

Password-Protected Key Retrieval with(out) HSM Protection*

Sebastian Faller
sebastian.faller@ibm.com
IBM Research Europe / ETH Zurich
Zurich, Switzerland

Tobias Handirk
tobias.handirk@uni-wuppertal.de
Bergische Universität Wuppertal
Wuppertal, Germany

Julia Hesse
juliahesse2@gmail.com
IBM Research Europe
Zurich, Switzerland

Máté Horváth
horvath@uni-wuppertal.de
Bergische Universität Wuppertal
Wuppertal, Germany

Anja Lehmann
anja.lehmann@hpi.de
Hasso Plattner Institute
Potsdam, Germany

Abstract

Password-protected key retrieval (PPKR) enables users to store and retrieve high-entropy keys from a server securely. The process is bootstrapped from a human-memorizable password only, addressing the challenge of how end-users can manage cryptographic key material. The core security requirement is protection against a corrupt server, which should not be able to learn the key or offline-attack it through the password protection. PPKR is deployed at a large scale with the WhatsApp Backup Protocol (WBP), allowing users to access their encrypted messaging history when switching to a new device. Davies et al. (Crypto’23) formally analyzed the WBP, proving that it satisfies most of the desired security. The WBP uses the OPAQUE protocol for password-based key exchange as a building block and relies on the server using a hardware security module (HSM) for most of its protection. In fact, the security analysis assumes that the HSM is incorruptible – rendering most of the heavy cryptography in the WBP obsolete.

In this work, we explore how provably secure and efficient PPKR can be built that either relies strongly on an HSM – but then takes full advantage of that – or requires less trust assumption for the price of more advanced cryptography. To this end, we expand the definitional work by Davies et al. to allow the analysis of PPKR with *fine-grained HSM corruption*, such as leakage of user records or attestation keys. For each scenario, we aim to give *minimal* PPKR solutions. For the strongest corruption setting, namely a fully corrupted HSM, we propose a protocol with a simpler design and better efficiency than the WBP. We also fix an attack related to client authentication that was identified by Davies et al.

1 Introduction

Secure messaging apps like Signal and WhatsApp have made end-to-end encrypted (E2EE) communication accessible to millions of end-users. With the deployment of E2EE on user-managed devices, new challenges arise. One is how to ensure that users maintain access to their data when they lose their devices or switch to a new phone. Protocols aiming at the secure backup of a user’s chat history – or a user’s cryptographic key in general – have been proposed by Signal [Lun19], Apple [Krs16], Google [Wal18] and WhatsApp [Wha21]. The common idea in all approaches is to back

up the user’s high-entropy (encryption) key on a server and enable retrieval when the user correctly authenticates via a human-memorizable secret, like a password or PIN. In particular, when the backup service is offered by the same provider that handles the E2EE communication, extra care is necessary to not undermine the encryption security. This is done by relying on trusted hardware enclaves, such as hardware security modules (HSMs) that protect the user’s key material and ensure protection even when the server that runs the HSM gets corrupted.

From the aforementioned approaches, the WhatsApp Backup protocol (WBP) has enjoyed the most attention due to its enhanced protocol design and widespread usage. The WBP deploys OPAQUE [JKX18], a strong asymmetric password-based key exchange protocol (saPAKE) [BM93, BMP00, GMR06], to ensure that no offline-attackable information is leaked as part of the protocol and security can be guaranteed even in the presence of a malicious server.

Password-Protected Key Retrieval. The security of WBP has been formally analyzed by Davies et al. [DFG⁺23b]. Therein the problem is abstracted as *password-protected key retrieval* (PPKR) and the desirable security guarantees are expressed through an ideal functionality in the universal composability (UC) [Can01] model. A secure PPKR guarantees a multitude of security properties against malicious users and a corrupt server. For the sake of simplicity, we focus on the properties most relevant for our work in the remainder of our paper and refer to [DFG⁺23b] for a full exposition of PPKR. The core properties (relevant to our work) are as follows:

Secret & Pseudorandom key: Honest users create pseudorandom keys, even if the server acts maliciously. Further, the key remains hidden from the server, as long as it does not correctly guess the password.

Security against Offline and Online Attacks: The user’s key is protected through a human-memorizable password, which naturally has low entropy. Thus, it is crucial that the protocol ensures security against offline attacks on the user’s password. In fact, PPKR even enforces an upper bound (set to 10 in WBP and [DFG⁺23b]) that strictly limits the number of wrong password guesses through an online attack. If the limit has been reached through consecutive wrong retrieval attempts, the user’s key gets securely deleted.

Authenticated Retrieval: As the user might use the key recovered via PPKR for further encryption, or other cryptographic operations, she should not be tricked into accepting a wrong key from the retrieval.

* An extended abstract [FHH⁺24] of this article appears in the proceedings of the 31st ACM Conference on Computer and Communications Security (ACM CCS 2024). This is the full version.

Davies et al. [DFG⁺23b] have proven that the original WBP securely instantiates (most of) such an ideal PPKR and provides the desired strong security even when the server is corrupted – assuming that the core cryptographic parts of the server are handled by a trusted HSM.

The role of the HSM. The WBP protocol and analysis crucially rely on an HSM that is run within the server as an incorruptible entity. This was manifested in the ideal functionality of PPKR [DFG⁺23b] that maintains all its essential security guarantees even if the server is corrupt. (We will refer to this security level as Lev-1). Consequently, the WBP protocol has been proven secure under the assumption that the HSM is fully trusted, i.e., all keys used for the internally run OPAQUE protocol, as well as all user-specific files must be protected through the HSM.

While this correctly captures the initial design choice made by the WBP protocol, this is somewhat unsatisfactory from a security and protocol design perspective. First, if an incorruptible HSM can be assumed, the protocol could take more advantage of that – the core of the WBP protocol is the OPAQUE protocol that provides strong security guarantees even when the cryptographic state gets compromised, which wouldn’t be needed if the assumption is that such an event can never occur. Second, relying on a perfectly secure sub-entity is a risky assumption. In fact, also trusted hardware modules have a history of getting breached or having to lower their security claims [VBMW⁺18, VBPS17, BC19, SRW22]. Thus, it would be desirable to clearly express and analyze the impact a partial or full corruption of the HSM has on the expected security guarantees of PPKR.

1.1 Our Contributions

In our work we revisit the design of password-protected key retrieval (PPKR), focusing on the impact of HSM-assurance. We first extend the security model of [DFG⁺23b] to include more corruption capabilities and express their impact on the expected security guarantees. Second, we present three provably secure PPKR protocols, each having a dedicated trade-off between the strength of the HSM assumption and simplicity. In more detail, we provide the following.

Security Model with Fine-Grained (HSM) Corruption. We extend the ideal functionality of [DFG⁺23b] to include two further corruption settings: The first is *file leakage*, expressing that the user-specific files that are stored on the server’s side to provide the password-authenticated key retrieval can be leaked (repeatedly) to the adversary. We consider this as a partial corruption, and give the dynamically created key files to the adversary but still assume the HSM to be trusted in its core. In practice, an HSM’s main task is secure attestation which we consider to be unaffected by this first type of attack. That is, the HSM only lost its protected files but still behaves honestly. The second setting of *full corruption* then models that all security of the HSM is compromised. This translates to an adversary getting access to all keys and files of the HSM, and from then on can deviate arbitrarily from the honest protocol execution.

Despite these strong corruption settings, we aim at a *graceful degradation* of the overall security properties. As in [DFG⁺23b], our security definition ensures that no offline attacks on the password and protected key are possible, even when the server, i.e., the entity

controlling the interface to the HSM, is corrupt. When a user’s file gets leaked, this security guarantee is lost, and offline attacks are inevitable – but are also still necessary. That is, the adversary must offline-attack the password of every key it wants to learn, and strong user passwords will still provide a line of defense. The impact of file leakage must also be strictly contained to only the users whose files got compromised. Further, even for users whose key file got lost, the security of active recovery sessions should remain, meaning that the authenticity of the recovered keys must still hold. Upon full corruption of the server, i.e., a complete break of the HSM, this property can no longer be guaranteed but is not fully lost either: it again requires the adversary to correctly guess the user’s (login) password to make her accept an incorrect key.

We refer to the optimal security guarantees under server corruption and file leakage as Lev-2 security, and to Lev-3 security for the guarantees under full server corruption.

Table 1: Security achieved by our PPKR protocols and WBP. ✓= achieves optimal protection in that corruption setting, ✗= not secure, where Lev-1: server corrupt, Lev-2: server corrupt + HSM leaks, Lev-3: server (and HSM) fully corrupt. Gray are conjectures. WBP was proven secure in a model that is slightly weaker than our Lev-1 model, denoted by ✓*, with similar relaxations likely needed for the higher levels, too. The last column indicates if the protocol relies on standard building blocks only.

	Lev-1	Lev-2	Lev-3	Standard BB
encPw, Fig. 3	✓	✗	✗	✓
encPw+, Fig. 3	✓	✓	✗	✓
OPRF-PPKR, Fig. 4	✓	✓	✓	✗
WBP [DFG ⁺ 23b]	✓*	✓*	✓*	✗

Simple Protocols for Different HSM-Levels. Having a security model that explicitly captures different levels of server corruption now allows to integrate internal HSM-protection more thoroughly in the protocol design. We propose three PPKR protocols, all relying on a server-internal HSM to protect the user’s key files, but each aiming at a different trust level. We present two protocols that provide the same (or even stronger) *formally proven* security than WBP, yet are significantly simpler. Both fall short of achieving Lev-3 security, for which we design a protocol that uses an oblivious pseudorandom function (OPRF) as crucial building blocks, giving up on the reliance on basic primitives only.

Basic/enhanced encrypt-to-HSM: We start with a protocol that fully relies on a trusted HSM for most of its operation, and leverage this to simplify the design as much as possible. The resulting protocol merely uses standard symmetric and asymmetric encryption. We prove our simple protocol to be secure on Lev-1, which gives the same (or even slightly stronger, see discussion below) security guarantees as had been formally proven for WBP in [DFG⁺23b]. We further show how to upgrade the protocol to Lev-2 security, by relying on fresh encryption keys and salted hashing in the stored password files. The simplicity hinges on the trust in the HSM though, as the protocol loses its security if the HSM gets fully corrupted. Interestingly, the best attacks against all long-term user files are still offline attacks, but the

login passwords that users send during an active retrieval session are now revealed in plain when the HSM is fully corrupted, which violates the Lev-3 guarantees.

OPRF-based: Our third protocol provides Lev-3 security, i.e., guarantees optimal protection even in the presence of full server/HSM corruption. This protocol partly resembles WBP, as it also relies on an OPRF as a core primitive. Recall that WBP relies on OPAQUE, an asymmetric PAKE that is built from an OPRF. By building our protocol directly from OPRFs and the clearly specified security guarantees as concrete target, our protocol enjoys a much cleaner design. In particular, we can omit all OPAQUE parts that are not needed for PPKR (which does not aim at fresh session keys as aPAKE). Overall, we propose an OPRF-based PPKR that has a simpler protocol design and better efficiency than WBP, and that fixes several attacks related to user authentication that were identified by Davies et al. [DFG⁺23b].

Strengthening Security. We also use the opportunity to remove unwanted (in)security artifacts that were present in the WBP and had to be included in the ideal functionality for a provable security treatment of that real-world protocol. These artifacts comprise the re-routing of logins to different files, whenever the server is corrupt. All our three protocols prevent that attack. In the Lev-2/3 setting, WBP has further undesirable – albeit not fatal – behavior, such as potentially leaking the information whether several honest users have the same password. Our protocols for these levels also improve on these shortcomings.

1.2 Related Works

The closest work to ours is the work by Davies et al. [DFG⁺23b], which we extend with a more fine-grained analysis of the HSM corruption settings. As we now distinguish between the "front-end" corruption of the server's networking part and further internal corruption (of the HSM), this might yield the question whether this is just a special 2-party version of the more general concept of *password-protected (t, n)-secret sharing* (PPSS) [BJS11, JKKX16, CLLN14]. In PPSS, the user can store a password-protected secret (such as the backup key) with the help of n servers. The user can later recover the secret if she provides the correct password to at least t servers. Secure PPSS solutions ensure security against offline attacks on the user's password as long as less than the threshold gets corrupted. While a (2, 2) version of PPSS is similar in spirit to PPKR, the main difference is in the overall setting: PPKR captures the single-server environment that is currently deployed by WhatsApp, and envisioned by other major E2EE providers too [Lun19, Krs16, Wal18]. Therein the user communicates with a single server only, which might run internal measures (such as an HSM) to enhance security. In contrast, PPSS requires the user to individually talk to all n parties, which must also run the equivalent code. In our HSM-supported setup, such direct communication between the user and the HSM is not desirable, and the server's networking part and HSM are clearly distinct. Thus, while PPSS appears more general, it also excludes the most common real-world deployment setting that we are targeting directly with PPKR.

The dedicated one-server setting of PPKR also distinguishes it from *password-protected encryption* protocols [AMMR18, WH20, CGMS21, DHL22] which work in a multi-server setting with n

independent parties too. Further, these works target symmetric encryption directly, and not key-recovery that can be used for any key and subsequent cryptographic usage of it.

The recent works on Credential-less Secret Recovery [Sca19, OSV23] are for a somewhat similar single-server setting, as PPKR. Therein, a user stores a secret on some cloud storage and uses the additional power of a trusted execution environment (TEE) for secure recovery. In contrast to PPKR however, [OSV23] relies on a publicly accessible blockchain instead of passwords to authenticate users.

2 Preliminaries

In our protocols, we make use of standard cryptographic primitives, namely strongly EUF-CMA secure digital signatures (Sig), and IND-CPA and IND-CCA secure public key encryption (PKE), see e.g., [BS23]. We will also use *authenticated encryption* (AE), i.e., symmetric encryption (SE) that is IND-CPA secure and has ciphertext integrity in the following sense:

Definition 1 (Ciphertext integrity). The advantage of an adversary \mathcal{A} against the *integrity of ciphertexts* (INT-CTXT) of a symmetric encryption scheme $SE = (\text{KeyGen}, \text{Enc}, \text{Dec})$ is defined as

$$\text{Adv}_{\mathcal{A}, SE}^{\text{INT-CTXT}}(\lambda) := \Pr[\text{Dec}(k, c^*) \neq \perp : c^* \xleftarrow{\$} \mathcal{A}^{\text{Enc}(k, \cdot)}(1^\lambda)],$$

where $k \xleftarrow{\$} \text{KeyGen}(1^\lambda)$ and c^* is fresh in the sense that it has never been output by the encryption oracle $\text{Enc}(k, \cdot)$.

For simplicity, we will assume that KeyGen samples $k \xleftarrow{\$} \{0, 1\}^\lambda$. We define two additional properties of an authenticated encryption scheme that will be required in this work, namely *equivocability* and *random-key robustness*. Equivocability means that a simulator is able to produce ciphertexts of the scheme without committing to the plaintext. If the ciphertext must be decrypted then the simulator can output an appropriate key to open the ciphertext to a message of choice. Random-key robustness means that an adversary cannot come up with a ciphertext that simultaneously decrypts under *two* keys if the keys are chosen uniformly at random. According to [JKX18], this can be instantiated in the standard model with encrypt-then-MAC using HMAC or in the ROM by adding $H(k, c)$ to an authenticated ciphertext c .

Definition 2. A symmetric encryption scheme SE is *equivocable* if for any efficient algorithm \mathcal{A} , there exists an efficient stateful simulator $\text{SIM}_{EQV}^{\mathcal{A}, SE}$ s.t. the distinguishing advantage $\text{Adv}_{\mathcal{A}, SE}^{\text{EQV}}$ of \mathcal{A} 's view in the following two games is a negligible function in λ :

- The real game: \mathcal{A} outputs message m and computes its final output given (c, k) , where $k \xleftarrow{\$} \{0, 1\}^\lambda$ and $c \xleftarrow{\$} \text{Enc}(k, m)$.
- The ideal game: \mathcal{A} outputs message m and computes its final output given (c, k) , where $c \xleftarrow{\$} \text{SIM}_{EQV}(|m|)$ and $k \xleftarrow{\$} \text{SIM}_{EQV}(c, m)$.

Definition 3. A symmetric encryption scheme SE is *random-key robust* if for any efficient adversary \mathcal{A}

$$\text{Adv}_{\mathcal{A}, SE}^{\text{rkr}}(\lambda) := \Pr_{k, k' \xleftarrow{\$} \{0, 1\}^\lambda} \left[\text{Dec}(k, c) \neq \perp \wedge \text{Dec}(k', c) \neq \perp : c \xleftarrow{\$} \mathcal{A}(k, k') \right]$$

is a negligible function of λ .

$\mathcal{F}_{\text{PPKR}}$ is parameterized with a security parameter λ . $\mathcal{F}_{\text{PPKR}}$ talks to a server S where S is encoded in sid . S can have corruption state `HONEST`, `CORRUPT`, or `FULLYCORRUPT`. $\mathcal{F}_{\text{PPKR}}$ further allows file leakage via the `LEAKFILE` corruption command. $\mathcal{F}_{\text{PPKR}}$ also talks to the adversary \mathcal{A} , and arbitrary clients ID_C . If the functionality tries to retrieve a record that does not exist, it ignores the incoming message. We write $\text{tx}[\cdot]$ for a list of counters. For brevity, we omit session identifier sid from all inputs, outputs, and records. $\mathcal{F}_{\text{PPKR}}$ ignores repeated inputs with the same ssid and in that case activates \mathcal{A} .

Corruption interfaces

On corruption command (`LEAKFILE`) from \mathcal{A} :

- LS.1 Create an empty list ID_C .
- LS.2 For every record $\langle \text{FILE}, [\text{ID}_C], [\text{pw}], [K], [\text{ctr}], * \rangle$:
 - (1) Append $\langle \text{ID}_C, \text{ctr} \rangle$ to ID_C .
 - (2) If no record $\langle \text{LEAKED}, \text{ID}_C, *, *, * \rangle$ exists, add a record $\langle \text{LEAKED}, \text{ID}_C, \text{pw}, K, 1 \rangle$.
 - (3) Otherwise, add a record $\langle \text{LEAKED}, \text{ID}_C, \text{pw}, K, i+1 \rangle$, where $i \in \mathbb{N}$ is the largest number, such that a record $\langle \text{LEAKED}, \text{ID}_C, *, *, i \rangle$ exists.
- LS.3 Output ID_C to \mathcal{A} .

On corruption command (`CORRUPT`, P) from \mathcal{A} :

- C.1 Mark P as `CORRUPT`.

On corruption command (`FULLYCORRUPT`, S) from \mathcal{A} :

- FC.1 Mark S as `FULLYCORRUPT` and run LS.1-LS.2
- FC.2 For each $\langle \text{ID}_C, \cdot \rangle \in \text{ID}_C$, set $\text{tx}[\text{ID}_C] \leftarrow i$, where i is the largest $i \in \mathbb{N}$ such that a record $\langle \text{LEAKED}, \text{ID}_C, \text{pw}, K, i \rangle$ exists.
- FC.3 In every record $\langle \text{FILE}, *, *, *, [\text{ctr}], * \rangle$ overwrite ctr with ∞ .
- FC.4 Output ID_C to \mathcal{A} .

Initialization phase

On input (`INITC`, ssid , pw) from ID_C : // Client ID_C starts initialization of a password-protected key

- IC.1 Choose $K \xleftarrow{\$} \{0, 1\}^\lambda$
- IC.2 If a record $\langle \text{INIT}, \text{ssid}, \text{ID}_C, \perp, \perp, \text{srvOk} \rangle$ exists, overwrite $\langle \perp, \perp \rangle$ in it with $\langle \text{pw}, K \rangle$
- IC.3 Otherwise record $\langle \text{INIT}, \text{ssid}, \text{ID}_C, \text{pw}, K, \perp \rangle$
- IC.4 Send $\langle \text{INITC}, \text{ssid}, \text{ID}_C \rangle$ to \mathcal{A}

On input (`INITS`, ssid , ID_C) from S : // Server agrees to initialize ID_C

- IS.1 If a record $\langle \text{INIT}, \text{ssid}, \text{ID}_C, *, *, \perp \rangle$ exists, overwrite \perp with `srvOk`
- IS.2 Otherwise record $\langle \text{INIT}, \text{ssid}, \text{ID}_C, \perp, \perp, \text{srvOk} \rangle$
- IS.3 Send $\langle \text{INITS}, \text{ssid}, \text{ID}_C \rangle$ to \mathcal{A}

On input (`COMPLETEINITC`, ssid , b_C) from \mathcal{A} : // Complete initialization for client

- CIC.1 Retrieve $\langle \text{INIT}, \text{ssid}, [\text{ID}_C], *, [K], \text{srvOk} \rangle$ and drop the query if $K = \perp$
- CIC.2 If $b_C = 1$, output $\langle \text{INITRES}, \text{ssid}, K \rangle$ to ID_C
- CIC.3 If $b_C = 0$, output $\langle \text{INITRES}, \text{ssid}, \text{FAIL} \rangle$ to ID_C

On input (`COMPLETEINITS`, ssid , b_S) from \mathcal{A} : // Complete initialization for server

- CIS.1 Retrieve $\langle \text{INIT}, \text{ssid}, [\text{ID}_C], [\text{pw}], [K], \text{srvOk} \rangle$ and drop the query if $K = \perp$
- CIS.2 If $b_S = 0$, output $\langle \text{INITRES}, \text{ssid}, \text{ID}_C, \text{FAIL} \rangle$ to S . Else continue.
- CIS.3 Store record $\langle \text{FILE}, \text{ID}_C, \text{pw}, K, \text{ctr}, [\text{HONEST}] \rangle$ if ID_C is honest and $\langle \text{FILE}, \text{ID}_C, \text{pw}, K, \text{ctr}, [\text{MALICIOUS}] \rangle$ otherwise, where $\text{ctr} \leftarrow \infty$ if S is `FULLYCORRUPT`, and $\text{ctr} \leftarrow 10$ else, overwriting any existing record $\langle \text{FILE}, \text{ID}_C, *, *, *, * \rangle$
- CIS.4 If S is `FULLYCORRUPT`, record $\langle \text{LEAKED}, \text{ID}_C, \text{pw}, K, i+1 \rangle$, where $i \in \mathbb{N}$ is the largest number, such that a record $\langle \text{LEAKED}, \text{ID}_C, *, *, i \rangle$ exists, or $i = 1$ if no such record exists. If $\text{tx}[\text{ID}_C]$ is undefined set $\text{tx}[\text{ID}_C] \leftarrow 1$.
- CIS.5 Send $\langle \text{INITRES}, \text{ssid}, \text{Succ} \rangle$ to S

Figure 1: Ideal functionality $\mathcal{F}_{\text{PPKR}}$ for password-protected key retrieval, offline attacks, and initialization interfaces. The boxed code can be dropped to strengthen $\mathcal{F}_{\text{PPKR}}$ (see Section 3.2).

