, M. Hamburg, and S. Turner. Elliptic Curves for Security. RFC 7748 (Informational), January 2016. URL: <https://www.rfc-editor.org/rfc/rfc7748.txt>, doi:10.17487/RFC7748.
- [LL24] You Lyu and Shengli Liu. Hybrid password authentication key exchange in the UC framework. *Cryptology ePrint Archive*, Paper 2024/1630, 2024. URL: <https://eprint.iacr.org/2024/1630>.
- [LLH24] You Lyu, Shengli Liu, and Shuai Han. Universal composable password authenticated key exchange for the post-quantum world. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology – EUROCRYPT 2024*, pages 120–150, Cham, 2024. Springer Nature Switzerland.

- [LLH25] You Lyu, Shengli Liu, and Shuai Han. Efficient asymmetric pake compiler from kem and ae. In Kai-Min Chung and Yu Sasaki, editors, *Advances in Cryptology – ASIACRYPT 2024*, pages 34–65, Singapore, 2025. Springer Nature Singapore.
- [LP11] Richard Lindner and Chris Peikert. Better key sizes (and attacks) for LWE-based encryption. In Aggelos Kiayias, editor, *CT-RSA 2011*, volume 6558 of *LNCS*, pages 319–339. Springer, Heidelberg, February 2011. doi:10.1007/978-3-642-19074-2_21.
- [MRR20] Ian McQuoid, Mike Rosulek, and Lawrence Roy. Minimal symmetric PAKE and 1-out-of-N OT from programmable-once public functions. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020*, pages 425–442. ACM Press, November 2020. doi:10.1145/3372297.3417870.
- [PW17] David Pointcheval and Guilin Wang. VTBPEKE: Verifier-based two-basis password exponential key exchange. In Ramesh Karri, Ozgur Sinanoglu, Ahmad-Reza Sadeghi, and Xun Yi, editors, *ASIACCS 17*, pages 301–312. ACM Press, April 2017.
- [PZ23] Jiaxin Pan and Runzhi Zeng. A generic construction of tightly secure password-based authenticated key exchange. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part VIII*, volume 14445 of *LNCS*, pages 143–175. Springer, Heidelberg, December 2023. doi:10.1007/978-981-99-8742-9_5.
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005. doi:10.1145/1060590.1060603.
- [RG22] Peixin Ren and Xiaozhuo Gu. Practical post-quantum password-authenticated key exchange based-on module-lattice. In Jong Hwan Park and Seung-Hyun Seo, editors, *Information Security and Cryptology – ICISC 2021*, pages 137–156, Cham, 2022. Springer International Publishing.
- [RX23] Lawrence Roy and Jiayu Xu. A universally composable PAKE with zero communication cost - (and why it shouldn't be considered UC-secure). In Alexandra Boldyreva and Vladimir Kolesnikov, editors, *PKC 2023, Part I*, volume 13940 of *LNCS*, pages 714–743. Springer, Heidelberg, May 2023. doi:10.1007/978-3-031-31368-4_25.
- [SGJ23] Bruno Freitas Dos Santos, Yanqi Gu, and Stanislaw Jarecki. Randomized half-ideal cipher on groups with applications to UC (a)PAKE. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part V*, volume 14008 of *LNCS*,

pages 128–156. Springer, Heidelberg, April 2023. doi:10.1007/978-3-031-30589-4_5.

- [Sho20] Victor Shoup. Security analysis of itSPAKE2+. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part III*, volume 12552 of *LNCS*, pages 31–60. Springer, Heidelberg, November 2020. doi:10.1007/978-3-030-64381-2_2.
- [Wes24] Bas Westerbaan. The state of the post-quantum Internet, March 2024. The Cloudflare Blog. URL: <https://blog.cloudflare.com/pq-2024>.

A Proofs

A.1 Proof of Theorem 1

We show the theorem for an intact P_1 PAKE protocol, i.e., for the protocol $\text{ParComb}[\mathcal{F}_{\text{PAKE}}^{(1)}, P_2]$. Because the UC framework follows a sequential model of execution where only one machine is active at any given point in time, we cannot execute both $\mathcal{F}_{\text{PAKE}}^{(1)}$ and P_2 in parallel. We hence analyze a protocol version where parties first receive their output from $\mathcal{F}_{\text{PAKE}}^{(1)}$ and only then send the first message in P_2 . Because the execution of P_2 does not depend on $\mathcal{F}_{\text{PAKE}}^{(1)}$, when instantiating $\mathcal{F}_{\text{PAKE}}^{(1)}$ with an interactive protocol, its messages can be sent together with the P_2 messages.

Proof. Game G_0 : The real execution. This is the real execution as in Figure 9, with a dummy adversary \mathcal{A} . W.l.o.g. we replace $P_1^{\text{sid}||1}$ by a hybrid functionality $\mathcal{F}_{\text{PAKE}}^{(1)}$.

Game G_1 : Change layout. We change the previous game as follows:

- We move the whole execution into a single machine and call it the simulator Sim .
- We add all the record-keeping of the simulator as in Figures 7 and 8.
- In between \mathcal{Z} and Sim we add one dummy party for each real party.
- In between the dummy parties and Sim , we add the ideal functionality $\mathcal{F}_{\text{PAKE}}$ as in Figure 1, but relaying passwords of honest parties to Sim and relaying outputs keys provided by Sim to the dummy parties. We call that functionality \mathcal{F} .

The changes are only syntactical since the real execution runs on same inputs and produces outputs the same way as in the previous game. We hence have

$$\Pr[G_0] = \Pr[G_1].$$

Game G_2 : Simulate the P_2 message of honest parties. In this game we change the simulation to compute P_2 messages using $(x, \tau_x) := \text{Sim}^{(2)}. \text{Start}()$, and compute the P_2 output keys as $\text{Sim}^{(2)}. \text{ComputeKey}(x, \tau_x, y, \text{pw})$ for incoming message y . Because P_2 has password hiding and trapdoor correctness, the output distribution is statistically close to the previous game. We hence have

$$|\Pr[G_2] - \Pr[G_1]| \leq \eta$$

for a function η that is negligible in λ .

Game G_3 : Functionality aligns keys in honest sessions. We change the functionality to ignore the simulator's output key for honest parties

```

On (NewSession, sid,  $\mathcal{P}, \mathcal{P}'$ ) from  $\mathcal{F}_{\text{PAKE}}$ 
  • (G10) Record (P1Session,  $\mathcal{P}, \mathcal{P}', \perp$ ) and mark it fresh
  • (G10) Send (NewSession, sid||1,  $\mathcal{P}, \mathcal{P}'$ ) to  $\mathcal{Z}$  from  $\mathcal{F}_{\text{PAKE}}^{(1)}$ 

On (NewSession, sid||1,  $\mathcal{P}, \mathcal{P}', \text{pw}$ ) from corrupt  $\mathcal{P}$  to  $\mathcal{F}_{\text{PAKE}}^{(1)}$ 
  • (G1) Record (P1Session, sid,  $\mathcal{P}, \mathcal{P}', \text{pw}$ ) and mark it fresh

On (TestPwd, sid||1,  $\mathcal{P}, \text{pw}$ ) from  $\mathcal{Z}$  to  $\mathcal{F}_{\text{PAKE}}^{(1)}$  or as internal call // w.l.o.g.  $\mathcal{P}$  is honest
  • (G6) Retrieve (P1Session, sid,  $\mathcal{P}, \cdot, \perp$ ) marked fresh
  • (G6) Send (TestPwd, sid,  $\mathcal{P}, \text{pw}$ ) to  $\mathcal{F}_{\text{PAKE}}$ 
  • (G6) Upon response “correct”:
    – (G6) Mark P1Session as compromised
    – (G6) Record (CorrectTestPwd, sid,  $\mathcal{P}, \text{pw}$ )
    – (G6) Return “correct”
  • (G6) Upon response “wrong”:
    – (G6) Mark P1Session as interrupted
    – (G6) Return “wrong”

On (NewKey, sid||1,  $\mathcal{P}, K'_1$ ) from  $\mathcal{Z}$  to  $\mathcal{F}_{\text{PAKE}}^{(1)}$ 
  • (G1) Retrieve (P1Session, sid,  $\mathcal{P}, [\mathcal{P}'], [\text{pw}]$ ) marked  $m \neq \text{completed}$ 
  • (G1) If  $m = \text{compromised}$ :  $K_1 := K'_1$ 
  • (G1) Elif  $m = \text{fresh}$  and  $\exists(\text{P1Key}, \text{sid}, \mathcal{P}', [K''_1], \text{fresh})$ : //  $\mathcal{P}$  finishes last
    – (G7) If  $\exists(\text{TestedPwd}, \text{sid}, \mathcal{P}', \text{pw}, \text{correct})$ :  $K_1 := K''_1$  //  $\mathcal{Z} \neq \perp \Rightarrow \mathcal{P}$  corrupt  $\Rightarrow$ 
      w.l.o.g,  $\mathcal{P}'$  is honest, and we already ran TestPwd when  $\mathcal{P}'$  finished.
    – (G10) Elif  $\mathcal{P}$  and  $\mathcal{P}'$  are honest:  $K_1 := \perp$  //  $\mathcal{F}_{\text{PAKE}}$  controls the final  $K$  anyway
    – (G1) Else:  $K_1 \leftarrow \{0, 1\}^\lambda$  //  $\mathcal{P}$  corrupt, wrong password
  • (G1) Elif  $\exists(\text{P1Session}, \text{sid}, \mathcal{P}', \mathcal{P}, \cdot)$  then set  $K_1 \leftarrow \{0, 1\}^\lambda$  //  $\mathcal{P}$  finishes first or interrupted, or  $\mathcal{P}'$  finished non-fresh.
  • (G1) Else ignore the query // Isolated  $\mathcal{P}$  does not output a key
  • (G7) If  $m = \text{fresh}$ :
    – (G7) If  $\mathcal{P}$  is honest and  $\mathcal{P}'$  is corrupt: // Now that TestPwd can no longer be called
      on this session, we can finally call it
      * (G7) Retrieve (P1Session, sid,  $\mathcal{P}', \mathcal{P}, [\text{pw}']$ )
      * (G7) Run code of (TestPwd, sid||1,  $\mathcal{P}, \text{pw}'$ )
  • (G1) Record (P1Key, sid,  $\mathcal{P}, K_1, m$ )
  • (G1) Mark P1Session as completed
  • (G1) If  $\mathcal{P}$  is corrupt:
    – (G1) Send (sid,  $K_1$ ) to  $\mathcal{P}$ 
    – (G1) Return
  • (G2) Compute  $x, \tau_x \leftarrow \text{Sim}^{(2)}. \text{Start}(1^\lambda)$ 
  • (G2) Record (P2State, sid,  $\mathcal{P}, x, \tau_x$ )
  • (G1) Record (Sent, sid,  $\mathcal{P}, \mathcal{P}', x$ )
  • (G1) Send ( $\mathcal{P} \rightarrow \mathcal{P}', \text{sid}, x$ ) to  $\mathcal{Z}$ 

