

OPAQUE: An Asymmetric PAKE Protocol Secure Against Pre-Computation Attacks

Stanislaw Jarecki¹, Hugo Krawczyk², and Jiayu Xu¹

¹ University of California, Irvine. Email: {stasio@ics.,jiayux@}uci.edu.

² IBM Research. Email: hugo@ee.technion.ac.il.

Abstract. Password-Authenticated Key Exchange (PAKE) protocols allow two parties that only share a password to establish a shared key in a way that is immune to offline attacks. *Asymmetric* PAKE (aPAKE) strengthens this notion for the more common client-server setting where the server stores a mapping of the password and security is required even upon server compromise, that is, the only allowed attack in this case is an (inevitable) offline exhaustive dictionary attack against individual user passwords. Unfortunately, most suggested aPAKE protocols (that dispense with the use of servers’ public keys) allow for *pre-computation attacks* that lead to the *instantaneous compromise* of user passwords upon server compromise, thus forgoing much of the intended aPAKE security. Indeed, these protocols use – in essential ways – deterministic password mappings or use random “salt” transmitted *in the clear* from servers to users, and thus are vulnerable to pre-computation attacks.

We initiate the study of *Strong aPAKE* protocols that are secure as aPAKE’s but *are also secure against pre-computation attacks*. We formalize this notion in the Universally Composable (UC) settings and present two modular constructions using an Oblivious PRF as a main tool. The first builds a Strong aPAKE from *any* aPAKE (which in turn can be constructed from any PAKE [18]) while the second builds a Strong aPAKE from *any* authenticated key-exchange protocol secure against reverse impersonation (a.k.a. KCI). Using the latter transformation, we show a *practical instantiation of a UC-secure Strong aPAKE* in the Random Oracle model. The protocol (“OPAQUE”) consists of 2 messages (3 with mutual authentication), requires 3 and 4 exponentiations for server and client, respectively (2 to 4 of which can be fixed-base depending on optimizations), provides forward secrecy, is PKI-free, supports user-side hash iterations, and allows a user-transparent server-side threshold implementation.

1 Introduction

Passwords constitute the most ubiquitous form of authentication in the Internet, from the most mundane to the most sensitive applications. The almost universal password authentication method in practice relies on TLS/SSL and consists of the user sending its password to the server under the protection of a client-to-server confidential TLS channel. At the server, the password is decrypted and

verified against a one-way image typically computed via hash iterations applied to the password and a random “salt” value. Both the password image and salt are stored for each user in a so-called “password file.” In this way, an attacker who succeeds in stealing the password file is forced to run an exhaustive *offline dictionary attack* to find users’ passwords given a set (“dictionary”) of candidate passwords. The two obvious disadvantages of this approach are: (i) the password appears in cleartext at the server during login; and (ii) security breaks if the TLS channel is established with a compromised server’s public key (a major concern given today’s too-common PKI failures³).

Password protocols have been extensively studied in the crypto literature – including in the above client-server setting where the user is assumed to possess an authentic copy of the server’s public key [19,20], but the main focus has been on *password-only* protocols where the user does not need to rely on any outside keying material (such as public keys). The basic setting considers two parties that share the same low-entropy password with the goal of establishing shared session keys secure against *offline dictionary attacks*, namely, against an active attacker that possesses a small dictionary from which the password has been chosen. The only viable option for the attacker should be the inevitable *online impersonation attack* with guessed passwords. Such model, known as *password-authenticated key exchange (PAKE)*, was first studied by Bellare and Merritt [5] and later formalized by Bellare et al. [4] in the game-based indistinguishability approach. Canetti et al. [12] formalized PAKE in the Universally Composable (UC) framework [11], which better captures PAKE security issues such as the use of arbitrary password distributions, the inputting of wrong passwords by the user, and the common use in practice of related passwords for different services.

Whereas the cryptographic literature on PAKE’s focuses on the above basic setting, in practice the much more common application of password protocols is in the client-server setting. However, sharing the same password between user and server would mean that a break to the server leaks plaintext passwords for all its users. Thus, what’s needed is that upon a server compromise, and the stealing of the password file, an attacker is forced to perform an exhaustive offline dictionary attack as in the above TLS scenario. No other attack, except for an inevitable online guessing attack, should be feasible. In particular, *the two main shortcomings of password-over-TLS* mentioned earlier - reliance on public keys and exposure of the password to the server - *need to be eliminated*. This setting, known as *aPAKE*, for *asymmetric PAKE* (also called *augmented* or *verifier-based*), was introduced by Bellare and Merritt [6], later formalized in the simulation-based approach by Boyko et al. [10], and in the UC framework by Gentry et al. [18]. Early protocols proven in the simulation-based model include [10,28,29]. Later, Gentry et al. [18] presented a compiler that transforms any UC-PAKE protocol into a UC-aPAKE (adding an extra round of communication and

³ PKI failures include stealing of server private keys, software that does not verify certificates correctly, users that accept invalid or suspicious certificates, certificates issued by rogue CAs, rogue CAs accepted as roots of trust, servers that share their TLS keys with others, e.g. CDN providers or security monitoring software; and more.

a client’s signature). This was followed by [24] who show the first simultaneous one-round adaptive UC-aPAKE protocol. In addition, several aPAKE protocols targeting practicality have been proposed, most with ad-hoc security arguments, and some have been (and are being) considered for standardization (see below).

A common unfortunate property of all these aPAKE protocols, including those being proposed for practical use and regardless of their underlying formalism, is that they are *all vulnerable to pre-computation attacks*. Namely, the attacker \mathcal{A} can *pre-compute* a table of values based on a passwords dictionary D , so as soon as \mathcal{A} succeeds in compromising a server it can *instantly* find a user’s password. This significantly weakens the benefits of security against server compromise that motivate the aPAKE notion in the first place. Moreover, while current definitions require that the attacker cannot exploit a server compromise without incurring a workload proportional to the dictionary size $|D|$, these definitions allow all this workload to be spent *before* the actual server compromise happens. Indeed, this weakening in the existing aPAKE security definition [18] is needed to accommodate aPAKE protocols that store a one-way *deterministic* mapping of the user’s password at the server, say $H(\text{pw})$. Such protocols trivially fall to a pre-computation attack as the attacker \mathcal{A} can build a table of $(H(\text{pw}), \text{pw})$ pairs for all $\text{pw} \in D$, and once it compromises the server, it finds the value $H(\text{pw})$ associated with a user and immediately, in $\log(|D|)$ time, finds that user’s password. Such devastating attack can be mitigated by “personalizing” the password map, e.g., hashing the password together with the user id. This forces \mathcal{A} to pre-compute separate tables for individual users, yet all this effort can still be spent before the actual server compromise. Note that in the case of passwords transmitted over TLS, pre-computation is prevented since password are hashed with a random salt visible to the server only. In contrast, existing aPAKE protocols that do not rely on PKI, either don’t use salt or if they do, the salt is transmitted from server to user during login *in the clear*⁴. Given that password stealing via server compromise is the main avenue for collecting billions of passwords by attackers, the above vulnerability of existing aPAKE protocols to pre-computation attacks is a serious flaw, and in this aspect password-over-TLS is more secure than all known aPAKE schemes.

Our contributions

We initiate the study of *Strong aPAKE (SaPAKE)* protocols that strengthen the aPAKE security notion by *disallowing pre-computation attacks*. We formalize this notion in the Universally Composable (UC) model by modifying the aPAKE functionality from [18] to eliminate an adversarial action which allowed such pre-computation attacks. As we explain above, allowing pre-computation attacks was indeed necessary to model the security of existing aPAKE protocols.

⁴ While aPAKE protocols are not intended to run over TLS, we point out that even in such a case, the transmitted salt would be open to a straightforward active attack.

The next contribution is building Strong aPAKE (SaPAKE) protocols. For this we present two generic constructions. The first builds the SaPAKE protocol from any aPAKE protocol (namely one that satisfies the original definition from [18]) so that one can “salvage” existing aPAKE protocols. To do so we resort to Oblivious PRF (OPRF) functions [17, 22], namely, a PRF with an associated two-party protocol that in our case is run between a server S that stores a PRF key k and a user U with a password pw . At the end of the interaction, U learns the PRF output $F_k(\text{pw})$ and S learns nothing (in particular, nothing about pw). We show that by preceding any aPAKE protocol with an OPRF interaction in which U computes the value $\text{rw} = F_k(\text{pw})$ with the help of S and uses rw as the password in the aPAKE protocol, one obtains a Strong aPAKE protocol. We show that if the OPRF and the given aPAKE protocol are, respectively, UC realizations of the OPRF functionality (defined in [22]) and the original aPAKE functionality from [18], the resultant scheme realizes our UC functionality $\mathcal{F}_{\text{SaPAKE}}$.

Our second transformation consists of the composition of an OPRF as above with a regular authenticated key exchange protocol AKE. We require UC security for the AKE protocol as well as a property known as resistance to KCI attacks. The latter means that an attacker that learns the secret keys of one party P , but does not actively control P , cannot use this information to impersonate another party P' to P . KCI resistance is a common property of most AKE protocols. In our SaPAKE construction, U first runs the OPRF with S to compute $\text{rw} = F_k(\text{pw})$; then it runs the AKE protocol with S using a private key stored, encrypted using an *authenticated* encryption under rw , at S who sends it to U . Crucial to the security of the protocol is the use of authenticated encryption with a “random-key robustness” property, which is achieved naturally by some schemes or otherwise can be easily ensured, e.g., by adding an HMAC to a symmetric encryption scheme. Under these conditions we show that the composed scheme realizes our UC functionality $\mathcal{F}_{\text{SaPAKE}}$.

Next, we use the above second transformation to instantiate a Strong aPAKE protocol with a very efficient OPRF and any efficient AKE with the KCI property. The OPRF scheme we use, essentially a Chaum-type blinded DH computation, has been proven UC-secure by Jarecki et al. [21, 22]. We show that this OPRF scheme, which we call DH-OPRF (called 2HashDH in [21, 22]), remains secure in spite of changes to the OPRF functionality that we introduce for supporting a stronger OPRF notion needed in our setting. We call the result of this instantiation, the *OPAQUE protocol*.

OPAQUE combines the best properties of existing aPAKE protocols and of the standard password-over-TLS. As any aPAKE-secure protocol, it offers two fundamental advantages over the TLS-based solution: It does not rely on PKI and the plaintext password is never in the clear at the server. The only way for an attacker that observes (or actively controls) a session at a server to learn the password is via an exhaustive offline dictionary attack. Watching or participating in a session with the user does not help the attacker. At the same time, OPAQUE resolves the major flaw of existing aPAKE protocols relative to password-over-TLS, namely, their vulnerability to pre-computation attacks.

In addition to the above fundamental properties, OPAQUE enjoys important properties for use in practice. Its modularity allows for its use with different key-exchange schemes that can provide different features and performance tradeoffs. When implemented with a 2-message implicit-authentication KE protocol (e.g., HMQV [27]), OPAQUE takes only 2 messages (or 3 with mutual explicit authentication). The computational cost (using the DH-OPRF scheme from Section A) is one exponentiation for the server and two for the client⁵ in addition to the KE protocol cost (with HMQV, this cost is 2.17 exponentiations per party). OPAQUE offers forward secrecy (a particularly crucial property for password protocols) if the KE does. OPAQUE further supports password hardening for increasing the cost of offline dictionary attacks (upon server compromise) through user-side iterated hashing without the need to transmit salt from S to U. In Figure 7 in Section 6 we show an instantiation of OPAQUE in the RO model with HMQV as the AKE.

Compared to the practical aPAKE protocols that have been and are being considered for standardization (cf., [1, 32]), OPAQUE fares clearly better on the security side as the only scheme that offers resistance to pre-computation attacks while all others are vulnerable. Performance-wise, OPAQUE is competitive with the more efficient among these protocols (see Sec. 6). Additional advantages of OPAQUE include its ability to store and retrieve user’s secrets such as a bitcoin wallet, authentication credentials, encrypted backup keys, etc., and to support a user-transparent server-side threshold implementation [23] (where the only exposure of the user password - or any stored secrets - is in case a threshold of servers is compromised and even then a full dictionary attack is required). Finally, we comment that while OPAQUE can completely replace password authentication in TLS, it can also be used in conjunction with TLS, either for bootstrapping client authentication (via an OPAQUE-retrieved client signing key) or as an hedge against PKI failures. In other words, while we are accustomed to use TLS to protect passwords, OPAQUE can be used to protect TLS.