3 Security model

Our ideal functionality is based on the definition from Davies et al. [DFG⁺23b], and models how a user – identified through a client identifier ID_C – can store and retrieve a strong cryptographic key

Recovery Phase

On input $\langle \text{REC}_C, \text{ssid}, \text{pw}' \rangle$ from ID_C : // Client starts key recovery from password pw'

- RC.1 If a record $\langle \text{REC}, \text{ssid}, \text{ID}_C, \perp, [\text{pw}], * \rangle$ exists, overwrite \perp with pw' and set $\text{match} \leftarrow (\text{pw} \stackrel{?}{=} \text{pw}')$.
- RC.2 Otherwise record $\langle \text{REC}, \text{ssid}, \text{ID}_C, \text{pw}' \rangle$ and set $\text{match} \leftarrow \perp$
- RC.3 Send $\langle \text{REC}_C, \text{ssid}, \text{ID}_C, \text{match} \rangle$ to \mathcal{A}

On input $\langle \text{REC}_S, \text{ssid}, \text{ID}_C, \text{pw}^*, K^*, i \rangle$ from S : // Server starts recovery session for ID_C

- RS.1 Retrieve $\langle \text{REC}, \text{ssid}, \text{ID}_C, [\text{pw}'] \rangle$, otherwise set $\text{pw}' \leftarrow \perp$ and record $\langle \text{REC}, \text{ssid}, \text{ID}_C, \text{pw}' \rangle$
- RS.2 If S is `FULLYCORRUPT` and $\text{pw}^* \neq \perp$: set $\overline{\text{pw}} \leftarrow \text{pw}^*$, $\overline{K} \leftarrow K^*$
- RS.3 If S is `FULLYCORRUPT`, $\text{pw}^* = \perp$, and $i \geq \text{tx}[\text{ID}_C]$, retrieve the record $\langle \text{LEAKED}, \text{ID}_C, [\text{pw}], [K], i \rangle$ and set $\text{ctr} \leftarrow \infty$. Otherwise, retrieve the record $\langle \text{FILE}, \text{ID}_C, [\text{pw}], [K], [\text{ctr}], * \rangle$. Set $\overline{\text{pw}}, \overline{K}$ as follows:
 - (1) If no such record exists, set $\overline{\text{pw}} \leftarrow \perp, \overline{K} \leftarrow \text{FAIL}$.
 - (2) If $\text{ctr} = 0$, set $\overline{\text{pw}} \leftarrow \perp, \overline{K} \leftarrow \text{DELREC}$ and delete $\langle \text{FILE}, \text{ID}_C, *, *, *, * \rangle$.
 - (3) Else, set $\overline{\text{pw}} \leftarrow \text{pw}, \overline{K} \leftarrow K$ and overwrite ctr with $\text{ctr} - 1$ in the `FILE` record.
- RS.4 Append $\overline{\text{pw}}$ and \overline{K} to record $\langle \text{REC}, \text{ssid}, \text{ID}_C, \text{pw}' \rangle$
- RS.5 If $\text{pw}' = \perp$, set $\text{match} \leftarrow \perp$, else set $\text{match} \leftarrow (\overline{\text{pw}} \stackrel{?}{=} \text{pw}')$
- RS.6 Send $\langle \text{REC}_S, \text{ssid}, \text{ID}_C, \text{match} \rangle$ to \mathcal{A} .

On input $\langle \text{COMPLETEREC}_C, \text{ssid}, b_C \rangle$ from \mathcal{A} : // Complete recovery for client

- CRC.1 Retrieve record $\langle \text{REC}, \text{ssid}, [\text{ID}_C], [\text{pw}'], [\text{pw}], [K] \rangle$ and drop the query if $\text{pw}' = \perp$.
- CRC.2 Determine the output K' as follows:
 - (1) If $K \in \{\text{FAIL}, \text{DELREC}\}$, set $K' \leftarrow K$
 - (2) If $\text{pw} = \text{pw}'$ and $b_C = 1$, set $K' \leftarrow K$
 - (3) In all other cases, set $K' \leftarrow \text{FAIL}$
- CRC.3 Send $\langle \text{RECRES}, \text{ssid}, K' \rangle$ to ID_C .

On input $\langle \text{COMPLETEREC}_S, \text{ssid}, b_S \rangle$ from \mathcal{A} : // Complete recovery for server

- CRS.1 Retrieve record $\langle \text{REC}, \text{ssid}, [\text{ID}_C], [\text{pw}'], [\text{pw}], [K] \rangle$ and drop the query if $\text{pw}' = \perp$.
- CRS.2 Determine the output result as follows:
 - (1) If $K \in \{\text{FAIL}, \text{DELREC}\}$, set $\text{result} \leftarrow K$.
 - (2) If $b_S = 1$ and either $\text{pw} = \text{pw}'$ or a record $\langle \text{FILE}, \text{ID}_C, *, *, *, \text{MALICIOUS} \rangle$ exists, set $\text{result} \leftarrow \text{Succ}$.
 - (3) In all other cases, set $\text{result} \leftarrow \text{FAIL}$.
- CRS.3 If $\text{result} = \text{Succ}$ and there exists a record $\langle \text{FILE}, \text{ID}_C, \text{pw}, K, [\text{ctr}], * \rangle$, overwrite ctr with 10 if S is not `FULLYCORRUPT` or with ∞ if S is `FULLYCORRUPT`.
- CRS.4 Send $\langle \text{RECRES}, \text{ssid}, \text{result} \rangle$ to S .

Attack interfaces

On input $\langle \text{MALICIOUSINIT}, \text{ID}_C, \text{pw}^*, K^* \rangle$ from \mathcal{A} : // Malicious initialization with adversarial password pw^* and key K^*

- MI.1 If S is `HONEST`, ignore this input.
- MI.2 Record $\langle \text{FILE}, \text{ID}_C, \text{pw}^*, K^*, 10, [\text{MALICIOUS}] \rangle$, overwriting any existing record $\langle \text{FILE}, \text{ID}_C, *, *, *, * \rangle$

On input $\langle \text{MALICIOUSREC}, \text{ID}_C, \text{pw}^* \rangle$ from \mathcal{A} : // Malicious server guesses ID_C 's password, subject to non-zero file counters

- MR.1 If S is `HONEST`, ignore this input.
- MR.2 If no record $\langle \text{FILE}, \text{ID}_C, [\text{pw}], [K], [\text{ctr}], * \rangle$ exists, output `FAIL` to \mathcal{A} .
- MR.3 If $\text{ctr} = 0$, delete record $\langle \text{FILE}, \text{ID}_C, \text{pw}, K, \text{ctr}, * \rangle$ and output $\langle \text{DELREC}, \text{ID}_C \rangle$ to \mathcal{A}
- MR.4 If $\text{pw}^* = \text{pw}$, overwrite ctr in the record with 10 and output K to \mathcal{A} . Otherwise, overwrite ctr with $\text{ctr} - 1$ and output `FAIL` to \mathcal{A}

On input $\langle \text{OFFLINEATTACK}, \text{ID}_C, \text{pw}^*, i \rangle$ from \mathcal{A} : // Offline attack on leaked key files - unlimited number of guesses allowed

- OA.1 If a record $\langle \text{LEAKED}, \text{ID}_C, \text{pw}^*, [K], i \rangle$ exists, output K to \mathcal{A} . Otherwise, output `FAIL` to \mathcal{A} .

Figure 2: Ideal functionality $\mathcal{F}_{\text{PPKR}}$, recovery interfaces. The boxed code can be dropped to strengthen $\mathcal{F}_{\text{PPKR}}$ (see Section 3.2).

K with the help of an external server. The server S is encoded in the session id sid and is the central entity that interacts with many clients. At initialization, the client ID_C creates a random key K that it stores with the server, protected under a password pw . Later, the client can recover the key K from the server, but only if it correctly authenticates with the same password again.

The definition of Davies et al. [DFG⁺23b] provides very strong guarantees in case of server compromise: even if the server is corrupt, the client’s key and password remain secure unless the adversary correctly guesses the password in a few attempts. In the realization, this essentially requires the server to internally run an *incorruptible* entity, like an HSM. The HSM is then responsible to protect the client’s key and enforce strict access control to it.

In our work, we model stronger types of server corruption to avoid the need to assume an incorruptible sub-part in the constructions. Beyond the server compromise from [DFG⁺23b], our model captures *file leakage* of the stored client files and a *fully corrupt* server. The latter translates to a corruption of all parts of the server, including an internally hosted HSM. Beyond stronger corruption settings, we further make the definition more general to allow for a broader set of protocol flows and also remove some artifacts that had to be included in [DFG⁺23b] to capture the security of the concrete WhatsApp protocol. For an overview and explanation of these changes, we refer to Appendix A.

3.1 Ideal PPKR Functionality

We start by recalling the main functional interfaces of $\mathcal{F}_{\text{PPKR}}$ (when all parties are honest) and then explain the different corruption settings and their impact on the guaranteed security properties.

Initialization. To use the key recovery service, the client ID_C and server S must first engage to create a password-protected account. This can be initiated by either party through the INITC and INITS interfaces, respectively. The client’s interface takes the user’s password pw as input and lets $\mathcal{F}_{\text{PPKR}}$ create a random key K that gets stored along with ID_C and pw . If completion towards both parties is signaled by the adversary through the COMPLETEINITC and COMPLETEINITS interfaces, a key file $\langle \text{FILE}, sid, ID_C, pw, K, ctr \rangle$ is stored in the functionality. The counter is set to $ctr \leftarrow 10$, which determines how many incorrect password recoveries will be tolerated before the file is deleted.

The initialization can be called repeatedly by the same client ID_C , which allows the renewal of the key and possibly the replacing of the password.

Recovery. If a client ID_C wishes to recover its key, it triggers the RECC interface for a password attempt pw' . When the server initiates a recovery for ID_C too, and a key file $\langle \text{FILE}, sid, ID_C, pw, K, ctr \rangle$ for ID_C is stored within $\mathcal{F}_{\text{PPKR}}$, the counter ctr is decremented by 1 and the adversary \mathcal{A} learns whether $pw \stackrel{?}{=} pw'$. This bit is considered unavoidable leakage in password-based protocols, as the observable protocol behavior usually strongly differs depending on whether the password authentication was successful or not. If the counter reaches 0, the file FILE for ID_C gets deleted, and the recovery “key” is set to DELREC .

If the adversary signals the completion of the recovery attempt towards the client via COMPLETERECC and the file hasn’t been

deleted, the client will receive the correct key K from its record if $pw = pw'$ or a failure notification FAIL otherwise. Likewise, if the adversary completes the session for ID_C towards the server via the COMPLETERECS interface, the server finally learns whether the password attempt was correct (or the file was deleted). Further, the counter ctr is reset to 10 when the password was correct and the session got completed.

Core Security Properties. In summary, the functionality $\mathcal{F}_{\text{PPKR}}$ ensures that only the legitimate client ID_C can initiate sessions to setup or recover a key (*implicit client-to-server authentication*), and all generated keys are fresh and chosen at random (*pseudorandomness of K*). The server S neither learns the password pw nor key K (*secrecy of pw and K*), only if a recovery was run on an incorrect password attempt pw' or not. If the password attempt was correct and the client recovered a key K , she can be ensured that it is the correct key (*authenticity of K*). An adversary has at most 10 guessing attempts against the client’s password (*no offline attacks, limited online attacks*). We give an overview of the core security properties of $\mathcal{F}_{\text{PPKR}}$ – and their degradation with corruption – in Table 2.

3.2 Modelling Server (and HSM) Corruption

The central entity in our system is the server S , which provides the password-protected key retrieval service to its clients, and which we model to be corruptible in several ways. Note that the HSM does not appear as an explicit entity in a PPKR functionality, but is rather considered an artifact on how the server’s code is deployed. The original PPKR definition [DFG⁺23b] required the server code to be split into a corruptible networking part, and an incorruptible cryptography part, e.g., deployed on a permanently secure HSM. In the protocol, the HSM was then responsible for performing the cryptographic operations and securely storing the users’ password-protected keys. Our model additionally allows fine-grained corruption of the cryptography part and distinguishes whether the (HSM-protected) key files are leaked, or even all of the additional protection is lost.

We start with a brief overview of the three corruption settings, and then explain their impact on the guaranteed security properties. For the sake of clarity, we explain how the additional corruption settings translate to an HSM-supported protocol setup.

Server Corrupt: This is the original corruption status of [DFG⁺23b], which gives a corrupt server some, rather benign, attack capabilities, but still maintains most of the security guarantees. In the construction, this requires that the server is mainly a connecting interface between the clients and a secure HSM, where the key material and client records remain protected during this corruption thanks to being executed in an HSM, that can securely attest all outgoing messages it sends (via the server) to the clients.

File Leakage: We additionally model that the client files, containing the information to verify the recovery password and retrieve the client’s key, can be leaked to the adversary. In the context of an HSM-supported setup, this means that the adaptively stored and maintained information by the HSM can get compromised. The file leakage can happen repeatedly, with the adversary obtaining snapshots of all the stored client files. This does neither

imply that the server nor HSM are corrupt though. In particular, the HSM still behaves honestly and its attestation remains secure.

We believe this realistically models that the core attestation (key) of the HSM enjoys particularly strong protection, whereas the HSM-protected database of possibly millions of client records is less secure and can be vulnerable to leakage attacks.

Server (and HSM) Fully Corrupt: We also want to capture the security (loss) when the server gets fully corrupted, meaning that all parts of the server – including the HSM – are under the control of the adversary. This implies the previous two corruption settings, i.e., normal server corruption and *continuous* compromise of all files, and also goes beyond as we now assume that the adversary is in full control of the HSM. That is, the attestation is no longer trusted, and the adversary can arbitrarily deviate from the original HSM protocol.

We now explain how each corruption impacts the security guarantees of $\mathcal{F}_{\text{PPKR}}$. Note that the status of server corruption and file leakage are independent of each other, i.e., they can occur separately or jointly. In the latter case, the adversary’s capabilities are simply the combination of both individual attacks.

Server Honest & No Leakage. If the server is honest and no files have been leaked, the functionality guarantees full protection of the clients’ information. In particular, only the legitimate client ID_C can initiate the recovery of a previously stored key. This ensures that the account and key of an honest client cannot be deleted by the adversary through repeated incorrect authentication attempts.

An initially honest client can get corrupted through the corruption command ($\text{CORRUPT}, \text{ID}_C$) though. In that case, the adversary gets up to 10 guesses on the client’s password and learns the key K when it provides the correct password.

An unavoidable attack even if all parties are honest, are network attacks, i.e., interrupting an honestly initiated session. This is modeled through the DoS-bit b_C and b_S , in the COMPLETEINITC and COMPLETEINITS interfaces for the initialization and the COMPLETEREC , COMPLETERECS interfaces for the recovery respectively. By providing $b_C = 0$ or $b_S = 0$, the adversary can make the session fail for either party. If it does so towards the server, this leads to a failed initialization, i.e., no file gets recorded in the functionality; or a failed recovery, i.e., the counter does not get reset to 10, even if the client provided the correct password.

Server Corrupt. To corrupt the server, the adversary sends the corruption command ($\text{CORRUPT}, S$). From then on, the functionality no longer enforces proper client authentication – which is the main impact of this corruption type. As a consequence, the adversary can now replace records of honest clients with a malicious password and key (through the MALICIOUSINIT interface). It still does not learn the passwords and keys of any honestly created or existing records but can try to recover a client’s key from a stored record via password guessing (through the MALICIOUSREC interface). This password guessing is limited by the counter ctr though (which will be enforced through the honest HSM in the protocol), i.e., the adversary has at most 10 guesses for each honest account. Thus, there are still no unlimited offline attacks possible on any accounts.

Difference to [DFG⁺ 23b]: In the original functionality, a corrupt server was able to re-route sessions from an honest client ID_C to a different client ID_C^* . This was modeled because the WhatsApp protocol allowed such an attack. In general, this can be prevented easily, and we therefore remove this weakness in our definition.

File Leakage. The adversary can repeatedly compromise the existing records, by sending (LEAKFILE) to $\mathcal{F}_{\text{PPKR}}$. The functionality then provides the adversary with a list of all clients ID_C for which files are recorded. It creates LEAKED copies of all records and now permits *unlimited* offline attacks on these records via the OFFLINEATTACK interface. As the initialization interface can be used by the client to renew its key and/or password, there can be different files for the same client ID_C at different times. We therefore attach a counter i to each leaked record and allow offline attacks on each version.

Note that the adversary can start using the OFFLINEATTACK on records only *after* it compromised them. This resembles the notion of pre-computation attack resistance which was introduced as a strengthening of aPAKE with the OPAQUE protocol [JKX18].

Apart from allowing offline attacks on leaked key files, all other security guarantees of $\mathcal{F}_{\text{PPKR}}$ are still intact. In particular, key authenticity is guaranteed and the deletion counter is still set honestly as instructed by our functionality.

Server (and HSM) Fully Corrupt. The strongest corruption level considers the server (and HSM) to be fully corrupted, which can be triggered through the ($\text{FULLYCORRUPT}, S$) interface. This marks the server as FULLYCORRUPT and leaks all existing and newly created key files to the adversary. This full corruption gives \mathcal{A} all capabilities described above for the corrupt server and file leakage settings. In addition, the functionality now gives up on key authenticity and the strict enforcement of the deletion counter.

The latter is done in the FULLYCORRUPT interface, which sets all counters of stored and newly created files to ∞ , such that they no longer get deleted automatically. When a client recovers her key, \mathcal{A} can now freely decide if it wants to signal DELREC or not.

The full corruption also enables the adversary to online attack recovery sessions from honest clients, in two ways. The first is an *key planting and online guessing* attack, where the adversary can provide an arbitrary key K^* along with a password guess pw^* as optional inputs in the RECS interface. It then learns whether $\text{pw}^* \stackrel{?}{=} \text{pw}'$. If it guessed the password attempt correctly, K^* is returned to the client instead of the proper key. Note that this planting requires the correct guess of the client’s password attempt, and also allows only a single password guess per recovery session. However, weak passwords might be known to the adversary through its offline attack capability on all key files.

The second enabled online attack is to re-use any of the previously leaked key files, which might have been replaced in the meantime. Here, \mathcal{A} provides an additional input i in the RECS interface, that allows it to run the honest clients recovery against the old key file of version i . If the client’s password attempt is pw' matches the pw contained in the old key file, the client gets tricked into accepting its old key.

Summary of Security Guarantees & Security Levels. Note that even in the full corruption setting, $\mathcal{F}_{\text{PPKR}}$ still guarantees two important security properties: (1) pseudorandomness of keys – as the keys generated by honest clients are still chosen at random by the functionality and (2) the honest client’s passwords and keys are never directly accessible to the adversary. For each key K that \mathcal{A} wants to compromise, it has to correctly guess the user’s password through an offline attack. Thus, users who have chosen a strong password can still hope that their key remains safe. Even the third core property of key authenticity is not entirely lost when S is fully corrupt, as this still requires the adversary to correctly guess the honest user’s password. We give a summary and comparison of the main security properties in the different corruption settings in Table 2 below.

For brevity, we refer to security that is maintained when the server S is corrupt, but no files got leaked as Lev-1 security. Lev-2 security covers the optimal guarantees up to joint server corruption and file leakage, considering both attack capabilities combined. Finally, Lev-3 security refers to schemes that maintain the optimal security guarantees up to full server corruption.

Flaws in $\mathcal{F}_{\text{PPKR}}$ from [DFG⁺23b]. In the WBP, if a corrupt party initializes and later runs a recovery, it can always reset the counter to 10, even if it used different passwords in initialization and recovery¹. This weakness also exists in our protocol $\pi_{\text{OPRF-PPKR}}$ aiming at Lev-3 security that we present in Section 4.3. The weakness was not captured in the original functionality in [DFG⁺23b] and we fix this in Figures 1 to 2. We remark that weakening $\mathcal{F}_{\text{PPKR}}$ in this way gives only very limited additional capabilities as the corrupt party can always reset its file by simply using the correct password.

Interestingly, our protocols π_{encPw} (see Section 4.1) and $\pi_{\text{encPw+}}$ (see Section 4.2), which only aim at Lev-1, resp. Lev-2, security, do not have this property. For this reason, we mark the corresponding code in Figures 1 to 2 in boxes, which allows us to strengthen $\mathcal{F}_{\text{PPKR}}$ by dropping the boxed code.

Table 2: Overview of security properties guaranteed by $\mathcal{F}_{\text{PPKR}}$ under different corruption settings. We refer to the following level: Lev-1 = S is corrupt, Lev-2 = S is corrupt and files are leaked (both attacks combined), Lev-3 = fully corrupt S .

Property	Honest S	Lev-1	Lev-2	Lev-3
Pseudorandom key K	✓	✓	✓	✓
No direct leakage of pw, K	✓	✓	✓	✓
Limit file access to ID_C	✓	✗	✗	✗
No offline attacks on pw, K	✓	✓	✗	✗
No precomputation for offline attacks	n.a.	n.a.	✓	✓
No offline attacks on pw’ (pw-attempt)	✓	✓	✓	✓
Upper limit on incorrect recoveries	✓	✓	✓	✗
Key authenticity	✓	✓	✓	✗

4 Constructions

In this section, we discuss three protocols for password-protected key retrieval that securely realize the functionality $\mathcal{F}_{\text{PPKR}}$. To ensure security in the single-server setting, we adapt the approach from

¹The corrupt client can remember the AKE secret keys from initialization and disregard the recovered keys, leading to a successful KE. In our protocol, this corresponds to remembering the original signing key.

WhatsApp [Wha21] and let all our protocols rely on a Hardware Security Module (HSM). One can think of the HSM as a particularly secure part of the server or a dedicated module that only the server can access. We make the assumption that the code of the HSM is executed honestly and messages from the HSM are attested, i.e., signed such that anyone can verify the message’s origin.

Each of the protocols in this section protects against a different level of attacks against the HSM. The first protocol, basic encrypt-to-HSM (Section 4.1), is based solely on public-key and symmetric encryption and is secure if the server gets corrupted, but the HSM is incorruptible and can be completely trusted (Lev-1). The second protocol, enhanced encrypt-to-HSM (Section 4.2), is a minor modification of the first protocol that provides additional protection of clear-text passwords and keys when the HSM leaks client account data (Lev-2). The third construction, OPRF-based PPKR (Section 4.3), is secure in the full corruption scenario (Lev-3), i.e., even when all files and all secret keys of the HSM are given to the adversary, an offline attack against the users’ password is still necessary.

Like in [DFG⁺23b], we also assume client-to-server authenticated channels, e.g., through an SMS/email passcode authentication. This is merely to prevent an adversary from tricking an honest server into deleting accounts of honest users through repeated incorrect retrievals. Thus, this is used as a basic protection against deletion only, but not for actual user authentication – this is what the password is for. Importantly, however, the client-to-HSM channel is not authenticated, which is what enabled the attack from [DFG⁺23b].

Further, as we work in the UC-framework, all messages contain globally unique session- and subsession identifiers. This is a typical UC-artefact and there are standard ways to handle these ids when implementing the protocol in the real world, see [Can00, Sec. 3.3.2].

Modeling the HSM. All our protocols internally rely on the concept of an HSM (or similar trusted execution environments). We assume the HSM to attest every message, i.e., sign it under a public verification key that is known to all parties. Such verification key could be published on the HSM vendor’s website, certified by some trusted authority.

In our model of leakage from the HSM, covered through Lev-2 security, HSM attestation is the *only* power kept from the adversary. Only in our strongest model, Lev-3 that considers full HSM corruption, we will provide the adversary with the attestation key. This distinction between the permanent attestation key and client-specific files is justified because attestation is a central capability of HSMs. Therefore, an HSM manufacturer might turn special attention to the protection of the single attestation key, such as more expensive hardware protection mechanisms.

For simplicity, we also assume that the HSM has some persistent memory where it stores file records. In practice, HSMs might securely outsource such storage as done, e.g. in the case of WhatsApp’s solution [DFG⁺23b], but the actual realization of storage is irrelevant for our work. This HSM-protected storage is still considered secret upon server corruption (Lev-1) but entirely leaked to the adversary up from Lev-2.

When proving security in the UC-framework [Can01], we model the HSM as a subroutine of the server, similar to [PST17], instead of modeling it as an individual party. This captures the exclusive

access that the server has to the HSM. Concretely, the code of the HSM in Figure 3 and Figure 4 is the code of an entity that cannot be corrupted by the adversary (because it is not a protocol party) and that does not accept input from anybody else than the server. For the protocols that can handle certain types of HSM corruption, we equip our HSM functionalities with the respective corruption interfaces. Concretely, if we allow leakage of client files and all the protocol-specific long-term state, the HSM functionality will expose a `LEAKFILE` interface to the adversary that returns the client files. For the OPRF PPKR protocol, we add a `FULLYCORRUPT` interface to the HSM functionality that additionally gives the adversary the HSM’s secret attestation key. A formal description of the respective HSM functionalities used by our protocols can be found in Figure 9 and Figure 10.