```

Figure 7: Simulator for $\text{ParComb}[\mathcal{F}_{\text{PAKE}}^{(1)}, \text{P}_2]$ realizing $\mathcal{F}_{\text{PAKE}}$. $\text{Sim}^{(2)}$ represents the simulator for the statistically password-hiding property of PAKE P_2 . Random oracle H is simulated as in the real execution. Label (**G**₁) indicates that the line of code is activated in game **G**₁.

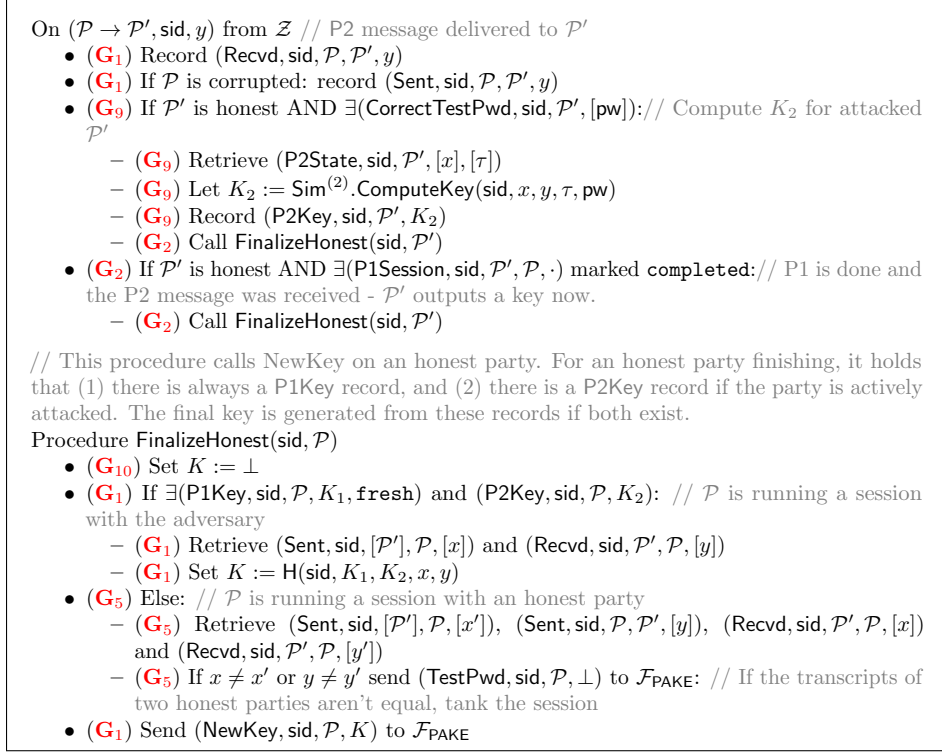


Figure 8: Simulator (cont.)

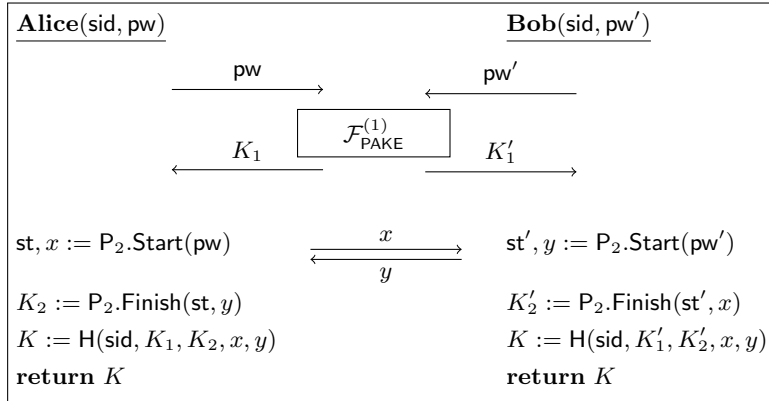


Figure 9: UC execution of the ParComb combiner with a hybrid functionality $\mathcal{F}_{\text{PAKE}}^{(1)}$ and P_2 being a statistically password-hiding one-round PAKE. When instantiating $\mathcal{F}_{\text{PAKE}}^{(1)}$, messages can be sent in parallel with x, y . The random oracle H outputs values in $\{0, 1\}^\lambda$.

finishing last in honest sessions, with matching passwords, and where both records were **fresh** at the time of output generation (i.e., no `TestPwd` was queried by the simulator on any of the two parties). In this case, the functionality also outputs the session key of the first party to the second party.

This and the previous game are not distinguishable by the (trivial) correctness of the protocol, and hence we have

$$\Pr[\mathbf{G}_2] = \Pr[\mathbf{G}_3].$$

Game \mathbf{G}_4 : Randomize output key of parties in an honest interaction.

We change the functionality \mathcal{F} to ignore the output keys from the simulation in case an honest party finishes in an unattacked session, and instead uses the output key from its internal `Key` record.

Note that the change (a) randomizes the output key of the honest user who finish first in an unattacked session, (b) randomizes the output key of the honest user who finish last in an unattacked session, and (c) aligns keys of the honest session in case of matching passwords. In this game, parties with tampered \mathcal{P}_2 transcripts still receive their output key from the simulation, which is uncorrelated to their honest partner's key (which is randomized in this game) due to the \mathcal{P}_2 transcript that is included in the hash.

Because $\mathcal{F}_{\text{PAKE}}^{(1)}$ draws K_1 uniformly at random and the adversary does not query $\mathcal{F}_{\text{PAKE}}^{(1)}$ in any way that it generates output to the adversary that depends on K_1 , \mathcal{Z} only sees a difference in this and the previous game if it queries the random oracle with $H(\text{sid}, x, y, K_1, K_2)$. We hence have

$$|\Pr[\mathbf{G}_4] - \Pr[\mathbf{G}_3]| \leq \frac{q_H}{2^\lambda}.$$

Game \mathbf{G}_5 : Randomize output key upon message tampering.

We change the simulator to send $(\text{TestPwd}, \text{sid}, \mathcal{P}, \perp)$ to \mathcal{F} (guaranteeing the session is **interrupted**) when \mathcal{Z} tampers with the \mathcal{P}_2 message sent to \mathcal{P} in an otherwise honest interaction (in particular, when there are no `TestPwd` queries against \mathcal{P} to $\mathcal{F}_{\text{PAKE}}^{(1)}$). At the same time, we let the functionality ignore the output keys from the simulator of the session of \mathcal{P} in case such a `TestPwd` happened, and instead output the key from its `Key` record.

The argument is very similar to \mathbf{G}_4 : \mathcal{Z} can only notice a difference if it queries the final hash of \mathcal{P} to the random oracle. We again have

$$|\Pr[\mathbf{G}_5] - \Pr[\mathbf{G}_4]| \leq \frac{q_H}{2^\lambda}.$$

As of this game, all **fresh** records in \mathcal{F} produce output keys.

Game G_6 : Relay TestPwd against P_1 . We change the code of the simulator upon \mathcal{Z} sending $(\text{TestPwd}, \text{sid}, \mathcal{P}, \text{pw})$ to $\mathcal{F}_{\text{PAKE}}^{(1)}$ for an honest \mathcal{P} . Sim now also sends this query to \mathcal{F} and uses that response as its own response to \mathcal{Z} . Because both functionality instances work on the same inputs and compromised or interrupted markings in \mathcal{F} do not yet affect the outputs, the changes are only syntactical and we have

$$\Pr[G_5] = \Pr[G_6].$$

Game G_7 : Extract a password from a corrupt party. We change the simulator to send $(\text{TestPwd}, \text{sid}, \mathcal{P}, \text{pw})$ to \mathcal{F} upon $(\text{NewKey}, \text{sid}, \mathcal{P}, \cdot)$ from \mathcal{Z} for an honest \mathcal{P} that is in a session with a corrupt \mathcal{P}' that had previously send input $(\text{NewSession}, \text{sid}, \mathcal{P}', \mathcal{P}, \text{pw})$ to $\mathcal{F}_{\text{PAKE}}^{(1)}$. The simulator then sends the output key of \mathcal{P} to \mathcal{F} via NewKey , and \mathcal{F} uses that key for compromised sessions.