We stress that variants of OPAQUE have been studied in prior work in several settings but none of these works presents a formal analysis of the protocol as an aPAKE, let alone as a Strong aPAKE, a notion that we introduce here for the first time. While our treatment frames OPAQUE in the context of Oblivious PRFs [21, 22], its design can be seen as an instantiation of the Ford-Kaliski paradigm for password hardening and credential retrieval using Chaum’s blinded exponentiation. Boyen [9] specifies and studies the protocol (called HPAKE) in the setting of client-side *halting KDF* [8]. Jarecki et al. [21, 22] study a threshold version (also using the OPRF abstraction) in the context of *password-protected secret sharing (PPSS)* protocols. Because of the relation between PPSS and Threshold PAKE protocols [21], this analysis implies security of OPAQUE as a PAKE protocol in the BPR model [4] but not as an aPAKE (let alone as a strong aPAKE).

⁵ A variant of the protocol discussed in Section 6.2 allows one or both of the client’s exponentiations to be fixed-base and offline.

2 The Strong aPAKE Functionality

We present the ideal UC Strong aPAKE functionality, $\mathcal{F}_{\text{SaPAKE}}$, that will serve as our definition of Strong aPAKE security; namely, we call a protocol a secure *Strong aPAKE* if it realizes $\mathcal{F}_{\text{SaPAKE}}$. Functionality $\mathcal{F}_{\text{SaPAKE}}$ is a simple but significant variant of the UC aPAKE functionality $\mathcal{F}_{\text{aPAKE}}$ from [18] (it was denoted $\mathcal{F}_{\text{apwKE}}$ in [18]) which we recall in Fig. 1.

The aPAKE functionality of [18] is based on the UC PAKE functionality from [12], and it includes extensions needed for taking care of the asymmetric nature of the aPAKE setting. First, in an aPAKE scheme the server and the user run different programs: The user runs an aPAKE session on a password (via command `USRSESSION`) while the server runs it on a “password file” `file[sid]` that represents server’s user-specific state corresponding to the user’s password, e.g., a password hash, which the server creates on input the user’s password during aPAKE initialization, via command `STOREPWDFILE`. Furthermore, $\mathcal{F}_{\text{aPAKE}}$ models a possible compromise of a server, via command `STEALPWDFILE`, from which the attacker obtains `file[sid]`. Such compromise subsequently allows the attacker to (1) impersonate the server to the user, via command `IMPERSONATE`, and (2) find the password via an offline dictionary attack, via command `OFFLINETESTPWD`. The way functionality $\mathcal{F}_{\text{aPAKE}}$ of [18] handles the offline dictionary attack is the focus of the Strong aPAKE functionality we propose, and we discuss them below.

Strong aPAKE vs. aPAKE. Our functionality $\mathcal{F}_{\text{SaPAKE}}$ is almost identical to $\mathcal{F}_{\text{aPAKE}}$ except that the text with the gray background in Fig. 1 is omitted. That is, the only difference between $\mathcal{F}_{\text{SaPAKE}}$ and $\mathcal{F}_{\text{aPAKE}}$ are in the actions upon the stealing of the password file; specifically, $\mathcal{F}_{\text{SaPAKE}}$ omits recording the $(\text{OFFLINE}, \text{pw})$ pairs and does not allow for `OFFLINETESTPWD` queries made before the `STEALPWDFILE` query. Let us explain. Let’s consider first the definition of $\mathcal{F}_{\text{SaPAKE}}$, i.e., with the gray text omitted. In this case, the actions upon server compromise, i.e., `STEALPWDFILE`, are simple. First, a flag is defined to mark that the password file has been compromised. Second, once this event happens, the adversary is allowed to submit password guesses and be informed if a guess was correct. Note that each guess “costs” the attacker one `OFFLINETESTPWD` query. This together with the restriction that these queries can only be made after the password file is compromised ensure that shortcuts in finding the password after such compromise are not possible, namely that the attacker needs to pay with one `OFFLINETESTPWD` query for each password it wants to test. Thus, pre-computation attacks are made infeasible.

Now, consider the $\mathcal{F}_{\text{aPAKE}}$ functionality from [18] which includes the text in gray too. This functionality allows the attacker, via $(\text{OFFLINE}, \text{pw})$ records, to make guess queries against the password even before the password file is compromised. The restriction is that the responses to whether a guess was correct or not are provided to the attacker only after a `STEALPWDFILE` event. But note that if one of these guesses was correct, the attacker learns it *immediately* upon server compromise. This provision was necessary in [18]

In the description below, we assume $P \in \{U, S\}$.

Password Registration

- On (STOREPWDFILE, sid, U, pw) from S , if this is the first STOREPWDFILE message, record (FILE, U, S, pw) and mark it UNCOMPROMISED.

Stealing Password Data

- On (STEALPWDFILE, sid) from \mathcal{A}^* , if there is no record (FILE, U, S, pw), return “no password file” to \mathcal{A}^* . Otherwise, if the record is marked UNCOMPROMISED, mark it COMPROMISED; regardless,
 - If there is a record (OFFLINE, pw), send pw to \mathcal{A}^* .
 - Else Return “password file stolen” to \mathcal{A}^* .
- On (OFFLINETESTPWD, sid, pw^*) from \mathcal{A}^* , do:
 - If there is a record (FILE, U, S, pw) marked COMPROMISED, do: if $pw^* = pw$, return “correct guess” to \mathcal{A}^* ; else return “wrong guess.”
 - Else record (OFFLINE, pw).

Password Authentication

- On (USRSESSION, $sid, ssid, S, pw'$) from U , send (USRSESSION, $sid, ssid, U, S$) to \mathcal{A}^* . Also, if this is the first USRSESSION message for $ssid$, record (SSID, U, S, pw') and mark it FRESH.
- On (SVRSESSION, $sid, ssid$) from S , retrieve (FILE, U, S, pw), and send (SVRSESSION, $sid, ssid, U, S$) to \mathcal{A}^* . Also, if this is the first SVRSESSION message for $ssid$, record (SSID, S, U, pw) and mark it FRESH.

Active Session Attacks

- On (TESTPWD, $sid, ssid, P, pw^*$) from \mathcal{A}^* , if there is a record (SSID, P, P', pw') marked FRESH, do: if $pw^* = pw'$, mark it COMPROMISED and return “correct guess” to \mathcal{A}^* ; else mark it INTERRUPTED and return “wrong guess.”
- On (IMPERSONATE, $sid, ssid$) from \mathcal{A}^* , if there is a record (SSID, U, S, pw') marked FRESH, do: if there is a record (FILE, U, S, pw) marked COMPROMISED and $pw' = pw$, mark (SSID, U, S, pw') COMPROMISED and return “correct guess” to \mathcal{A}^* ; else mark it INTERRUPTED and return “wrong guess.”

Key Generation and Authentication

- On (NEWKEY, $sid, ssid, P, SK$) from \mathcal{A}^* where $|SK| = \tau$, if there is a record (SSID, P, P', pw') not marked COMPLETED, do:
 - If the record is marked COMPROMISED, or either P or P' is corrupted, send (sid, $ssid, SK$) to P .
 - If the record is marked FRESH, a (sid, $ssid, SK'$) tuple was sent to P' , and at that time there was a record (SSID, P', P, pw') marked FRESH, send (sid, $ssid, SK'$) to P .
 - Else pick $SK'' \leftarrow_R \{0, 1\}^\tau$ and send (sid, $ssid, SK''$) to P .
 Finally, mark (SSID, P, P', pw') COMPLETED.
- On (TESTABORT, $sid, ssid, P$) from \mathcal{A}^* , if there is a record (SSID, P, P', pw') not marked COMPLETED, do:
 - If it is marked FRESH and record (SSID, P', P, pw') exists, send SUCC to \mathcal{A}^* .
 - Else send FAIL to \mathcal{A}^* and (ABORT, $sid, ssid$) to P , and mark (SSID, P, P', pw') COMPLETED.

Fig. 1: Functionalities $\mathcal{F}_{\text{aPAKE}}$ (full text) and $\mathcal{F}_{\text{SaPAKE}}$ (shadowed text omitted)

because the $\text{file}[sid]$ in their aPAKE construction contains a deterministic publicly-computable hash of the password, thus allowing for a pre-computation attack which lets the adversary instantaneously identify the password with a single table lookup upon server compromise. Indeed, one can think of the pairs $(\text{OFFLINE}, \text{pw})$ in the original $\mathcal{F}_{\text{aPAKE}}$ functionality as a pre-computed table that the attacker builds overtime and which it can use to identify the password as soon as the server is compromised. By eliminating the ability to get guesses $(\text{OFFLINE}, \text{pw})$ answered before server compromise in our $\mathcal{F}_{\text{SaPAKE}}$ functionality, we make such pre-computation attacks infeasible in the case of a Strong aPAKE.

Modeling Server Compromise and Offline Dictionary Queries. As in [18], we specify that STEALPWDFILE and OFFLINETESTPWD messages from \mathcal{A}^* to $\mathcal{F}_{\text{SaPAKE}}$ are accounted for by the environment. This is consistent with the UC treatment of adaptive corruption queries and is crucial to our modeling. Note that if the environment does not observe adaptive corruption queries then the ideal model adversary, i.e., the simulator, could immediately corrupt all parties at the beginning of the protocol, learning their private inputs and thus making the work of simulation easier. By making the player-corruption queries, modeled by STEALPWDFILE command in our context, observable by the environment, we ensure that the environment’s view of both the ideal and the real execution includes the same player-corruption events. This way we keep the simulator “honest,” because it can only corrupt a party if the environment accounts for it.

The same concern pertains to offline dictionary queries OFFLINETESTPWD , because if they were not observable by the environment, the ideal adversary could make such queries even if the real adversary does not. In particular, without environmental accounting for these queries the $\mathcal{F}_{\text{aPAKE}}$ and $\mathcal{F}_{\text{SaPAKE}}$ functionalities would be equivalent because the simulator could internally gather all the offline dictionary attack queries made by the real-world adversary before server corruption, and it would send them all via the OFFLINETESTPWD query to $\mathcal{F}_{\text{SaPAKE}}$ after server corruption via the STEALPWDFILE query. Such simulator would make the ideal-world view indistinguishable from the real-world view to the environment *if* the environment does not observe the sequence of OFFLINETESTPWD and STEALPWDFILE queries.

Finally, we note that the functionality $\mathcal{F}_{\text{SaPAKE}}$, like $\mathcal{F}_{\text{aPAKE}}$, has effectively two separate notions of a server corruption. Formally, it considers a *static* adversarial model where all entities, including users and servers, are either honest or corrupt throughout the life-time of the scheme. In addition, it allows for an *adaptive* server compromise of an honest server, via the STEALPWDFILE , which leaks to the adversary the server’s private state corresponding to a particular password file, but it does not give the adversary full control over the server’s entity. In particular, the accounts on the same server for which the adversary does not explicitly issue the STEALPWDFILE command must remain unaffected. We adopt this convention from [18] and we call a server “corrupted” if it is (statically) corrupt and adversarially

controlled, and we call an aPAKE instance “compromised” if the adversary steals its password file from the server.

Non-black-box Assumptions. Note that the aPAKE functionality requires the simulator, playing the role of the ideal-model adversary, to detect offline password guesses made by the real-world adversary. As pointed out by [18], this seems to require a non-black-box hardness assumption on some cryptographic primitive, e.g., the Random Oracle Model (ROM), which would allow the simulator to extract a password guess from adversary’s *local* computation, e.g., a local execution of aPAKE interaction on a password guess and a stolen password file.

Server Initialization. We note that while $\mathcal{F}_{\text{aPAKE}}$ defines password registration as an internal action of server S , with the user’s password as a local input, one can modify it to support an interactive procedure between user and server, e.g., to prevent S from ever learning the plaintext password. To that end one needs to assume that during the Password Registration phase there is an *authenticated channel* from server to user, so the user can verify that it is registering the password with the correct server. (Functionality $\mathcal{F}_{\text{aPAKE}}$ effectively also assumes such authenticated channel because otherwise the user’s password cannot be safely transported to S .) In practice, the server also needs to verify the user’s identity, and the password file could be created by the user and transported to the server. However, this is beyond the scope of the formal aPAKE functionality.

3 Oblivious Pseudorandom Function

Oblivious Pseudorandom Functions (OPRF) are a central tool in all our constructions. An OPRF consists of a pseudorandom function family F with an associated two-party protocol run between a server that holds a key k for F and a user with an input x . At the end of the interaction, the user learns the PRF output $F_k(x)$ and nothing else, and the server learns nothing (in particular, nothing about x). The notion of OPRF was introduced in [17]. The first UC formulation of it was given in [21], including a verifiability property that lets the user check the correct behavior of the server during the OPRF execution. Later [22] gave an alternative UC definition of OPRF which dispensed with the verifiability property, allowing for more efficient instantiations. The main idea in the OPRF formulations of [21, 22] is the use of a *ticketing mechanism* $\text{tx}(\cdot)$ that ensures that the number of input values on which anyone can compute the OPRF on a key held by an honest server S is no more than the number of executions of the OPRF recorded by S . This mechanism dispenses with the need to extract users’ inputs as is typically needed in UC simulations and it leads to much more efficient OPRF instantiations.