4.1 Lev-1 Protocol: Basic Encrypt-to-HSM

The security analysis of the WBP protocol of [DFG⁺23b] assumed that the HSM cannot be corrupted and never leaks information. A natural question is if this strong assumption allows for a simpler PPKR protocol. The basic encrypt-to-HSM protocol π^{encPw} , depicted in Figure 3, answers the question in the affirmative.

The central observation is that if no information is leaked by the HSM, then one can store the client’s passwords and keys in the clear at the HSM. For recovery, the HSM is handed the clear-text password attempt and performs the comparison. To protect against network attackers and the potentially corrupt server, the client encrypts all messages to the HSM under a long-term encryption key pk_{Enc} . Like in WBP, all clients have access to pk_{Enc} , e.g., because the key is hard-coded into the client’s source code.

To prevent the replay of messages by network attackers or corrupt servers, we let the client encrypt session identifier `ssid` and their own identity `IDC`, and also include these items in the clear. By comparing the decrypted identifiers with the clear-text ones, the HSM will notice when a malicious server changes the identifiers of honest client requests, and abort. Hence, with this little check, π^{encPw} is stronger (on Lev-1) than the WBP, which admits such replay attacks [DFG⁺23b].

We show that π^{encPw} is a secure PPKR when only considering server corruptions, i.e., it provides Lev-1 security.

Theorem 1. *The protocol π^{encPw} from Figure 3 UC-realizes $\mathcal{F}_{\text{PPKR}}$ in the $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$ -hybrid model (i.e., assuming no corruption or leakage of the HSM), assuming HSM attested messages, and assuming that PKE is IND-CCA secure and SE is IND-CPA secure.*

More precisely, let $\text{Dist}_{\mathcal{Z}}^{A,B}$ denote the advantage of \mathcal{Z} to distinguish distributions A and B . We get that for every efficient real-world adversary \mathcal{A} against π^{encPw} , there is an efficient simulator $\text{Sim}^{\text{encPw}}$ (see Figures 11 to 12) that interacts with $\mathcal{F}_{\text{PPKR}}$ such that for every efficient environment \mathcal{Z} there exists an efficient adversaries \mathcal{B}_1 against the IND-CCA security of PKE and \mathcal{B}_2 against the IND-CPA security of SE such that

$$\begin{aligned} \text{Dist}_{\mathcal{Z}}^{\pi^{\text{encPw}}, \{\mathcal{F}_{\text{PPKR}}, \text{Sim}^{\text{encPw}}\}}(\lambda) &\leq (q_{\text{INIT}} + q_{\text{REC}}) \text{Adv}_{\text{PKE}, \mathcal{B}_1}^{\text{IND-CCA}}(\lambda) \\ &\quad + q_{\text{REC}} \text{Adv}_{\text{SE}, \mathcal{B}_2}^{\text{IND-CPA}}(\lambda) \\ &\quad + \text{Adv}_{\text{Sig}, \mathcal{B}_4}^{\text{SEUF-CMA}}(\lambda), \end{aligned}$$

where q_{INIT} is the number of initializations and q_{REC} is the number of recoveries.

PROOF SKETCH. The gist of the proof is that the incorruptible HSM holds the secret key sk_{Enc} for all encryptions. As the HSM is modeled as a hybrid functionality, the simulator has access to the secret key. This allows for the following proof strategy:

To simulate *honest clients* without the password, the simulator can replace the PKE ciphertexts C by encryptions of \perp . The environment cannot detect this change by the IND-CCA security of the encryption scheme. Indeed, we require CCA security, as the HSM’s `FAIL` responses on mismatching `IDC` and `ssid` give the adversary a very limited decryption oracle. The ciphertext C' is again replaced by an encryption of \perp . Note that this time we only need to reduce to IND-CPA security because the integrity of C' is ensured by the HSM attestation. Finally, the simulator can use $\mathcal{F}_{\text{PPKR}}$ to provide clients with their correct output.

To extract the inputs of *corrupt clients* the simulator can use the secret key sk_{Enc} of the HSM. The simulator knows this key because the HSM is modeled as a hybrid functionality. That means concretely that in the ideal-world execution, the simulator plays the role of the HSM (and thus, chooses sk_{Enc} by itself). Using sk_{Enc} , the simulator can decrypt the PKE ciphertext C provided by the corrupt client and provide the used password pw^* to $\mathcal{F}_{\text{PPKR}}$. If the password guess was correct the functionality will give the simulator the recovered backup key K . The simulator can also extract the corrupt client’s symmetric key k_{sym} by using the secret key sk_{Enc} of the HSM and respond to the client with an encrypted version of its backup key C' .

Simulating the server amounts to executing the rest of the protocol as usual because the server holds no private information. \square

The full proof can be found in Appendix D.1.

4.2 Lev-2 Protocol: Enhanced Encrypt-to-HSM

When the HSM leaks permanently stored protocol data such as account information of clients, also called “password files”, one cannot hope to prevent an adversary from, e.g., offline-attacking a PPKR protocol. That is because the file must contain enough information for the HSM to decide whether a recovery attempt is successful or not. Nonetheless, one can demand that an adversary still needs to guess a user’s password and cannot read the password and/or the backup key immediately from the file. In other words, clients that choose very strong passwords should still have a certain level of security even if the HSM’s files get leaked.

Clearly, the basic encrypt-to-HSM protocol from Section 4.1 falls short of this goal, and hence cannot reach Lev-2 security: The passwords and backup keys are stored in the clear and a compromise of the HSM’s long-term state immediately gives them away. Further, the leakage includes all protocol-specific long-term state, i.e., the adversary also gets the HSM’s decryption key sk_{Enc} , which leaks all the password (attempts) from sessions happening *after* the compromise to the attacker.

To reach Lev-2 security, we strengthen π^{encPw} in two ways: First, we store the password only in hashed form at the HSM, with a user-specific salt to prevent pre-computation attacks; and encrypt the backup key with another salted password hash. The salts are stored

in the password file. Second, we let the HSM use ephemeral encryption keys for each session. Figure 3 including grayboxes and skipping [dashed boxes] shows our enhanced encrypt-to-HSM protocol $\pi^{\text{encPw}+}$. We can formally prove in Theorem 2 that these relatively simple (although in the case of the encryption keys significantly more expensive) measures are enough to restrict an adversary in Lev-2 back to offline guessing, after getting access to the long-term storage of the HSM.

We stress that our HSM leakage model at Lev-2 returns all permanently stored user files and protocol-specific *long-term* state to the adversary, but not the temporary values that occur during execution of the protocol.

Theorem 2. *The protocol $\pi^{\text{encPw}+}$ from Figure 3 UC-realizes $\mathcal{F}_{\text{PPKR}}$ in the $\mathcal{F}_{\text{HSM}}^{\text{encPw}+}$ -hybrid model (i.e., assuming no corruption but state leakage of the HSM) and assuming that H is modeled as random oracle, messages are HSM-attested, PKE is IND-CCA secure, and SE is IND-CPA secure.*

More precisely, let $\text{Dist}_{\mathcal{Z}}^{A,B}$ denote the advantage of \mathcal{Z} to distinguish distributions A and B . We get that for every efficient real-world adversary \mathcal{A} against $\pi^{\text{encPw}+}$, there is an efficient simulator $\text{SIM}^{\text{encPw}+}$ (see Figure 13) that interacts with $\mathcal{F}_{\text{PPKR}}$ such that for every efficient environment \mathcal{Z} there exist efficient adversaries \mathcal{B}_1 against the IND-CCA security of PKE and \mathcal{B}_2 against the IND-CPA security of SE such that

$$\begin{aligned} \text{Dist}_{\mathcal{Z}}^{\pi^{\text{encPw}+}, \{\mathcal{F}_{\text{PPKR}}, \text{SIM}^{\text{encPw}+}\}}(\lambda) &\leq (q_{\text{INIT}} + q_{\text{REC}}) \text{Adv}_{\text{PKE}, \mathcal{B}_1}^{\text{IND-CCA}}(\lambda) \\ &\quad + q_{\text{REC}} \text{Adv}_{\text{SE}, \mathcal{B}_2}^{\text{IND-CPA}}(\lambda) + \text{Adv}_{\text{Sig}, \mathcal{B}_4}^{\text{SEUF-CMA}} \\ &\quad + \frac{q_{\text{INIT}}(2q_{\text{INIT}} - 1)}{2^\lambda} + \frac{2q_{\text{INIT}}q_H}{2^\lambda}, \end{aligned}$$

where q_{INIT} is the number of initializations, q_{REC} is the number of recoveries, and q_H is the number of H queries.

PROOF SKETCH. Because the formal proof of Theorem 2 shares most of its steps with the proof of Theorem 1, we focus on the changes to the proof that are required to also achieve security under leaked HSM files. We show the modified and additional interfaces that the simulator provides in Figure 13. Now we have to simulate leakage of HSM files, because $\mathcal{F}_{\text{HSM}}^{\text{encPw}+}$ allows the adversary to call its LEAKFILE interface. To that end, the simulator must observe and program the random oracle H and use the OFFLINEATTACK interface of $\mathcal{F}_{\text{PPKR}}$. The simulator now also has to generate fresh encryption keys $(\text{pk}_{\text{Enc}}, \text{sk}_{\text{Enc}})$ for every initialization and recovery and store them until they are used. Note that we can keep the rest of the simulation as in the proof of Theorem 1. That is because the leaked client file does not contain values that would allow impersonation of the HSM (as will be the case when we consider full corruption in Section 4.3).

Before the state of the HSM is leaked, an adversary has only a negligible chance to query $H(s_2, \text{pw})$ to the random oracle. That is because s_2 has high entropy. Therefore, the simulator can store a uniformly random value c as the encoding of the backup key. But once the HSM state is leaked, the adversary knows s_1, s_2 and h and thus, can guess passwords and verify its guess using h and s_1 . If one of the guesses is correct, the adversary will be able to check if c is indeed an encoding of K by using s_2 . The simulator can use the OFFLINEATTACK interface of $\mathcal{F}_{\text{PPKR}}$ to see if the adversary's guess was correct. In case of a successful guess, the simulator learns K .

Then, it can program $H(s_2, \text{pw}) \leftarrow c \oplus K$ to equivocate c after the fact, and similarly SIM can program $H(s_1, \text{pw}) \leftarrow h$. \square

The full proof can be found in Appendix D.2.

Missing Lev-3 Security of $\pi^{\text{encPw}+}$. Clearly, our enhanced encrypt-to-HSM protocol cannot satisfy Lev-3 security: if the adversary gets the HSM's attestation key, it has full control over the public keys for the encryption scheme under which the client encrypts the user's password. Thus, for all users that still use the PPKR service after the full compromise happened, the adversary immediately learns their plaintext passwords from the request – and consequently can recover their keys too.

To achieve Lev-3 security, we need to securely communicate some password-dependent data to a possibly entirely malicious party for authentication, and cannot rely on any server-held secret for the password's protection. This requires more advanced cryptographic techniques, for which we revert to an oblivious pseudorandom function in our third construction.

4.3 Lev-3 Protocol: OPRF-Based PPKR

In this section, we build PPKR that provides protection against full server and HSM corruption. That is, even when all keys and files are leaked, the best an adversary can do is an offline attack against each user's password. Our goal is to propose a simpler protocol than WBP that reaches such Lev-3 security.

Oblivious Pseudorandom Function. At the core of our construction is an oblivious pseudorandom function (OPRF). Such a function allows to deterministically compute $y = \text{PRF}(k, x)$ through an interactive protocol between an evaluator and requester. The evaluator knows the PRF key k , but learns nothing about the input x or output y it computes. We will use such an OPRF to deterministically compute password-dependent key material, which allows the client to authenticate towards the HSM, and also derive the encryption key under which K gets wrapped.

High-Level Idea of $\pi^{\text{OPRF-PPKR}}$. For initialization, the client executes the OPRF with the HSM (through the server) to obtain a value $\rho = \text{PRF}(k_{\text{OPRF}}, (\text{pw}, \text{ID}_C))$, depending on the user's password and ID_C . The HSM chooses a client-specific key k_{OPRF} . The derived ρ then serves as the key for an authenticated encryption (AE) scheme, under which the client's randomly chosen key K gets encrypted. The client also generates a signature key pair $(\text{sk}_C, \text{pk}_C)$ for future re-authentication. It encrypts both sk_C and K under ρ , obtaining a ciphertext c . The HSM gets c and pk_C and stores them with $\text{ID}_C, k_{\text{OPRF}}$, and a counter value ctr , initially set to 10, forming the password file for ID_C .

For recovery, the client receives c from the HSM and runs the OPRF again. If the client uses the same password as in initialization, it obtains the same ρ as earlier. The client then uses ρ to decrypt c , obtaining (K, sk_C) . Now the client can prove knowledge of the right password towards the HSM, by using the sk_C to sign the transcript of previous messages exchanged with the HSM. The HSM decreases the counter ctr at the beginning of every recovery, and if it receives such a signature that verifies under the pk_C in the stored user file, it considers the client as correctly authenticated and resets the ctr value to 10 again.

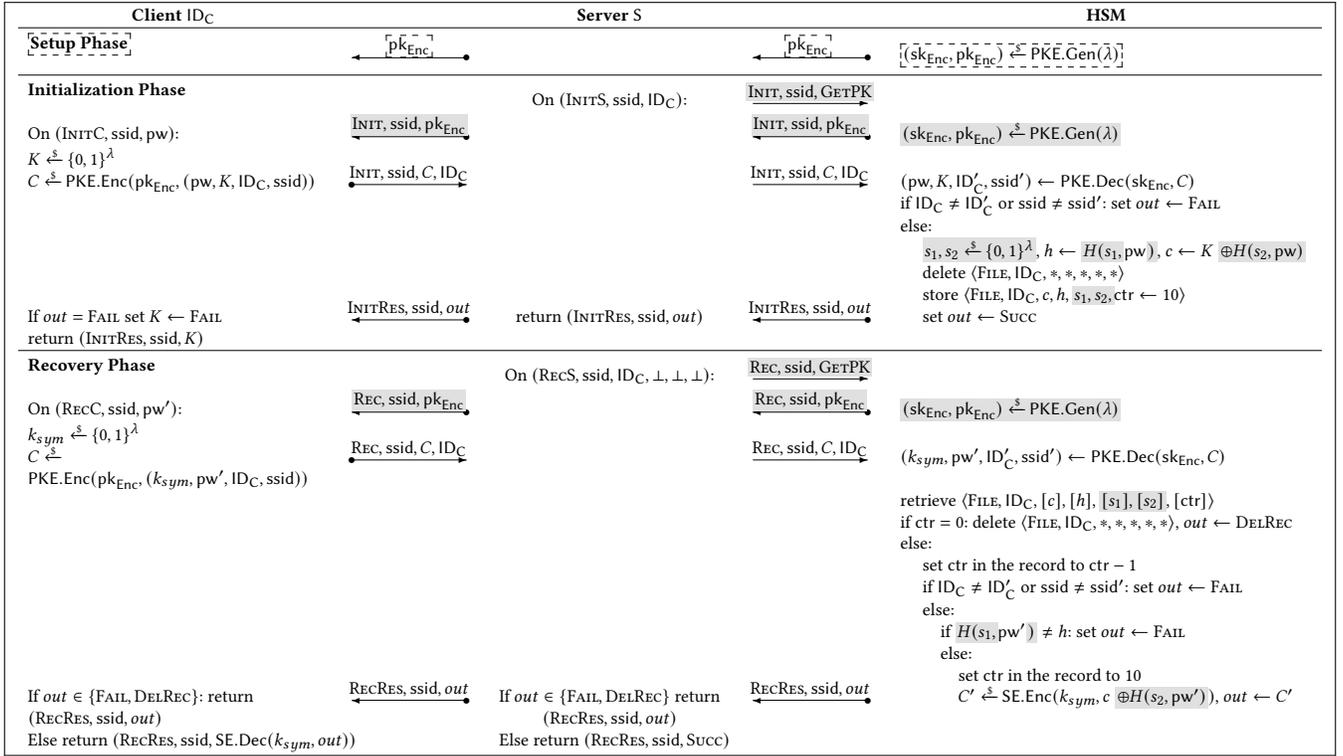


Figure 3: Protocols π^{encPw} and $\pi^{\text{encPw+}}$. The code in [] is only executed in π^{encPw} , the code in [] is only executed in $\pi^{\text{encPw+}}$. \xleftarrow{x} is the HSM-attested transmission of x . \xrightarrow{y} is the client-to-server authenticated transmission of y . We implicitly assume that each message contains $ssid$ and the random oracle H takes $ssid$ as first input. ID_C and S output (INITRES, ssid, FAIL) or (RECRES, ssid, FAIL), whenever they receive an unexpected message (i.e. with mismatching $ssid$ or ID_C) or when attestation verification fails. If any party receives the same type of message twice with the same $ssid$, it ignores the second one. Init and Rec are strings indicating which phase the client wishes to initiate. The server receives the input (RECS, ssid, ID_C, ⊥, ⊥, ⊥) instead of just (RECS, ssid, ID_C) only for technical reasons as the syntax needs to match $\mathcal{F}_{\text{PPKR}}$.

Interestingly, we still need a fresh encryption key pair for communication towards the HSM – but only for the initialization. Therein, the client’s values (pk_C, c) and the $ssid$ get encrypted under the freshly chosen key. The purpose of this encryption is to prevent a malicious server from combining the ciphertext c of an honest user with a different public key pk_C . Such a mixing attack would allow the server to plant the honest c with some authentication information where it could easily circumvent the strict limit of failed retrievals imposed by the HSM. Thus, the purpose of the public-key encryption here is to bind the honest user’s information together.

Comparison to WBP. Our protocol shares a lot of similarities with the WBP as stated by Davies et al. [DFG⁺23b], but simplifies the design enabling both a more efficient and more secure variant. We start by sketching the WBP core ideas and then explain the main differences to our protocol.

The overall idea of the WBP is to let the client and the HSM execute an *asymmetric PAKE* (aPAKE) protocol to exchange a fresh symmetric key K^{session} from a clear-text password of the client, and a password file stored by the HSM. K^{session} is subsequently used to prove the correctness of the client’s password to the HSM using

standard key confirmation techniques, to reset the file counter. The client’s key K is encrypted under a static (i.e., depending only on the password) key K^{export} and stored by the HSM. The protocol is instantiated with OPAQUE [JKX18], which in turn crucially relies on an OPRF to produce the static export key.

In our protocol, we start from an OPRF instead of aPAKE, and implement the proof of password knowledge separately through digital signatures. This yields a simpler protocol layout and cleaner security proof. In terms of similarity, the AE ciphertext stored in the client’s password files corresponds to the OPAQUE password files, and the OPRF is used to deterministically derive a key to decrypt this file and perform the re-authentication. How this authentication is done differs though: we save one message by using signatures instead of authenticated key exchange. Note that this exploits that we are not aiming at the key exchange, which was the goal of OPAQUE – and thus is more than what is needed for PPKR.

Apart from improving efficiency, our protocol provides better security than WBP. While the following attacks can all be considered minor, they are still in conflict with the desired security properties – and easily preventable as shown by our protocol. The security improvements are as follows:

- (1) In the WBP, a corrupt server can prevent old password files of honest clients from being overwritten upon a client re-initializing with a new password. The corrupt server could then still use up all remaining password guesses against old password files, to recover previous keys of the client. We prevent this by having the HSM attest the client identity ID_C for each session ssid, and letting the client abort when it receives a mismatching ID_C . This enforces an agreement between an honest client and honest HSM when the server is corrupt (needed for Lev-1 security).
- (2) We fix a “rerouting” attack on the WBP that was discovered by [DFG⁺23b], and which allows a corrupt server to reroute honest recovery attempts to wrong password files. This results in honest Alice recovering Bob’s key if both use the same password. We therefore invoke the OPRF not only on the user’s password but also append the unique client identity ID_C to the input. This simple measure ensures that users derive unique wrapping keys ρ , which protects against this attack (again needed for Lev-1 security).
- (3) The above two measures additionally prevent two attacks on the WBP that a fully corrupt server in the Lev-3 setting² can mount: (I) checking whether two honest clients use the same password, (II) resetting the counter in an honest client’s password file if another client having the same password runs a recovery. Both attacks work by routing a recovery attempt of Alice to Bob’s password file, which in the WBP goes unnoticed by Alice, and can succeed because the passwords of honest clients are not enforced to be unique. Both is prevented through measures (1) and (2).
- (4) Considering the same HSM leakage model that we use (everything except the HSM’s attestation key gets leaked on Lev-2), in the WBP, a malicious server can run offline password guessing attack not only against all leaked files but also against files that are created *after* the compromise. This is because the WBP relies on a long-term encryption key at the HSM for initialization. After a Lev-2 compromise, the malicious server can decrypt an honest user’s initialization request, and re-encrypt the user’s wrapped key K together with a maliciously chosen (AKE) public key. This allows the malicious server to get unlimited password guesses against the user’s real password-wrapped key, as it can correctly complete key confirmation towards the HSM, even when performing the retrieval with the wrong passwords. Our protocol prevents that attack by using fresh encryption keys in initialization, achieving the necessary security for Lev-2.

Concrete OPRF for Efficiency. While our protocol can be securely realized with any 2-round OPRF, we build $\pi^{\text{OPRF-PPKR}}$ from the 2HashDH OPRF [JKKX16] in a non-black-box way. This was mainly done for efficiency reasons. A generic approach relying on an ideal OPRF functionality would not allow binding the OPRF in- and outputs directly to other protocol values, nor do this efficiently in an HSM-attested way. However, even though our protocol is non-generic, our proof actually provides some modularity: we use the 2HashDH simulator as a step in our proof, and from then on

²We note that the fully corrupt server case was not analyzed in [DFG⁺23b], and hence they did not claim the protocol to prevent these attacks. It is nonetheless easy to verify that a fully corrupt server in the WBP can use the same OPRF key for Alice and Bob and perform these attacks.

rely on the abstracted security properties of a UC-secure OPRF. Nevertheless, our proof additionally relies on the internals of the 2HashDH simulator at several points, making the switch to another OPRF non-trivial.

Theorem 3. *The protocol $\pi^{\text{OPRF-PPKR}}$ from Figure 4 UC-realizes $\mathcal{F}_{\text{PPKR}}$ in the $\mathcal{F}_{\text{HSM}}^{\text{OPRF-PPKR}}$ -hybrid model (i.e., allowing file leakage and full corruption of the server), assuming adaptive corruptions outside of ongoing initialization and recovery phases, HSM-attested messages, client-to-server authenticated channels, that AE is equivocal and random-key robust and has INT-CTXT-security, Sig is sEUF-CMA-secure, PKE is CCA-secure, H_1 and H_2 are modeled as random oracles, and that the (N, Q) -OMDH assumption holds in \mathbb{G} .*

More precisely, let $\text{Dist}_{\mathcal{Z}}^{A,B}$ denote the advantage of \mathcal{Z} to distinguish distributions A and B . Then, for every efficient real-world adversary \mathcal{A} against $\pi^{\text{OPRF-PPKR}}$, there is an efficient simulator $\text{Sim}^{\text{OPRF-PPKR}}$ (see Figures 14 to 17) that interacts with $\mathcal{F}_{\text{PPKR}}$ such that for every efficient environment \mathcal{Z} there exists efficient adversaries $\mathcal{B}_1, \dots, \mathcal{B}_7$ such that

$$\begin{aligned} \text{Dist}_{\mathcal{Z}}^{\pi^{\text{OPRF-PPKR}}, \{\mathcal{F}_{\text{PPKR}}, \text{Sim}^{\text{OPRF-PPKR}}\}}(\lambda) &\leq \text{Adv}_{\mathcal{B}_1, \text{Sig}}^{\text{sEUF-CMA}}(\lambda) \\ &+ q_{\text{INIT}} \text{Adv}_{\mathcal{B}_2, \mathbb{G}}^{(q_E + q_H \cdot q_E) - \text{OMDH}}(\lambda) + q_E^2 \text{Adv}_{\mathcal{B}_3, \text{AE}}^{\text{rkr}}(\lambda) \\ &+ q_{\text{INIT}} \text{Adv}_{\mathcal{B}_4, \text{AE}}^{\text{INT-CTXT}}(\lambda) + \text{Adv}_{\mathcal{B}_5, \text{AE}}^{\text{EQV}}(\lambda) + q_{\text{INIT}} \text{Adv}_{\mathcal{B}_6, \text{Sig}}^{\text{sEUF-CMA}}(\lambda) \\ &+ q_{\text{INIT}} \text{Adv}_{\mathcal{B}_7, \text{PKE}}^{\text{IND-CCA}}(\lambda) + (q_E + q_H)^2 / q \end{aligned}$$

where q_{INIT} is the number of initializations, q_{REC} is the number of recoveries, $q_E = q_{\text{INIT}} + q_{\text{REC}}$, q_H is the number of queries to H_1 and q is the order of \mathbb{G} .