While in G_6 \mathcal{F} just outputted all keys from the simulation, in this game it still does the same but receives the ones for compromised sessions through the simulator's NewKey queries. Hence, the changes are only syntactical and we have

$$\Pr[G_6] = \Pr[G_7].$$

Game G_8 : Randomize output keys of attacked users with wrong password guess. We change the functionality to ignore the simulator's key for parties who finish on **interrupted** sessions, and instead use the output key from the Key record.

Again, the argument is very similar to the one in G_4 : in G_7 , upon a wrong password guess, $\mathcal{F}_{\text{PAKE}}^{(1)}$ issues a random K_1 to \mathcal{P} of which \mathcal{Z} only sees the output $H(\text{sid}, K_1, K_2, x, y)$ of \mathcal{P} . We can now randomize this output unnoticed by \mathcal{Z} except with the negligible probability that \mathcal{Z} guesses K_1 , i.e.,

$$|\Pr[G_8] - \Pr[G_7]| \leq \frac{q_H}{2^\lambda}.$$

As of this game, all **interrupted** records in \mathcal{F} produce output keys.

Game G_9 : Simulate compromised parties using TestPwd guesses. We change the simulator upon sending $(\text{TestPwd}, \text{sid}, \mathcal{P}, \text{pw})$ to \mathcal{F} for an honest \mathcal{P} , as detailed in G_6 and G_7 . If \mathcal{F} replies with "correct", we now let the simulator use pw in the computation of $\text{Sim}^{(2)}. \text{ComputeKey}$.

The change is only syntactical because the passwords used in ComputeKey are identical in this and the previous game. We have

$$\Pr[G_8] = \Pr[G_9].$$

Game \mathbf{G}_{10} : Remove the passwords from the simulation. We modify the simulator to use \perp as password in all `NewSession` inputs to the internally simulated $\mathcal{F}_{\text{PAKE}}^{(1)}$ for honest parties. At the same time, we change the functionality not to forward input passwords of honest parties to the simulator, but relay (`NewSession`, `sid`, \mathcal{P} , \mathcal{P}') from \mathcal{F} to \mathcal{Z} as coming from $\mathcal{F}_{\text{PAKE}}^{(1)}$.

We need to argue that outputs of $\mathcal{F}_{\text{PAKE}}^{(1)}$ that depend on input passwords of honest parties, i.e., any K_1 that was output to an honest party and chosen by $\mathcal{F}_{\text{PAKE}}^{(1)}$, are not relevant to the execution anymore. This concerns **fresh** and **interrupted** records (**compromised** records get an adversarially chosen key instead of a secure one picked by $\mathcal{F}_{\text{PAKE}}^{(1)}$). This is immediate to see in \mathbf{G}_9 : \mathcal{F} computes all output keys as of \mathbf{G}_9 , and the \mathbf{P}_2 execution does not depend on K_1 .

Because forwarding of input passwords and relaying of simulator output keys was the only difference when we introduced \mathcal{F} in \mathbf{G}_1 , the functionality \mathcal{F} in \mathbf{G}_{10} is equal to $\mathcal{F}_{\text{PAKE}}$ and we have

$$\Pr[\mathbf{G}_9] = \Pr[\mathbf{G}_{10}].$$

The theorem thus follows with the simulator depicted in Figures 7 and 8. \square

A.2 Proof of Theorem 2

The proof of this theorem is split into two cases, namely analyzing \mathbf{P} relying on the UC security of \mathbf{P}_1 , and \mathbf{P} relying on the UC security of \mathbf{P}_2 . The corresponding protocol relying on a secure \mathbf{P}_1 is depicted in Figure 13. The protocol relying on a secure \mathbf{P}_1 is depicted in Figure 10, relying on a \mathbf{P}_2 with arbitrary communication pattern.

Theorem 2 then follows from Lemmas 2 and 3. Since the proofs each require the transcript of either the first or the second PAKE to be included in the final hash, for the combiner to be secure, the full protocol transcript needs to be included in the final hash.

Lemma 2. *Let \mathbf{P}_1 be a one-round statistically password-hiding PAKE protocol. Then protocol $\text{SeqComb}[\mathbf{P}_1, \mathcal{F}_{\text{PAKE}}^{(2)}]$ (Figure 10) UC-realizes $\mathcal{F}_{\text{PAKE}}$ in the $(\mathcal{F}_{\text{RO}}, \mathcal{F}_{\text{PAKE}})$ -hybrid model, where $\mathbf{H}_0, \mathbf{H}_1$ are modeled as random oracles (i.e., as calls to \mathcal{F}_{RO}), and $\mathcal{F}_{\text{PAKE}}^{(2)} = \mathcal{F}_{\text{PAKE}}$ (i.e., the superscript is only added to differentiate between the two instances of $\mathcal{F}_{\text{PAKE}}$ in the statement).*

For proving the lemma we make the following simplifying assumptions.

1. \mathcal{Z} corrupts at most one party in a session
2. \mathcal{Z} never issues a `TestPwd` query against a corrupt party

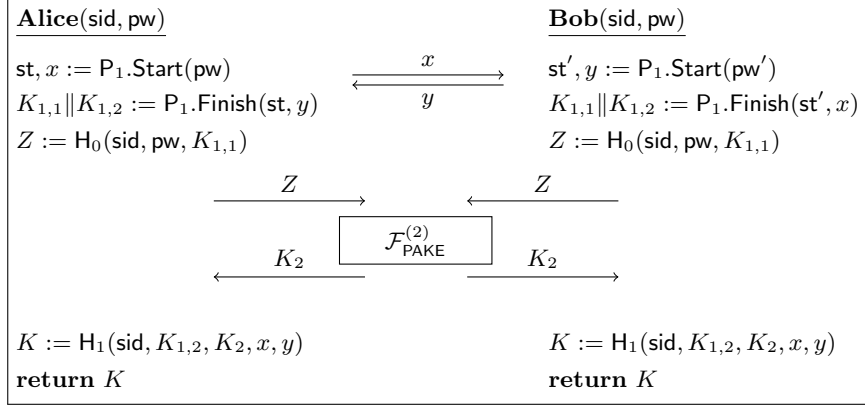


Figure 10: Protocol $\text{SeqComb}[P_1, \mathcal{F}_{\text{PAKE}}^{(2)}]$ relying on an intact second PAKE protocol.

3. \mathcal{Z} never issues a `TestPwd` query against an honest party if the other party is corrupt

These simplifications are without loss of generality, i.e., our proof still implies that there is no distinguishing environment \mathcal{Z} . This is because (1) in the fully corrupt setting the simulator knows all secrets and hence the simulation is straightforward, (2) a `TestPwd` query against a corrupt party compares two passwords that were both given by \mathcal{Z} , and (3) the password guess can be issued via a corrupt `NewSession` instead of `TestPwd`.

Game G_0 : The real execution. This is the real execution as in Figure 10, with a dummy adversary \mathcal{A} and a hybrid functionality $\mathcal{F}_{\text{PAKE}}^{(2)}$ in place of P_2 .

Game G_1 : Change layout. We change the previous game as follows:

- We move the whole execution into a single machine and call it the simulator `Sim`.
- We add all the record-keeping of the simulator as in Figures 11 and 12.
- In between \mathcal{Z} and `Sim` we add one dummy party for each real party.
- In between the dummy parties and `Sim`, we add the ideal functionality $\mathcal{F}_{\text{PAKE}}$ as in Figure 1, but relaying passwords of honest parties to `Sim` and relaying outputs keys provided by `Sim` to the dummy parties. We call that functionality \mathcal{F} .

The changes are only syntactical since the real execution runs on same inputs and produces outputs the same way as in the previous game. We hence have

$$\Pr[G_0] = \Pr[G_1].$$

On (NewSession, sid, \mathcal{P} , \mathcal{P}') from $\mathcal{F}_{\text{PAKE}}$ // Simulate the P1 message

- (**G₂**) Compute $x, \tau_x \leftarrow \text{Sim}^{(1)}$
- (**G₁**) Record (P1State, sid, \mathcal{P} , x, τ_x)
- (**G₁**) Record (Sent, sid, \mathcal{P} , \mathcal{P}' , x)
- (**G₁**) Send ($\mathcal{P} \rightarrow \mathcal{P}'$, sid, x) to \mathcal{Z}

On ($\mathcal{P} \rightarrow \mathcal{P}'$, sid, y) from \mathcal{Z} // P1 message delivery

- (**G₁**) If $\exists(\text{Recvd}, \text{sid}, \mathcal{P}, \mathcal{P}', \cdot)$: ignore this query
- (**G₁**) Record (Recvd, sid, \mathcal{P} , \mathcal{P}' , y)
- If \mathcal{P} is corrupted: // Record the sent message because it's adversarially generated
 - (**G₁**) Record (Sent, sid, \mathcal{P} , \mathcal{P}' , y)
- If \mathcal{P}' is honest:// If the receiver is honest, use this as the time to start its P2 session
 - (**G₈**) Record (P2Session, sid, \mathcal{P} , \mathcal{P}' , \perp) marked **fresh**
 - (**G₁**) Send (NewSession, sid||2, \mathcal{P} , \mathcal{P}') to \mathcal{Z}

On (NewSession, sid||2, \mathcal{P}' , Z) from corrupt \mathcal{P} to $\mathcal{F}_{\text{PAKE}}^{(2)}$

- (**G₁**) Record (P2Session, sid, \mathcal{P} , \mathcal{P}' , Z)
- (**G₁**) Send (NewSession, sid||2, \mathcal{P} , \mathcal{P}') to \mathcal{Z} //Wait with TestPwd until other party received input