Here we adopt the formulation from [22] as the basis for our definition of functionality $\mathcal{F}_{\text{OPRF}}$ presented in Fig. 2. We refer to [22] for detailed rationale, but we note that it requires PRF outputs to be pseudorandom even to the

<p><u>Public Parameters:</u> PRF output-length ℓ, polynomial in security parameter τ.</p> <p><u>Convention on F and tx:</u> For every (sid, S) value $\text{tx}(sid, S)$ is initially set to 0, and for every x value $F_{sid,S}(x)$ is initially undefined. Whenever $F_{sid,S}(x)$ is referenced below and it is undefined, pick $F_{sid,S}(x) \leftarrow_{\text{R}} \{0, 1\}^{\ell}$.</p> <p><u>Initialization</u></p> <ul style="list-style-type: none"> – On (INIT, sid, x) from server S, send $(\text{INIT}, sid, F_{sid,S}(x))$ to S. <p><u>Server Compromise</u></p> <ul style="list-style-type: none"> – On $(\text{COMPROMISE}, sid)$ from \mathcal{A}^* for a server S, mark S COMPROMISED. – On $(\text{OFFLINEEVAL}, sid, S, x)$ from \mathcal{A}^*, if S is corrupted or marked COMPROMISED, send $(\text{OFFLINEEVAL}, sid, F_{sid,S}(x))$ to \mathcal{A}^*. <p><u>Evaluation</u></p> <ul style="list-style-type: none"> – On $(\text{EVAL}, sid, ssid, S, x)$ from party $P \in \{U, \mathcal{A}^*\}$, record $\langle ssid, P, S, x \rangle$ and send $(\text{EVAL}, sid, ssid, P, S)$ to \mathcal{A}^*. – On $(\text{SNDRCOMPLETE}, sid, ssid)$ from S, send $(\text{SNDRCOMPLETE}, sid, ssid, S)$ to \mathcal{A}^*. If this is the first SNDRCOMPLETE message for $ssid$, set $\text{tx}(sid, S)++$. – On $(\text{RCVCOMPLETE}, sid, ssid, S^*)$ from \mathcal{A}^*, retrieve $\langle ssid, P, S, x \rangle$ (where $P \in \{U, \mathcal{A}^*\}$); abort if (i) such record does not exist, or (ii) S is honest and not marked COMPROMISED and $S^* \neq S$, or (iii) $\text{tx}(sid, S^*) = 0$. Otherwise set $\text{tx}(sid, S^*)--$ and send $(\text{EVAL}, sid, ssid, F_{sid,S^*}(x))$ to P.

Fig. 2: Functionality $\mathcal{F}_{\text{OPRF}}$ with Adaptive Compromise

owner of the PRF key k . This does not seem achievable under non-black-box assumptions, but it is achievable, indeed very efficiently, in the Random Oracle Model (ROM). Note that the reliance on non-black-box assumptions like ROM is called for in the aPAKE context, see Section 2.

Changes from OPRF functionality of [22]. To use UC OPRF in our application(s) we need to make some changes to the way functionality $\mathcal{F}_{\text{OPRF}}$ was defined in [22], as described below. Changes (2), (3) and (4) are essentially syntactic and require only cosmetic changes in the security argument. Change (1) is the only one which influences the security argument in a more essential way. Fortunately, the DH-OPRF protocol that we use for OPRF instantiation in our protocols, shown in [22] to realize their version of the OPRF functionality $\mathcal{F}_{\text{OPRF}}$, also realizes our modified $\mathcal{F}_{\text{OPRF}}$ functionality. We recall the DH-OPRF protocol in Figure 9 in Appendix A, adapting its syntax to our changes in $\mathcal{F}_{\text{OPRF}}$, and we argue that the security proof of [22] which shows that it realizes $\mathcal{F}_{\text{OPRF}}$ defined by [22] extends to the modified functionality $\mathcal{F}_{\text{OPRF}}$ presented here.

- (1) We extend the OPRF functionality to allow the *adaptive compromise* of a server holding the PRF key via a COMPROMISE message. Such action is needed in the aPAKE setting where the attacker \mathcal{A}^* can compromise a server’s password file that contains the server’s OPRF key. After the compromise, \mathcal{A}^* is allowed to compute that server’s PRF function by itself on any value of its choice using OFFLINEEVAL and without the restrictions of the ticketing mechanism. We note that functionality $\mathcal{F}_{\text{OPRF}}$ distinguishes between (statically) *corrupted servers* and (adaptively) *compromised sessions* (the latter representing different OPRF keys at the same server). This distinction allows for a granular separation between compromised and uncompromised OPRF keys held by the same server. We adopt this distinction for consistency with the aPAKE functionality from Fig. 1 that distinguishes between an entirely corrupted server and particular aPAKE instances that can be adaptively compromised by an adversary.
- (2) We change the SDRCOMPLETE message such that it is sent from S instead of \mathcal{A} , thus restricting the number of OPRF invocations per *ssid* to one. This enforces a single password guess per aPAKE sub-session which is crucial for aPAKE security.
- (3) We change the session-id syntax used in [22] to model the use of multiple OPRF keys by the same server. In the formulation of [22] each PRF key was identified with a server identity making a one-to-one correspondence between OPRF keys and servers. Here, we allow multiple OPRF keys to be associated with one server. Each such key is identified with a tag *sid* and a server can be associated with multiple such tags. In the context of our application to aPAKE protocols, each aPAKE session is associated with a unique OPRF key used by the server for a particular user, so the session-id *sid* corresponds to a user account at that server. Any *sid* can include *sub-sessions*, denoted by *ssid*, corresponding to different runs of the OPRF protocol between a user and a server.
- (4) We add an Initialization phase to the functionality, which models a server picking an OPRF key and, in addition, computing the OPRF value on any input. This interface simplifies the usage of OPRF in our applications to aPAKE, where the server will pick an OPRF key for a new user and evaluate the OPRF on the user’s password (for generating an encryption key). This modeling differs from [22] who framed OPRF initialization as an interactive procedure through an EVAL call while here it is performed locally by the server.

4 A Compiler from aPAKE to Strong aPAKE via OPRF

In Fig. 3 we specify a compiler that transforms any OPRF and any aPAKE into a Strong aPAKE protocol. In UC terms the Strong aPAKE protocol is defined in the $(\mathcal{F}_{\text{OPRF}}, \mathcal{F}_{\text{aPAKE}})$ -hybrid world, for $\mathcal{F}_{\text{OPRF}}$ with the output length parameter $\ell = 2\tau$. The compiler is simple. First, the user transforms its password pw into a randomized value rw by interacting with the server in an OPRF protocol where the user inputs pw and the server inputs the OPRF key. Nothing is learned at the

server about \mathbf{pw} (i.e., \mathbf{rw} is indistinguishable from random as long as the input \mathbf{pw} is not queried as input to the OPRF). Next, the user sets \mathbf{rw} as its password in the given aPAKE protocol. Note that since the password \mathbf{rw} is taken from a pseudorandom set, then even if the size of this set is the same as the original dictionary D from which \mathbf{pw} was taken, the pseudorandom set is unknown to the attacker (the attacker can only learn this set via OPRF queries which require an online dictionary attack). Thus, any previous ability to run a pre-computation attack against the aPAKE protocol based on dictionary D is now lost.

We assume that \mathcal{A} always simultaneously sends queries ($\text{COMPROMISE}, sid$) and ($\text{STEALPWDFILE}, sid$) for the same sid , resp. to $\mathcal{F}_{\text{OPRF}}$ to $\mathcal{F}_{\text{aPAKE}}$, because in any instantiation of this scheme the server's OPRF-related state and aPAKE-related state would be part of the same file[sid]. Consequently, for a single sid , S 's status (COMPROMISED or not) in $\mathcal{F}_{\text{OPRF}}$ and $\mathcal{F}_{\text{aPAKE}}$ is always the same.

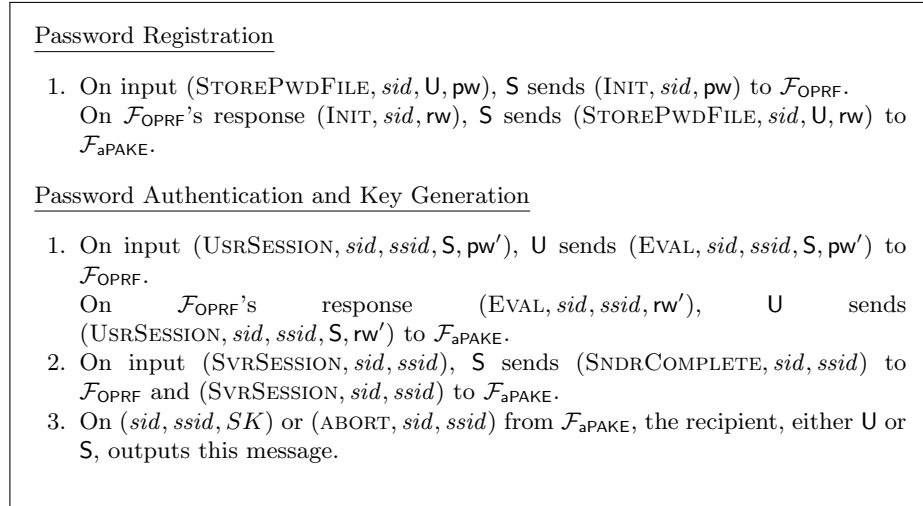


Fig. 3: Strong aPAKE Protocol in the $(\mathcal{F}_{\text{OPRF}}, \mathcal{F}_{\text{aPAKE}})$ -Hybrid World

4.1 Proof of Security

Theorem 1. *The protocol in Fig. 3 UC-realizes the $\mathcal{F}_{\text{SaPAKE}}$ functionality assuming access to the OPRF functionality $\mathcal{F}_{\text{OPRF}}$ and aPAKE functionality $\mathcal{F}_{\text{aPAKE}}$.*

Concretely, for any adversary \mathcal{A} against the protocol, there is a simulator SIM that produces a view in the simulated ideal world (henceforth simulated world) such that the advantage that an environment has in distinguishing between this view and the view in the $(\mathcal{F}_{\text{OPRF}}, \mathcal{F}_{\text{aPAKE}})$ -hybrid real world

(henceforth real world) is no more than $(q_F^2 + 2q_O + 6)/2^{2\tau+1}$, where τ is the security parameter, q_F is the number of `EVAL` and `OFFLINEEVAL` messages aimed at $\mathcal{F}_{\text{OPRF}}$ from \mathcal{A} , and q_O is the number of `OFFLINETESTPWD` messages aimed at $\mathcal{F}_{\text{aPAKE}}$ from \mathcal{A} . (In the real world, \mathcal{A} sends the messages to $\mathcal{F}_{\text{OPRF}}$ and $\mathcal{F}_{\text{aPAKE}}$. In the simulated world, \mathcal{A} sends the messages to `SIM` acting as both $\mathcal{F}_{\text{OPRF}}$ and $\mathcal{F}_{\text{aPAKE}}$.)

Due to lack of space, we leave the proof to the full version of this paper.

5 A Compiler from AKE-KCI to Strong aPAKE via OPRF

Our second transformation for building a Strong aPAKE protocol composes an OPRF with an Authenticated Key Exchange (AKE) protocol, “glued” together using authenticated encryption. We require the AKE to be secure in the UC model, namely, to realize the UC KE functionality of [14], but we also require it to be “KCI secure,” a property which we call here “security against reverse impersonation.” The notion of AKE-KCI security has been formalized with a game-based approach in [27], but to the best of our knowledge it was not formalized in UC setting, and we present such formalization in Section 5.1.

5.1 UC Definition of AKE-KCI

The KCI notion for KE protocols, which stands for “key-compromise impersonation,” captures the property we call “security against reverse impersonation,” which concerns an attacker \mathcal{A} who learns party P ’s long-term keys but otherwise does not actively control P . Resistance to KCI attacks, or “KCI security” for short, postulates that even though \mathcal{A} can impersonate P to other parties, sessions which P itself runs with honest peers need to remain secure. A game-based definition of this notion appears in [27], and here we formalize it in the UC model through functionality $\mathcal{F}_{\text{AKE-KCI}}$ presented in Fig. 4. We specialize functionality $\mathcal{F}_{\text{AKE-KCI}}$ to our user-server setting where only servers can be compromised, but it can be extended to allow for compromise of any protocol party.