PROOF SKETCH. The proof heavily relies on the security of the 2HashDH OPRF. We use the simulator SIM_{OPRF} from Davies et al. [DFG⁺23b] (restated in Appendix B), which demonstrates that 2HashDH UC-realizes $\mathcal{F}_{\text{OPRF}}$, in a non-black-box way throughout the proof. Whenever we have to simulate a message a, b, a' , or b' for some honest party, we “outsource” the simulation to SIM_{OPRF} . This also means that we let SIM_{OPRF} choose the key k_{OPRF} , which we then need to obtain from SIM_{OPRF} whenever there is a LEARFILE query from \mathcal{A} . Furthermore, from the random oracle queries to H_2 by \mathcal{A} , SIM_{OPRF} is able to extract the password chosen by a corrupt party in an initialization, which allows us to install a corresponding file in $\mathcal{F}_{\text{PPKR}}$.

A challenge that arises from allowing the LEARFILE query is that \mathcal{A} can do offline password guessing using k_{OPRF} obtained from leaked files. If \mathcal{A} guesses the correct password pw of some ID_C , it can decrypt the ciphertext c that was encrypted under the output ρ of the OPRF. Hence, c has to be simulated such that it decrypts to the key K chosen randomly by $\mathcal{F}_{\text{PPKR}}$ for ID_C , as otherwise, the simulation would be distinguishable from the real protocol. However, when we simulate c in the initialization phase, K is unknown, as $\mathcal{F}_{\text{PPKR}}$ has not even given it to ID_C , yet. To solve this, we require AE to be equivocal, which allows us to produce a simulated c and only later decide to which values c decrypts. More precisely, whenever \mathcal{A} queries H_2 on an input of the form $(\text{pw}' \parallel ID_C, H_1(\text{pw}' \parallel ID_C)^{k_{\text{OPRF}}})$, where k_{OPRF} is some OPRF key chosen by SIM_{OPRF} , we submit pw' to the OFFLINEATTACK interface of $\mathcal{F}_{\text{PPKR}}$. If $\text{pw} = \text{pw}'$, then we obtain the key K and can equivocate c to obtain a ρ such that c decrypts to K under ρ and program the output of H_2 to ρ .

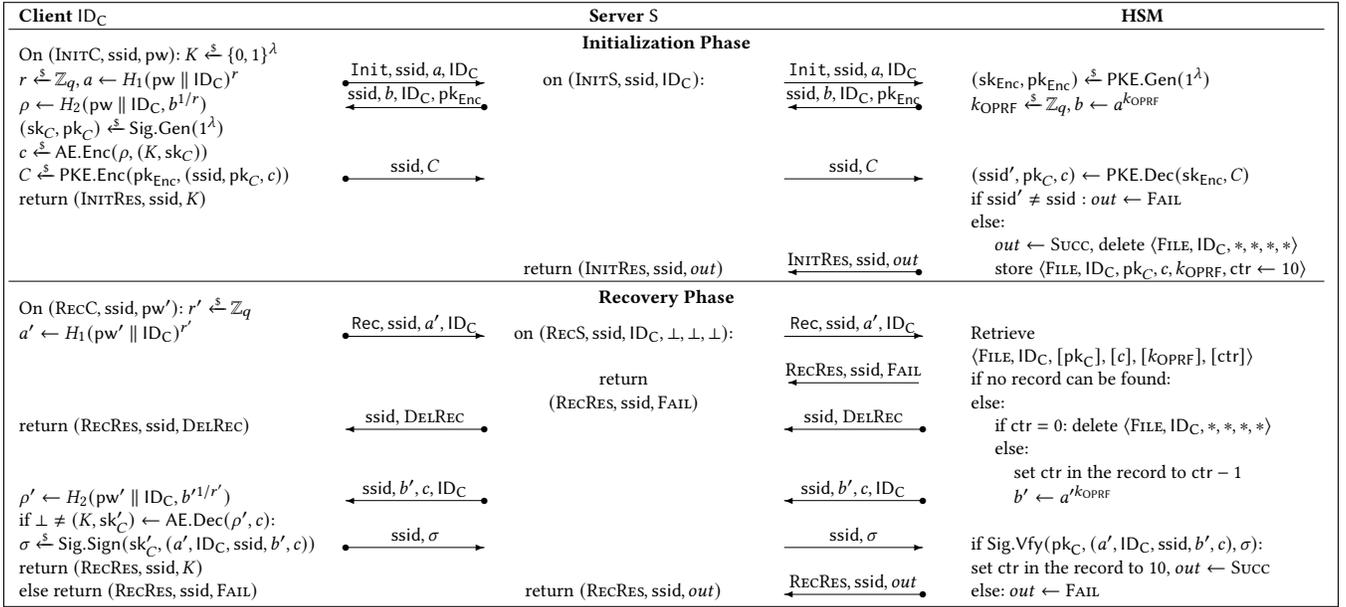


Figure 4: Protocol $\pi^{\text{OPRF-PPKR}}$. \xrightarrow{x} indicates that the message x is signed by the HSM; and \xrightarrow{y} denotes the client-to-server authenticated transmission of y . ID_C outputs $(INITRES, ssid, FAIL)$ or $(RECRES, ssid, FAIL)$, whenever it receives an unexpected message (i.e. with mismatching $ssid$ or ID_C) or when signature verification fails for the received signed message. If any party receives the same type of message twice with the same $ssid$ it ignores the second one. $Init$ and Rec are strings indicating which phase the client wishes to initiate. The server receives the input $(REC_S, ssid, ID_C, \perp, \perp, \perp)$ instead of just $(REC_S, ssid, ID_C)$ only for technical reasons as the syntax needs to match $\mathcal{F}_{\text{PPKR}}$.

A similar challenge arises when S is fully corrupt. Then, S can make the critical query to H_2 even before we simulate c as it chooses the OPRF key k_{OPRF} used in the initialization. Thus, we have to check for all previous H_2 queries whether it was this critical query via the `OFFLINEATTACK` interface and if so, encrypt K under the corresponding output of H_2 instead of outputting an equivocable c .

The last major challenge is that a fully corrupt S can do key-planting attacks and can mix-and-match different files during a recovery, e.g., when ID_C recovers, S could use k_{OPRF} of some $ID'_C \neq ID_C$ and an adversarial c . To deal with this, we check whether the password $pw'' \parallel ID'_C$ and key k'_{oprf} used to derive the ρ , under which the adversarial c was encrypted, are the same as the password $pw \parallel ID_C$ and key k_{OPRF} used in the recovery by ID_C . In the real world, ID_C would then output the key K decrypted from c . To simulate this, we use the key-planting capabilities in the `REC_S` interface. We can again find the password $pw'' \parallel ID'_C$ used to derive ρ via the H_2 queries by \mathcal{A} and check whether the same OPRF key was used with the help of `SIMOPRF`. We can then decrypt c to get K'' and submit pw'' and K'' to the `REC_S` interface. If $pw = pw''$, ID_C then outputs K'' . We can proceed similarly if c instead is equivocable. \square

The full proof can be found in Appendix D.3.

5 Evaluation & Discussion

In this section, we give an overview of the concrete efficiency of our protocols, compare them to WBP, and also discuss the respective advantages of our protocols.

Instantiations of Building Blocks. For the efficiency overview, we choose concrete instantiations of the required primitives, such that we can count the number of operations performed in each protocol. We only chose group-based public key primitives to keep the numbers comparable. But of course, one could use any other secure instantiation of the primitives, e.g., based on RSA or lattices.

Concretely, we chose Schnorr-Signatures, HMAC, HKDF, ElGamal encryption as a CPA secure encryption and DHIES as CCA secure encryption. For simplicity, we assumed that all hash evaluations cost a uniform unit “1 Hash”³ and similarly that one AE encryption or decryption costs “1 AES”. A detailed overview of the computation costs of each primitive can be found in Table 4. We did not list the HSM’s attestation signatures as they are automatically produced by the HSM anyway.

Efficiency Comparison. As Table 3 shows, OPRF-PPKR is more efficient than the WBP in both, initialization and recovery. This comes mostly from WBP performing an authenticated key exchange where OPRF-PPKR uses a digital signature. In particular, in the recovery phase, which will be the more time-critical phase in deployment,

³Ignoring e.g., exponentiations that might be needed to hash into a group.

Table 3: Efficiency of PPKR realizations expressed in terms of the number of exponentiations (Exp), multiplications (Mult), hash evaluations (Hash) and AES encryptions/decryptions (AES). Since the server mostly just relays messages, we ignore its costs in the comparison. In case of encPw and WBP, we assume that the static encryption key of the HSM is hardcoded into the client, therefore no communication is needed for key sharing in Init and Rec and we neglect the costs of the one-time generation. For more details on the instantiations of the primitives used for the comparison, see Table 4.

		encPw (Sec. 4.1)	encPw+ (Sec. 4.2)	OPRF-PPKR (Sec. 4.3)	WBP [DFG+23b]
Init	Client	2 Exp, 3 Hash, 1 AES	2 Exp, 3 Hash, 1 AES	5 Exp, 5 Hash, 2 AES	7 Exp, 12 Hash, 1 AES, 1 Mult
	HSM	3 Hash, 1 AES	2 Exp, 5 Hash, 1 AES	2 Exp, 3 Hash, 1 AES	3 Exp, 3 Hash
	no. rounds	2	3	3	3
Rec	Client	2 Exp, 3 Hash, 2 AES	2 Exp, 3 Hash, 2 AES	3 Exp, 3 Hash, 1 AES	8 Exp, 27 Hash, 2 AES, 1 Mult
	HSM	3 Hash, 2 AES	2 Exp, 5 Hash, 2 AES	2 Exp, 1 Hash, 1 Mult	6 Exp, 15 Hash, 1 AES
	no. rounds	2	3	3	4

Table 4: Costs of concrete building blocks for the efficiency evaluation in Section 5.

	KeyGen	Enc	Dec
CPA Enc (ElGamal)	1 Exp	2 Exp, 1 Hash	1 Exp, 1 Hash
CCA Enc (DHIES)	1 Exp	2 Exp, 3 Hash, 1 AES	1 Exp, 3 Hash, 1 AES
AE (AES-GCM)	-	1 AES	1 AES

	KeyGen	Sign	Vfy
Signature (Schnorr)	1 Exp	1 Exp, 1 Hash	2 Exp, 1 Mult, 1 Hash
MAC (HMAC)	-	2 Hash	2 Hash

KDF (n keys, HKDF)	$(2n + 2)$ Hash
-----------------------	-----------------

as it will be run more often than initialization, our protocol outperforms WBP: OPRF-PPKR reduces the round⁴ complexity from 4 to 3, and uses roughly one-third of the operations required by WBP, for both the client and HSM.

If we compare the two enhanced encrypt-to-HSM protocols to OPRF-PPKR, they are more efficient and they save one round of communication, as they strongly rely on the HSM security, with the weakest protocol having the lowest computational requirements. Interestingly, in bare numbers, they are not significantly more efficient though.

Thus, considering that OPRF-PPKR provides much better security for almost the same costs, this might raise the question of whether there are any advantages in using our simpler protocols – which is what we answer next.

Advantages of Standard Primitives. The core benefit of our two basic/enhanced encrypt-to-HSM protocols is that they explore how the extended trust in the HSM can be traded for simplicity in the protocol design. Both protocols rely on standard and well-understood primitives only, which are public-key and symmetric encryption, and hash functions. In contrast, WBP and our OPRF-PPKR require (dedicated discrete-log-based) OPRFs.

In particular, when internally using an HSM, reliance on simple and standard building blocks is an advantage – established primitives are implemented in many well-tested frameworks and have been studied for resistance against side-channel attacks. OPRFs are

⁴We count one round as a message from party A to party B (and not as a full round-trip A to B to A).

still a somewhat more modern primitive that just recently came into focus of practitioners, i.e., it might require developers to implement low-level cryptographic procedures instead of merely invoking APIs of trusted libraries. Thus, the operations needed for OPRF-PPKR and WBP might be more prone to implementation errors or be simply not available or accessible through the shielded HSM APIs when using “off-the-shelf” HSMs. The clear downside of both simple protocols is that they lose security when the HSM is (fully) compromised. However, given that they require only well-understood operations by the HSM, such a simpler HSM might be easier to protect, making a security breach less likely.

Further, realizing an efficient *quantum-safe* OPRF is still an open problem. This is in contrast to the other standard building blocks for which quantum-safe options exist. Here the basic/enhanced encrypt-to-HSM protocols again have the advantage over OPRF-PPKR and WBP, as they rely on standard primitives only, and can easily benefit from a post-quantum “upgrade”.

Lev-3 Security for Offline Users in enhanced encrypt-to-HSM. While we prove the enhanced encrypt-to-HSM protocol to satisfy at most Lev-2 security, it actually does preserve strong guarantees if the HSM is fully corrupted, but only for “offline” clients. Recall that we assumed the HSM’s attestation key to be the single value that enjoys particularly strong protection – and thus is the only additional information the adversary gets upon full server corruption. Consequently, the impact of such a compromise on the protocol’s security is then rather limited. In fact, for users who never use the PPKR *after* the full corruption happened, our enhanced encrypt-to-HSM protocol provides the same protection as OPRF-PPKR: their leaked files must still be cracked through individual offline attacks against each password. This might be a sufficient guarantee in reality. Especially in settings where it will be known which HSM is used by the server, and assuming that a breach of the HSM most critical operation would become public. Then either the service is stopped, and the server updates to a secure HSM; or the particularly cautious users would no longer log into the PPKR service – and their password and key would be just as secure as with the OPRF-PPKR protocol. Only the users who still engage in active sessions would lose their security. However, their security is fully compromised in enhanced encrypt-to-HSM, as the malicious server will learn their password and key in plain as soon as they start a new init or

recovery session. Here, OPRF-PPKR still ensures the secrecy of the user's data.

What is the best protocol? Overall, there is no clear answer to what can be considered the "best" protocol. We believe that both our enhanced encrypt-to-HSM and OPRF-PPKR protocols have their individual strengths that can make each the right choice for a dedicated deployment setting. The enhanced encrypt-to-HSM is clearly superior to its unsalted variant, yet preserves all simplicity advantages. Thus, here we do not see a strong reason to favor the π^{encPw} protocol (Lev-1 security) and would recommend opting for the $\pi^{\text{encPw+}}$ version (Lev-2 security) whenever the advantages of the simple and standard construction outweigh the concerns of a full HSM corruption. As just discussed, the guarantees of $\pi^{\text{encPw+}}$ in case of full corruption are actually not lost entirely. For those users who never use the retrieval again after the full HSM corruption occurred, it provides the same security as $\pi^{\text{OPRF-PPKR}}$. Nevertheless, for applications with very high-security requirements, the $\pi^{\text{OPRF-PPKR}}$ construction provides the strongest (Lev-3) guarantees, but requires more care in the implementation.

Acknowledgments

The authors would like to thank the anonymous referees for their valuable comments and helpful suggestions.

Tobias Handirk has been supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme, grant agreement 802823. Julia Hesse was supported by the Swiss National Science Foundation (SNSF) under the AMBIZIONE grant "Cryptographic Protocols for Human Authentication and the IoT". Máté Horváth has been supported by the German Research Foundation (DFG), project JA2445/6-1.

References

- [AMMR18] Shashank Agrawal, Payman Mohassel, Pratyay Mukherjee, and Peter Rindal. DiSE: Distributed symmetric-key encryption. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018: 25th Conference on Computer and Communications Security*, pages 1993–2010, Toronto, ON, Canada, October 15–19, 2018. ACM Press.
- [BC19] Jean-Baptiste Bedrune and Gabriel Campana. Everybody be cool, this is a robbery! BlackHat USA 2019, 2019. <http://i.blackhat.com/USA-19/Thursday/us-19-Campana-Everybody-Be-Cool-This-Is-A-Robbery.pdf>, Accessed: 24.04.2024.
- [BJS11] Ali Bagherzandi, Stanislaw Jarecki, Nitesh Saxena, and Yanbin Lu. Password-protected secret sharing. In Yan Chen, George Danezis, and Vitaly Shmatikov, editors, *ACM CCS 2011: 18th Conference on Computer and Communications Security*, pages 433–444, Chicago, Illinois, USA, October 17–21, 2011. ACM Press.
- [BM93] Steven M. Bellovin and Michael Merritt. Augmented encrypted key exchange: A password-based protocol secure against dictionary attacks and password file compromise. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *ACM CCS 93: 1st Conference on Computer and Communications Security*, pages 244–250, Fairfax, Virginia, USA, November 3–5, 1993. ACM Press.
- [BMP00] Victor Boyko, Philip D. MacKenzie, and Sarvar Patel. Provably secure password-authenticated key exchange using Diffie-Hellman. In Bart Preneel, editor, *Advances in Cryptology – EUROCRYPT 2000*, volume 1807 of *Lecture Notes in Computer Science*, pages 156–171, Bruges, Belgium, May 14–18, 2000. Springer, Heidelberg, Germany.
- [BS23] Dan Boneh and Victor Shoup. A graduate course in applied cryptography. <http://toc.cryptobook.us/>, 2023.
- [Can00] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. Cryptology ePrint Archive, Report 2000/067, 2000. <https://eprint.iacr.org/2000/067>.
- [Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd Annual Symposium on Foundations of Computer Science*, pages 136–145, Las Vegas, NV, USA, October 14–17, 2001. IEEE Computer Society Press.
- [CGMS21] Mihai Christodorescu, Sivanarayana Gaddam, Pratyay Mukherjee, and Rohit Sinha. Amortized threshold symmetric-key encryption. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021: 28th Conference on Computer and Communications Security*, pages 2758–2779, Virtual Event, Republic of Korea, November 15–19, 2021. ACM Press.
- [CLLN14] Jan Camenisch, Anja Lehmann, Anna Lysyanskaya, and Gregory Neven. Memento: How to reconstruct your secrets from a single password in a hostile environment. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology – CRYPTO 2014, Part II*, volume 8617 of *Lecture Notes in Computer Science*, pages 256–275, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Heidelberg, Germany.
- [DFG+23a] Gareth T. Davies, Sebastian Faller, Kai Gellert, Tobias Handirk, Julia Hesse, Máté Horváth, and Tibor Jager. Security analysis of the whatsapp end-to-end encrypted backup protocol. Cryptology ePrint Archive, Paper 2023/843, 2023. <https://eprint.iacr.org/2023/843>.
- [DFG+23b] Gareth T. Davies, Sebastian F. Faller, Kai Gellert, Tobias Handirk, Julia Hesse, Máté Horváth, and Tibor Jager. Security analysis of the WhatsApp end-to-end encrypted backup protocol. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023, Part IV*, volume 14084 of *Lecture Notes in Computer Science*, pages 330–361, Santa Barbara, CA, USA, August 20–24, 2023. Springer, Heidelberg, Germany.
- [DHL22] Poulami Das, Julia Hesse, and Anja Lehmann. DPaSE: Distributed password-authenticated symmetric-key encryption, or how to get many keys from one password. In Yuji Suga, Kouichi Sakurai, Xuhua Ding, and Kazuo Sako, editors, *ASIACCS 22: 17th ACM Symposium on Information, Computer and Communications Security*, pages 682–696, Nagasaki, Japan, May 30 – June 3, 2022. ACM Press.
- [FHH+24] Sebastian Faller, Tobias Handirk, Julia Hesse, Máté Horváth, and Anja Lehmann. Password-protected key retrieval with(out) HSM protection. In *Proceedings of the 2024 ACM SIGSAC Conference on Computer and Communications Security (CCS '24)*, October 14–18, 2024, Salt Lake City, UT, USA. ACM, 2024.
- [GMR06] Craig Gentry, Philip MacKenzie, and Zulfikar Ramzan. A method for making password-based key exchange resilient to server compromise. In Cynthia Dwork, editor, *Advances in Cryptology – CRYPTO 2006*, volume 4117 of *Lecture Notes in Computer Science*, pages 142–159, Santa Barbara, CA, USA, August 20–24, 2006. Springer, Heidelberg, Germany.
- [JKKX16] Stanislaw Jarecki, Aggelos Kiayias, Hugo Krawczyk, and Jiayu 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, Saarbrücken, Germany, March 21–24, 2016*, pages 276–291. IEEE, 2016.
- [JKX18] Stanislaw Jarecki, Hugo Krawczyk, and Jiayu Xu. OPAQUE: An asymmetric PAKE protocol secure against pre-computation attacks. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018, Part III*, volume 10822 of *Lecture Notes in Computer Science*, pages 456–486, Tel Aviv, Israel, April 29 – May 3, 2018. Springer, Heidelberg, Germany.
- [Krs16] Ivan Krstic. Behind the scenes with ios security, 2016. <https://www.blackhat.com/docs/us-16/materials/us-16-Krstic.pdf>, Accessed: 18.04.2024.
- [Lun19] Joshua Lund. Technology preview for secure value recovery, 2019. <https://signal.org/blog/secure-value-recovery/>, Accessed: 18.04.2024.
- [OSV23] Chris Orsini, Alessandra Scafuro, and Tanner Verber. How to recover a cryptographic secret from the cloud. Cryptology ePrint Archive, Paper 2023/1308, 2023. <https://eprint.iacr.org/2023/1308>.
- [PST17] Rafael Pass, Elaine Shi, and Florian Tramèr. Formal abstractions for attested execution secure processors. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *Advances in Cryptology – EUROCRYPT 2017, Part I*, volume 10210 of *Lecture Notes in Computer Science*, pages 260–289, Paris, France, April 30 – May 4, 2017. Springer, Heidelberg, Germany.
- [Sca19] Alessandra Scafuro. Break-glass encryption. In Dongdai Lin and Kazuo Sako, editors, *PKC 2019: 22nd International Conference on Theory and Practice of Public Key Cryptography, Part II*, volume 11443 of *Lecture Notes in Computer Science*, pages 34–62, Beijing, China, April 14–17, 2019. Springer, Heidelberg, Germany.
- [SRW22] Alon Shokevsky, Eyal Ronen, and Avishai Wool. Trust dies in darkness: Shedding light on samsung's TrustZone keymaster design. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 251–268, Boston, MA, August 2022. USENIX Association.
- [VBMW+18] Jo Van Bulck, Marina Minkin, Ofir Weiss, Daniel Genkin, Baris Kasikci, Frank Piessens, Mark Silberstein, Thomas F Wenihs, Yuval Yarom, and Raoul Strackx. Foreshadow: Extracting the keys to the intel {SGX} kingdom with transient {Out-of-Order} execution. In *27th USENIX Security Symposium (USENIX Security 18)*, pages 991–1008, 2018.

- [VBPS17] Jo Van Bulck, Frank Piessens, and Raoul Strackx. Sgx-step: A practical attack framework for precise enclave execution control. In *Proceedings of the 2nd Workshop on System Software for Trusted Execution*, pages 1–6, 2017.
- [Wal18] Shabsi Walfish. Google cloud key vault service, 2018. <https://developer.android.com/about/versions/pie/security/ckv-whitepaper>, Accessed: 18.04.2024.
- [WH20] Xunhua Wang and Ben Huson. Robust distributed symmetric-key encryption. Cryptology ePrint Archive, Report 2020/1001, 2020. <https://eprint.iacr.org/2020/1001>.
- [Wha21] WhatsApp. Security of End-to-End Encrypted Backups, September 2021. https://www.whatsapp.com/security/WhatsApp_Security_Encrypted_Backups_Whitepaper.pdf, Accessed: 18.04.2024.