On (TestPwd, sid||2, \mathcal{P} , Z) from \mathcal{Z} to $\mathcal{F}_{\text{PAKE}}^{(2)}$ or as internal call // w.l.o.g., \mathcal{P} is honest

- (**G₁**) Retrieve (P2Session, sid, \mathcal{P} , \cdot , \perp) marked **fresh**
- // Z is the correct password iff it is what \mathcal{P} would have computed. This is the case iff both the pw and $K_{1,1}$ used in the computation of Z are correct. So we test those.
- (**G₄**) Retrieve (H_0 , sid, [pw], [K_1^*], Z) or GOTO **wrong**
- (**G₄**) Retrieve (P1State, sid, \mathcal{P} , [x], [τ_x]), (Recvd, sid, \cdot , \mathcal{P} , [y])
- (**G₄**) Let $[K_{1,1}]||[K_{1,2}] := \text{Sim}^{(1)}. \text{ComputeKey}(\text{sid}, x, y, \tau_x, \text{pw})$
- (**G₄**) Record (P1Key, \mathcal{P} , $K_{1,1}$, $K_{1,2}$)
- (**G₄**) If $K_1^* \neq K_{1,1}$ GOTO **wrong** // Early-fail if the P1 key disagrees
- (**G₄**) Send (TestPwd, sid||2, \mathcal{P} , pw) to $\mathcal{F}_{\text{PAKE}}$
- (**G₁**) If “correct”: GOTO **correct**
- (**G₁**) If “wrong”: GOTO **wrong**
- (**G₁**) Label **wrong**:
 - (**G₁**) Mark retrieved P2Session **interrupted**
 - (**G₁**) Record (TestedPwd, sid, \mathcal{P} , pw, **wrong**)
 - (**G₁**) Return “wrong”
- (**G₁**) Label **correct**:
 - (**G₁**) Mark retrieved P2Session **compromised**
 - (**G₁**) Record (TestedPwd, sid, \mathcal{P} , pw, **correct**)
 - (**G₁**) Return “correct”

Figure 11: Simulator for $\text{SeqComb}[\text{P}_1, \mathcal{F}_{\text{PAKE}}^{(2)}]$ realizing $\mathcal{F}_{\text{PAKE}}$, Lemma 2. Random oracles H_0, H_1 are simulated as in the real execution. Label (**G₄**) indicates that the line of code is activated in game **G₄**.

```

On (NewKey, sid||2,  $\mathcal{P}$ ,  $K_2^*$ ) from  $\mathcal{Z}$  to  $\mathcal{F}_{\text{PAKE}}^{(2)}$ 
• (G1) Retrieve (P2Session, sid,  $\mathcal{P}$ ,  $\mathcal{P}'$ , [ $Z$ ]) with marking  $m \neq \text{completed}$ 
• // Pick  $K_2$  the same way a normal PAKE functionality would pick the key
• (G1) If  $m = \text{compromised}$ :  $K_2 := K_2^*$ 
• (G1) Else if  $m = \text{fresh}$  AND  $\exists(\text{P2Key}, \text{sid}, \mathcal{P}', [K_2'], \text{fresh})$ : //  $\mathcal{P}$  finishes last
  - (G5) If  $\exists(\text{TestedPwd}, \text{sid}, \mathcal{P}', Z, \text{correct})$ :  $K_2 := K_2'$  //  $Z \neq \perp \Rightarrow \mathcal{P}$  corrupt  $\Rightarrow$ 
    w.l.o.g,  $\mathcal{P}'$  is honest, and we already ran TestPwd when  $\mathcal{P}'$  finished.
  - (G8) Elif  $\mathcal{P}$  and  $\mathcal{P}'$  are honest:  $K_2 := \perp$  //  $\mathcal{F}_{\text{PAKE}}$  controls the final  $K$  anyway
  - (G1) Else:  $K_2 \leftarrow \{0, 1\}^\lambda$  //  $\mathcal{P}$  corrupt, wrong password
• (G1) Elif  $\exists(\text{P2Session}, \text{sid}, \mathcal{P}', \cdot)$  then set  $K_2 \leftarrow \{0, 1\}^\lambda$  //  $\mathcal{P}$  finishes first or interrupted, or  $\mathcal{P}'$  finished non-fresh.
• (G1) Else ignore the query // Isolated  $\mathcal{P}$  does not output a key
• // Same as PAKE functionality, record the key and return it to the party (if corrupt)
• (G1) Record (P2Key, sid,  $\mathcal{P}$ ,  $K_2$ ,  $m$ )
• (G1) Mark P2Session completed
• (G1) If  $\mathcal{P}$  is corrupt:
  - (G1) Send (sid,  $K_2$ ) to  $\mathcal{P}$  // No need to call  $\mathcal{F}_{\text{PAKE}}$ .NewKey on a corrupt party
• // Now we compute the final key  $H_2(\text{sid}, K_{1,2}, K_2, \text{tr})$  for honest  $\mathcal{P}$  if we can find  $K_1$ .
• (G4) If  $\mathcal{P}'$  is corrupt AND P2Session is fresh: // Attacked  $\mathcal{P}$  finishes. Opportunistically run TestPwd in the honest-corrupt setting.
  - (G4) Retrieve (P2Session, sid,  $\mathcal{P}'$ ,  $\mathcal{P}$ , [ $Z'$ ])
  - (G4) Run code of (TestPwd, sid||2,  $\mathcal{P}$ ,  $Z'$ ) // Ensure P1Key is set in attacked session
• (G4) If  $\mathcal{P}'$  is corrupt OR  $\exists(\text{TestedPwd}, \text{sid}, \mathcal{P}, \cdot, \text{correct})$ :
  - (G4) Retrieve (P1Key, sid,  $\mathcal{P}$ ,  $\cdot$ , [ $K_{1,1}, K_{1,2}$ ]) // Must exist in a non-tanked session because, above, Sim always calls TestPwd on an honest-corrupt setting
  - (G4)  $K := H_2(\text{sid}, K_{1,2}, K_2, \text{tr})$ 
• Else:
  - (G8)  $K := \perp$  // Honest-honest setting:  $\mathcal{F}_{\text{PAKE}}$  decides
  - (G3) Retrieve (Sent, sid, [ $\mathcal{P}'$ ],  $\mathcal{P}$ , [ $x'$ ]), (Sent, sid,  $\mathcal{P}$ ,  $\mathcal{P}'$ , [ $y$ ]), (Recv, sid,  $\mathcal{P}'$ ,  $\mathcal{P}$ , [ $x$ ]) and (Recv, sid,  $\mathcal{P}'$ ,  $\mathcal{P}$ , [ $y'$ ])
  - (G3) If  $x \neq x'$  or  $y \neq y'$  send (TestPwd, sid,  $\mathcal{P}$ ,  $\perp$ ) to  $\mathcal{F}_{\text{PAKE}}$ : // DoS attack on  $\mathcal{P}_1$ 
• Send (NewKey, sid,  $\mathcal{P}$ ,  $K$ ) to  $\mathcal{F}_{\text{PAKE}}$ 

```

Figure 12: Cont. simulator.

Game \mathbf{G}_2 : Simulate the \mathcal{P}_1 message of honest parties. In this game we change the simulation to compute \mathcal{P}_1 messages using $(x, \tau_x) := \text{Sim}^{(1)}. \text{Start}()$, and compute the \mathcal{P}_1 output keys as $\text{Sim}^{(1)}. \text{ComputeKey}(x, \tau_x, y, \text{pw})$ for incoming message y . Because \mathcal{P}_1 has password hiding and trapdoor correctness, the output distribution is statistically close to the previous game. We hence have

$$|\Pr[\mathbf{G}_2] - \Pr[\mathbf{G}_1]| \leq \eta$$

for a function η that is negligible in λ .

Game \mathbf{G}_3 : Randomize outputs keys in case of DoS attack in \mathcal{P}_1 . We modify the simulator to send (TestPwd, sid, \mathcal{P} , \perp) for an honest party \mathcal{P} in case \mathcal{P}' is also honest, and there was no TestPwd query against \mathcal{P} by \mathcal{Z} , and \mathcal{Z} did not deliver both x, y untampered. This can only be decided

once \mathcal{Z} sends `NewKey` for \mathcal{P} to $\mathcal{F}_{\text{PAKE}}^{(2)}$, and hence the simulator performs this check at the very end of the `NewKey` interface.

Because `interrupted` markings in \mathcal{F} do not yet affect output keys, the changes do not affect the output distribution and we have

$$\Pr[\mathbf{G}_2] = \Pr[\mathbf{G}_3].$$

Game \mathbf{G}_4 : Extract password from corrupt party. We change the simulator for an honest \mathcal{P} , either upon `(TestPwd, sid, \mathcal{P} , Z)` from \mathcal{Z} where w.l.o.g. \mathcal{P}' is honest, or upon a corrupt \mathcal{P}' receiving input `(NewSession, sid, \mathcal{P} , Z)`, where $Z = \text{H}_0(\text{sid}, [\text{pw}], [K_{1,1}^*])$. The simulator computes $[K_{1,1}] \parallel [K_{1,2}] := \text{Sim}^{(1)}. \text{ComputeKey}(\text{sid}, x, y, \tau_x, \text{pw})$ and if $K_{1,1} = K_{1,1}^*$ sends `(TestPwd, sid, \mathcal{P} , pw)` to $\mathcal{F}_{\text{PAKE}}$. In all other cases, the simulator keeps letting the internally emulated $\mathcal{F}_{\text{PAKE}}^{(2)}$ instance handle the guess. When \mathcal{P} is supposed to generate output, if the guess came back “correct”, the simulator sets the key of \mathcal{P} to be $\text{H}_2(\text{sid}, K_{1,2}, K_2^*, \text{tr})$, where K_2^* is taken from the `NewKey` query from \mathcal{Z} to $\mathcal{F}_{\text{PAKE}}^{(2)}$ for \mathcal{P} . The simulator sends this key through `NewKey` to \mathcal{F} , which outputs it for `compromised` records.