Functionality $\mathcal{F}_{\text{AKE-KCI}}$ extends the standard KE functionality of [14] with two adversarial actions. The first, `COMPROMISE`, is targeted at a server and captures the compromise of the server’s keys. The second is `IMPERSONATE` which is borrowed from the aPAKE functionality of [18] shown in Fig. 1. This action can only be targeted at users’ sessions, and only for sessions with servers compromised via the `COMPROMISE` action, and it marks such session as `COMPROMISED`, which implies that the attacker can determine the session key this session outputs, via the `NEWKEY` action. This models the fact that user’s sessions with a compromised S as a peer cannot be assumed to be secure since they could have been run with the adversary who has stolen S ’s keys. However, sessions at S itself must not be affected by the `IMPERSONATE` action, and they remain secure. All other elements in $\mathcal{F}_{\text{AKE-KCI}}$ are the same as in the basic UC

In the description below, we assume $P \in \{U, S\}$.

- On (USRSESSION, $sid, ssid, S$) from U , send (USRSESSION, $sid, ssid, U, S$) to \mathcal{A}^* . If there is no record $(ssid, U, \cdot)$, record $(ssid, U, S)$ and mark it FRESH.
- On (SVRSESSION, $sid, ssid, U$) from S , send (SVRSESSION, $sid, ssid, U, S$) to \mathcal{A}^* . If there is no record $(ssid, S, \cdot)$, record $(ssid, S, U)$ and mark it FRESH.
- On (COMPROMISE, sid) from \mathcal{A}^* , mark S COMPROMISED.
- On (IMPERSONATE, $sid, ssid$) from \mathcal{A}^* , if S is marked COMPROMISED and there is a record $(ssid, U, S)$ marked FRESH, mark the record COMPROMISED.
- On (NEWKEY, $sid, ssid, P, SK$) from \mathcal{A}^* where $|SK| = \tau$, if there is a record $(ssid, P, P')$ not marked COMPLETED, do:
 - If the record is marked COMPROMISED, or either P or P' is corrupted, send $(sid, ssid, SK)$ to P .
 - If the record is marked FRESH, a $(sid, ssid, SK')$ tuple was sent to P' , and at that time there was a record $(ssid, P', P)$ marked FRESH, send $(sid, ssid, SK')$ to P .
 - Else pick $SK'' \leftarrow_{\mathcal{R}} \{0, 1\}^\tau$ and send $(sid, ssid, SK'')$ to P .
 Finally, mark $(ssid, P, P')$ COMPLETED.

Fig. 4: Functionality $\mathcal{F}_{\text{AKE-KCI}}$

KE functionality, except of some syntactic specialization to the user-server setting.

AKE-KCI security of HMQV. A concrete instantiation of protocol OPAQUE shown in Fig. 7 in Section 6, which instantiates the generic Strong aPAKE protocol shown in Section 5.2 below, using HMQV [27] as the AKE-KCI protocol. The KCI property of HMQV was proved in [27] in the game-based Canetti-Krawczyk model [13] extended to include KCI security. Here we require UC security, namely, a protocol that realizes functionality $\mathcal{F}_{\text{AKE-KCI}}$. Fortunately, [14] proves the equivalence of the game-based definition of [13] and their UC AKE formulation. Thanks to this equivalence, HMQV, as a basic KE, is secure in the UC model. More precisely, this applies to the three-message HMQV with client authentication (which satisfies the “ACK” property required for the equivalence in [14]). For the 2-message version of HMQV, the equivalence still holds using the notion of *non-information oracle* [14] that holds for HMQV under Computational Diffie-Hellman (CDH) assumption in the RO model. For our purposes, however, we need HMQV to realize the extended AKE-KCI functionality of Fig. 4. Luckily, the equivalence with the game-based definition extends to this case. Indeed, since the original equivalence from [14] holds even in the case of adaptive party corruptions, the COMPROMISE and IMPERSONATE actions introduced here – which constitute a *limited* form of adaptive corruptions – follow as a special case. Finally, we note

that the equivalence between the above models also preserves forward secrecy, so this property (proved in the game-based Canetti-Krawczyk model in [27]) holds in the UC too. We note that by the results in [27], the 3-message HMQV enjoys full PFS while the 2-message only weak PFS (against passive attackers only). The above security of HMQV (without including security against the leakage of ephemeral exponents) is based on the CDH assumption in the RO model [27].

5.2 Strong aPAKE Construction from OPRF and AKE-KCI

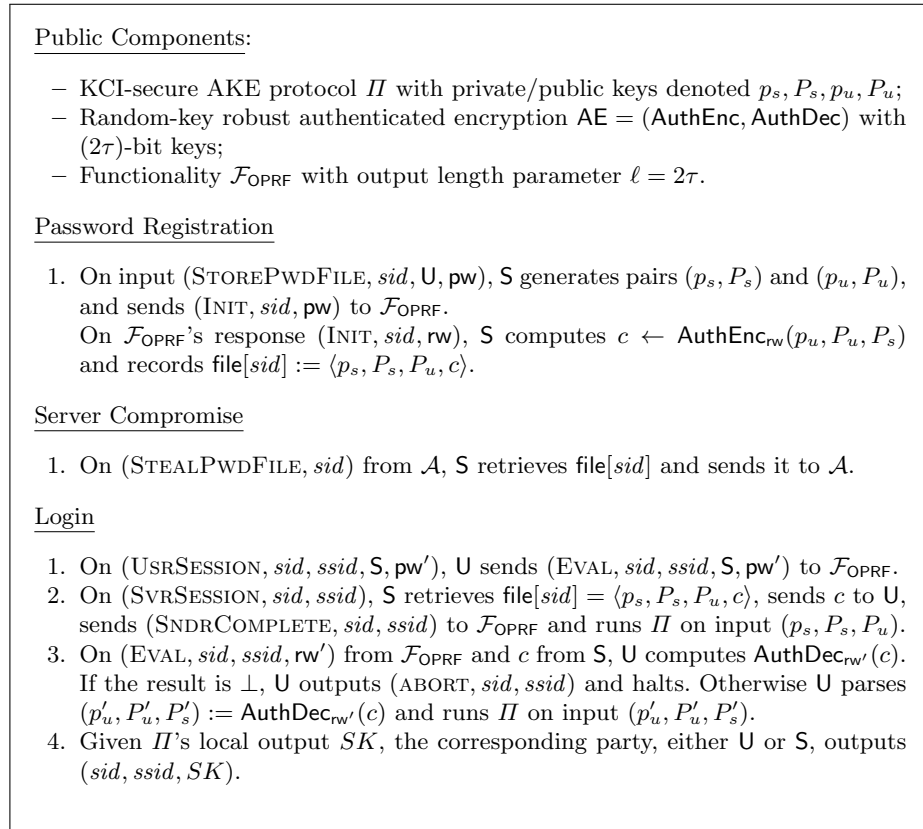


Fig. 5: Strong aPAKE based on AKE-KCI in the $\mathcal{F}_{\text{OPRF}}$ -Hybrid World

Our Strong aPAKE protocol based on OPRF and AKE-KCI is shown in Fig. 5. The protocol uses the same OPRF tool as the Strong aPAKE construction of Section 4, for length parameter $\ell = 2\tau$, which defines the “randomized password” value $\text{rw} = F_k(\text{pw})$ for user \mathbf{U} 's password pw and OPRF key k held by server \mathbf{S} .

We assume that in the AKE-KCI protocol Π each party holds a (private,public) key pair, and that the each party runs the Login subprotocol using its key pair and the public key of the counterparty as inputs. In Password Registration phase, server S generates the user U 's keys, and S 's password file contains S 's key pair p_s, P_s ; U 's public key P_u ; and a ciphertext c of U 's private key p_u , and the public keys P_u and P_s created using an Authenticated Encryption scheme using $\text{rw} = F_k(\text{pw})$ as the key. After creating the password file, value p_u is erased at S . In Login phase, S runs OPRF with U , which lets U compute $\text{rw} = F_k(\text{pw})$, it sends c to U , who can decrypt it under rw and retrieves its key-pair p_u, P_u together with the server's key P_s , at which point both parties have appropriate inputs to the AKE-KCI protocol Π to compute the session key.

Role of Authenticated Encryption. The Strong aPAKE protocol of Fig. 5 utilizes an *Authenticated Encryption* scheme $\text{AE} = (\text{AuthEnc}, \text{AuthDec})$ to encrypt and authenticate U 's AKE “credential” $m = (p_u, P_u, P_s)$. We encrypt the whole payload m for simplicity, because unlike U 's private key p_u , values P_u, P_s could be public and need to be only authenticated, not encrypted. However, the authentication property of AE must apply to the whole payload. Intuitively, U must authenticate S 's public key P_s , but if U derived even its key pair (p_u, P_u) using just the secrecy of $\text{rw} = F_k(\text{pw})$, e.g., using rw as randomness in a key generation, and U then executed AKE on such (p_u, P_u) pair, the resulting protocol would already be insecure. To see an example, if an AKE leaks U 's public key input P_u (note that AKE does not guarantee privacy of the public key) then an adversary \mathcal{A} who engages U in a single protocol instance can find U 's password pw via an offline dictionary attack by running the OPRF with U on some key k^* , and then given P_u leaked in the subsequent AKE it finds pw s.t. the key generation outputs P_u as a public key on randomness $\text{rw} = F_{k^*}(\text{pw})$.

Thus the role of the authentication property in authenticated encryption is to commit \mathcal{A} to a single guess of rw and consequently, given the OPRF key k^* , to a single guess pw . (Note that our UC OPRF notion implies that F is collision-resistant.) To that end we need the authenticated encryption to satisfy the following property which we call *random-key robustness*:⁶ For any efficient algorithm \mathcal{A} there is a negligible probability that \mathcal{A} on input (k_1, k_2) for two random keys k_1, k_2 outputs c s.t. $\text{AuthDec}_{k_1}(c) \neq \perp$ and $\text{AuthDec}_{k_2}(c) \neq \perp$. In other words, it must be infeasible to create an authenticated ciphertext that successfully decrypts under two different randomly generated keys. This property can be achieved in the standard model using e.g. encrypt-then-MAC with a MAC that is collision resistant with respect to the message and key, a property enjoyed by HMAC with full hash output. In the RO model used by our aPAKE application one can also enforce it for any authenticated encryption scheme by attaching to its ciphertext c a hash $H(k, c)$ for a RO hash H with 2τ -bit outputs.

⁶ This notion is a weakening of *full robustness (FROB)* from [16] where the attacker is allowed to choose k_1, k_2 (in our case these keys are random). An even weaker notion, *Semi-FROB*, is defined in [16] where k_1, k_2 are random but only k_1 is provided to \mathcal{A} .

Note on not utilizing $\mathcal{F}_{\text{AKE-KCI}}$. In Fig. 5 we abstract the OPRF protocol as functionality $\mathcal{F}_{\text{OPRF}}$, but we use the real-world AKE-KCI protocol Π , rather than functionality $\mathcal{F}_{\text{AKE-KCI}}$. The reason for this presentation is that in the KE functionality of [14], of which $\mathcal{F}_{\text{AKE-KCI}}$ is an extension, it is not clear how to support a usage of the KE protocol on keys which are computed via some other mechanism than the intended KE key generation. The KE functionality of [14] assumes that each entity keeps its private key as a permanent state, authenticates to a counterparty given its identity, and a KE party cannot specify any bitstring as one’s own private key and a counterparty’s public key. This is not how we use AKE in our Strong aPAKE of Fig. 5 precisely because U does not keep state and has to reconstruct its keys from a password (via OPRF). However, we can still use the real-world protocol Π , which UC-realizes $\mathcal{F}_{\text{AKE-KCI}}$, giving it the OPRF-computed information as input. In the proof of security we utilize the simulator SIM_{AKE} , which shows that Π UC-realizes $\mathcal{F}_{\text{AKE-KCI}}$, in our simulator construction, but we rely on its correctness only if U runs Π on the correctly reconstructed (p_u, P_s, P_s) , and if the adversary causes U to reconstruct a different string we interpret this as a successful attack on U ’s login session.

5.3 Proof of Security

In Theorem 2 below we state security of the Strong aPAKE protocol of Fig. 5.