Appendix

A Comparison with $\mathcal{F}_{\text{PPKR}}$ of [DFG⁺23b]

In this part, we summarize the changes that we introduced compared to the ideal functionality $\mathcal{F}_{\text{PPKR}}$ in Figures 5 and 6, defined by [DFG⁺23b]. We can group the differences into the following categories.

Extension with HSM corruption.

- Added LEAKFILE interface to capture leakage of password files. The functionality keeps records $\langle \text{LEAKED}, \text{ID}_C, \text{pw}, K, i \rangle$ of leaked files (see LS.2), where the counter value i distinguishes between different leakages that affected the data of ID_C . We note that running Init with a fully corrupt S also has the effect of leaking the password file (see CIS.4).
- Added OFFLINEATTACK interface to model both offline password guessing against leaked records, and offline password guessing against currently and previously stored records by a fully corrupt HSM.
- Earlier the server had two corruption states, HONEST, and CORRUPT, that we extended with a third state called FULLYCORRUPT to model the case when the server has full control over the HSM (see the (FULLYCORRUPT, S) interface). This new corruption state also requires changes of the Init and Rec interfaces (see CIS.3, CIS.4, and RS.2, RS.3 respectively). For ease of expression, we set all ctr values to ∞ when the server gets FULLYCORRUPT.
- Added key planting in recovery in RS.2. A FULLYCORRUPT server may submit a password guess and key to be planted. If it submits a correct password guess, the client recovers the planted key. If it submits a wrong password guess, the client fails. If S is not FULLYCORRUPT or it doesn't submit a password guess, recovery works as before.

Changes affecting the security of PPKR.

- [DFG⁺23b] artificially weakened the PPKR functionality to be able to analyze WBP that allows a malicious server to prevent the erasure of password files when a client re-runs Init (see boxed code on figs. 5 to 6). We strengthened our functionality by disabling this attack.
- We identified and fixed a bug in functionality from Davies et al. [DFG⁺23b]. In particular, WBP and our OPRF-PPKR always enable resetting the counter in recoveries for files, stored in initialization by corrupt client or corrupt server. However, this was not captured in the ideal functionality of [DFG⁺23b] (see CRC.2 and CRS.3). We fix this problem by marking the file records either HONEST or MALICIOUS (see CIS.3) and allow for resetting the counter in case of a file record marked MALICIOUS (see CRS.2).
- In the original PPKR functionality, S receives an output after invoking the INITC and REC interfaces (see IC.3 and RC.2 respectively). When dealing with a corrupt client, the simulator needs early-extract the password to be able to send the proper input to the functionality. However, early extraction is impossible in any OPRF-based PPKR protocol that is instantiated with an OPRF that information-theoretically hides the PRF input (i.e. the password), such as 2HashDH, that is used by both WBP and our OPRF-PPKR. The problem is that we cannot extract the password from the client's first OPRF message (the security proof in [DFG⁺23b] missed this problem) so the simulation runs into a problem. We fix this issue by modifying the functionality so that S does not get output anymore (see IC.4 and RC.3 respectively), which is in line with our other modification that sessions can be initiated by both parties.

Broadening the applicability.

- Allowed arbitrary order of inputs in Init (see IC.2, IS.2) and Recovery (see RC.1, RS.1), instead of sessions always being initiated by the client (see IS.1 and IR.2 respectively).

Syntactical changes.

- Integrated former DoS interfaces called COMPLETEINITC-DoS, COMPLETEINITS-DoS, COMPLETEREC-DoS and COMPLETERECS-DoS into adversarial complete interfaces (see CIC.3, CIS.2, CRC.2 and CRS.2 respectively) via b_C, b_S bits indicating whether a DoS attack is successful (1) or not (0).

Recovery Phase	
On input $(\text{REC}, \text{sid}, \text{ID}_C, \text{pw}')$ from ID_C (or \mathcal{A} if ID_C is corrupt):	
RC.1	Record $(\text{REC}, \text{sid}, \text{ID}_C, \text{pw}')$, overwriting any existing record $(\text{REC}, \text{sid}, \text{ID}_C, *)$ // Storing ID_C 's current init state; a client can only be in one recovery session.
RC.2	Send $(\text{REC}, \text{sid}, \text{ID}_C)$ to \mathcal{A} and S . // S learns which clients started recovery, with the guarantee that attempts by honest clients cannot be faked.
On input $(\text{REC}, \text{sid}, \text{ID}_C, \text{ID}_C^*)$ from S (or \mathcal{A} if S is corrupt): // Server agrees to assist ID_C in recovery. If S is corrupt, \mathcal{A} can reroute the recovery to a different ID_C^* . Note that ID_C^* does not have to be the identity of an existing client and can be an arbitrary identity	
RS.1	If S is honest, set $\text{ID}_C' \leftarrow \text{ID}_C$, otherwise set $\text{ID}_C' \leftarrow \text{ID}_C^*$
RS.2	Retrieve record $(\text{REC}, \text{sid}, \text{ID}_C, [\text{pw}'])$ // Continue only if ID_C started recovery already
RS.3	If there exists no record $(\text{FILE}, \text{ID}_C', \text{sid}, [\text{pw}], [K])$ marked STORED , send $(\text{RECRES}, \text{aid}, \text{FAIL})$ to S (or \mathcal{A} if S is corrupt). Else retrieve the record. // The currently stored K and pw (for ID_C') are used. If ID_C' re-initiates afterwards, it has no effect on this recovery session.
RS.4	If $\text{tx}_{\text{sid}}[\text{ID}_C'] = 0$, delete record $(\text{FILE}, \text{sid}, \text{ID}_C', \text{pw}, K)$ marked STORED and send $(\text{DELREC}, \text{sid}, \text{ID}_C')$ to S and \mathcal{A} Else continue.
RS.5	Set $\text{tx}_{\text{sid}}[\text{ID}_C'] \leftarrow \text{tx}_{\text{sid}}[\text{ID}_C'] - 1$
RS.6	Append pw and K to record $(\text{REC}, \text{sid}, \text{ID}_C, \text{pw}')$, overwriting any existing record $(\text{REC}, \text{sid}, \text{ID}_C, *, *, *)$ // The recovery session of ID_C is used
RS.7	Send $(\text{REC}, \text{sid}, \text{ID}_C', \text{pw} \stackrel{?}{=} \text{pw}')$ to \mathcal{A} // A ppKR protocol may not hide whether recovery was successful or not
On input $(\text{COMPLETEREC}, \text{sid}, \text{ID}_C)$ from \mathcal{A} :	
CRC.1	Retrieve record $(\text{REC}, \text{sid}, \text{ID}_C, [\text{pw}'], [\text{pw}], [K])$. If record is marked RECOVERED , delete it. Otherwise, mark it RECOVERED . // Ensures that record can be retrieved twice before deletion.
CRC.2	Determine the output as follows: (1) If $\text{pw} = \text{pw}'$ then set $K' \leftarrow K$ // Recovering the key! (2) In all other cases, set $K' \leftarrow \text{FAIL}$
CRC.3	Send $(\text{RECRES}, \text{sid}, K')$ to ID_C (or \mathcal{A} if ID_C is corrupt).
On input $(\text{COMPLETERECS}, \text{sid}, \text{ID}_C)$ from \mathcal{A} : // Server finishes recovery session by learning whether the password was correct or not:	
CRS.1	Retrieve record $(\text{REC}, \text{sid}, \text{ID}_C, [\text{pw}], [\text{pw}'], [K])$. If record is marked RECOVERED , delete it. Otherwise, mark it RECOVERED . // Ensures that record can be retrieved twice before deletion.
CRS.2	If $\text{pw} = \text{pw}'$, set $\text{tx}_{\text{sid}}[\text{ID}_C] \leftarrow 10$ and send $(\text{RECRES}, \text{sid}, \text{ID}_C, \text{Succ})$ to S .
CRS.3	If $\text{pw} \neq \text{pw}'$, then send $(\text{RECRES}, \text{sid}, \text{ID}_C, \text{FAIL})$ to S (or \mathcal{A} if S is corrupt).
<u>Attacks on Recovery Phase</u>	
On input $(\text{COMPLETEREC-DoS}, \text{sid}, \text{ID}_C)$ from \mathcal{A} // Network attacker or malicious server can always make the client fail:	
CRCD.1	Retrieve record $(\text{REC}, \text{sid}, \text{ID}_C, *, *, *)$ and delete it.
CRCD.2	Send $(\text{RECRES}, \text{sid}, \text{FAIL})$ to ID_C (or \mathcal{A} if ID_C is corrupt).
On input $(\text{COMPLETERECS-DoS}, \text{sid}, \text{ID}_C)$ from \mathcal{A} : // Server finishes with failure. In particular, it never learns if the password was correct	
CRSC.1	Retrieve record $(\text{REC}, \text{sid}, \text{ID}_C, *, *, *)$ and delete it.
CRSD.2	Output $(\text{RECRES}, \text{sid}, \text{ID}_C, \text{FAIL})$ to S (or \mathcal{A} if S is corrupt).

Figure 6: Ideal functionality $\mathcal{F}_{\text{PPKR}}$, original version from [DFG⁺23b], cont'd (recovery interfaces). Boxed code reflects an attack on WBP.

$\mathcal{F}_{\text{PPKR}}$ is parameterized with a security parameter λ . $\mathcal{F}_{\text{PPKR}}$ talks to a server S where S is encoded in sid . $\mathcal{F}_{\text{PPKR}}$ also talks to the adversary \mathcal{A} , and arbitrary clients ID_C . If the functionality tries to retrieve a record that does not exist, it ignores the incoming message. We write $\text{tx}_{\text{sid}}[\cdot]$ for a list of counters.	
<u>Offline attacks</u>	
On input $(\text{MALICIOUSINIT}, \text{sid}, \text{ID}_C, \text{pw}^*, K^*)$ from \mathcal{A} : // A corrupt server can impersonate an either honest or corrupt ID_C and initialize on his behalf.	
MI.1	If S is honest ignore this input.
MI.2	Record $(\text{FILE}, \text{sid}, \text{ID}_C, \text{pw}^*, K^*)$, overwriting any existing record $(\text{FILE}, \text{sid}, \text{ID}_C, *, *)$. Set $\text{tx}_{\text{sid}}[\text{ID}_C] \leftarrow 10$
On input $(\text{MALICIOUSREC}, \text{sid}, \text{ID}_C, \text{pw}^*)$ from \mathcal{A} : // Attacking an honest client's stored key: bury the key after 10 subsequent wrong password guesses.	
MR.1	If S is honest ignore this input. // Server needs to be corrupt to mount an offline attack.
MR.2	Retrieve record $(\text{FILE}, \text{sid}, \text{ID}_C, [\text{pw}], [K])$ marked STORED .
MR.3	If $\text{tx}_{\text{sid}}[\text{ID}_C] = 0$, delete record $(\text{FILE}, \text{sid}, \text{ID}_C, \text{pw}, K)$ and output $(\text{DELREC}, \text{sid}, \text{ID}_C)$ to \mathcal{A} // The key is buried if zero guesses remain.
MR.4	If $\text{pw}^* = \text{pw}$, set $\text{tx}_{\text{sid}}[\text{ID}_C] \leftarrow 10$ and output (sid, K) to \mathcal{A} . Otherwise, set $\text{tx}_{\text{sid}}[\text{ID}_C] \leftarrow \text{tx}_{\text{sid}}[\text{ID}_C] - 1$ an output $(\text{sid}, \text{FAIL})$ to \mathcal{A}
<u>Initialization phase</u>	
On input $(\text{INITC}, \text{sid}, \text{ID}_C, \text{pw})$ from ID_C (or \mathcal{A} if ID_C is corrupt): // Client always starts initialization	
IC.1	Choose $K \xleftarrow{\$} \{0, 1\}^\lambda$
IC.2	Record $(\text{INITC}, \text{sid}, \text{ID}_C, \text{pw}, K)$, overwriting any existing record $(\text{INITC}, \text{sid}, \text{ID}_C, *, *)$ // Storing ID_C 's current init state; a client can only be in one initialization session
IC.3	Send $(\text{INITC}, \text{sid}, \text{ID}_C)$ to \mathcal{A} and to S
On input $(\text{INITS}, \text{sid}, \text{ID}_C, [\text{ID}_C^*])$ from S (or \mathcal{A} if S is corrupt): // Server agrees to assist ID_C in initialization. If S is corrupt, \mathcal{A} can reroute the initialization to a different ID_C^* . Note that ID_C^* does not have to be the identity of an existing client and can be an arbitrary identity	
IS.1	Retrieve $(\text{INITC}, \text{sid}, \text{ID}_C, [\text{pw}], [K])$ // Continue only if ID_C started initialization already
IS.2	If S is honest, record $(\text{FILE}, \text{sid}, \text{ID}_C, \text{pw}, K)$, overwriting any existing record $(\text{FILE}, \text{sid}, \text{ID}_C, *, *)$, and send $(\text{INITS}, \text{sid}, \text{ID}_C)$ to \mathcal{A} // Storing S 's current init state; the server can only be in one initialization session. Invariant: There is only one key stored
IS.3	Otherwise, record $(\text{FILE}, \text{sid}, \text{ID}_C^*, \text{pw}, K)$, overwriting any existing record $(\text{FILE}, \text{sid}, \text{ID}_C^*, *, *)$.
On input $(\text{COMPLETEINITC}, \text{sid}, \text{ID}_C)$ from \mathcal{A} : // Client completes the protocol and outputs a key	
CIC.1	Retrieve record $(\text{INITC}, \text{sid}, \text{ID}_C, *, [K])$ and delete it
CIC.2	Output $(\text{INITRES}, \text{sid}, K)$ to ID_C (or \mathcal{A} if ID_C is corrupt)
On input $(\text{COMPLETEINITS}, \text{sid}, \text{ID}_C)$ from \mathcal{A} : // Server concludes initialization with file storage	
CIS.1	Retrieve record $(\text{FILE}, \text{sid}, \text{ID}_C, *, *)$ not marked STORED and mark it STORED // Note: there is only one such record thanks to overwriting in INITS interface. This becomes the stored key now!
CIS.2	Set $\text{tx}_{\text{sid}}[\text{ID}_C] \leftarrow 10$
CIS.3	Send $(\text{INITRES}, \text{sid}, \text{ID}_C, \text{Succ})$ to S (or \mathcal{A} if S is corrupt)
<u>Attacks on Initialization Phase</u>	
On input $(\text{COMPLETEINITC-DoS}, \text{sid}, \text{ID}_C)$ from \mathcal{A} : // DoS attack against ID_C , who concludes the initialization session with failure.	
CICD.1	Retrieve record $(\text{INITC}, \text{sid}, \text{ID}_C, *, *)$ and delete it
CICD.2	Output $(\text{INITRES}, \text{sid}, \text{FAIL})$ to ID_C (or \mathcal{A} if ID_C is corrupt)
On input $(\text{COMPLETEINITS-DoS}, \text{sid}, \text{ID}_C)$ from \mathcal{A} : // DoS attack against the server, such that it cannot store a file	
CISD.1	Delete any record $(\text{FILE}, \text{sid}, \text{ID}_C, *, *)$ // Server's state in current initialization session no longer needed
CISD.2	Send $(\text{INITRES}, \text{sid}, \text{ID}_C, \text{FAIL})$ to S (or \mathcal{A} if S is corrupt)

Figure 5: Ideal functionality $\mathcal{F}_{\text{PPKR}}$ for password-protected key retrieval, offline attacks and initialization interfaces, original version from [DFG⁺23b]. Boxed code reflects an attack on WBP.

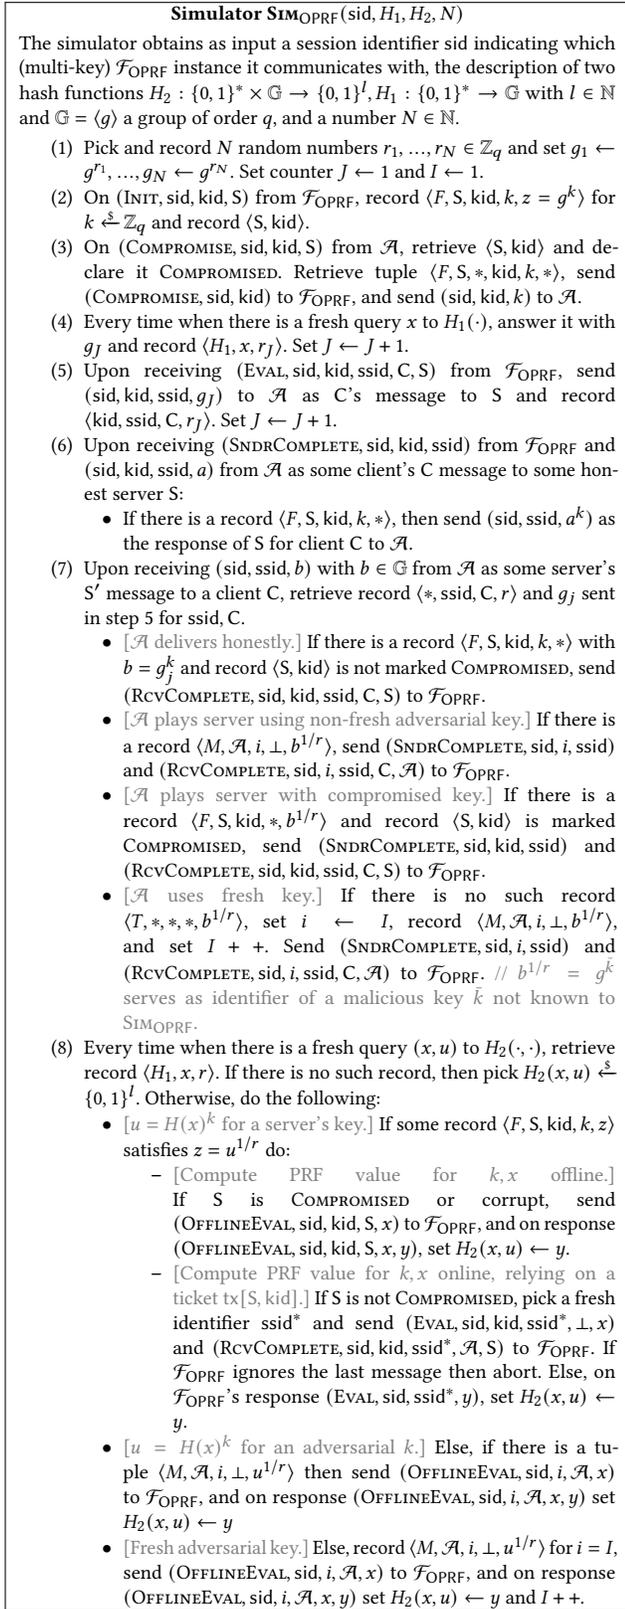


Figure 8: The simulator that demonstrates that “multi-key” 2HashDH UC-realizes our “multi-key” $\mathcal{F}_{\text{OPRF}}$, restated from [DFG⁺23b].

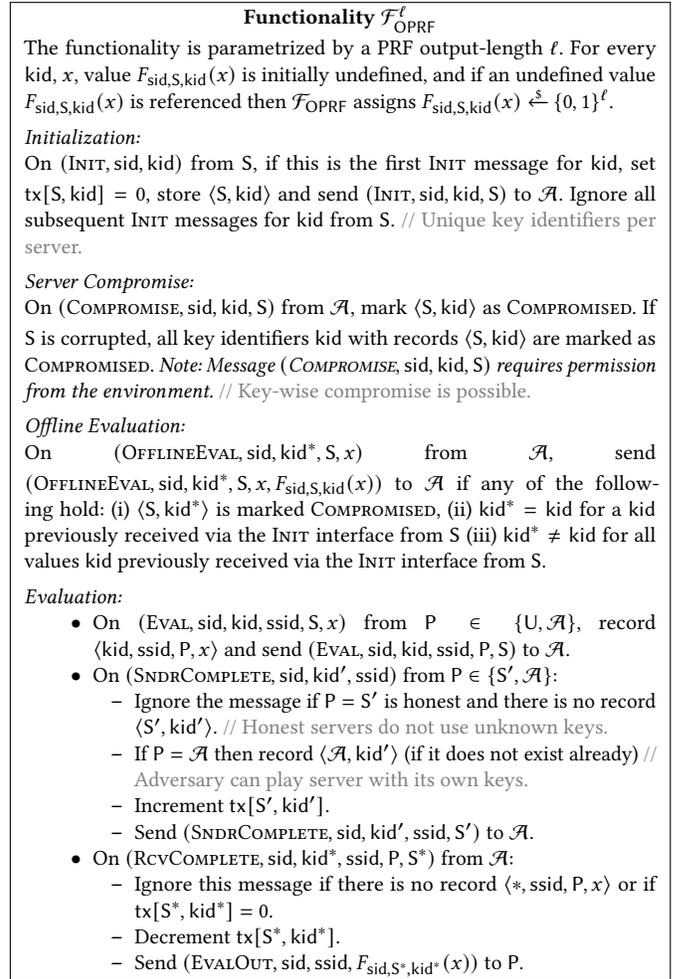


Figure 7: A multi-key version of the ideal functionality $\mathcal{F}_{\text{OPRF}}$ [DFG⁺23b].

B The 2HashDH simulator

We restate a result of Davies et al. [DFG⁺23b] about the multi-session security of the 2HashDH OPRF protocol of Jarecki et al [JKKX16]. The security is proven under the following assumption.

Definition 4 ((N, Q) one-more DH assumption [JKKX16]). The (N, Q) one-more Diffie–Hellman (DH) assumption holds in a cyclic group $\mathbb{G} = \langle g \rangle$ if for any polynomial-time adversary \mathcal{A} ,

$$\text{Adv}_{\mathcal{A}, \mathbb{G}}^{(N, Q)\text{-OMDH}}(\lambda) := \Pr_{k \xleftarrow{\$} \mathbb{Z}_q, g_i \xleftarrow{\$} \mathbb{G}} \left[\mathcal{A}^{(\cdot)^k, \text{DDH}(\cdot, \cdot, \cdot)}(g, g^k, g_1, \dots, g_N) = S \right]$$

is negligible, where $S = \{(g_{j_s}, g_{j_s}^k) \mid s = 1, \dots, Q + 1\}$ with $j_s \in [N]$ for $s \in [Q + 1]$ and

- $(\cdot)^k$ is an exponentiation oracle that \mathcal{A} can query Q times and on input $h \in \mathbb{G}$ it returns h^k ;

- $\text{DDH}(\cdot, \cdot, \cdot, \cdot)$ is a Diffie–Hellman oracle that takes as input (g, g^k, g^x, g^y) with $g \in \mathbb{G}$ and returns 1 if $y = kx$ and 0 otherwise.

Theorem 4. Let $H_1 : \{0, 1\}^* \rightarrow \mathbb{G}$ be a hash function into a group of order $q \in \mathbb{N}$, $H_2 : \{0, 1\}^* \times \mathbb{G} \rightarrow \{0, 1\}^\ell$ with $\ell \in \mathbb{N}$ be another hash function, and let $k \xleftarrow{\$} \mathbb{Z}_q$. Suppose the (N, Q) one-more DH assumption holds for \mathbb{G} , where $Q := q_E$ is the maximum number of $(\text{EVAL}, *, \text{kid}, S, *)$ queries over all tuples (kid, S) made by the environment \mathcal{Z} , $N := q_E + q_H$, and q_H is the total number of H_1 queries made by \mathcal{Z} . Then the “multi-key” protocol 2HashDH UC-realizes

the “multi-key” functionality $\mathcal{F}_{\text{OPRF}}$ of Figure 7, with hash functions H_1, H_2 modeled as random oracles.