The changes of key computation of \mathcal{P} in this game are only syntactical because the correct `TestPwd` guess guarantees that `pw` is the input password of \mathcal{P} . We hence have

$$\Pr[\mathbf{G}_3] = \Pr[\mathbf{G}_4].$$

Game \mathbf{G}_5 : Removing two password equality checks from the simulation. We change the functionality to ignore the simulator’s output key for honest parties finishing last in honest sessions, with matching passwords, and where both records were `fresh` at the time of output generation (i.e., no `TestPwd` was queried by the simulator on any of the two parties). In this case, the functionality outputs the output key of the first party also to the second party.

At the same time we make a similar change in the simulation of $\mathcal{F}_{\text{PAKE}}^{(2)}$, for a corrupt party finishing last in a session with an honest party, where simulator’s `TestPwd` query to $\mathcal{F}_{\text{PAKE}}$ returned “correct” for that honest party, for Z input by the corrupt party. Instead of comparing the honest party’s password with the corrupt party’s password, we let the simulator output the key K_2^* that was previously output to the corrupt party.

This and the previous game are not distinguishable by the (trivial) correctness of the protocol and the correctness of $\mathcal{F}_{\text{PAKE}}$. We note that the correctness of $\mathcal{F}_{\text{PAKE}}$ can be used as an argument only in case of no adversarial interference, as is the case in this game, but does not hold in general as demonstrated by Roy and Xu [RX23]. Hence we have

$$\Pr[\mathbf{G}_4] = \Pr[\mathbf{G}_5].$$

Game \mathbf{G}_6 : Randomize output keys of interrupted sessions. We change the functionality to ignore the simulator's output keys for **interrupted** sessions and instead output the key stored in the **Key** record, i.e., a randomly chosen key.

In this game, **interrupted** sessions got marked as such due to either an incorrect password guess, or \mathcal{Z} tampering with the transcript between two honest users. \mathcal{Z} can only notice a difference if it queries what the honest user with the interrupted session outputs, namely $H_1(\text{sid}, K_{1,2}, K_2, x, y)$. In case of message tampering without an active attack on $\mathcal{F}_{\text{PAKE}}^{(2)}$, this happens at most with probability $q_H/2^\lambda$ because $\mathcal{F}_{\text{PAKE}}^{(2)}$ chooses K_2 uniformly at random and does not leak any information about it since the other party is also honest and outputs a hash $H_1(\text{sid}, K_{1,2}, K_2, x', y')$ where $(x, y) \neq (x', y')$. In case of an incorrect password guess, $\mathcal{F}_{\text{PAKE}}^{(2)}$ chooses a uniformly random K_2 for the honest party with the interrupted session and outputs it only to that party. Hence, we have

$$|\Pr[\mathbf{G}_6] - \Pr[\mathbf{G}_5]| \leq \frac{q_H}{2^\lambda}.$$

Game \mathbf{G}_7 : Randomize keys of an honest party finishing first with a fresh record. We change the functionality to ignore the simulator's output keys for an honest \mathcal{P} finishing on a **fresh** session, and instead output to \mathcal{P} the key from the **Key** record, i.e., a randomly chosen key.

Note that, as of \mathbf{G}_5 , the functionality aligns the output key in case of matching passwords, which implies that the randomized key chosen of \mathcal{P} in this game might get repeated to the other party.

Because the simulation of \mathbf{G}_7 always asks **TestPwd** queries if an honest \mathcal{P} either runs with a corrupt party, or with an honest party but is attacked by either a message tampering adversary or an active adversary through **TestPwd** on $\mathcal{F}_{\text{PAKE}}^{(2)}$, \mathcal{P} finishing on a fresh record is not under any attack. In \mathbf{G}_7 , $\mathcal{F}_{\text{PAKE}}^{(2)}$ chooses a fresh K_2 for \mathcal{P} that is only given to the honest \mathcal{P}' (if passwords match; otherwise \mathcal{P}' receives no information about K_2). We hence have

$$|\Pr[\mathbf{G}_7] - \Pr[\mathbf{G}_6]| \leq \frac{q_H}{2^\lambda}.$$

Game \mathbf{G}_8 : Remove the passwords from the simulation. We modify the simulator to use \perp as password in all **NewSession** inputs to the internally simulated $\mathcal{F}_{\text{PAKE}}^{(1)}$ for honest parties. We let it omit **ComputeKey** runs on the input passwords of honest parties. At the same time, we change the functionality not to forward input passwords of honest parties to the simulator, but relay $(\text{NewSession}, \text{sid}, \mathcal{P}, \mathcal{P}')$ from \mathcal{F} to \mathcal{Z} as coming from $\mathcal{F}_{\text{PAKE}}^{(1)}$.

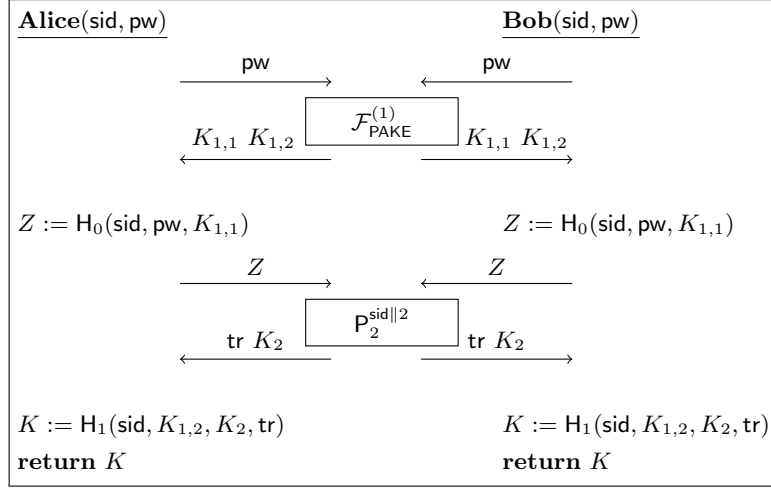


Figure 13: Protocol $\text{SeqComb}[\mathcal{F}_{\text{PAKE}}^{(1)}, \mathcal{P}_2]$ relying on an intact first PAKE protocol.

As of \mathbf{G}_7 , outputs are fully determined by \mathcal{F} except for compromised records where \mathcal{F} forwards the key from the simulator. A careful inspection of the simulator code of \mathbf{G}_8 shows that it does not use honest parties' input passwords anymore: the protocol transcript is generated without passwords of honest parties as of \mathbf{G}_2 , with `ComputeKey` taking extracted passwords as input as of \mathbf{G}_4 . The internal simulation of $\mathcal{F}_{\text{PAKE}}^{(2)}$ also works without passwords since we replace password equality checks in \mathbf{G}_5 . We hence have

$$\Pr[\mathbf{G}_7] = \Pr[\mathbf{G}_8].$$

Because forwarding of input passwords and relaying of simulator output keys was the only difference when we introduced \mathcal{F} in \mathbf{G}_1 , the functionality \mathcal{F} in \mathbf{G}_8 is equal to $\mathcal{F}_{\text{PAKE}}$. The theorem thus follows with the simulator depicted in Figures 11 and 12.

Lemma 3. *Let \mathcal{P}_2 be any PSK equality hiding PAKE protocol. Then protocol $\text{SeqComb}[\mathcal{F}_{\text{PAKE}}^{(1)}, \mathcal{P}_2]$ (Figure 13) UC-realizes $\mathcal{F}_{\text{lePAKE}}$ in the $(\mathcal{F}_{\text{RO}}, \mathcal{F}_{\text{PAKE}})$ -hybrid model with respect to static party corruptions, where H_0, H_1 are modeled as random oracles (i.e., as calls to \mathcal{F}_{RO}), and $\mathcal{F}_{\text{PAKE}}^{(1)} = \mathcal{F}_{\text{PAKE}}$ (i.e., the superscript is only added to differentiate between the two instances of $\mathcal{F}_{\text{PAKE}}$ in the statement).*

We again make the following simplifying assumptions.