Theorem 2. *If protocol Π UC-realizes functionality $\mathcal{F}_{\text{AKE-KCI}}$ then protocol in Fig. 5 UC-realizes functionality $\mathcal{F}_{\text{SaPAKE}}$ in the $\mathcal{F}_{\text{OPRF}}$ -hybrid model.*

Concretely, suppose that there is a simulator SIM_{AKE} such that the distinguishing advantage of an environment \mathcal{Z} between the real execution of Π and \mathcal{Z} ’s interaction with SIM_{AKE} is at most $\text{Adv}_{\text{SIM}_{\text{AKE}}, \mathcal{Z}}^{\text{DIST}}(\tau)$, where τ is the security parameter. Then for any adversary \mathcal{A} with running time T against the protocol, there is a simulator SIM that produces a view in the simulated world such that the advantage that \mathcal{Z} has in distinguishing between this view and the view in the real world is no more than $\text{Adv}_{\text{AE}, T}^{\text{AUTH}}(\tau) + q_F^2 \cdot \text{Adv}_{\text{AE}, T}^{\text{RK-RBST}}(\tau) + 2\text{Adv}_{\text{SIM}_{\text{AKE}}, \mathcal{Z}}^{\text{DIST}}(\tau)$, where q_F is the number of EVAL and OFFLINEEVAL messages aimed at $\mathcal{F}_{\text{OPRF}}$ from \mathcal{A} , and $\text{Adv}_{\text{AE}, T}^{\text{AUTH}}(\tau)$ and $\text{Adv}_{\text{AE}, T}^{\text{RK-RBST}}(\tau)$ are the probabilities that any algorithm in running time T breaks the authenticity of AE and the random-key robustness of AE, respectively.


Proof. For any adversary \mathcal{A} , we construct a simulator SIM as in Fig. 6. While interacting with SIM_{AKE} , SIM plays the role of both $\mathcal{F}_{\text{AKE-KCI}}$ and \mathcal{A} .

Following [11], without loss of generality, we may assume that \mathcal{A} is a “dummy” adversary that merely passes all its messages and computations to the environment \mathcal{Z} . We omit all interactions with corrupted U and S where SIM acts as $\mathcal{F}_{\text{OPRF}}$, since the simulation is trivial (SIM gains all information needed and simply follows the code of $\mathcal{F}_{\text{OPRF}}$). To keep notation brief we denote functionality $\mathcal{F}_{\text{SaPAKE}}$ as \mathcal{F} .

In order to account for the advantage of the environment \mathcal{Z} in distinguishing between its views in the real world and the simulated world, we compare between



(a) real world



(b) simulated world

these two settings in the different simulator actions and derive the distinguishing advantages in cases where the simulation is not perfect. Below we assume that \mathcal{Z} issues the $(\text{STOREPWDFILE}, sid, U, pw)$ command to S for some pw ; otherwise any subsequent server-side commands of \mathcal{Z} will not have any effect.

- $\text{file}[sid] = \langle p_s, P_s, P_u, c \rangle$ (from \mathcal{A}): In both worlds, \mathcal{Z} receives this message after \mathcal{A} sends $(\text{COMPROMISE}, sid)$ aimed at $\mathcal{F}_{\text{OPRF}}$ and $(\text{STEALPWDFILE}, sid)$ to S , provided that \mathcal{Z} input $(\text{STOREPWDFILE}, sid, U, pw)$ to S previously.
In both worlds, p_s , P_s and P_u are generated in the same way, and c is computed as $\text{AuthEnc}_{rw}(p_u, P_u, P_s)$. The only difference is that rw is $F_{sid,S}(pw)$ in the real world, while it is chosen from random in the simulated world. There is no way for \mathcal{Z} to distinguish unless and until it queries $F_{sid,S}(pw)$ by letting \mathcal{A} send $(\text{OFFLINEEVAL}, sid, S, pw)$ aimed at $\mathcal{F}_{\text{OPRF}}$. However, once \mathcal{A} sends such message, SIM sets $F_{sid,S}(pw)$ to rw . Therefore, in both worlds, $F_{sid,S}(pw) = rw$ and \mathcal{Z} cannot distinguish.
- $(\text{OFFLINEEVAL}, sid, \rho)$ (from \mathcal{A}): In both worlds, \mathcal{Z} receives this message after \mathcal{A} sends $(\text{OFFLINEEVAL}, sid, S, x)$ to $\mathcal{F}_{\text{OPRF}}$, provided that S is corrupted or marked COMPROMISED . The selection of ρ is the same in the two worlds, except that in the simulated world, if $x = pw$, ρ is set to rw which was chosen from random in advance, while in the real world, ρ is always chosen from random directly. There is no way to distinguish between these two cases.
- $(\text{EVAL}, sid, ssid, U, S)$ (from \mathcal{A}): In both worlds, \mathcal{Z} receives this message after inputting $(\text{USRSESSION}, sid, ssid, S, pw')$ to U .
- c and $(\text{SNDRCOMPLETE}, sid, ssid, S)$ (from \mathcal{A}): In both worlds, \mathcal{Z} receives these two messages after inputting $(\text{SVRSESSION}, sid, ssid)$ to S . As argued above, \mathcal{Z} cannot distinguish the two c 's in the two worlds.
- $(\text{ABORT}, sid, ssid)$ (from U): In both worlds, \mathcal{Z} may receive this message after \mathcal{A} sends $(\text{RCVCOMPLETE}, sid, ssid, S^*)$ aimed at $\mathcal{F}_{\text{OPRF}}$ and c' aimed at U , provided that (i) there is a record $\langle ssid, U, S, pw' \rangle$ in $\mathcal{F}_{\text{OPRF}}$ (or a

For every sid and every server S , initialize $\text{tx}(sid, S)$ to 0.

Stealing Password Data and Offline Queries

1. On $(\text{COMPROMISE}, sid)$ from \mathcal{A} aimed at $\mathcal{F}_{\text{OPRF}}$ and $(\text{STEALPWDFILE}, sid)$ from \mathcal{A} aimed at S , send $(\text{STEALPWDFILE}, sid)$ to \mathcal{F} .
If \mathcal{F} returns “password file stolen,” mark S COMPROMISED, generate two key pairs (p_s, P_s) and (p_u, P_u) , pick $rw \leftarrow_{\mathcal{R}} \{0, 1\}^{2\tau}$, compute $c \leftarrow \text{AuthEnc}_{rw}(p_u, P_u, P_s)$, record $\text{file}[sid] := \langle p_s, P_s, P_u, c \rangle$, and send $\text{file}[sid]$ to \mathcal{A} as a message from S .
2. On $(\text{OFFLINEEVAL}, sid, S, x)$ from \mathcal{A} aimed at $\mathcal{F}_{\text{OPRF}}$, if S is marked COMPROMISED or corrupted, send $(\text{OFFLINETESTPWD}, sid, x)$ to \mathcal{F} . If \mathcal{F} returns “correct guess,” set $F_{sid, S}(x) := rw$. Regardless, send $(\text{OFFLINEEVAL}, sid, F_{sid, S}(x))$ to \mathcal{A} as a message from $\mathcal{F}_{\text{OPRF}}$ (if $F_{sid, S}(x)$ is undefined, pick $\rho \leftarrow_{\mathcal{R}} \{0, 1\}^{2\tau}$ and set $F_{sid, S}(x) := \rho$).

Password Authentication

1. On $(\text{USRSESSION}, sid, ssid, U, S)$ from \mathcal{F} , send $(\text{EVAL}, sid, ssid, U, S)$ to \mathcal{A} as a message from $\mathcal{F}_{\text{OPRF}}$. Also, if this is the first USRSESSION message for $ssid$, record $\langle ssid, U, S, \cdot \rangle$.
2. On $(\text{SVRSESSION}, sid, ssid, U, S)$ from \mathcal{F} , retrieve $\text{file}[sid] = \langle p_s, P_s, P_u, c \rangle$, send c and $(\text{SNDRCOMPLETE}, sid, ssid, S)$ to \mathcal{A} as a message from S to U and from $\mathcal{F}_{\text{OPRF}}$, respectively, and send $(\text{SVRSESSION}, sid, ssid, U, S)$ to SIM_{AKE} as a message from $\mathcal{F}_{\text{AKE-KCI}}$. Also, if this is the first SVRSESSION message for $ssid$, set $\text{tx}(sid, S)++$.
3. On $(\text{RCVCOMPLETE}, sid, ssid, S^*)$ from \mathcal{A} aimed at $\mathcal{F}_{\text{OPRF}}$, retrieve $\langle ssid, U, S, \cdot \rangle$; ignore this message if (i) such record does not exist, or (ii) S is honest and not marked COMPROMISED and $S^* \neq S$, or (iii) $\text{tx}(sid, S^*) = 0$. Else set $\text{tx}(sid, S^*)--$, augment $\langle ssid, U, S, \cdot \rangle$ to $\langle ssid, U, S, S^*, \cdot \rangle$ and mark $(ssid, U)$ COMPLETED.

Key Generation and Authentication

1. As soon as $(ssid, U)$ is marked COMPLETED and a c' is sent from \mathcal{A} aimed at U , retrieve $\text{file}[sid] = \langle p_s, P_s, P_u, c \rangle$ and $\langle ssid, U, S, S^*, \cdot \rangle$, and proceed as follows:
 - If $c' = c$ and $S^* = S$, send $(\text{TESTABORT}, sid, ssid, U)$ to \mathcal{F} .
If \mathcal{F} returns SUCC, send $(\text{USRSESSION}, sid, ssid, U, S)$ to SIM_{AKE} as a message from $\mathcal{F}_{\text{AKE-KCI}}$. Mark this case (1).
 - Else for every x such that $F_{sid, S^*}(x)$ is defined (denote it y), check whether $\text{AuthDec}_y(c') \neq \perp$.
 - If there are more than one such x 's, output HALT and abort.
 - If there is a unique such x , send $(\text{TESTPWD}, sid, ssid, U, x)$ to \mathcal{F} .
If \mathcal{F} returns “correct guess,” parse $(p'_u, P'_u, P'_s) := \text{AuthDec}_y(c')$. Mark this case (2).
If \mathcal{F} returns “wrong guess,” send $(\text{TESTABORT}, sid, ssid, U)$ to \mathcal{F} and halt.
 - If there is no such x , send $(\text{TESTABORT}, sid, ssid, U)$ to \mathcal{F} and halt.
2. In case (1): (i) On $(\text{IMPERSONATE}, sid, ssid)$ from SIM_{AKE} , if S is marked COMPROMISED, pass this message to \mathcal{F} ; (ii) While SIM_{AKE} simulates the execution of Π , pass messages between it and \mathcal{A} ; (iii) On $(\text{NEWKEY}, sid, ssid, P, SK)$ from SIM_{AKE} , pass this message to \mathcal{F} .
3. In case (2): (i) On \mathcal{A} 's message as from S to U , run U 's algorithm in Π (henceforth Π_u) on (p'_u, P'_u, P'_s) ; (ii) On \mathcal{A} 's message as from U to S , pass it to SIM_{AKE} as a message from \mathcal{A} , and pass SIM_{AKE} 's response to \mathcal{A} as from S to U ; (iii) When Π_u is completed with output SK , send $(\text{NEWKEY}, sid, ssid, U, SK)$ to \mathcal{F} ; (iv) On $(\text{NEWKEY}, sid, ssid, S, SK)$ from SIM_{AKE} , send $(\text{NEWKEY}, sid, ssid, S, 0^\tau)$ to \mathcal{F} .

Fig. 6: The Simulator SIM

record $\langle ssid, U, S, \cdot \rangle$ in SIM), (ii) if S is honest and not marked COMPROMISED, then $S^* = S$, and (iii) $\text{tx}(sid, S^*) > 0$.

Note that \mathcal{Z} may see a HALT message from SIM at this time. HALT occurs when there exists $x_1 \neq x_2$ such that $\text{AuthDec}_{y_1}(c') \neq \perp$ and $\text{AuthDec}_{y_2}(c') \neq \perp$, where $y_1 = F_{sid, S^*}(x_1)$ and $y_2 = F_{sid, S^*}(x_2)$. Since $F_{sid, S^*}(\cdot)$ is a random function onto $\{0, 1\}^{2\tau}$, y_1 and y_2 are independent random strings in $\{0, 1\}^{2\tau}$; thus, for fixed y_1 and y_2 , the probability that \mathcal{A} finds c' such that $\text{AuthDec}_{y_1}(c') \neq \perp$ and $\text{AuthDec}_{y_2}(c') \neq \perp$ is at most $\mathbf{Adv}_{\text{AE}, T}^{\text{RK-RBST}}(\tau)$ due to the random-key robustness of AE. Since \mathcal{A} queries F q_F times, there are q_F independent y 's; using a polynomial reduction, we have $\Pr[\text{HALT}] \leq q_F^2 \cdot \mathbf{Adv}_{\text{AE}, T}^{\text{RK-RBST}}(\tau)$.