More precisely, for any adversary against 2HashDH , there is a simulator SIM_{OPRF} that interacts with $\mathcal{F}_{\text{OPRF}}$ and produces a view that no environment \mathcal{Z} can distinguish with advantage better than

$$\Pr[\text{FAIL}] \leq q_I \text{Adv}_{\mathcal{A}, \mathbb{G}}^{(q_E + q_H, q_E)\text{-OMDH}}(\lambda) + (q_E + q_H)^2 / q,$$

where q_I is the number of honestly initialized keys in the system.

The full proof can be found in [DFG⁺23b]. Because in our work we will make use of the simulator of this statement, we restate it in Figure 8.

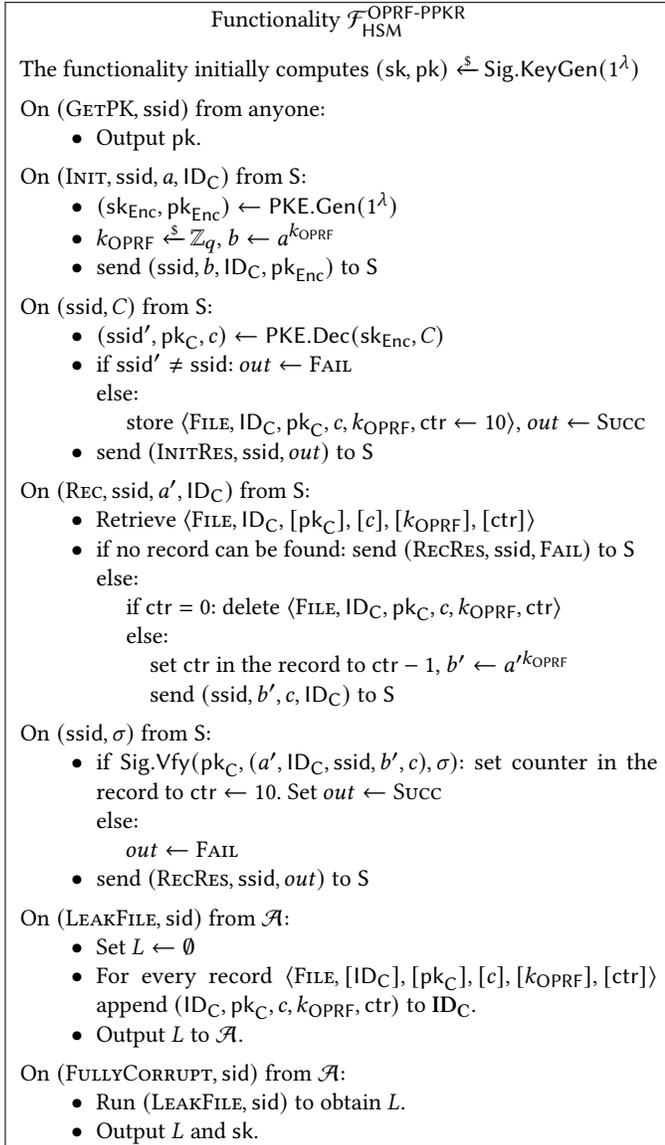


Figure 10: The ideal functionality $\mathcal{F}_{\text{HSM}}^{\text{OPRF-PPKR}}$.

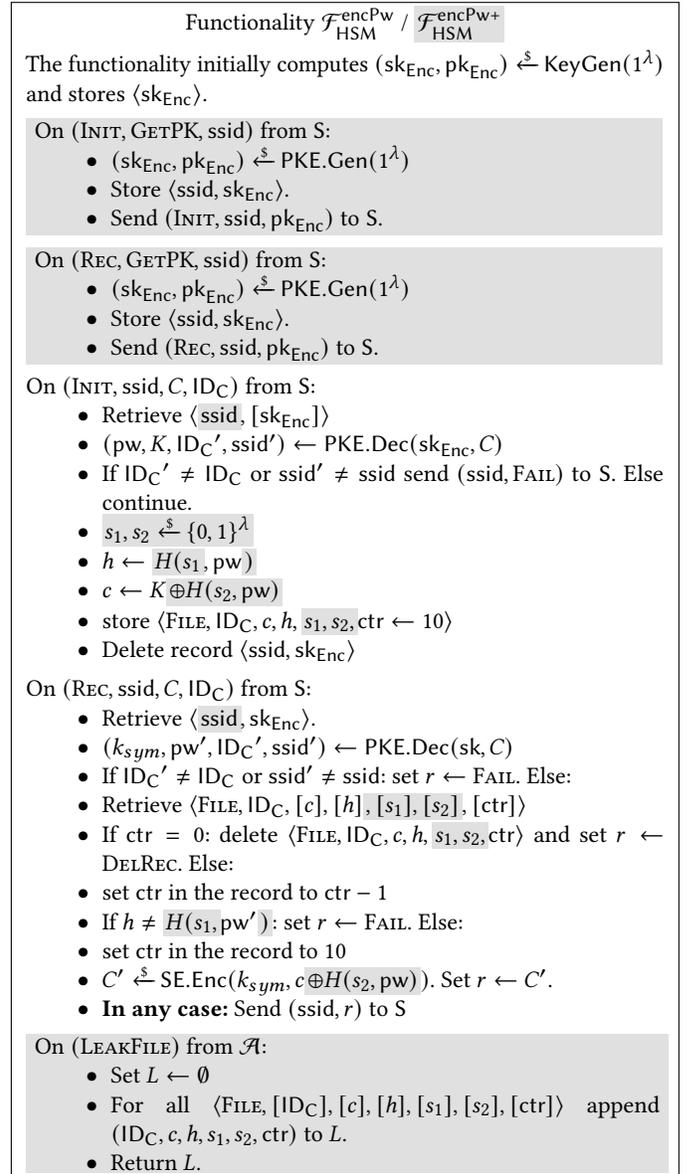


Figure 9: The ideal functionalities $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$ and $\mathcal{F}_{\text{HSM}}^{\text{encPw+}}$. The code in gray boxes is only executed in $\mathcal{F}_{\text{HSM}}^{\text{encPw+}}$.

C HSM Functionalities

We define the HSM functionalities used by our protocols in Figure 9 and Figure 10.

D Full proofs

In this section, we render the full proofs for Theorems 1 to 3.

D.1 Proof of Theorem 1

PROOF. We construct a sequence of hybrid games \mathbf{G}_0 to \mathbf{G}_{14} where we gradually change the real-world execution of the protocol π^{encPw} (interacting with the hybrid functionality $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$) to reach the ideal-world execution, where the environment interacts with the simulator from Figures 11 to 12 and the ideal functionality $\mathcal{F}_{\text{PPKR}}$. We write $\Pr[\mathbf{G}_i]$ to denote the probability that the environment outputs 1 in the hybrid game \mathbf{G}_i .

Game \mathbf{G}_0 : Real world. This is the real world.

Game \mathbf{G}_1 : Create simulator. In this game we create two new entities called the ideal functionality \mathcal{F} and the simulator SIM . Initially, \mathcal{F} just forwards the input of the dummy parties to SIM and outputs what SIM instructs it to output. In particular, \mathcal{F} has interfaces $(\text{INITC}, \text{ssid}, \text{pw})$, $(\text{INITS}, \text{ssid}, \text{ID}_C)$, $(\text{RECC}, \text{ssid}, \text{pw}')$, and $(\text{RECS}, \text{ssid}, \text{ID}_C, \text{pw}^*, K^*, i)$ that just forward the input to SIM . The simulator executes the code of all honest parties of the protocol internally on the input that it is provided by \mathcal{F} and it internally runs the code of the hybrid functionality $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$. Note that these are just syntactical changes and the protocol is still executed as in the real world. We have

$$\Pr[\mathbf{G}_1] = \Pr[\mathbf{G}_0].$$

Game \mathbf{G}_2 : Switch ciphertext in initialization. In this game, we change how SIM computes the ciphertext C produced by honest clients during initialization. Instead of computing $C \stackrel{\$}{\leftarrow} \text{PKE.Enc}(\text{pk}_{\text{Enc}}, (\text{pw}, K, \text{ID}_C, \text{ssid}))$, the simulator computes $C \stackrel{\$}{\leftarrow} \text{PKE.Enc}(\text{pk}_{\text{Enc}}, \perp)$, where \perp is an arbitrary string that is used nowhere else. Nonetheless, SIM still receives the secret input pw from \mathcal{F} and chooses $K \leftarrow \{0, 1\}^\lambda$, which it stores in a record $\langle \text{INIT}, \text{ssid}, \text{ID}_C, C, \text{pw}, K \rangle$ (see I.2). Note that storing pw and K in the record is only a temporary change and we remove this again in \mathbf{G}_{11} .

Further, we change how SIM acts when receiving a ciphertext C from an honest ID_C or from a corrupt server that honestly delivers C from an honest ID_C to $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$ in an initialization. In more detail, SIM keeps track of all ciphertexts that it computes for honest ID_C as $C \stackrel{\$}{\leftarrow} \text{PKE.Enc}(\text{pk}_{\text{Enc}}, \perp)$ in the $\langle \text{INIT}, \text{ssid}, \text{ID}_C, C, \text{pw}, K \rangle$ records. When SIM receives a message $(\text{INIT}, \text{ssid}, C, \text{ID}_C)$, then SIM tries to retrieve $\langle \text{INIT}, \text{ssid}, \text{ID}_C, C, [\text{pw}], [K] \rangle$, i.e., it checks if C was computed by SIM for the honest client ID_C as an encryption of \perp in subsession ssid . If that is the case, then SIM does not use sk_{Enc} to decrypt C (that would yield \perp anyways) but instead stores a record $(\text{FILE}, \text{ID}_C, \text{pw}, K, 10)$ and sends $(\text{INITRES}, \text{ssid}, \text{Succ})$ to the server (see IC.2 and IC.6). Again, creating this record is only a temporary change that we remove again in games \mathbf{G}_8 , \mathbf{G}_{11} and \mathbf{G}_{14} . If C was not computed by SIM , indicated by the fact that no record $\langle \text{INIT}, \text{ssid}, \text{ID}_C, C, *, * \rangle$ exists, then SIM continues as in \mathbf{G}_1 . Note that this may lead to SIM decrypting C to \perp , e.g., if \mathcal{A} replays some C that was computed by SIM on behalf of some honest ID_C . However, in that case, in \mathbf{G}_1 , SIM would output FAIL to \mathcal{S} as there would be a mismatch between

ID_C and ID'_C or ssid and ssid' . Thus, we let SIM output FAIL to \mathcal{S} if C is decrypted to \perp (see IC.3(b) and IC.7).

It is easy to see that the outputs of the server and the client are just as in \mathbf{G}_1 . Hence, the only difference is the distribution of C . If \mathcal{Z} can distinguish \mathbf{G}_2 from \mathbf{G}_1 , then we can construct an adversary \mathcal{B}_1 against the IND-CCA security of PKE as follows: First, \mathcal{B}_1 does not compute $(\text{sk}_{\text{Enc}}, \text{pk}_{\text{Enc}})$ itself but uses the pk^* provided by its challenger. Let $q_{\text{INIT}} \in \mathbb{N}$ be the number of initializations. We construct a sequence of games $\mathbf{G}_1^{(0)}, \dots, \mathbf{G}_1^{(q_{\text{INIT}})}$, where in $\mathbf{G}_1^{(i)}$ the first i ciphertexts are computed as encryptions of \perp if simulated for an honest ID_C and the remaining ciphertexts are encrypted as in \mathbf{G}_1 (except that pk^* is used). We have $\mathbf{G}_1 = \mathbf{G}_1^{(0)}$ and $\mathbf{G}_2 = \mathbf{G}_1^{(q_{\text{INIT}})}$. Because \mathcal{Z} can distinguish \mathbf{G}_1 from \mathbf{G}_2 , there must be an index $i^* \in [q_{\text{INIT}}]$ such that \mathcal{Z} has a non-negligible advantage in distinguishing $\mathbf{G}_1^{(i^*)}$ and $\mathbf{G}_1^{(i^*-1)}$. Now, in the i^* -th initialization the reduction \mathcal{B}_1 gives $m_0 := (k_{\text{sym}}, \text{pw}, \text{ID}_C, \text{ssid})$ and $m_1 := \perp$ to the challenger and uses the returned C^* as ciphertext for the i^* -th initialization. Additionally, whenever \mathcal{B}_1 receives a message $(\text{INIT}, \text{ssid}, C^*, \text{ID}_C)$ from \mathcal{A} to $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$ on behalf of a corrupt server such that there is no record $\langle \text{INIT}, \text{ssid}, C^*, \text{ID}_C, [\text{pw}], [K] \rangle$, then \mathcal{B}_1 uses its decryption oracle to decrypt C^* . The oracle answers with $(\text{pw}^*, K^*, \text{ID}_C^*, \text{ssid}^*)$ and \mathcal{B}_1 then proceeds as in \mathbf{G}_1 .

Now, if C^* encrypts m_0 , the game is distributed as in $\mathbf{G}_1^{(i^*-1)}$ and if it encrypts m_1 , then the game is distributed as in $\mathbf{G}_1^{(i^*)}$. We get

$$|\Pr[\mathbf{G}_2] - \Pr[\mathbf{G}_1]| \leq q_{\text{INIT}} \text{Adv}_{\text{PKE}, \mathcal{B}_1}^{\text{IND-CCA}}(\lambda).$$

Game \mathbf{G}_3 : Inform \mathcal{F} of malicious initializations. In this game, we add the MALICIOUSINIT interface as in $\mathcal{F}_{\text{PPKR}}$ to \mathcal{F} . We also let SIM use MALICIOUSINIT to make \mathcal{F} create records for malicious initializations. More precisely, when \mathcal{A} instructs a corrupted server to send a message $(\text{INITC}, \text{ssid}, C^*, \text{ID}_C)$ to $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$, where C^* was never sent by an honest client before, then SIM first proceeds just as in \mathbf{G}_2 , i.e., it executes the code of $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$, but additionally gives input $(\text{MALICIOUSINIT}, \text{ssid}, \text{ID}_C, \text{pw}^*, K^*)$ to \mathcal{F} (see IC.7).

Note that on receiving the MALICIOUSINIT input \mathcal{F} only responds to SIM . Therefore, \mathcal{Z} 's view did not change and we get

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

Game \mathbf{G}_4 : Provide \mathcal{F} with input of corrupted clients in initialization. In this game we change SIM such that when it receives a message $(\text{INIT}, \text{ssid}, C, \text{ID}_C)$ from a corrupted ID_C to an honest \mathcal{S} , in addition to decrypting $(\text{pw}, K, \text{ID}'_C, \text{ssid}')$, checking $\text{ID}_C = \text{ID}'_C, \text{ssid} = \text{ssid}'$ and recording $\langle \text{FILE}, \text{ID}_C, \text{pw}, K, 10 \rangle$, it also gives input $(\text{INITC}, \text{ssid}, \text{pw})$ to \mathcal{F} (see IC.3(a)). Since a corrupted ID_C may replay ciphertexts C that were computed by SIM as an encryption of \perp , SIM sets $\text{pw} = \perp$ for its INITC query if C decrypts to \perp . This cannot modify the view of \mathcal{Z} as SIM outputs FAIL to the corrupt client if such a replay happens

Initially, compute $(sk_{Enc}, pk_{Enc}) \leftarrow \text{KeyGen}(1^\lambda)$ and store record $\langle sid, sk_{Enc}, pk_{Enc} \rangle$. *Wait for X* means that $\text{SIM}^{\text{encPw}}$ does not proceed to the next instruction before receiving X and meanwhile gives back activation to \mathcal{A} . Once it receives X, it first proceeds with the instructions on input X and then continues at the instruction, where it waited for X. If a record cannot be retrieved, the query is ignored. For brevity we omit session identifier sid from all inputs, outputs, and records.

On $(\text{INITC}, ssid, ID_C)$ from $\mathcal{F}_{\text{PPKR}}$:

I.1 Retrieve record $\langle sid, *, [pk_{Enc}] \rangle$ (**G₁**)

I.2 Compute $C \stackrel{\$}{\leftarrow} \text{PKE.Enc}(pk_{Enc}, \perp)$ and record $\langle \text{INIT}, ssid, ID_C, C \rangle$ (**G₂**). Send $(\text{INIT}, ssid, C, ID_C)$ to S on behalf of ID_C (**G₁**).

On $(\text{INITS}, ssid, ID_C)$ from $\mathcal{F}_{\text{PPKR}}$:

I.3 If S is honest, wait for $(\text{INIT}, ssid, C, ID_C)$ from ID_C to S. (**G₁**)

On $(\text{INIT}, ssid, C, ID_C)$ from \mathcal{A} to S on behalf of ID_C :

IC.1 Wait for $(\text{INITS}, ssid, ID_C)$ from $\mathcal{F}_{\text{PPKR}}$. (**G₁**)

IC.2 If ID_C is honest, retrieve $\langle \text{INIT}, ssid, ID_C, C \rangle$. (**G₂**)

IC.3 If ID_C is corrupt:

(a) Retrieve $\langle sid, [sk_{Enc}], * \rangle$, compute $(pw, K, ID'_C, ssid') \leftarrow \text{PKE.Dec}(sk_{Enc}, C)$ (**G₁**) and give input $(\text{INITC}, ssid, pw)$ to $\mathcal{F}_{\text{PPKR}}$ on behalf of ID_C , where $pw = \perp$ if the decryption resulted in \perp . On response $(\text{INITC}, ssid, ID_C)$ continue below. (**G₄**)

(b) If the decryption resulted in \perp (**G₂**), $ID_C \neq ID'_C$, or $ssid \neq ssid'$, send $(\text{INITRES}, ssid, \text{FAIL})$ to ID_C (**G₁**) and $(\text{COMPLETEINITS}, ssid, 0)$ to $\mathcal{F}_{\text{PPKR}}$ (**G₅**) // 0 meaning fail.

(c) Otherwise, store (FILE, ID_C, K) , overwriting any existing $\langle \text{FILE}, ID_C, * \rangle$. (**G₁**) // No need to store pw and ctr, as $\mathcal{F}_{\text{PPKR}}$ takes care of this.

IC.4 If a record $\langle \text{INIT}, ssid, ID_C, C \rangle$ exists, store record $\langle \text{FILE}, ID_C, \text{ctr} \leftarrow 10 \rangle$. (**G₁**) // SIM must keep counter for DELREC in honest ID_C , honest server case.

IC.5 Send $(\text{INITRES}, ssid, \text{Succ})$ to ID_C (**G₁**) and give input $(\text{COMPLETEINITS}, ssid, 1)$ to $\mathcal{F}_{\text{PPKR}}$ (**G₅**) // 1 meaning no fail.

On $(\text{INIT}, ssid, C, ID_C)$ from \mathcal{A} to $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$ on behalf of corrupt S:

IC.6 If there is a record $\langle \text{INIT}, ssid, ID_C, C \rangle$ (**G₂**), then give input $(\text{INITS}, ssid, ID_C)$ to $\mathcal{F}_{\text{PPKR}}$ (**G₅**). // Honest C from ID_C is used On response $(\text{INITS}, ssid, ID_C)$, send $(\text{INITRES}, ssid, \text{Succ})$ to S (**G₂**). Send $(\text{COMPLETEINITS}, ssid, 1)$ to $\mathcal{F}_{\text{PPKR}}$ (**G₅**). // $b_S = 1$ for successful Init.

IC.7 Else if there is no such record, then retrieve record $\langle sid, [sk_{Enc}], * \rangle$ and compute $(pw^*, K^*, ID_C^*, ssid^*) \leftarrow \text{Dec}(sk_{Enc}, C)$ (**G₁**). If the decryption results in \perp (**G₂**), $ID_C^* \neq ID_C$, or $ssid^* \neq ssid$, then return $(\text{RECRES}, ssid, \text{FAIL})$ to S (**G₁**). Else give input $(\text{MALICIOUSINIT}, ID_C, pw^*, K^*)$ to \mathcal{F}_{HSM} (**G₃**), delete any existing record $\langle \text{FILE}, ID_C, * \rangle$ (**G₁₂**), and send $(\text{INITRES}, ssid, \text{Succ})$ to S (**G₁**). // Maliciously chosen key and pwd.

On $(\text{INITRES}, ssid, out)$ from \mathcal{A} to ID_C on behalf of S:

IO.1 If $(\text{INITRES}, ssid, out)$ was never output by $\text{SIM}^{\text{encPw}}$ on behalf of S or $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$ (**G₁₀**), send $(\text{COMPLETEINITC}, ssid, 0)$ to $\mathcal{F}_{\text{PPKR}}$ (**G₁₁**).

IO.2 Else, give input $(\text{COMPLETEINITC}, ssid, 1)$ to $\mathcal{F}_{\text{PPKR}}$ (**G₁₁**) // 1 meaning no success.

Figure 11: The initialization part of the simulator $\text{SIM}^{\text{encPw}}$ for the protocol π^{encPw} .

(cf. **IC.3(b)**). Note that the INITC interface of \mathcal{F} is still a dummy interface, so this change does not alter the view of \mathcal{Z} , so

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

Game G₅: Let \mathcal{F} generate server output in initialization.

In this game, we change the initialization interfaces INITC and INITS , and add the interface COMPLETEINITS to \mathcal{F} . We change INITC and INITS to be exactly as in $\mathcal{F}_{\text{PPKR}}$ except that INITC still provides SIM with the input pw of ID_C . Further, we change SIM such that it uses COMPLETEINITS to produce output for honest S in the initialization phase. Now, when SIM receives $(\text{INITS}, ssid, ID_C)$ from \mathcal{F} and a message $(\text{INITC}, ssid, C, ID_C)$ to the server, where either C was produced by SIM itself or $(pw, K, ID'_C, ssid') \leftarrow \text{PKE.Dec}(sk_{Enc}, C)$ and $ID'_C = ID_C$ and $ssid' = ssid$, then SIM does not directly make the server output $(\text{INITRES}, ssid, ID_C, \text{Succ})$ but instead gives $(\text{COMPLETEINITS}, ssid, ID_C, 1)$ to \mathcal{F} (**IC.5**). If the decrypted

ID'_C is different from ID_C or $ssid'$ is different from $ssid$, then SIM gives input $(\text{COMPLETEINITS}, ssid, ID_C, 0)$ to \mathcal{F} (**IC.3(b)**).

Furthermore, SIM uses the COMPLETEINITS interface, whenever a corrupt S honestly delivers a message from some honest ID_C to $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$. Even though in that case the interface gives its output to SIM , this query is necessary to ensure that \mathcal{F} creates a FILE record. Note that \mathcal{F} only creates a FILE record, if it received the INITC and INITS input. Thus, since S is corrupt, SIM first gives the input $(\text{INITS}, ssid, ID_C)$ to \mathcal{F} (see **IC.6**).