1. \mathcal{Z} corrupts at most one party in a session
2. \mathcal{Z} never issues a `TestPwd` query against a corrupt party

```

On (NewSession, sid,  $\mathcal{P}$ ,  $\mathcal{P}'$ ) from  $\mathcal{F}_{\text{PAKE}}$ 
  • Record (P1Session, sid,  $\mathcal{P}$ ,  $\mathcal{P}'$ ,  $\perp$ ) and mark it fresh
  • (G1) Send (NewSession, sid||1,  $\mathcal{P}$ ,  $\mathcal{P}'$ ) to  $\mathcal{Z}$ 

On (NewSession, sid||1,  $\mathcal{P}$ ,  $\mathcal{P}'$ , pw) from corrupt  $\mathcal{P}$  to  $\mathcal{F}_{\text{PAKE}}^{(1)}$ 
  • (G1) Record (P1Session, sid,  $\mathcal{P}$ ,  $\mathcal{P}'$ , pw) and mark it fresh
  • (G1) Send (NewSession, sid||1,  $\mathcal{P}$ ,  $\mathcal{P}'$ ) to  $\mathcal{Z}$ 

On (NewKey, sid,  $\mathcal{P}$ ,  $K^*$ ) from  $\mathcal{Z}$  to  $\mathcal{F}_{\text{PAKE}}^{(1)}$ 
// Follow the logic of the  $\mathcal{F}_{\text{PAKE}}$  NewKey interface to compute the output key of  $\mathcal{P}_2$ :
  • (G1) Retrieve (P1Session, sid,  $\mathcal{P}$ , [ $\mathcal{P}'$ ], [pw]) with mark  $m \neq$  completed
  • (G1) If  $m =$  compromised:  $K_{1,1}, K_{1,2} := K^*$ 
  • (G8) Elif  $m =$  fresh AND  $\exists(\text{P1Key}, \text{sid}, \mathcal{P}', [K'], \text{fresh})$ : //  $\mathcal{P}$  finishes last
    - If  $\exists(\text{TestedPwd}, \text{sid}, \mathcal{P}', \text{pw}, \text{correct})$ :  $K_{1,1}, K_{1,2} := K'$  //  $Z \neq \perp \Rightarrow \mathcal{P}$  corrupt  $\Rightarrow$  w.l.o.g,  $\mathcal{P}'$ 
      is honest, and we already ran TestPwd when  $\mathcal{P}'$  finished.
    - (G1) Elif  $\mathcal{P}$  and  $\mathcal{P}'$  are honest:  $K_{1,1}, K_{1,2} := \perp$  //  $\mathcal{F}_{\text{PAKE}}$  controls the final  $K$  anyway
    - (G1) Else:  $K_{1,1}, K_{1,2} \leftarrow \{0, 1\}^\lambda$  //  $\mathcal{P}$  corrupt, wrong password
  • (G1) Elif  $\exists(\text{P1Session}, \text{sid}, \mathcal{P}', \cdot)$  then set  $K_{1,1}, K_{1,2} \leftarrow \{0, 1\}^\lambda$  //  $\mathcal{P}$  finishes first or interrupted,
    or  $\mathcal{P}'$  finished non-fresh.
  • (G1) Else ignore the query // Isolated  $\mathcal{P}$  does not output a key
  • (G2) If P1Session is fresh,  $\mathcal{P}$  is honest and  $\mathcal{P}'$  is corrupt: // Now that TestPwd can no longer be
    called on this session, we can finally call it
    - (G2) Retrieve (P1Session, sid,  $\mathcal{P}'$ ,  $\mathcal{P}$ , [pw'])
    - (G2) Run code of (TestPwd, sid||1,  $\mathcal{P}$ , pw')
  • (G1) Record (P1Key, sid,  $\mathcal{P}$ ,  $K_{1,1}$ ,  $K_{1,2}$ ,  $m$ )
  • (G1) Mark the P1Session completed
  • If  $\mathcal{P}$  is corrupt:
    - (G1) Send (sid,  $K_{1,1}$ ,  $K_{1,2}$ ) to  $\mathcal{P}$ 
    - (G1) Return
  • Else: // Now compute  $Z$  for honest  $\mathcal{P}$  if we can, and run the rest of the protocol
    - (G2) If  $\exists(\text{TestedPwd}, \text{sid}, \mathcal{P}$ , [pw'], "correct"):  $Z := \text{H}_0(\text{sid}, \text{pw}', K_{1,1})$  //  $K_{1,1} \neq \perp$  because it's
      only  $\perp$  when both parties are honest and  $\mathcal{P}$ 's session is uncompromised. But when that's the
      case, (TestedPwd, sid,  $\mathcal{P}$ , ...) doesn't exist
    - (G9) & (G10) Else:  $Z \leftarrow \{0, 1\}^\lambda$ 
    - (G1) Store (P2Input, sid,  $\mathcal{P}$ ,  $Z$ )
    - (G1) Run  $\mathcal{P}_2$  on behalf of  $\mathcal{P}$  on input  $Z$ 

On (TestPwd, sid,  $\mathcal{P}$ , pw) from  $\mathcal{Z}$  to  $\mathcal{F}_{\text{PAKE}}^{(1)}$  or as internal call // W.l.o.g.,  $\mathcal{P}$  is honest
  • (G2) Retrieve (P1Session, sid,  $\mathcal{P}$ ,  $\cdot$ ,  $\cdot$ ) marked fresh
  • (G2) Send (TestPwd, sid,  $\mathcal{P}$ , pw) to  $\mathcal{F}_{\text{PAKE}}$ , get result  $b$ , and mark P1Session accordingly
  • (G2) Record (TestedPwd, sid,  $\mathcal{P}$ , pw,  $b$ )
  • (G2) Return  $b$ 

On  $\mathcal{Z}$  sending the final message in  $\mathcal{P}_2$  to honest (sid,  $\mathcal{P}$ )
  • (G1) Ignore if simulated  $\mathcal{P}$  is not ready to receive the final  $\mathcal{P}_2$  message
  • (G1) Retrieve (P1Key, sid,  $\mathcal{P}$ ,  $\cdot$ , [ $K_{1,1}, K_{1,2}$ ],  $\cdot$ ) // Exists because  $\mathcal{P}$  starts  $\mathcal{P}_2$  after P1Key was
    recorded
  • (G1) If  $K_{1,1}, K_{1,2} \neq \perp$ : // If we can compute  $K_2$  and  $K$  then do so
    - (G1) Compute  $\mathcal{P}_2$  protocol output  $\text{tr}$ ,  $K_2$  on behalf of  $\mathcal{P}$ 
    - (G1) Record (P2Key, sid,  $\mathcal{P}$ ,  $K_2$ ,  $\text{tr}$ )
    - (G1) Compute  $K := \text{H}_1(\text{sid}, K_{1,2}, K_2, \text{tr})$ 
  • Else:  $K := \perp$  //  $\mathcal{P}_1$  is uncompromised and interacting with honest party: Sim's key will be ignored.
  • (G3) If  $\exists$  record (P2Key, sid,  $\mathcal{P}'$ , *, [tr']) with  $\text{tr} \neq \text{tr}'$ , send (TestPwd, sid,  $\mathcal{P}$ ,  $\perp$ ) to  $\mathcal{F}_{\text{PAKE}}$  // DoS
    on  $\mathcal{P}_2$ 
  • (G1) Send (NewKey, sid,  $\mathcal{P}$ ,  $K$ ) to  $\mathcal{F}_{\text{PAKE}}$ 

```

Figure 14: Simulator for $\text{SeqComb}[\mathcal{F}_{\text{PAKE}}^{(1)}, \mathcal{P}_2]$ realizing $\mathcal{F}_{\text{PAKE}}$, Lemma 3. Random oracles H_0, H_1 are simulated as in the real execution. Label (**G**₉) indicates that the line of code is activated in game **G**₉.

3. \mathcal{Z} never issues a `TestPwd` query against an honest party if the other party is corrupt

Game G_0 : The real execution.

Game G_1 : Change layout. We change the previous game as follows:

- We move the whole execution into a single machine and call it the simulator `Sim`.
- We add all the record-keeping of the simulator as in Figure 14.
- In between \mathcal{Z} and `Sim` we add one dummy party for each real party.
- In between the dummy parties and `Sim`, we add the ideal functionality $\mathcal{F}_{\text{PAKE}}$ as in Figure 1, but relaying passwords of honest parties to `Sim` and relaying outputs keys provided by `Sim` to the dummy parties. We call that functionality \mathcal{F} .

The changes are only syntactical since the real execution runs on same inputs and produces outputs the same way as in the previous game. We hence have

$$\Pr[G_0] = \Pr[G_1].$$

Game G_2 : Extract password from corrupt party We change the simulator upon a corrupt \mathcal{P} sending input $(\text{NewSession}, \text{sid}, \mathcal{P}, \mathcal{P}', \text{pw})$ to $\mathcal{F}_{\text{PAKE}}^{(1)}$, or upon \mathcal{Z} sending $(\text{TestPwd}, \text{sid}, \mathcal{P}', \text{pw})$. W.l.o.g., \mathcal{P}' from these queries is honest. Upon `TestPwd`, the simulator immediately forwards the query to \mathcal{F} . For the `NewSession`, the simulator first waits until \mathcal{P}' is supposed to output a key through $\mathcal{F}_{\text{PAKE}}^{(1)}$ (i.e., \mathcal{Z} sends a `NewKey` query for \mathcal{P}' to $\mathcal{F}_{\text{PAKE}}^{(1)}$), and then sends the password guess to \mathcal{F} . If the response from $\mathcal{F}_{\text{PAKE}}$ is “correct”, the simulator uses `pw` in the computation of Z for \mathcal{P} . Because of the input password of \mathcal{P} is the same in the simulation and in \mathcal{F} , the switch to `pw` upon “correct” is only syntactical and we have

$$\Pr[G_1] = \Pr[G_2].$$

Game G_3 : Randomize output key upon message tampering in P_2 . We change the simulator to send $(\text{TestPwd}, \text{sid}, \mathcal{P}, \perp)$ to \mathcal{F} in case of $\mathcal{P}, \mathcal{P}'$ both honest and \mathcal{Z} tampering with a message from \mathcal{P}' . At the same time, we let the functionality ignore the simulator’s output key for \mathcal{P} in case such a `TestPwd` happened, and instead take the key from its `Key` record (i.e., a randomly chosen one).