Next we assume that HALT does not occur. In the real world, \mathcal{Z} receives (ABORT, $sid, ssid$) from U if and only if $\text{AuthDec}_{rw'}(c') = \perp$; that is, \mathcal{Z} does not receive this message if and only if $\text{AuthDec}_{rw'}(c') \neq \perp$. There are only three possibilities:

- (1) $(pw', S^*, c') = (pw, S, c)$: Then $rw' = rw = F_{sid, S}(pw)$, thus $\text{AuthDec}_{rw'}(c') = \text{AuthDec}_{rw}(c) = (p_u, P_u, P_s)$.
- (2) \mathcal{A} queries $rw' = F_{sid, S^*}(pw')$ previously, and $\text{AuthDec}_{rw'}(c') \neq \perp$: If \mathcal{A} learns rw' , then it can compute an AuthEnc instance on rw' and any message to find a c' such that $\text{AuthDec}_{rw'}(c') \neq \perp$.
- (3) Other cases where \mathcal{A} finds a c' such that $\text{AuthDec}_{rw'}(c') \neq \perp$, while rw' is independently random of everything else in \mathcal{Z} 's view (since \mathcal{A} does not query $F_{sid, S^*}(pw')$), and \mathcal{Z} does not query $\text{AuthEnc}_{rw'}(p'_u, P'_u, P'_s)$ (\mathcal{Z} queries $\text{AuthEnc}_{rw'}(p'_u, P'_u, P'_s)$ by setting $pw' = pw$ [thus making $rw' = rw$] and receiving $c = \text{AuthEnc}_{rw}(p_u, P_u, P_s)$ from S). Since AE is an authenticated encryption, the probability of (3) is at most $\mathbf{Adv}_{\text{AE}, T}^{\text{AUTH}}(\tau)$.

In the simulated world, \mathcal{Z} does not receive this message if and only if either of the following two conditions holds:

- (1) $c' = c$, $S^* = S$ and \mathcal{F} returns SUCC on (TESTABORT, $sid, ssid, U$) from SIM. The last condition holds if and only if there are two records $\langle ssid, U, S, pw' \rangle$ and $\langle ssid, S, U, pw'' \rangle$, the former marked FRESH and $pw' = pw''$. Note that no TESTPWD, IMPERSONATE or NEWKEY message has been issued yet, so the record must be FRESH. According to the syntax of SVRSESSION, we have $pw'' = pw$. Therefore, the last condition is equivalent to $pw' = pw$, thus this case is equivalent to case (1) in the real world.
- (2) There exists x s.t. $y = F_{sid, S^*}(x)$ is defined in SIM, $\text{AuthDec}_y(c') \neq \perp$ and \mathcal{F} returns “correct guess” on (TESTPWD, $sid, ssid, x$) from SIM. The last condition is equivalent to $x = pw'$; thus, the three conditions combined are equivalent to $rw' = F_{sid, S^*}(pw')$ is defined in SIM and $\text{AuthDec}_{rw'}(c') \neq \perp$. SIM defines $F_{sid, S^*}(pw')$ only when receiving (OFFLINEEVAL, sid, S^*, pw') from \mathcal{A} . Therefore, this case is equivalent to case (2) in the real world.

Hence, \mathcal{Z} receives this message in the two worlds under the same conditions, except for case (3) in the real world.

- Messages sent from U and S while executing Π (in the real world), or messages sent from SIM (in the simulated world) (from \mathcal{A}): In case (1) and

messages sent from S in case (2), they are simulated by SIM who in turn receives them from SIM_{AKE} . Since SIM_{AKE} generates \mathcal{A} 's view indistinguishable from \mathcal{A} 's view in the real world, SIM , who merely passes messages between SIM_{AKE} and \mathcal{A} , can also achieve that; the distinguishing advantage of \mathcal{Z} is at most $\mathbf{Adv}_{SIM_{AKE}, \mathcal{Z}}^{DIST}(\tau)$. For messages sent from U in case (2), they are the results of Π_u on (p'_u, P'_u, P'_s) , and are simulated perfectly.

- $(sid, ssid, SK')$ (from U): In both worlds, \mathcal{Z} receives this message when Π is completed and sends output to U . In the real world, there are two cases:
 - $(p'_u, P'_u, P'_s) = (p_u, P_u, P_s)$, i.e., the input of U to Π is correct. This corresponds to case (1) above. There are two subcases regarding Π :
 - * S is not compromised. Then according to the syntax of $\mathcal{F}_{AKE-KCI}$, SK' is a random string in $\{0, 1\}^\tau$ (independent of everything else, or the same with S 's output if S already output previously). In the simulated world, the record $\langle ssid, U, S, pw' \rangle$ in \mathcal{F} is marked FRESH, so SK' is also a random string in $\{0, 1\}^\tau$.
 - * S is compromised (then \mathcal{A} may impersonate S while interacting with U in the execution of Π and set U 's output). In the simulated world, SIM_{AKE} sends $(\text{IMPERSONATE}, sid, ssid)$ to SIM , who transfers this message to \mathcal{F} , which makes the record $\langle ssid, U, S, pw' \rangle$ marked COMPROMISED (note that we have $pw' = pw$ here since this is a condition of case (1)). Therefore, SK' chosen by SIM_{AKE} (which is the same with the SK' output by Π in the real world except for probability at most $\mathbf{Adv}_{SIM_{AKE}, \mathcal{Z}}^{DIST}(\tau)$) is the value output to U .
 - $(p'_u, P'_u, P'_s) \neq (p_u, P_u, P_s)$, i.e., the input of U to Π is incorrect. This may occur only in cases (2) and (3) above. As argued above, the probability of (3) is at most $\mathbf{Adv}_{AE, T}^{AUTH}(\tau)$. (2) is equivalent to case (2) in the simulated world, where SIM sends $(\text{TESTPWD}, sid, ssid, U, x)$ to \mathcal{F} and \mathcal{F} returns “correct guess” (meaning that $x = pw'$). After this, the record $\langle ssid, U, S, pw' \rangle$ is marked COMPROMISED. Therefore, SK' , which is computed by SIM as Π_u 's output on (p'_u, P'_u, P'_s) , is the value output to U . In the real world, U also outputs SK' .
- $(sid, ssid, SK)$ (from S): In both worlds, \mathcal{Z} receives this message when Π is completed and sends output to S .
In the real world, SK is always a random string in $\{0, 1\}^\tau$ (independent of everything else, or the same with U 's output if U already output previously). Note that in the simulated world, the record $\langle ssid, S, U, pw' \rangle$ is always marked FRESH. Therefore, SK is also random string in $\{0, 1\}^\tau$.

It remains to show that SIM_{AKE} 's view while interacting with SIM is the same as interacting with $\mathcal{F}_{AKE-KCI}$ and \mathcal{A} . When SIM acts as \mathcal{A} , the interaction is trivial since SIM merely passes messages between SIM_{AKE} and the real \mathcal{A} . Consider when SIM acts as $\mathcal{F}_{AKE-KCI}$, and note that SIM engages with SIM_{AKE} only in cases (1) and (2):

- (1) U 's input is correct: Same effect as honest U and S executing Π ;

- (2) U’s input is incorrect: Same effect as corrupted U and honest S executing Π . Note that SIM engages with SIM_{AKE} on the side of S only, so SIM_{AKE} ’s view is again the same.

We conclude that \mathcal{Z} ’s view in the real world and the simulated world is the same, except for (1) (ABORT, $sid, ssid$) or HALT after \mathcal{A} sends (RCVCOMPLETE, $sid, ssid, S^*$) and c' , (2) messages sent during the execution of Π , and (3) ($sid, ssid, SK'$) output from U. The probabilities that (1), (2) and (3) are different in the two worlds are no more than $\text{Adv}_{\text{AE},T}^{\text{AUTH}}(\tau) + q_F^2 \cdot \text{Adv}_{\text{AE},T}^{\text{RK-RBST}}(\tau)$, $\text{Adv}_{\text{SIM}_{\text{AKE}},\mathcal{Z}}^{\text{DIST}}(\tau)$ and $\text{Adv}_{\text{SIM}_{\text{AKE}},\mathcal{Z}}^{\text{DIST}}(\tau)$, respectively. Using a hybrid argument, we can see that \mathcal{Z} ’s advantage is no more than $\text{Adv}_{\text{AE},T}^{\text{AUTH}}(\tau) + q_F^2 \cdot \text{Adv}_{\text{AE},T}^{\text{RK-RBST}}(\tau) + 2\text{Adv}_{\text{SIM}_{\text{AKE}},\mathcal{Z}}^{\text{DIST}}(\tau)$.

6 OPAQUE: A Strong Asymmetric PAKE Instantiation

Figure 7 shows OPAQUE, a concrete instantiation of the generic OPRF+AKE protocol from Fig. 5. An illustration is presented in Figure 8.

The OPRF is instantiated with the DH-OPRF scheme from [22] recalled in Appendix A, while the AKE protocol can be instantiated with any UC-secure 2-message implicitly-authenticated AKE-KCI; in Fig. 7 this is illustrated with HMV [27]. Fortunately, the two messages of DH-OPRF and the two messages from HMV (or a similar protocol) can be run “in parallel” hence obtaining a 2-message SaPAKE.

By Theorem 2 on the security of the generic OPRF+AKE construction, by Lemma 1 in Appendix A on the security of DH-OPRF, and by security of HMV (see below), we get that protocol OPAQUE realizes functionality $\mathcal{F}_{\text{SaPAKE}}$, hence it is a provably-secure Strong aPAKE, under the One-More Diffie-Hellman assumption [3, 22] in ROM.

6.1 Protocol Details and Properties

We expand on the specification of OPAQUE and the protocol’s properties.

- *Password registration.* Password registration is the only part of the protocol assumed to run over secure channels where parties can authenticate each other. We note that while OPAQUE is presented with S doing all the registration operations, in practice one may want to avoid that. Instead, we can let S choose an OPRF key k_s and U choose pw, and then run the OPRF protocol between U and S so only U learns its secrets ($\text{pw}, \text{rw}, p_u$) and only S learns p_s . A problem arises with this approach if S’s policy is to check the user’s password for compliance with some rules. A possible workaround is to adapt techniques from [26] that present zero-knowledge proofs for proving compliance without disclosing the password.
- *Authenticated encryption.* As specified in Section 5.2, the scheme AuthEnc used in the protocol needs to satisfy the key-committing property defined there. In practice, using an encrypt-then-mac scheme with HMAC-256 (or larger) as

Public Parameters and Components

- Security parameter τ
- Group G of prime order q , $|q| = 2\tau$ and generator g (G^* denotes $G \setminus \{1\}$).
- Hash functions $H(\cdot, \cdot)$, $H'(\cdot)$ with ranges $\{0, 1\}^{2\tau}$ and G , respectively.
- Pseudorandom function (PRF) $f(\cdot)$ with range $\{0, 1\}^{2\tau}$.
- OPRF function defined as $F_k(x) = H(x, (H'(x))^k)$ for key $k \in \mathbb{Z}_q$.
- Key-committing authenticated encryption scheme ($\text{AuthEnc}, \text{AuthDec}$).
- Key exchange formula KE defined below.

Password Registration

1. ($\text{STOREPWDFILE}, sid, \mathbf{U}, \text{pw}$): \mathbf{S} computes $k_s \leftarrow_{\mathbb{R}} \mathbb{Z}_q$, $\text{rw} := F_{k_s}(\text{pw})$, $p_s \leftarrow_{\mathbb{R}} \mathbb{Z}_q$, $p_u \leftarrow_{\mathbb{R}} \mathbb{Z}_q$, $P_s := g^{p_s}$, $P_u := g^{p_u}$, $c \leftarrow \text{AuthEnc}_{\text{rw}}(p_u, P_u, P_s)$; it records $\text{file}[sid] := \langle k_s, p_s, P_s, P_u, c \rangle$.