We can see that the output of the honest server did not change. In **G₄** the honest server outputs $(\text{INITRES}, ssid, \text{Succ})$ if it received a ciphertext that decrypted to the correct $ID_C, ssid$. If C contains a different ID'_C or $ssid'$, the server aborts. The simulator makes \mathcal{F} provide the corresponding output by choosing the bit b_S

On $(\text{RECC}, \text{ssid}, \text{ID}_C, \text{match})$ from $\mathcal{F}_{\text{PPKR}}$:

- R.1** Retrieve record $\langle \text{sid}, [\text{sk}_{\text{ENC}}], [\text{pk}_{\text{ENC}}] \rangle$. (**G₁**)
- R.2** Choose $k_{\text{sym}} \xleftarrow{\$} \{0, 1\}^\lambda$ (**G₁**). Compute $C \xleftarrow{\$} \text{PKE.Enc}(\text{pk}_{\text{ENC}}, \perp)$ and record $\langle \text{REC}, \text{ssid}, \text{ID}_C, C, k_{\text{sym}} \rangle$ (**G₆**). Send $(\text{REC}, \text{ssid}, C, \text{ID}_C)$ to S on behalf of ID_C . (**G₁**)

On $(\text{RECS}, \text{ssid}, \text{ID}'_C, \text{match})$ from $\mathcal{F}_{\text{PPKR}}$:

- R.3** If S is honest, wait for $(\text{REC}, \text{ssid}, C, \text{ID}_C)$ from ID_C to S . (**G₁**)

On $(\text{REC}, \text{ssid}, C, \text{ID}_C)$ from \mathcal{A} to S on behalf of ID_C :

- RC.1** Wait for $(\text{RECS}, \text{ssid}, \text{ID}_C, \text{match})$ from $\mathcal{F}_{\text{PPKR}}$. (**G₁**)
- RC.2** If ID_C is honest:
 - (a) Retrieve $\langle \text{REC}, \text{ssid}, \text{ID}_C, C, [k_{\text{sym}}] \rangle$. (**G₆**)
 - (b) Retrieve record $\langle \text{FILE}, \text{ID}_C, [\text{ctr}] \rangle$. If $\text{ctr} = 0$ set $C' \leftarrow \text{DELREC}$. Else if $\text{match} = 1$ (**G₁₄**), compute $C' \xleftarrow{\$} \text{SE.Enc}(k_{\text{sym}}, \perp)$ (**G₉**) and set ctr in the record to 10. Else set $C' \leftarrow \text{FAIL}$ (**G₁**) and ctr in the record to $\text{ctr} - 1$.
- RC.3** If ID_C is corrupt:
 - (a) Retrieve $\langle \text{sid}, [\text{sk}_{\text{ENC}}], * \rangle$, compute $(k_{\text{sym}}, \text{pw}', \text{ID}'_C, \text{ssid}') \leftarrow \text{PKE.Dec}(\text{sk}_{\text{ENC}}, C)$ (**G₁**), and give input $(\text{RECC}, \text{ssid}, \text{pw}')$ to $\mathcal{F}_{\text{PPKR}}$ on behalf of ID_C , where $\text{pw}' = \perp$ if the decryption resulted in \perp . On response $(\text{RECC}, \text{ssid}, \text{ID}_C, \text{match})$ continue below. (**G₇**)
 - (b) If the decryption resulted in \perp (**G₆**), $\text{ID}_C \neq \text{ID}'_C$, or $\text{ssid} \neq \text{ssid}'$, send $(\text{RECRES}, \text{ssid}, \text{FAIL})$ to ID_C . (**G₁**)
 - (c) Send $(\text{COMPLETEREC}, \text{ssid}, 1)$ to $\mathcal{F}_{\text{PPKR}}$. On response $(\text{RECRES}, \text{ssid}, K)$, if $K \in \{\text{FAIL}, \text{DELREC}\}$, set $C' \leftarrow K$. (**G₁₂**)
 - (d) If $K \notin \{\text{FAIL}, \text{DELREC}\}$ and a record $\langle \text{FILE}, \text{ID}_C, [K'] \rangle$ exists, compute $C' \xleftarrow{\$} \text{SE.Enc}(k'_{\text{sym}}, K')$ (**G₁₂**). // Ignore, K , it is different from the extracted K'
 - (e) If no such record exists, compute $C' \xleftarrow{\$} \text{SE.Enc}(k'_{\text{sym}}, K)$. (**G₁**) // ID_C was honest at Init
- RC.4** Send $(\text{COMPLETEREC}, \text{ssid}, 1)$ to $\mathcal{F}_{\text{PPKR}}$ (**G₈**) and $(\text{RECRES}, \text{ssid}, C')$ to ID_C . (**G₁**)

On $(\text{REC}, \text{ssid}, \text{ID}_C, C)$ from \mathcal{A} to $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$ on behalf of corrupt S :

- RC.5** If a record $\langle \text{REC}, \text{ssid}, \text{ID}_C, C, [k_{\text{sym}}] \rangle$ exists, then give input $(\text{RECS}, \text{ssid}, \text{ID}_C, \perp, \perp, \perp)$ to $\mathcal{F}_{\text{PPKR}}$. On response $(\text{RECS}, \text{ssid}, \text{ID}_C, \text{match})$, give input $(\text{COMPLETEREC}, \text{ssid}, 1)$ to $\mathcal{F}_{\text{PPKR}}$ (**G₈**) and on the subsequent response $(\text{RECRES}, \text{ssid}, K')$ execute the step **RC.2**(b). (**G₁₄**)
- RC.6** If no record $\langle \text{REC}, \text{ssid}, \text{ID}_C, C, * \rangle$ exists, then retrieve $\langle \text{sid}, \text{sk}_{\text{ENC}}, \text{pk}_{\text{ENC}} \rangle$ and compute $(k_{\text{sym}}, \text{pw}', \text{ID}'_C, \text{ssid}') \leftarrow \text{Dec}(\text{sk}_{\text{ENC}}, C)$. (**G₁**)
- RC.7** If the decryption results in \perp (**G₆**), $\text{ID}'_C \neq \text{ID}_C$, or $\text{ssid}' \neq \text{ssid}$, then send $(\text{RECRES}, \text{ssid}, \text{FAIL})$ to S . (**G₁**)
- RC.8** Give input $(\text{MALICIOUSREC}, \text{ssid}, \text{ID}_C, \text{pw}')$ to $\mathcal{F}_{\text{PPKR}}$ and on response K' determine the output as follows (**G₁₂**).
 - (a) If $K' \in \{\text{DELREC}, \text{FAIL}\}$, then send $(\text{RECRES}, \text{ssid}, K')$ to S . (**G₁₂**)
 - (b) If $K' \notin \{\text{DELREC}, \text{FAIL}\}$, then compute $C' \leftarrow \text{SE.Enc}(k_{\text{sym}}, K)$ if a record $\langle \text{FILE}, \text{ID}_C, [K] \rangle$ exists and $C' \leftarrow \text{SE.Enc}(k_{\text{sym}}, K')$ otherwise. Send $(\text{RECRES}, \text{ssid}, C')$ to S (**G₁₂**)

On $(\text{RECRES}, \text{ssid}, \text{out})$ from \mathcal{A} to ID_C on behalf of S :

- RO.1** If $(\text{RECRES}, \text{ssid}, \text{out})$ was never output by $\text{SIM}^{\text{encPw}}$ on behalf of S or $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$ (**G₁₀**), send $(\text{COMPLETEINITC}, \text{ssid}, 0)$ to $\mathcal{F}_{\text{PPKR}}$ (**G₁₁**).
- RO.2** Else, give input $(\text{COMPLETEREC}, \text{ssid}, 1)$ to $\mathcal{F}_{\text{PPKR}}$. (**G₁₁**) // $\text{out} = C'$ is attested by the HSM and \mathcal{A} didn't tamper with it.

Figure 12: The recovery part of the simulator $\text{SIM}^{\text{encPw}}$ for the protocol π^{encPw} .

accordingly for the **COMPLETEINITS** message. Hence,

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

Note that due to the changes introduced in Games **G₃**-**G₅**, \mathcal{F} now stores **FILE** records in all initializations, although \mathcal{F} does not use them in recoveries, yet.

Game **G₆: Switch ciphertext in recovery.** In this game, we change how SIM computes the ciphertext C produced by honest clients during recovery. Instead of computing the ciphertext as $C \xleftarrow{\$} \text{PKE.Enc}(\text{pk}_{\text{ENC}}, (k_{\text{sym}}, \text{pw}', \text{ID}_C, \text{ssid}))$ the simulator computes it as $C \xleftarrow{\$} \text{PKE.Enc}(\text{pk}_{\text{ENC}}, \perp)$. Still, it chooses a symmetric key $k_{\text{sym}} \xleftarrow{\$} \{0, 1\}^\lambda$ and stores a record $\langle \text{REC}, \text{ssid}, \text{ID}_C, C, k_{\text{sym}}, \text{pw}' \rangle$ to remember the symmetric key (see **R.2**). Note that SIM still gets the client's input pw' from \mathcal{F} , which we change in **G₁₄**.

We further change how SIM reacts on a message $(\text{REC}, \text{ssid}, C, \text{ID}_C)$ to $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$, where C was computed by SIM itself for the recovery of an honest client. Then, SIM does not retrieve sk_{ENC} to decrypt C (that would yield \perp anyways). Instead, SIM retrieves the record $\langle \text{REC}, \text{ssid}, \text{ID}_C, C, [k_{\text{sym}}], [\text{pw}'] \rangle$ (**RC.2**(a)) and proceeds as in **G₅** by executing the code of $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$. Similarly to **G₂**, the changes introduced here may lead to SIM decrypting C to \perp , e.g., if \mathcal{A} replays some C that was computed by SIM on behalf of some honest ID_C . Again, we let SIM output **FAIL** to S in that case (cf. **RC.3**(b) and **RC.7**).

Note that the output behavior of SIM did not change. That is because SIM essentially behaves like $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$ except that it does not encrypt and decrypt pw' , k_{sym} but it stores these values and uses the stored values later. However, the

distribution of C did change. Now, if the environment could distinguish G_6 from G_5 , we can construct an adversary \mathcal{B}_2 against the IND-CCA security of PKE as follows:

Let $q_{\text{REC}} \in \mathbb{N}$ be the number of recovery phases executed. First, \mathcal{B}_2 does not compute $(\text{sk}_{\text{Enc}}, \text{pk}_{\text{Enc}})$ itself anymore but uses the pk^* provided by its challenger. We construct a sequence of games $G_5^{(0)}, \dots, G_5^{(q_{\text{REC}})}$, where in $G_5^{(i)}$ the first i ciphertexts that honest clients produce in recovery are replaced by encryptions of \perp and the remaining ciphertexts stay as in G_5 (except that pk^* is used). If \mathcal{Z} can distinguish G_6 from G_5 then there is an index $i^* \in [q_{\text{REC}}]$ such that \mathcal{Z} distinguishes $G_5^{(i^*)}$ and $G_5^{(i^*-1)}$ with non-negligible advantage. \mathcal{B}_2 internally runs \mathcal{Z} and simulates $G_5^{(i^*-1)}$ except for the i^* -th recovery. In this recovery, \mathcal{B}_2 gives the messages $m_0 := (k_{\text{sym}}, \text{pw}', \text{ID}_C, \text{ssid})$ and $m_1 := \perp$ to its challenger. When the challenger responds with a ciphertext C^* then \mathcal{B}_2 uses C^* as the ciphertext in the message $(\text{REC}, \text{ssid}, C^*, \text{ID}_C)$ that the honest client sends to S . Further, if \mathcal{B}_2 receives at some point a message $(\text{REC}, \text{ssid}, C, \text{ID}_C)$ from a corrupted server to $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$, where there is no record $\langle \text{REC}, [\text{ssid}'], [\text{ID}'_C], C, [k_{\text{sym}}] \rangle$, i.e., the ciphertext towards the HSM is tampered with, then \mathcal{B}_2 gives C to the decryption oracle provided by the IND-CCA challenger. On the oracle's answer $(k_{\text{sym}}^*, \text{pw}^*, \text{ID}_C^*, \text{ssid}^*)$ the reduction checks if $\text{ID}_C^* = \text{ID}_C$ and $\text{ssid}^* = \text{ssid}$. If these checks, fail the reduction gives output $(\text{ssid}, \text{FAIL})$ to the server. Else, the reduction retrieves the record $(\text{FILE}, \text{sid}, \text{ID}_C, [\text{pw}], [K], [\text{ctr}])$. If $\text{ctr} = 0$, then the reduction deletes the record and sends $(\text{DELREC}, \text{ID}_C)$ to S . Else, if $\text{pw} = \text{pw}^*$, the reduction sets $\text{ctr} \leftarrow 10$, computes $C' \stackrel{\$}{\leftarrow} \text{SE.Enc}(k_{\text{sym}}^*, K)$, and sends (ssid, C') to S . Else, the reduction decrements ctr and sends $(\text{ssid}, \text{FAIL})$ to S . Finally, \mathcal{B}_2 outputs whatever \mathcal{Z} outputs. Note that if C^* encrypts m_0 , then the view of \mathcal{Z} is distributed exactly as in $G_5^{(i^*-1)}$, and if C^* encrypts m_1 , then the view of \mathcal{Z} is distributed exactly as in $G_5^{(i^*)}$. We get

$$|\Pr[G_6] - \Pr[G_5]| \leq q_{\text{REC}} \text{Adv}_{\text{PKE}, \mathcal{B}_2}^{\text{IND-CCA}}(\lambda).$$

Game G_7 : Provide \mathcal{F} with input of corrupted clients in recovery. In this game we change SIM such that when it receives a message $(\text{REC}, \text{ssid}, C, \text{ID}_C)$ from a corrupted ID_C the simulator now, in addition to decrypting $(k'_{\text{sym}}, \text{pw}', \text{ID}'_C, \text{ssid}') \leftarrow \text{PKE.Dec}(\text{sk}_{\text{Enc}}, C)$ and checking $\text{ID}_C = \text{ID}'_C, \text{ssid} = \text{ssid}'$, also gives input $(\text{RECC}, \text{ssid}, \text{ID}_C, \text{pw}')$ to \mathcal{F} . Similarly to G_4 , we use $\text{pw}' = \perp$ if C is decrypted to \perp (RC.3(a)). Note that the RECC interface of \mathcal{F} is still a dummy interface, so this change does not alter the view of \mathcal{Z} , so

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

Game G_8 : Let \mathcal{F} generate server output in recovery. In this game, we change the RECC and RECS interfaces of \mathcal{F} and we add the COMPLETEREC interface to \mathcal{F} .

We change RECC and RECS such that they are as in $\mathcal{F}_{\text{PPKR}}$ except that RECC still forwards the secret client input pw' to the simulator. Further, we change SIM such that it uses

the interfaces to produce outputs for honest servers in the recovery phase. That means, when SIM receives a message $(\text{REC}, \text{ssid}, C, \text{ID}_C)$ from an honest or corrupted ID_C , it first proceeds as in G_7 , and before sending its output to ID_C , it sends the message $(\text{COMPLETEREC}, \text{ssid}, \text{ID}_C, 1)$ to \mathcal{F} (RC.4).

Additionally, SIM uses the COMPLETEREC interface whenever the corrupted server honestly delivers a message from some honest ID_C to $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$. Similarly to G_5 , this query gives its output back to SIM , but this query is again necessary to appropriately update the internal state of \mathcal{F} , in particular this interface resets the counter of ID_C to 10 if the recovery is successful. Again, note that this requires the simulator to first give the input $(\text{RECS}, \text{ssid}, \text{ID}_C, \perp, \perp, \perp)$ to \mathcal{F} (RC.5).

First, note that by definition, the interface RECC , resp. RECS , is called whenever an honest client, resp. honest server, starts a recovery. Next, note that SIM also ensures that the RECC input is given to \mathcal{F} when a corrupted client performs a recovery. As SIM controls sk_{Enc} , it is able to extract pw' by decrypting C and can give this as input to \mathcal{F} . Thus, \mathcal{F} always retrieves a record $\langle \text{RECC}, \text{ssid}, \text{ID}_C, \text{pw}', \text{pw}, K \rangle$ with $\text{pw}' \neq \perp$ when SIM provides it with the COMPLETEREC input. Consequently, if the provided password was correct, \mathcal{F} outputs SUCC to S and else it outputs FAIL to S . Finally, note that \mathcal{F} maintains the counter exactly as $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$ did, i.e., the counter is set to 10 when a new FILE record is created, the counter is decremented when a password guess was wrong, and the counter is reset to 10 when a password guess was correct. Thus, from now on, SIM no longer keeps a counter in its FILE records. Overall, the output of an honest server did not change in this game and we get

$$\Pr[G_8] = \Pr[G_7].$$

Game G_9 : Switch symmetric ciphertext in recovery. In this game, we change how SIM computes the ciphertext C' produced by the HSM during recoveries of honest clients. Instead of computing $C' \stackrel{\$}{\leftarrow} \text{SE.Enc}(k_{\text{sym}}, K)$, where K is the recovered key, the simulator now computes $C' \stackrel{\$}{\leftarrow} \text{SE.Enc}(k_{\text{sym}}, \perp)$ (RC.2(b)). Note that k_{sym} is still chosen by SIM uniformly at random for every honest client that starts a recovery.

Thus, if \mathcal{Z} can distinguish G_8 and G_9 , then we can construct an adversary \mathcal{B}_3 against the IND-CPA security of SE . We construct a sequence of games $G_8^{(0)}, \dots, G_8^{(q_{\text{REC}})}$, where in $G_8^{(i)}$ the first i ciphertexts C' are replaced by encryptions of \perp as described above. If \mathcal{Z} can distinguish G_8 from G_9 , then there is an index i^* such that \mathcal{Z} has non-negligible advantage in distinguishing $G_8^{(i^*)}$ from $G_8^{(i^*-1)}$. The reduction \mathcal{B}_3 internally runs \mathcal{Z} and plays the role of \mathcal{F} and SIM in $G_8^{(i^*)}$ except for the i^* -th recovery. There, \mathcal{B}_3 does not choose a symmetric key k_{sym} but sends the messages $m_0 := K$ and $m_1 := \perp$ to its challenger, where K is the backup key that the HSM would use to compute C' . When the challenger returns a ciphertext C^* , \mathcal{B}_3 outputs the message $(\text{RECRES}, \text{ssid}, C^*)$.

Note that the view of \mathcal{Z} in the case that C^* encrypts m_0 is distributed exactly as in $\mathbf{G}_8^{(i^*-1)}$ and in the case that C^* encrypts m_1 the view is distributed exactly as in $\mathbf{G}_8^{(i^*)}$. Therefore we get

$$|\Pr[\mathbf{G}_9] - \Pr[\mathbf{G}_8]| \leq q_{\text{REC}} \text{Adv}_{\text{SE}, \mathcal{B}_3}^{\text{IND-CPA}}(\lambda).$$

Game \mathbf{G}_{10} : Output FAIL on tampered message from HSM.

In this game, we change how the simulator reacts when an honest client receives a message from the HSM that was tampered with. Whenever an honest client receives a message (INITRES, ssid, SUCC) where SIM never produced this message on behalf of S or $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$, then SIM makes the client output FAIL (see IO.1). Similarly, when an honest client receives a message (RECRES, ssid, out) that was never produced by SIM on behalf of the HSM, then SIM makes the client output FAIL (see RO.1).

This game is identical to \mathbf{G}_9 unless \mathcal{A} delivers a message to ID_C that was never output by SIM but the attestation signature is still valid. If an environment \mathcal{Z}^* exists that can distinguish between \mathbf{G}_{10} and \mathbf{G}_9 , we can construct an adversary \mathcal{B}_4 against the sEUF-CMA security of the signature scheme Sig used for attestation. \mathcal{B}_4 internally runs the whole experiment including \mathcal{Z}^* and outputs the above described message with the signature to its challenger. We get

$$|\Pr[\mathbf{G}_{10}] - \Pr[\mathbf{G}_9]| \leq \text{Adv}_{\text{Sig}, \mathcal{B}_4}^{\text{sEUF-CMA}}(\lambda).$$

Game \mathbf{G}_{11} : Generate output for honest clients by \mathcal{F} . In this game, we change \mathcal{F} and SIM such that SIM lets \mathcal{F} produce the output for an honest client in the initialization and recovery phase. To this end, SIM will make use of the interfaces COMPLETEINITC and COMPLETERECC that we add to \mathcal{F} exactly as in $\mathcal{F}_{\text{PPKR}}$. We change the following things:

- Whenever the simulator produces a message (INIT, ssid, C, ID_C) on behalf of some honest client ID_C , the simulator no longer draws a random backup key K . Further, when it then receives a message (INITRES, ssid, out) to ID_C , it uses the COMPLETEINITC interface with the bit b_C set accordingly (IO.2 and IO.1).
- Similarly, when SIM receives a message (RECRES, ssid, C') towards an honest client ID_C , the simulator now uses the COMPLETERECC interface with the bit b_C set depending on the value of out (RO.1 and RO.2).

As a consequence, SIM no longer needs to store the key K in the FILE records for honest initializations and the values pw and K in the INIT records.

Note that as of \mathbf{G}_5 , \mathcal{F} already chooses a uniformly random key K when it receives an INITC message, although that key is not used up to this game. Now, the COMPLETEINITC interface ensures that it is given as output to ID_C in the initialization phase. Similarly, in the recovery phase of \mathbf{G}_{10} the simulator retrieved its own stored backup key from the initialization phase to give it as output to ID_C in the recovery phase. Now, \mathcal{F} , and not SIM, stores the backup key

K in its FILE record, retrieves it when a recovery is successful, and gives it as output to ID_C or SIM when SIM sends a COMPLETERECC message for a successful subsession.

Overall, the distribution of \mathcal{Z} 's view did not change and we get

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

Game \mathbf{G}_{12} : Retrieve maliciously initialized records from \mathcal{F} .

In this game we change how SIM responds in recoveries by a corrupt party. When it receives a message (RECC, ssid, C^* , ID_C) from \mathcal{A} on behalf of a corrupted ID_C , after checking $\text{ID}'_C = \text{ID}_C$ and $\text{ssid}' = \text{ssid}$, SIM now additionally sends (COMPLETERECC, ssid, 1) to \mathcal{F} . If \mathcal{F} answers with (RECRES, ssid, K'), where $K \in \{\text{FAIL}, \text{DELREC}\}$, then SIM forwards this to ID_C on behalf of S (RC.3(c)). If $K' \notin \{\text{FAIL}, \text{DELREC}\}$, then SIM needs to produce the ciphertext C' . However, the key K' received from \mathcal{F} may be different from the key K that the corrupt ID_C chose in its last initialization, as \mathcal{F} always chooses a random key when receiving the input INITC even for a corrupted ID_C . Therefore, if a record $\langle \text{FILE}, \text{ID}_C, [K] \rangle$ exists, which implies that ID_C executed an initialization phase after being corrupted, then SIM encrypts K instead of K' to obtain C' (RC.3(d) and RC.3(e)). Note that this requires the simulator to delete any potentially existing FILE record of some ID_C whenever the corrupted S maliciously initializes for ID_C to ensure that it always returns the key of the most recent initialization (cf. IC.7).

Furthermore, we add the MALICIOUSRECOVER interface as in $\mathcal{F}_{\text{PPKR}}$ to \mathcal{F} , which SIM now uses in recoveries by a corrupt S. When \mathcal{A} instructs the corrupted server to send a message (RECC, ssid, C^* , ID_C) to $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$, where C^* was never sent by an honest client, which is again captured by checking if C^* decrypts to \perp (RC.7), SIM gives the input (MALICIOUSRECOVER, ssid, ID_C , pw*) to \mathcal{F} , where pw* is the password it obtained from decrypting C^* (RC.8). If the response of \mathcal{F} is (DELREC, ID_C) or FAIL, then SIM forwards the response to the server (RC.8(a)). Else, \mathcal{F} responds with K' . As before, we need to check whether the key returned by \mathcal{F} is the correct one by checking whether a FILE record exists for ID_C (RC.8(b)).