In an honest interaction, the only values depending on key $K_{1,2}$ are the outputs of the two honest parties, which because of the inclusion of `tr` in the final hash are different. Hence, this and the previous game are equally distributed unless \mathcal{Z} queries $(\text{sid}, K_{1,2}, K_2, \text{tr})$ to H_1 which happens only with probability $1/2^\lambda$ for each H_1 query of \mathcal{Z} , because $K_{1,2}$ is chosen uniformly at random. Hence we have

$$|\Pr[\mathbf{G}_3] - \Pr[\mathbf{G}_2]| \leq \frac{q_H}{2^\lambda}.$$

Game \mathbf{G}_4 : Abort if \mathcal{Z} queries a secret P_1 output key to H_0 . We abort the simulation if \mathcal{Z} queries $H_0(\text{sid}, \cdot, K_{1,1})$ for a $K_{1,1}$ that $\mathcal{F}_{\text{PAKE}}^{(1)}$ generated as output to an honest party \mathcal{P} , except when \mathcal{P} is attacked with a correct password guess.

Because $\mathcal{F}_{\text{PAKE}}^{(1)}$ generates output keys uniformly at random and, unless \mathcal{P} is subject to a successful active attack, these output keys are not given to anybody else than the honest \mathcal{P} , the probability for the abort to happen is $1/2^\lambda$ per hash query, and hence we have

$$|\Pr[\mathbf{G}_4] - \Pr[\mathbf{G}_3]| \leq \frac{q_H}{2^\lambda}.$$

Game \mathbf{G}_5 : Randomize output keys of interrupted sessions. We change \mathcal{F} to ignore the simulation's output keys for interrupted sessions and instead output the key from the **Key** record. This means that honest parties who received wrong password guesses now obtain a fresh random key from \mathcal{F} . The argument is the same as in \mathbf{G}_4 , this time relying on the negligible probability of \mathcal{Z} querying $K_{1,2}$ to H_1 . We again have

$$|\Pr[\mathbf{G}_5] - \Pr[\mathbf{G}_4]| \leq \frac{q_H}{2^\lambda}.$$

Game \mathbf{G}_6 : Functionality aligns keys In this game we let \mathcal{F} output the same key to the party who finishes last in a fresh session, with matching passwords, and thereby ignore the output key provided by the simulator.

Due to the (trivial) correctness of the protocol and of $\mathcal{F}_{\text{PAKE}}^{(1)}$ in an unattacked session, the changes are only syntactical and do not affect the output distribution. We have

$$\Pr[\mathbf{G}_5] = \Pr[\mathbf{G}_6].$$

Game \mathbf{G}_7 : Randomize output keys parties finishing first in fresh sessions We change \mathcal{F} to ignore the simulator's output keys for honest parties finishing first on a fresh session, and instead output the key from the **Key** record. I.e., these parties get a fresh uniform key generated by \mathcal{F} now.

Because of games \mathbf{G}_2 and \mathbf{G}_3 , the simulator queries **TestPwd** for honest parties under attack, and hence the fresh sessions modified in this game are guaranteed to belong to honest parties that are running with another honest party and an untampered transcript. Hence, $K_{1,2}$ chosen uniformly by $\mathcal{F}_{\text{PAKE}}^{(1)}$ are only given to the honest parties. \mathcal{Z} can only notice a difference

in the output distribution of \mathbf{G}_7 and \mathbf{G}_6 if it queries $K_{1,2}$ to \mathbf{H}_1 , which happens with negligible probability

$$|\Pr[\mathbf{G}_7] - \Pr[\mathbf{G}_6]| \leq \frac{q_H}{2^\lambda}.$$

Game \mathbf{G}_8 : Remove password equality check from simulation. We change the simulation to replace the password equality check in the `NewKey` interface of the internally simulated $\mathcal{F}_{\text{PAKE}}^{(1)}$: instead of comparing the input password of a finishing corrupt \mathcal{P} with its honest counterparty's password, the simulator compares with the correct password guess against the counterparty (if it happened).

Because `TestPwd` compares a password guess against an input password, the changes are only syntactical and we have

$$\Pr[\mathbf{G}_7] = \Pr[\mathbf{G}_8].$$

Game \mathbf{G}_9 : Randomize Z of honest unattacked parties with matching passwords. We change the simulation to use a randomly chosen $Z \leftarrow \{0, 1\}^\lambda$ as a password in \mathbf{P}_2 both honest parties if they have matching passwords and are not attacked, i.e., parties that received a session key from $\mathcal{F}_{\text{PAKE}}^{(1)}$ while their session at \mathcal{F} was marked `fresh`.

A distinguisher between this and the previous game breaks the pre-shared key equality hiding property of \mathbf{P}_2 . To see this, in \mathbf{G}_8 both parties were running on the same Z , which was randomly chosen by the random oracle. In this game, both parties are running on randomly chosen Z values. Further, \mathcal{Z} does not see the output keys K_2 since we randomized outputs of honest parties in \mathbf{G}_7 . We hence have

$$|\Pr[\mathbf{G}_9] - \Pr[\mathbf{G}_8]| \leq \text{Adv}_{\mathbf{P}_2}^{\text{PEH}}.$$

Game \mathbf{G}_{10} : Randomize Z of honest attacked parties and honest parties with mismatching passwords. We now change the simulation to use a random Z for an honest \mathcal{P} in all other cases except if \mathcal{P} 's session in \mathcal{F} is `compromised`, which the simulator can determine via the existence of a `(TestedPwd, sid, \mathcal{P} , \cdot , correct)` entry.

In all these cases, $\mathcal{F}_{\text{PAKE}}^{(1)}$ outputs a key $K_{1,1}$ to \mathcal{P} that is not given to anybody else. Hence, \mathcal{Z} can only notice a difference if it queries $K_{1,1}$ to \mathbf{H}_0 . We hence have

$$|\Pr[\mathbf{G}_{10}] - \Pr[\mathbf{G}_9]| \leq \frac{q_H}{2^\lambda}.$$

We are now in a situation where, in the simulation, honest parties whose sessions are not marked `compromised` in \mathcal{F} no longer use $K_{1,1}$ to compute Z .

The output keys of these parties are determined by \mathcal{F} and the corresponding output keys produced by the simulation are ignored by \mathcal{F} . Hence, the whole execution does not depend on the input passwords of honest parties anymore, and we can replace them by dummy values in the next game.

Game \mathbf{G}_{11} : Remove passwords from simulation. We let the simulator use \perp as input password for honest parties, and change \mathcal{F} to not forward passwords of honest parties to the simulator anymore.

As argued above, the changes go unnoticed by \mathcal{Z} because the distribution does not depend on the honest parties' passwords in the simulation as of \mathbf{G}_{10} . We hence have

$$\Pr[\mathbf{G}_{10}] = \Pr[\mathbf{G}_{11}].$$

Because forwarding of input passwords and relaying of simulator output keys was the only difference when we introduced \mathcal{F} in \mathbf{G}_1 , the functionality \mathcal{F} in \mathbf{G}_{11} is equal to $\mathcal{F}_{\text{PAKE}}$. The theorem thus follows with the simulator depicted in Figure 14.

A.3 Proof of Lemma 1

The overall idea of the proof is to let $\mathcal{F}_{\text{rlePAKE}}$ handle `LateTestPwd` queries by \mathcal{Z} . However, \mathcal{Z} issues those queries against the first sub-PAKE, i.e., $\mathcal{F}_{\text{lePAKE}}^{(1)}$, while the simulator has a `LateTestPwd` interface at $\mathcal{F}_{\text{rlePAKE}}$ which handles the overall key exchange. This causes a slight timing issue: if \mathcal{P}_1 completed but \mathcal{P}_2 is still ongoing, \mathcal{Z} can issue a late password guess against \mathcal{P}_1 while $\mathcal{F}_{\text{rlePAKE}}$ would reject such a query. The solution is to let our simulator translate late password guesses against \mathcal{P}_1 to online guesses against $\mathcal{F}_{\text{rlePAKE}}$ while the overall key exchange is not yet completed.

The proof is derived from the proof of Lemma 3 by changing all occurrences of $\mathcal{F}_{\text{PAKE}}^{(1)}$ to $\mathcal{F}_{\text{lePAKE}}^{(1)}$, changing all occurrences of $\mathcal{F}_{\text{PAKE}}$ to $\mathcal{F}_{\text{rlePAKE}}$, and modifying the games as follows.

- **Registering lazy extractions at $\mathcal{F}_{\text{rlePAKE}}$.** We add a new game $\mathbf{G}_{2.1}$ right after \mathbf{G}_2 where we change the simulation to forward a `(RegisterTest, sid, \mathcal{P})` query to $\mathcal{F}_{\text{rlePAKE}}$, but only at the very last moment before \mathcal{P} is producing an output key. That is, we mark the `P1Session` of \mathcal{P} with a `tested` flag and, right before sending `(NewKey, sid, \mathcal{P} , \cdot)` to $\mathcal{F}_{\text{rlePAKE}}$, the simulator sends `(RegisterTest, sid, \mathcal{P})` to $\mathcal{F}_{\text{rlePAKE}}$. Because the effect of `RegisterTest` is an **interrupted** record and these markings do not yet affect any output keys in this game, the changes do not affect the output distribution and the new game $\mathbf{G}_{2.1}$ is equally distributed to \mathbf{G}_2 .
- **Answering “online” `LateTestPwd` queries without passwords.** We add a new game $\mathbf{G}_{2.2}$ right after $\mathbf{G}_{2.1}$ where we change the simulation in case \mathcal{Z} sends `(LateTestPwd, sid, \mathcal{P} , pw)` for a `P1Session` that is completed

and flagged `tested`, but where \mathcal{P} has not yet produced an output key. The simulator sends $(\text{TestPwd}, \text{sid}, \mathcal{P}, \text{pw})$ to $\mathcal{F}_{\text{rePAKE}}$. Upon “correct”, it returns the `P1Key` record $K_{1,1}, K_{1,2}$ to \mathcal{Z} , otherwise a random key. Because the passwords in both $\mathcal{F}_{\text{rePAKE}}$ and $\mathcal{F}_{\text{lePAKE}}^{(1)}$ match, the changes are only syntactical and $\mathbf{G}_{2.1}$ and $\mathbf{G}_{2.2}$ are equally distributed.