Login

1. ($\text{USRSESSION}, sid, ssid, \mathbf{S}, \text{pw}$): \mathbf{U} picks $r, x_u \leftarrow_{\mathbb{R}} \mathbb{Z}_q$; sets $\alpha := (H'(\text{pw}))^r$ and $X_u := g^{x_u}$; sends α and X_u to \mathbf{S} .
2. ($\text{SVRSESSION}, sid, ssid$): On input α from \mathbf{U} , \mathbf{S} proceeds as follows:
 - (a) Checks that $\alpha \in G^*$. If not, outputs $(\text{ABORT}, sid, ssid)$ and halts;
 - (b) Retrieves $\text{file}[sid] = \langle k_s, p_s, P_s, P_u, c \rangle$;
 - (c) Picks $x_s \leftarrow_{\mathbb{R}} \mathbb{Z}_q$ and computes $\beta := \alpha^{k_s}$ and $X_s := g^{x_s}$;
 - (d) Computes $K := \text{KE}(p_s, x_s, P_u, X_u)$ and $SK := f_K(0)$;
 - (e) Sends β , X_s and c to \mathbf{U} ;
 - (f) Outputs $(sid, ssid, SK)$.
3. On β , X_s and c from \mathbf{S} , \mathbf{U} proceeds as follows:
 - (a) Checks that $\beta \in G^*$. If not, outputs $(\text{ABORT}, sid, ssid)$ and halts;
 - (b) Computes $\text{rw} := H(\text{pw}, \beta^{1/r})$;
 - (c) Computes $\text{AuthDec}_{\text{rw}}(c)$. If the result is \perp , outputs $(\text{ABORT}, sid, ssid)$ and halts. Otherwise sets $(p_u, P_u, P_s) := \text{AuthDec}_{\text{rw}}(c)$;
 - (d) Computes $K := \text{KE}(p_u, x_u, P_s, X_s)$ and $SK := f_K(0)$;
 - (e) Outputs $(sid, ssid, SK)$.

Key exchange formula KE with HMQV instantiation (if any of $X_u, P_u, X_s, P_s \notin G^*$ the receiving party outputs $(\text{ABORT}, sid, ssid)$ and halts)

$$\text{For } \mathbf{S}: \text{KE}(p_s, x_s, P_u, X_u) = H((X_u P_u^{e_u})^{x_s + e_s p_s})$$

$$\text{For } \mathbf{U}: \text{KE}(p_u, x_u, P_s, X_s) = H((X_s P_s^{e_s})^{x_u + e_u p_u})$$

where $e_u = H(X_u, \mathbf{S}) \bmod q$, $e_s = H(X_s, \mathbf{U}) \bmod q$.

Fig. 7: Protocol OPAQUE




Fig. 8: Schematic Representation of OPAQUE (see Fig. 7 for the details)

the MAC provides this property (if a scheme does not have this property then adding on top of it such a HMAC computed on the scheme's ciphertext will ensure this property).

- *Key exchange.* The generic AKE representation via the KE formula applies to any protocol whose session key is computed as a function of the long-term private-public key pair of each party and ephemeral session-specific private-public values. These values are represented as (p_s, P_s, x_s, X_s) for the server and (p_u, P_u, x_u, X_u) for the user. We note that while more general key-exchange protocols can be used with OPAQUE, this representation applies to many such protocols and, in particular, to HMQV [27] which we use here as our main instantiation.

- *Explicit mutual authentication.* The protocol as illustrated takes just two messages but does not provide explicit user authentication. With a third message the protocol achieves mutual authentication by simply adding the value $f_K(1)$ to the server's message and adding a third message where U sends $f_K(2)$ to S. Each party verifies that the value received from the other is computed correctly and if not it aborts.

- *Use of HMQV.* Recall that the security of OPAQUE depends on the KE protocol being AKE-secure in the UC model with the additional KCI property; namely, it should realize the AKE-KCI UC functionality from Fig. 4. As argued in Section 5.1, HMQV indeed realizes this functionality (under the CDH assumption in the RO model), hence it is appropriate for use in OPAQUE. Moreover, HMQV enjoys forward secrecy. Specifically, the 2-message protocol provides weak forward secrecy (i.e., forward secrecy is guaranteed for sessions where the user’s message delivered to the server came from the real U) while the 3-message variant with explicit client authentication provides full forward secrecy, namely, against arbitrary active attacks [27].
- *Forward secrecy.* This property (or lack of it) is inherited by OPAQUE from the key exchange component KE. In the case of HMQV, forward secrecy is achieved as stated above. *One cannot overstate the importance of forward secrecy in password protocols: it guarantees that past session keys remain secure upon the compromise of a user’s password (or server’s information).*
- *User iterated hashing.* OPAQUE can be strengthened by increasing the cost of a dictionary attack in case of server compromise. This is done by changing the computation of rw to $rw = H^n(F_k(\text{pw}))$, that is, the client applies n iterations of the function H on top of the result of the OPRF value $F_k(\text{pw})$. In practice, the iterations H^n would be replaced with one of the standard password-based KDFs, such as PBKDF2 [25] or bcrypt [31]. This forces an attacker that compromises the password file at the server to compute for *each* candidate password pw' the function $F_k(\text{pw}')$ as well as the additional n hash iterations. Note that n needs not be remembered by the user; it can be sent from S to U in the server’s message. Furthermore, one can follow Boyen’s design and apply the probabilistic Halting KDF function [8] as used in [9] so that the iterations count is hidden from the attacker and even from the server.
- *Performance.* OPAQUE takes two messages (three with explicit mutual authentication); one exponentiation for S, two and a hashing-into- G for U, plus the cost of KE. With HMQV, the latter cost is one offline fixed-base exponentiation and one multi-exponentiation (at the cost of 1.16 regular exponentiations) per party (about three exponentiations in total for the server and four for the user). All exponentiations are in regular DH groups, hence accommodating the fastest elliptic curves (e.g., no pairings). It is common in PAKE protocols to count number of group elements transmitted between the parties. In OPAQUE, U sends two while S sends three (one, P_u , can be omitted at the cost of one fixed-based exponentiation at the client).
- *Performance comparison.* The introduction presents background on OPAQUE and other password protocols. Here we provide a comparison with the more efficient among these protocols, particularly those that are being, or have been, considered for standardization. Clearly, OPAQUE is superior security-wise as the only one not subject to pre-computation attacks, but it also fares well in terms of performance.

AugPAKE [33, 34], is computationally very efficient with only 2.17 exponentiations per party; however, it uses 4 messages and does not provide

forward secrecy. In addition, the protocol has only been analyzed as a PAKE protocol, not aPAKE [34]. Another proposed aPAKE protocol, SPAKE2+ [2, 15], uses two messages only and 3 multi-exponentiations (or about 3.5 exponentiations) per party which is similar to OPAQUE cost. The security of the protocol has only been informally argued in [15] and to the best of our knowledge no formal analysis has appeared. We also mention SRP which has been included in TLS ciphersuites in the past but is considered outdated as it does not have an instantiation that works over elliptic curves (the protocol is defined over rings and uses both addition and multiplication). Its implementations over RSA moduli is therefore less efficient than those over elliptic curve; it also takes 4 messages.

We also mention two very recent schemes that have been formally analyzed as aPAKE protocols but, as the rest, are vulnerable to pre-computation. The protocol VTBPEKE in [30] uses 3 messages and 4 exponentiations per party and was proven secure in the non-UC aPAKE model of [7], while [24] shows a *simultaneous* one-round scheme that they prove secure in the UC aPAKE model of [18] augmented with adaptive security. The protocol works over bilinear groups and its computational cost includes 4 exponentiations and 3 pairing per party. We note that all of the above protocols require an initial message from server to user in order to transmit salt, which results in one or two added messages to the above message counts (except for VTBPEKE which already includes the salt transmission in its 3 messages). Also, all these protocols, like OPAQUE, work in the RO model.

- *Threshold implementation.* We comment on a simple extension of OPAQUE that can be very valuable in large deployments, namely, the ability to implement the OPRF phase as a Threshold OPRF [23]. In this case, an attacker needs to break into a threshold of servers to be able to impersonate the servers to the user or to run an offline dictionary attack. Such an implementation requires no user-side changes, i.e., the user does not need to know if the system is implemented with one or multiple servers.
- *Secret retrieval and hedging TLS.* Additional features of OPAQUE include the ability to store and retrieve user’s secrets (such as a bitcoin wallet, authentication credentials, encrypted backup keys, etc.) as part of the information encrypted and authenticated at the server under ciphertext c . In one particular use case such secret can be a client signature key for TLS. In this case, the key exchange part of OPAQUE can reuse that of TLS and a server’s certificate can be replaced with the server’s public key stored under the client-authenticated ciphertext c .

6.2 An OPAQUE variant: Multiplicative blinding

A variant of OPAQUE is obtained by replacing the user’s exponential blinding operation $\alpha := (H'(\text{pw}))^r$ with $\alpha := (H'(\text{pw})) \cdot g^r$. The server responds as before with $\beta = \alpha^{k_s}$. Assuming that U knows the value $y = g^{k_s}$ (previously stored or received from S), it can compute the same “hashed Diffie-Hellman” value $(H'(\text{pw}))^{k_s}$ as β/y^r . The advantage of this variant is that while the number

of client exponentiations remains the same, one is fixed-base (g^r) and the other (y^r) can also be fixed-base if \mathbf{U} caches y , a realistic possibility for accounts where the user logs in frequently (e.g., a personal email or social network). Computing y^r can also be done while waiting for the server’s response to reduce latency. Moreover, both exponentiations can be done offline although only short-term storage is recommended as the leakage of r exposes $H'(\text{pw})$. If \mathbf{U} does not store y , it needs to be transmitted to \mathbf{U} by \mathbf{S} together with the response β . This still allows for fixed-base optimization for computing g^r but not for y^r .

However, it turns out that this multiplicative mechanism results in an OPRF protocol that does *not* realize our OPRF functionality $\mathcal{F}_{\text{OPRF}}$. Thus, our analysis here does not imply the security of the multiplicative OPAQUE variant in general. If rw is redefined as $\text{rw} := H(\text{pw}, y, H'(\text{pw})^{k_s})$, i.e. if y is included under the hash, then the resulting OPRF does realize our functionality, and OPAQUE remains secure as SaPAKE under both blinding variants. This change, however, introduces a (slight) overhead of having to transmit y even if it is not strictly needed, e.g. if the client implements the exponential blinding operation. An alternative approach would be to replace the OPRF functionality $\mathcal{F}_{\text{OPRF}}$ with a weaker form $\mathcal{F}'_{\text{OPRF}}$ and to show that (i) $\mathcal{F}'_{\text{OPRF}}$ is realized by the multiplicative variant (even without hashing y) and (ii) $\mathcal{F}'_{\text{OPRF}}$ is sufficient for proving Theorem 2 hence implying the security of OPAQUE as SaPAKE. We intend to investigate this weakening of $\mathcal{F}_{\text{OPRF}}$.

References

1. CFRG, Crypto Forum Research Group, <https://datatracker.ietf.org/rg/cfrg/documents/>.
2. M. Abdalla and D. Pointcheval. Simple password-based encrypted key exchange protocols. In *Topics in Cryptology – CT-RSA 2005*, pages 191–208. Springer, 2005.
3. M. Bellare, C. Namprempe, D. Pointcheval, and M. Semanko. The one-more-RSA-inversion problems and the security of Chaum’s blind signature scheme. *Journal of Cryptology*, 16(3), 2003.
4. M. Bellare, D. Pointcheval, and P. Rogaway. Authenticated key exchange secure against dictionary attacks. In *Advances in Cryptology – EUROCRYPT 2000*, pages 139–155. Springer, 2000.
5. S. M. Bellovin and M. Merritt. Encrypted key exchange: Password-based protocols secure against dictionary attacks. In *IEEE Computer Society Symposium on Research in Security and Privacy – S&P 1992*, pages 72–84. IEEE, 1992.
6. S. M. Bellovin and M. Merritt. Augmented encrypted key exchange: a password-based protocol secure against dictionary attacks and password file compromise. In *ACM Conference on Computer and Communications Security – CCS 1993*, pages 244–250. ACM, 1993.
7. F. Benhamouda and D. Pointcheval. Verifier-based password-authenticated key exchange: New models and constructions. *IACR Cryptology ePrint Archive*, 2013:833, 2013.
8. X. Boyen. Halting password puzzles. In *Usenix Security Symposium – SECURITY 2007*, pages 119–134. The USENIX Association, 2007.