- **Answering “completed” `LateTestPwd` queries without passwords.** We add a new game $\mathbf{G}_{2.3}$ right after $\mathbf{G}_{2.2}$ where we change the simulation in case \mathcal{Z} sends $(\text{LateTestPwd}, \text{sid}, \mathcal{P}, \text{pw})$ for a \mathcal{P} that has a completed `P1` session flagged `tested`, and that has already produced an output. In case the guess comes back correct with a key $K \neq \perp$, the simulator programs $H_0[\text{sid}, \text{pw}^*, K_{1,1}] := Z$ and $H_1[\text{sid}, K_{1,2}, K_2, \text{tr}] := K$ for $K_{1,1}, K_{1,2}, Z, K_2$ from the simulation of \mathcal{P} . The changes are only syntactical because the simulator re-programs exactly what \mathcal{P} already computed.
- **Randomize output keys of interrupted sessions.** This is the original \mathbf{G}_5 , where we additionally let the simulator skip the output key computation of interrupted sessions. Because `LateTestPwd` queries also result in interrupted records, these output keys are also randomized by this game. The effect is that `LateTestPwd` responses are now also random keys chosen by $\mathcal{F}_{\text{rePAKE}}$, and hence the programming of the simulator in $\mathbf{G}_{2.3}$ becomes crucial as it reinstalls consistency between the simulated protocol values for \mathcal{P} and its $\mathcal{F}_{\text{rePAKE}}$ -determined output key. Hence, the switch to different keys goes unnoticed by the environment except with negligible probability as argued in \mathbf{G}_5 , despite the additional interrupted records through the `LateTestPwd` queries.

On \mathcal{Z} sending the final message in P_2 to honest $(\text{sid}, \mathcal{P})$

- Ignore if simulated \mathcal{P} is not ready to receive the final P2 message
- Retrieve $(P1Key, \text{sid}, \mathcal{P}, \cdot, [K_{1,1}, K_{1,2}], \cdot)$ // Exists because \mathcal{P} starts P2 after P1Key was recorded
- If $K_{1,1}, K_{1,2} \neq \perp$: // If we can compute K_2 and K then do so
 - Compute P_2 protocol output tr, K_2 on behalf of \mathcal{P}
 - Record $(P2Key, \text{sid}, \mathcal{P}, K_2, \text{tr})$
 - Compute $K := H_1(\text{sid}, K_{1,2}, K_2, \text{tr})$
- Else: $K := \perp$ // P1 is uncompromised and interacting with honest party: Sim's key will be ignored.
- If \exists record $(P2Key, \text{sid}, \mathcal{P}', \cdot, [\text{tr}'])$ with $\text{tr} \neq \text{tr}'$, send $(\text{TestPwd}, \text{sid}, \mathcal{P}, \perp)$ to $\mathcal{F}_{\text{lePAKE}}$ // DoS on P_2
- **(G_{2.1})** If \mathcal{P} is honest AND P1Session is flagged **tested** AND $\nexists(\text{TestedPwd}, \text{sid}, \mathcal{P}, \dots)$: send $(\text{RegisterTest}, \text{sid}, \mathcal{P})$ to $\mathcal{F}_{\text{lePAKE}}$ // Forward the RegisterTest at the last possible moment
- Send $(\text{NewKey}, \text{sid}, \mathcal{P}, K)$ to $\mathcal{F}_{\text{lePAKE}}$

On $(\text{RegisterTest}, \text{sid}||1, \mathcal{P})$ from \mathcal{Z} to $\mathcal{F}_{\text{lePAKE}}^{(1)}$

- **(G_{2.1})** Retrieve $(P1Session, \text{sid}, \mathcal{P}, \cdot, \cdot)$ marked **fresh**
- **(G_{2.1})** Mark it **interrupted** and flag it **tested**
- // Defer forwarding the RegisterTest until right before we do NewKey

On $(\text{LateTestPwd}, \text{sid}||1, \mathcal{P}, \text{pw}^*)$ from \mathcal{Z} to $\mathcal{F}_{\text{lePAKE}}^{(1)}$ // W.l.o.g., \mathcal{P} is honest

- **(G₁)** Retrieve $(P1Session, \text{sid}, \mathcal{P}, \cdot, \cdot)$ marked **completed** with flag **tested**
- **(G₁)** Remove the flag **tested**
- **(G₁)** Retrieve $(P1Key, \text{sid}, \mathcal{P}, [K_{1,1}, K_{1,2}], [m])$ // $K_{1,1}, K_{1,2}$ can't be \perp because that only happens for untested sessions
- **(G_{2.2})** If $m = \text{interrupted}$: // TestPwd was not queried on this session yet. It was marked interrupted by RegisterTest, so the normal TestPwd points were never triggered
 - **(G_{2.2})** Send $(\text{TestPwd}, \text{sid}, \mathcal{P}, \text{pw}^*)$ to $\mathcal{F}_{\text{lePAKE}}$, get result b
 - **(G_{2.2})** If “correct”: return $K_{1,1}, K_{1,2}$
 - **(G_{2.2})** Else: return $K'_1 \leftarrow_{\$} \{0, 1\}^{2\lambda}$
- Else: // The session ended. Need to make it such that $K_{1,2}$ explains the final K
 - **(G_{2.3})** Retrieve $(P2Key, \text{sid}, \mathcal{P}, K_2, \cdot)$
 - **(G_{2.3})** Retrieve $(P2Input, \text{sid}, \mathcal{P}, Z)$
 - **(G_{2.3})** Send $(\text{LateTestPwd}, \text{sid}, \mathcal{P}, \text{pw}^*)$ to $\mathcal{F}_{\text{lePAKE}}$ and get K
 - **(G_{2.3})** If $K = \perp$ return $K'_1 \leftarrow_{\$} \{0, 1\}^{2\lambda}$
 - Else:
 - * **(G_{2.3})** Set $H_0[\text{sid}, \text{pw}^*, K_{1,1}] := Z$ // Program the P2 input to the correct password.
 - * **(G_{2.3})** Set $H_1[\text{sid}, K_{1,2}, K_2, \text{tr}] := K$ // Program the final hash to the key K previously output by \mathcal{P} .
 - * **(G₁)** Return $K_{1,1}, K_{1,2}$

Figure 15: Simulator for Lemma 1 in terms of added/modified interfaces from Figure 14.

B Formal distinguisher for SeqComb with lazy extraction

To make SeqComb work with lazy extraction PAKEs, we would need to prove the following statement: if P_1 UC-realizes $\mathcal{F}_{\text{lePAKE}}$ and P_2 is PSK equality hiding, then P UC-realizes $\mathcal{F}_{\text{lePAKE}}$ (in the \mathcal{F}_{RO} -hybrid model). In this section we give a formal distinguisher for this statement.

The PSK equality hiding property does not rule out offline attacks against P_2 . We use that to distinguish the protocol from $\mathcal{F}_{\text{lePAKE}}$.

1. \mathcal{Z} starts honest \mathcal{P} with pw , sends RegisterTest to $\mathcal{F}_{\text{lePAKE}}^{(1)}$ and completes \mathcal{P} 's session to receive output key K .
2. \mathcal{Z} flips a coin b and sends pw_b as LateTestPwd guess to $\mathcal{F}_{\text{lePAKE}}^{(1)}$, where $\text{pw}_0 = \text{pw}$ and pw_1 is a different password.
3. \mathcal{Z} receives key $K_{1,1}, K_{1,2}$ from $\mathcal{F}_{\text{lePAKE}}^{(1)}$
4. \mathcal{Z} tests whether $H_0(\text{sid}, \text{pw}_b, K_{1,1})$ was used by \mathcal{P} in P_2 .
 - If $b = 0$ \mathcal{Z} expects that test to pass
 - If $b = 1$ \mathcal{Z} expects that test to fail.

A simulator Sim can react in the following way to this \mathcal{Z}

1. Sim has no information about \mathcal{P} 's password, so it chooses a random Z to run P_2
2. Sim submits pw_b via LateTestPwd to $\mathcal{F}_{\text{lePAKE}}$ and receives back K_b , where $K_b = K$ for $b = 0$ and random otherwise.
3. Sim chooses random $K_{1,1}, K_{1,2}$ to reply to \mathcal{Z} .
4. Sim obtains query $H_0(\text{sid}, \text{pw}_b, K_{1,1})$ and can now program it:
 - If Sim programs the hash to Z and $b = 1$, the test passes— \mathcal{Z} can distinguish
 - If Sim programs the hash to a different Z' and $b = 0$, the test fails— \mathcal{Z} can distinguish

C Changelog

Version 1.6

- Corrected preconditions for ParComb in results table.
- Clarified claim about the QROM analysis of SeqComb in [LL24].

Version 1.5

- Added footnote explaining SPAKE2 initiators behave slightly differently from receivers, and that this difference does not matter for our analysis.
- Added comparisons between our theorems about ParComb and SeqComb and those from [LL24], including a table of results.
- Removed CHIC and OCAKE from list of PSK-equality-hiding PAKEs, and added EKE-PRF. Also fixed proof of PSK equality hiding.
- Added proof that CPace is perfectly password hiding.
- Amended definition of password-hiding to be defined for maliciously chosen y .
- Amended Lemma 1 to include requirement of password hiding on P_1 .
- Relaxed requirement in theorems to *statistical* password hiding. Perfect password hiding was never necessary.

Version 1.0

Initial ePrint submission