9. X. Boyen. HPAKE: Password authentication secure against cross-site user impersonation. In *Cryptology and Network Security – CANS 2009*, pages 279–298. Springer, 2009.
10. V. Boyko, P. MacKenzie, and S. Patel. Provably secure password-authenticated key exchange using Diffie-Hellman. In *Advances in Cryptology – EUROCRYPT 2000*, pages 156–171. Springer, 2000.
11. R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *IEEE Symposium on Foundations of Computer Science – FOCS 2001*, pages 136–145. IEEE, 2001.
12. R. Canetti, S. Halevi, J. Katz, Y. Lindell, and P. MacKenzie. Universally composable password-based key exchange. In *Advances of Cryptology – EUROCRYPT 2005*, pages 404–421. Springer, 2005.
13. R. Canetti and H. Krawczyk. Analysis of key-exchange protocols and their use for building secure channels. In *Advances in Cryptology – EUROCRYPT 2001*, pages 453–474. Springer, 2001.
14. R. Canetti and H. Krawczyk. Universally composable notions of key exchange and secure channels. In *Advances in Cryptology – EUROCRYPT 2002*, pages 337–351. Springer, 2002.
15. D. Cash, E. Kiltz, and V. Shoup. The twin Diffie-Hellman problem and applications. In *Advances in Cryptology – EUROCRYPT 2008*, pages 127–145. Springer, 2008.
16. P. Farshim, C. Orlandi, and R. Rosie. Security of symmetric primitives under incorrect usage of keys. *IACR Transactions on Symmetric Cryptology*, 2017(1):449–473, 2017.
17. M. J. Freedman, Y. Ishai, B. Pinkas, and O. Reingold. Keyword search and oblivious pseudorandom functions. In *Theory of Cryptography – TCC 2005*, pages 303–324. Springer, 2005.
18. C. Gentry, P. MacKenzie, and Z. Ramzan. A method for making password-based key exchange resilient to server compromise. In *Advances in Cryptology – CRYPTO 2006*, pages 142–159. Springer, 2006.
19. L. Gong, M. A. Lomas, R. M. Needham, and J. H. Saltzer. Protecting poorly chosen secrets from guessing attacks. *IEEE Journal on Selected Areas in Communications*, 11(5):648–656, 1993.
20. S. Halevi and H. Krawczyk. Public-key cryptography and password protocols. *ACM Transactions on Information and System Security (TISSEC)*, 2(3):230–268, 1999.
21. S. Jarecki, A. Kiayias, and H. Krawczyk. Round-optimal password-protected secret sharing and T-PAKE in the password-only model. In *Advances in Cryptology – ASIACRYPT 2014*, pages 233–253. Springer, 2014.
22. S. Jarecki, A. Kiayias, H. Krawczyk, and J. Xu. Highly-efficient and composable password-protected secret sharing (or: how to protect your bitcoin wallet online). In *IEEE European Symposium on Security and Privacy – EuroS&P 2016*, pages 276–291. IEEE, 2016.
23. S. Jarecki, A. Kiayias, H. Krawczyk, and J. Xu. TOPPSS: Cost-minimal password-protected secret sharing based on threshold OPRF. In *Applied Cryptology and Network Security – ACNS 2017*, pages 39–58. Springer, 2017.
24. C. S. Jutla and A. Roy. Smooth NIZK arguments with applications to asymmetric UC-PAKE. *IACR Cryptology ePrint Archive*, 2016:233, 2016.
25. B. Kaliski. PKCS# 5: Password-based cryptography specification version 2.0. 2000.

26. F. Kiefer and M. Manulis. Zero-knowledge password policy checks and verifier-based PAKE. In *Computer Security – ESORICS 2014*, pages 295–312. Springer, 2014.
27. H. Krawczyk. HMQV: A high-performance secure Diffie-Hellman protocol (extended abstract). In *Advances in Cryptology – CRYPTO 2005*, page 546. Springer, 2005.
28. P. Mackenzie. More efficient password-authenticated key exchange. In *Topics in Cryptology – CT-RSA 2001*, pages 361–377. Springer, 2001.
29. P. MacKenzie, S. Patel, and R. Swaminathan. Password-authenticated key exchange based on RSA. In *Advances in Cryptology – ASIACRYPT 2000*, pages 599–613. Springer, 2000.
30. D. Pointcheval and G. Wang. VTB-peke: Verifier-based two-basis password exponential key exchange. In *ACM Asia Conference on Computer and Communications Security – AsiaCCS 2017*, pages 301–312. ACM, 2017.
31. N. Provos and D. Mazieres. A future-adaptable password scheme. In *USENIX Annual Technical Conference, FREENIX Track*, pages 81–91, 1999.
32. J. Schmidt. Requirements for password-authenticated key agreement (PAKE) schemes. Technical report, 2017.
33. S. Shin and K. Kobara. Augmented password-authenticated key exchange (AugPAKE). *draft-irtf-cfrg-augpake-08*.
34. S. Shin, K. Kobara, and H. Imai. Security proof of AugPAKE. *IACR Cryptology ePrint Archive*, 2010:334, 2010.

A The DH-OPRF Protocol Realizing Revised $\mathcal{F}_{\text{OPRF}}$

Figure 9 shows the DH-OPRF protocol of [22] (who calls it 2HashDH), syntactically modified to realize functionality $\mathcal{F}_{\text{OPRF}}$, see Fig. 2 in Section 3. Recall that the $\mathcal{F}_{\text{OPRF}}$ functionality we show in Section 3 is a revision of the OPRF functionality defined in [22], with the most important difference being modeling adaptive corruptions. The protocol shown below is essentially the same as in [22], and requires the same One-More Diffie-Hellman assumption [3, 22] for security.

We defer the proof of the following Lemma 1 to the full version because it is very similar to the proof of security given in [22].

Lemma 1. *The DH-OPRF protocol shown in Fig. 9 UC-realizes the OPRF functionality $\mathcal{F}_{\text{OPRF}}$ under the One-More Diffie Hellman assumption in ROM.*

Modifications in the Proof of [22]. We briefly discuss how our modifications to $\mathcal{F}_{\text{OPRF}}$ influence the security proof, and leave the detailed proof to the full version of this paper.

Since no message is sent to \mathcal{A}^* in the Initialization phase, adding Initialization has no impact on simulation. Allowing for sub-sessions (identified by *ssid*) results in adding *ssid* in the simulator whenever appropriate. The impact of changing SNDERCOMPLETE messages as sent from \mathcal{Z} , instead of from \mathcal{A}^* , is that no such messages are sent from SIM any more in steps 6 and 7; however, this does not influence the reduction that $\Pr[\text{HALT}]$ is negligible, since

Components: Hash functions $H(\cdot, \cdot)$, $H'(\cdot)$ with ranges $\{0, 1\}^\ell$ and G , respectively.

Initialization

- On input (INIT, sid, x) , \mathcal{S} picks $k \leftarrow_{\mathcal{R}} \mathbb{Z}_q$ and outputs $(\text{INIT}, sid, H(x, H'(x)^k))$.

Evaluation

- On input $(\text{EVAL}, sid, ssid, \mathcal{S}, x)$, \mathcal{U} proceeds as follows:
 - If there is a record $\langle \mathcal{S}, x, r, y \rangle$, outputs $(\text{EVAL}, sid, ssid, y)$ to \mathcal{Z} .
 - Else if there is a record $\langle \mathcal{S}', x, r, y \rangle$ (where $\mathcal{S}' \neq \mathcal{S}$), sends $a := H'(x)^r$ to \mathcal{S} .
 - Else picks $r \leftarrow_{\mathcal{R}} \mathbb{Z}_q$, records $\langle \mathcal{S}, x, r, \cdot \rangle$ and sends $a := H'(x)^r$ to \mathcal{S} .
- On input $(\text{SNDRCOMPLETE}, sid, ssid)$ and a from \mathcal{U} , \mathcal{S} sends $b := a^k$ to \mathcal{U} .
- On b from \mathcal{S} , if this is the first such message for $ssid$, \mathcal{U} retrieves record $\langle \mathcal{S}, x, r, \cdot \rangle$, replaces \cdot with $y := H_2(x, b^{1/r})$ and outputs $(\text{EVAL}, sid, ssid, y)$.

Fig. 9: Protocol DH-OPRF (for PRF output length ℓ)

the only SNDRCOMPLETE messages which count are those in step 5, which are still there (the only difference is that their issuers become \mathcal{Z} instead of \mathcal{SIM} , but they still have the effect of increasing the tx value).

The remaining change is that \mathcal{A} may compromise a server (for a specific sid) at any time; after that, \mathcal{A} can compute the server's function value on any valid input. \mathcal{SIM} is able to simulate this by sending OFFLINEEVAL messages to \mathcal{F} . Furthermore, note that HALT may only occur on servers who is not marked COMPROMISED at that time; therefore, the argument upper-bounding $\Pr[\text{HALT}]$ (in the setting where a server cannot be compromised) is not influenced.

1. Pick $r_1, \dots, r_N \leftarrow_{\mathbb{R}} \mathbb{Z}_m$, and compute $g_1 := g^{r_1}, \dots, g_N := g^{r_N}$. Record $(r_1, g_1), \dots, (r_N, g_N)$. Set counter $J := 1$.
2. Every time when there is a fresh query x to $H'(\cdot)$, answer it with g_J and record (x, r_J, g_J) . Finally, set $J++$.
3. On $(\text{COMPROMISE}, \text{sid})$ from \mathcal{A} as a message to \mathbf{S} , mark \mathbf{S} COMPROMISED, send $(\text{COMPROMISE}, \text{sid})$ to \mathcal{F} and do:
 - If there is no record $\langle \mathbf{S}, k, z \rangle$, pick $k \leftarrow_{\mathbb{R}} \mathbb{Z}_q$, compute $z := g^k$, record $\langle \mathbf{S}, k, z \rangle$ and send k to \mathcal{A} as \mathbf{S} 's response.
 - Else retrieve $\langle \mathbf{S}, k, z \rangle$ and send k to \mathcal{A} as \mathbf{S} 's response.
4. On $(\text{EVAL}, \text{sid}, \text{ssid}, \mathbf{U}, \mathbf{S})$ from \mathcal{F} , send g_J to \mathcal{A} as \mathbf{U} 's message to \mathbf{S} and record $\langle \text{ssid}, \mathbf{U}, \mathbf{S}, r_J, g_J \rangle$. Finally, set $J++$.
5. On $(\text{SNDRCOMPLETE}, \text{sid}, \text{ssid}, \mathbf{S})$ from \mathcal{F} and a from \mathcal{A} as some user \mathbf{U} 's message to \mathbf{S} , do:
 - If there is no record $\langle \mathbf{S}, k, z \rangle$, pick $k \leftarrow_{\mathbb{R}} \mathbb{Z}_q$, compute $z := g^k$, record $\langle \mathbf{S}, k, z \rangle$ and send a^k to \mathcal{A} as \mathbf{S} 's response to \mathbf{U} .
 - Else retrieve $\langle \mathbf{S}, k, z \rangle$ and send a^k to \mathcal{A} as \mathbf{S} 's response to \mathbf{U} .
6. On b from \mathcal{A} as some server \mathbf{S} 's message to a user \mathbf{U} , retrieve record $\langle \text{ssid}, \mathbf{U}, \cdot, r_j, g_j \rangle$ and do:
 - If there is a record $\langle \mathbf{S}', \cdot, z \rangle$ such that $b^{1/r_j} = z$, send $(\text{RCVCOMPLETE}, \text{sid}, \text{ssid}, \mathbf{S}')$ to \mathcal{F} .
 - Else create a new server \mathbf{S}' , record $\langle \mathbf{S}', \cdot, b^{1/r_j} \rangle$ and send $(\text{RCVCOMPLETE}, \text{sid}, \text{ssid}, \mathbf{S}')$ to \mathcal{F} .
7. Every time when there is a fresh query (x, u) to $H(\cdot, \cdot)$, do:
 - (a) If there is a record (x, r_j, g_j) , do:
 - (1) If there is a record $\langle \mathbf{S}, k, z \rangle$ such that $u^{1/r_j} = z$ and \mathbf{S} is marked COMPROMISED, send $(\text{OFFLINEEVAL}, \text{sid}, \mathbf{S}, x)$ to \mathcal{F} .
On \mathcal{F} 's response $(\text{OFFLINEEVAL}, \text{sid}, y)$, set $H(x, u) := y$.
 - (2) Else if there is a record $\langle \mathbf{S}, \cdot, z \rangle$ such that $u = z^{r_j}$ and \mathbf{S} is not marked COMPROMISED, send $(\text{EVAL}, \text{sid}, \text{ssid}, \mathbf{S}, x)$ and then $(\text{RCVCOMPLETE}, \text{sid}, \text{ssid}, \mathbf{S})$ to \mathcal{F} .
If \mathcal{F} ignores this message, output HALT and abort.
Otherwise on \mathcal{F} 's response $(\text{EVAL}, \text{sid}, \text{ssid}, y)$, set $H(x, u) := y$.
 - (b) In any other case, set $H(x, u) \leftarrow_{\mathbb{R}} \{0, 1\}^l$.

Fig. 10: The Simulator SIM for the DH-OPRF Protocol ($\mathcal{F}_{\text{OPRF}}$ Denoted \mathcal{F})