Note that the added interface does not output anything to the protocol parties but only to SIM. Further, note that \mathcal{F} now behaves exactly as $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$ did in \mathbf{G}_{11} when a corrupted server interacts with $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$. More precisely, on a (MALICIOUSINIT, ssid, ID_C , pw*, K^*) the functionality creates a record containing the password pw* and the backup key K^* together with a counter, initialized to 10. Then on a (MALICIOUSRECOVER, ssid, ID_C , pw*) the functionality tries to retrieve this record, checks the counter, and compares the provided password with the stored password.

$$\Pr[\mathbf{G}_{12}] = \Pr[\mathbf{G}_{11}].$$

Game \mathbf{G}_{13} : Add unused attack interfaces. In this game, we add the interfaces LEAKFILE, CORRUPT, FULLYCORRUPT, and OFFLINEATTACK to \mathcal{F} . Note that our simulator never uses these interfaces. That is because $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$ does not leak any key files and cannot be corrupted. Thus, the protocol even

realizes a version of $\mathcal{F}_{\text{PPKR}}$ that does not allow offline attacks. We get

$$\Pr[\mathbf{G}_{13}] = \Pr[\mathbf{G}_{12}].$$

Game \mathbf{G}_{14} : Ideal world. We change the ideal functionality such that no more private input pw, pw' is given to SIM . Note that SIM only used these inputs to determine if the password pw' used in a recovery is correct. Instead it can now use the output *match* from \mathcal{F} (RC.2(b) and RC.5). Hence, we finally remove pw' from the REC and FILE records. Also, we take away the ability of SIM to give output to parties. This is also not used anymore. Thus we get

$$\Pr[\mathbf{G}_{14}] = \Pr[\mathbf{G}_{13}].$$

After this change, we reached the ideal-world execution of the protocol encPw with the simulator $\text{SIM} = \text{SIM}^{\text{encPw}}$ as described in Figures 11 and 12 and the functionality $\mathcal{F} = \mathcal{F}_{\text{PPKR}}$ as described in Figures 1 and 2. In particular, we never had to change how SIM reacts on receiving a GETPK message when acting as $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$. \square

D.2 Proof of Theorem 2

PROOF. As in the proof of Theorem 1 we construct a sequence of hybrid games starting from the real world and ending in the ideal world. As the majority of game hops are identical to the previous proof, we only provide the games that are new. We write \mathbf{G}_x for $x \in (i, i+1) \subset \mathbb{R}$ to denote that the game is added between game \mathbf{G}_i and \mathbf{G}_{i+1} in the sequence of games from the proof of Theorem 1.

Game $\mathbf{G}_{1.1}$: Abort on salt collision. In this game, the simulator aborts if it randomly samples a salt value that is already used for some other ID_C . More precisely, when SIM simulates $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$ and draws s_1 and s_2 uniformly at random during initialization, SIM checks if it already sampled either value in some previous initialization. If that is the case, SIM aborts the execution (see LF.3(b)).

Let $q_{\text{INIT}} \in \mathbb{N}$ be the number of initializations. We get

$$|\Pr[\mathbf{G}_{1.1}] - \Pr[\mathbf{G}_1]| \leq \frac{q_{\text{INIT}}(2q_{\text{INIT}} - 1)}{2^\lambda}.$$

Game $\mathbf{G}_{5.1}$: Add OFFLINEATTACK and LEAKFILE interfaces to \mathcal{F} . In this game, we add the interfaces OFFLINEATTACK and LEAKFILE to \mathcal{F} exactly as they are in $\mathcal{F}_{\text{PPKR}}$. Because SIM does currently not use them, we get

$$\Pr[\mathbf{G}_{5.1}] = \Pr[\mathbf{G}_5].$$

Game $\mathbf{G}_{5.2}$: Use LEAKFILE interface of \mathcal{F} . In this game, we change how SIM responds to a (LEAKFILE) message from \mathcal{A} to $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$. SIM gives input (LEAKFILE) to \mathcal{F} and receives a list ID_C from \mathcal{F} (see LF.1). The simulator initializes $L \leftarrow \emptyset$. Now, for all $(\text{ID}_C, \text{ctr}) \in \text{ID}_C$ the simulator tries to retrieve a record $\langle \text{LEAKED}, \text{ID}_C, [c], [h], [s_1], [s_2], *, * \rangle$ that is not marked old.

If no such record exists then SIM chooses $s_1, s_2 \leftarrow \{0, 1\}^\lambda$ uniformly at random. Then, SIM computes $h \leftarrow H(s_1, \text{pw})$ and $c \leftarrow K \oplus H(s_2, \text{pw})$ and creates a record $\langle \text{LEAKED}, \text{ID}_C, c, h, s_1, s_2, \text{ctr}, 1 \rangle$. SIM appends $(\text{ID}_C, c, h, s_1, s_2, \text{ctr})$ to L .

Otherwise, SIM adds $(\text{ID}_C, c, h, s_1, s_2, \text{ctr})$ to L and records $\langle \text{LEAKED}, \text{ID}_C, c, h, s_1, s_2, \text{ctr}, i+1 \rangle$, where i is the biggest number such that a record $\langle \text{LEAKED}, \text{ID}_C, c, h, s_1, s_2, *, i \rangle$ exists.

SIM also does an additional step when an initialization was completed, i.e., in IC.6 or IC.8, depending on whether the server is corrupted or not. SIM goes through all so far leaked records and marks them as old. That means that they still can be offline attacked, but the next time when a record is leaked, SIM will simulate the creation of a new file. In particular, SIM keeps LEAKED records of already overwritten initializations. This will be important for programming H later. However, the records that are added to L are always the “current” records, i.e., records that correspond to FILE records of \mathcal{F} .

Also, note that SIM only appends leaked data for ID_C if $\text{ID}_C \in \text{ID}_C$. However, in the previous games, we ensured that \mathcal{F} creates a FILE record whenever $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$ would have done so. More precisely, when SIM receives a message $(\text{INIT}, \text{ssid}, C, \text{ID}_C)$ from \mathcal{A} to S on behalf of a corrupted ID_C and all the checks pass then SIM gives input $(\text{COMPLETEINITC}, \text{ssid}, \text{ID}_C, 1)$ to \mathcal{F} so \mathcal{F} creates a FILE record. Similarly, when SIM receives a message $(\text{INIT}, \text{ssid}, C, \text{ID}_C)$ from \mathcal{A} to $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$ on behalf of S and all the checks pass then SIM sends either a MALINIT message or $(\text{COMPLETEINITS}, \text{ssid}, \text{ID}_C, 1)$ to \mathcal{F} so \mathcal{F} creates a FILE record.

Overall, the view of \mathcal{Z} did not change.

$$\Pr[\mathbf{G}_{5.2}] = \Pr[\mathbf{G}_{5.1}].$$

Game $\mathbf{G}_{5.3}$: Simulate the password hash h . In this game, we let SIM simulate $h = H(s_1, \text{pw})$ without using pw and we let SIM program the random oracle accordingly.

First, when SIM creates a record $\langle \text{LEAKED}, \text{ID}_C, c, h, s_1, s_2, \text{ctr}, i \rangle$ it no longer computes h like in the real protocol but now chooses it uniformly at random, i.e., $h \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda$ (see LF.3(b)). Second, SIM observes the random oracle queries. If there is a query $H(s_1^*, \text{pw}^*)$ such that a record $\langle \text{LEAKED}, [\text{ID}_C], *, [h], s_1^*, *, *, [i] \rangle$ exists, then SIM gives input $(\text{OFFLINEATTACK}, \text{ID}_C, \text{pw}^*, i)$ to \mathcal{F} . If \mathcal{F} answers with FAIL , then SIM responds with a uniformly random value. If \mathcal{F} answers with $K \neq \text{FAIL}$, then SIM programs $H(s_1^*, \text{pw}^*)$ to h (see H.2). This ensures that if \mathcal{A} guesses the password of some client whose file was leaked previously, then the output of the random oracle query is consistent with the values that SIM output in the leaked file.

Since H is modeled as a random oracle, choosing h as a uniformly random value does not modify the distribution of h . However, observe that \mathcal{Z} could distinguish this game from the previous one if \mathcal{A} guesses the salt s_1^* that some ID_C uses together with the password pw^* before leaking the file of ID_C . If that happens, SIM would output a random value h^* on the oracle query (s_1^*, pw^*) . Then, upon \mathcal{A} leaking the file of ID_C , the simulator would not output h^* in the leaked file and instead choose a random value h since it cannot know that ID_C chose the password pw^* . As s_1 is chosen

<p>On (INITC, ssid, ID_C) from $\mathcal{F}_{\text{PPKR}}$:</p> <p>I.1 Wait for (INIT, ssid, pk_{Enc}) from S to ID_C. (G₁)</p> <p>On (INITS, ssid, ID_C) from $\mathcal{F}_{\text{PPKR}}$:</p> <p>I.2 Compute $(\text{sk}_{\text{Enc}}, \text{pk}_{\text{Enc}}) \xleftarrow{\\$} \text{PKE.Gen}(1^\lambda)$ and store $\langle \text{ssid}, \text{sk}_{\text{Enc}}, \text{pk}_{\text{Enc}} \rangle$. Send (INIT, ssid, pk_{Enc}) to ID_C on behalf of S. (G₁)</p> <p>On (INIT, ssid, pk_{Enc}) from \mathcal{A} to ID_C on behalf of S:</p> <p>IPK.1 Compute $C \xleftarrow{\\$} \text{PKE.Enc}(\text{pk}_{\text{Enc}}, \perp)$ and store $\langle \text{INIT}, \text{ssid}, \text{ID}_C, C \rangle$ (G₂). Send (INIT, ssid, C, ID_C) to S on behalf of ID_C. (G₁) // pk_{Enc} cannot be tampered with.</p> <p>On (RECC, ssid, ID_C, match) from $\mathcal{F}_{\text{PPKR}}$:</p> <p>R.1 Wait for (REC, ssid, pk_{Enc}) from S to ID_C. (G₁)</p> <p>On (RECS, ssid, ID_C', match) from $\mathcal{F}_{\text{PPKR}}$:</p> <p>R.2 Compute $(\text{sk}_{\text{Enc}}, \text{pk}_{\text{Enc}}) \xleftarrow{\\$} \text{PKE.Gen}(1^\lambda)$ and store $\langle \text{ssid}, \text{sk}_{\text{Enc}}, \text{pk}_{\text{Enc}} \rangle$. Send (REC, ssid, pk_{Enc}) to ID_C on behalf of S. (G₁)</p> <p>On (REC, ssid, pk_{Enc}) from \mathcal{A} to ID_C on behalf of S:</p> <p>RPK.1 Compute $C \xleftarrow{\\$} \text{PKE.Enc}(\text{pk}_{\text{Enc}}, \perp)$, choose $k_{\text{sym}} \xleftarrow{\\$} \{0, 1\}^\lambda$, and store $\langle \text{REC}, \text{ssid}, \text{ID}_C, C, k_{\text{sym}} \rangle$ (G₂). Send (REC, ssid, C, ID_C) to S on behalf of ID_C (G₁). // REC came before pk_{Enc}.</p> <p>On (INIT, ssid, GETPK) from \mathcal{A} to $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$ on behalf of S:</p> <p>G.1 Compute $(\text{sk}_{\text{Enc}}, \text{pk}_{\text{Enc}}) \xleftarrow{\\$} \text{PKE.Gen}(1^\lambda)$ and store $\langle \text{ssid}, \text{sk}_{\text{Enc}}, \text{pk}_{\text{Enc}} \rangle$. Send (INIT, ssid, pk_{Enc}) to S on behalf of $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$ (G₁). // attested</p> <p>On (REC, ssid, GETPK) from \mathcal{A} to $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$ on behalf of S:</p> <p>G.2 Compute $(\text{sk}_{\text{Enc}}, \text{pk}_{\text{Enc}}) \xleftarrow{\\$} \text{PKE.Gen}(1^\lambda)$ and store $\langle \text{ssid}, \text{sk}_{\text{Enc}}, \text{pk}_{\text{Enc}} \rangle$. Send (REC, ssid, pk_{Enc}) to S on behalf of $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$ (G₁). // attested</p> <p>On (INIT, ssid, C, ID_C) from \mathcal{A} to S on behalf of ID_C:</p> <p>IC.1 through IC.5 as in Figure 11.</p> <p>IC.6 Try to retrieve all records $\langle \text{LEAKED}, \text{ID}_C, [c], [h], [s_1], [s_2], *, * \rangle$ not marked old. If such records exist, then mark all of them as old. (G_{5.2}) // After init, the file contains fresh values.</p> <p>On (INIT, ssid, C, ID_C) from \mathcal{A} to $\mathcal{F}_{\text{HSM}}^{\text{encPw}+}$ on behalf of S:</p> <p>IC.6 and IC.7 as in Figure 11.</p> <p>IC.8 Try to retrieve all records $\langle \text{LEAKED}, \text{ID}_C, [c], [h], [s_1], [s_2], *, * \rangle$ not marked old. If such records exist, then mark all of them as old. (G_{5.2}) // After init, the file contains fresh values.</p> <p>On a query (LEAKFILE) from \mathcal{A} to $\mathcal{F}_{\text{HSM}}^{\text{encPw}+}$:</p> <p>LF.1 Give input (LEAKFILE) to $\mathcal{F}_{\text{PPKR}}$. $\mathcal{F}_{\text{PPKR}}$ returns ID_C. (G_{5.2})</p> <p>LF.2 $L \leftarrow \emptyset$ (G_{5.2})</p> <p>LF.3 For each $(\text{ID}_C, \text{ctr}) \in \text{ID}_C$:</p> <ol style="list-style-type: none"> Try to retrieve a record $\langle \text{LEAKED}, \text{ID}_C, [c], [h], [s_1], [s_2], *, [i] \rangle$ not marked old. (G_{5.2}) If no such record exists, choose $c, h, s_1, s_2 \xleftarrow{\\$} \{0, 1\}^\lambda$ (G_{5.3} and G_{5.4}). If s_1 or s_2 was used before as salt by SIM, then abort (G_{1.1}). Else store $\langle \text{LEAKED}, \text{ID}_C, c, h, s_1, s_2, \text{ctr}, 1 \rangle$. Append $(\text{ID}_C, c, h, s_1, s_2, \text{ctr})$ to L. (G_{5.2}) Else store a record $\langle \text{LEAKED}, \text{ID}_C, c, h, s_1, s_2, \text{ctr}, i + 1 \rangle$, where i is the biggest number such that a record $\langle \text{LEAKED}, \text{ID}_C, [c], [h], [s_1], [s_2], *, i \rangle$ exists. Append $(\text{ID}_C, c, h, s_1, s_2, \text{ctr})$ to L. (G_{5.2}) <p>LF.4 Send L to \mathcal{A} on behalf of $\mathcal{F}_{\text{HSM}}^{\text{encPw}+}$. (G₁)</p> <p>On a query $H(s^*, \text{pw}^*)$:</p> <p>H.1 If a record $\langle H, s^*, \text{pw}^*, [y] \rangle$ exists, output y. (G₁)</p> <p>H.2 Try to retrieve the record $\langle \text{LEAKED}, [\text{ID}_C], *, [h], s^*, *, *, [i] \rangle$. If such a record exists, give input (OFFLINEATTACK, ID_C, pw[*], i) to $\mathcal{F}_{\text{PPKR}}$. If $\mathcal{F}_{\text{PPKR}}$ responds with $K \neq \text{FAIL}$, record $\langle H, s^*, \text{pw}^*, h \rangle$ and output h. If $K = \text{FAIL}$, choose $y \xleftarrow{\\$} \{0, 1\}^\lambda$, record $\langle H, s^*, \text{pw}^*, y \rangle$ and output y. (G_{5.3})</p> <p>H.3 If no such LEAKED record exists, try to retrieve the record $\langle \text{LEAKED}, [\text{ID}_C], [c], *, *, s^*, *, [i] \rangle$. If such a record exists, give input (OFFLINEATTACK, ID_C, pw[*], i) to $\mathcal{F}_{\text{PPKR}}$. If $\mathcal{F}_{\text{PPKR}}$ responds with $K \neq \text{FAIL}$, record $\langle H, s^*, \text{pw}^*, c \oplus K \rangle$ and output $c \oplus K$. If $K = \text{FAIL}$, choose $y \xleftarrow{\\$} \{0, 1\}^\lambda$, record $\langle H, s^*, \text{pw}^*, y \rangle$ and output y. (G_{5.4})</p> <p>H.4 If both such LEAKED do not exist (G_{5.4}), choose $y \xleftarrow{\\$} \{0, 1\}^\lambda$, record $\langle H, s^*, \text{pw}^*, y \rangle$ and output y (G₁).</p>
--

Figure 13: Differences from the simulator $\text{SIM}^{\text{encPw}}$ (Figures 11 to 12) to the simulator $\text{SIM}^{\text{encPw}+}$.

uniformly at random from $\{0, 1\}^\lambda$, the probability of this happening can be bounded by $q_H 2^{-\lambda}$, where $q_H \in \mathbb{N}$ is the number of queries to H . Note that this even holds if other clients use the same pw^* because due to $\mathbf{G}_{1,1}$ they will use a different salt s'_1 and $H(s'_1, \text{pw}^*)$ is independent of $H(s_1^*, \text{pw}^*)$. Thus, we have

$$|\Pr[\mathbf{G}_{5,3}] - \Pr[\mathbf{G}_{5,2}]| \leq \frac{q_{\text{INIT}} q_H}{2^\lambda}.$$

Game $\mathbf{G}_{5,4}$: Simulate the password hash c . In this game, we let SIM simulate the leaked records without using K and pw and we let SIM program the random oracle accordingly. First, when SIM creates a record $\langle \text{LEAKED}, \text{ID}_C, c, h, s_1, s_2, \text{ctr}, i \rangle$ it no longer computes c like in the real protocol but now chooses it uniformly at random, i.e., $c \leftarrow \{0, 1\}^\lambda$ (see [LF.3\(b\)](#)). Second, SIM observes the random oracle queries. If there is a query $H(s_2^*, \text{pw}^*)$ such that a record $\langle \text{LEAKED}, [\text{ID}_C], [c], *, *, s_2^*, *, [i] \rangle$ exists, then SIM gives input $(\text{OFFLINEATTACK}, \text{ID}_C, \text{pw}^*, i)$ to \mathcal{F} . If \mathcal{F} answers with FAIL, then SIM responds with a uniformly random value. If \mathcal{F} answers with $K \neq \text{FAIL}$, then, SIM programs $H(s_2^*, \text{pw}^*) \leftarrow c \oplus K$. Note that for honest ID_C , SIM uses the key K it chose on behalf of ID_C in the initialization instead of the one returned by \mathcal{F} . Once we let \mathcal{F} generate the output of honest clients in \mathbf{G}_{11} , we change this to use the key returned from \mathcal{F} . This again ensures that the leaked files are consistent with the random oracle outputs ([H.3](#)).

Again, it holds that the distribution of c does not change by choosing it as uniformly random value since H is a random oracle. However, similarly to $\mathbf{G}_{5,3}$, the environment could distinguish this game from the previous one if \mathcal{A} guesses the salt s_2^* of some ID_C^* . Following the same arguments as in $\mathbf{G}_{5,3}$, we have

$$|\Pr[\mathbf{G}_{5,4}] - \Pr[\mathbf{G}_{5,3}]| \leq \frac{q_{\text{INIT}} q_H}{2^\lambda}.$$

Changed Reductions. Also, we have to change the reductions on IND-CCA security in \mathbf{G}_2 and \mathbf{G}_6 because of the difference in computing $(\text{sk}_{\text{ENC}}, \text{pk}_{\text{ENC}})$:

Game \mathbf{G}_2 : Let $q_{\text{INIT}} \in \mathbb{N}$ be the number of initializations. We construct a sequence of games $\mathbf{G}_1^{(0)}, \dots, \mathbf{G}_1^{(q_{\text{INIT}})}$, where in $\mathbf{G}_1^{(i)}$ the first i ciphertexts are computed as encryptions of \perp and the remaining ciphertexts are encrypted as in \mathbf{G}_1 . We have $\mathbf{G}_1 = \mathbf{G}_1^{(0)}$ and $\mathbf{G}_2 = \mathbf{G}_1^{(q_{\text{INIT}})}$. Because \mathcal{Z} can distinguish \mathbf{G}_1 from \mathbf{G}_2 , there must be an index $i^* \in [q_{\text{INIT}}]$ such that \mathcal{Z} has a non-negligible advantage in distinguishing $\mathbf{G}_1^{(i^*)}$ and $\mathbf{G}_1^{(i^*-1)}$. Now, in the i^* -th initialization the reduction \mathcal{B}_1 does not choose a new key-pair (sk, pk) when it receives a GETPK or INIT_S message, but uses the public key that it is provided by the challenger. When it simulates ID_C in the i^* -th simulation, it gives $m_0 := (k_{\text{sym}}, \text{pw}', \text{ID}_C, \text{ssid})$ and $m_1 := \perp$ to the challenger and uses the returned C^* as the ciphertext for the initialization. Further, if \mathcal{B}_1 receives a message $(\text{INIT}, \text{ssid}, C, \text{ID}_C)$ from \mathcal{A} to S or to $\mathcal{F}_{\text{HSM}}^{\text{encPW}}$ such that $C \neq C^*$, then \mathcal{B}_1 uses its decryption oracle to decrypt C^* and proceeds as in \mathbf{G}_1 .

Now, if C^* encrypts m_0 the game is distributed as in $\mathbf{G}_1^{(i^*-1)}$ and if it encrypts m_1 then the game is distributed as in $\mathbf{G}_1^{(i^*)}$.

Game \mathbf{G}_6 : Let $q_{\text{REC}} \in \mathbb{N}$ be the number of recovery phases executed. We construct a sequence of games $\mathbf{G}_5^{(0)}, \dots, \mathbf{G}_5^{(q_{\text{REC}})}$, where in $\mathbf{G}_5^{(i)}$ the first i ciphertexts that honest clients produce in recovery are replaced by encryptions of \perp and the remaining ciphertexts stay as in \mathbf{G}_5 . If \mathcal{Z} can distinguish \mathbf{G}_6 from \mathbf{G}_5 , then there is an index $i^* \in [q_{\text{REC}}]$ such that \mathcal{Z} distinguishes $\mathbf{G}_5^{(i^*)}$ and $\mathbf{G}_5^{(i^*-1)}$ with non-negligible advantage. \mathcal{B}_2 internally runs \mathcal{Z} and simulates $\mathbf{G}_5^{(i^*-1)}$ except for the i^* -th recovery. In this recovery, \mathcal{B}_2 does not choose a fresh pair $(\text{pk}_{\text{ENC}}, \text{sk}_{\text{ENC}})$ when it receives a GETPK message, but uses the public key pk_{ENC}^* provided by the challenger. \mathcal{B}_2 gives the messages $m_0 := (\text{pw}, K, \text{ID}_C, \text{ssid})$ and $m_1 := \perp$ to its challenger. When the challenger responds with a ciphertext C^* , then \mathcal{B}_2 uses C^* as the ciphertext in the message $(\text{REC}, \text{ssid}, C^*, \text{ID}_C)$ that the honest client sends to S . Further, if \mathcal{B}_2 receives at some point a message $(\text{REC}, \text{ssid}, C, \text{ID}_C)$ to S or $\mathcal{F}_{\text{HSM}}^{\text{encPW}}$, where $C \neq C^*$, then \mathcal{B}_2 gives C to the decryption oracle provided by the IND-CCA challenger and proceeds as in \mathbf{G}_5 . Finally, \mathcal{B}_2 outputs whatever \mathcal{Z} outputs.

Note that if C^* encrypts m_0 then the view of \mathcal{Z} is distributed exactly as in $\mathbf{G}_5^{(i^*-1)}$ and if C^* encrypts m_1 then the view of \mathcal{Z} is distributed exactly as in game $\mathbf{G}_5^{(i^*)}$. \square

D.3 Proof of Theorem 3

We depict our simulator in Figures 14 to 17. In the below, when referring to blue-colored boxes such as [L.1](#), we always mean the ones from these figures. The gray boxes [CIS.4](#) refer to $\mathcal{F}_{\text{PPKR}}$ instructions from Figures 1 and 2. The simulator works with session state records

$$\langle \text{INIT}, \text{ID}_C, \text{kid}, \text{ssid}, a, \text{ID}_C^*, \text{kid}, a^*, \text{sk}_{\text{ENC}}, \text{sk}_C, b, C \rangle$$

for initialization. All values are initialized to \perp and potentially updated through an initialization. At the end of an initialization by ID_C , the record has the following semantics.

- INIT₁ $\in \{\text{ID}_C, \perp\}$ is either the ID_C of an honest client running initialization or it is set to \perp if the client is corrupt or if the corrupt S impersonates ID_C ;
- INIT₂ $\in \{\text{kid}, \perp\}$ is either the number of initializations started by ID_C , as seen by the HSM, or it is set to \perp if the client is corrupt or if the corrupt S impersonates ID_C ;
- INIT₃ = ssid indicates the sub-session id used by ID_C ;
- INIT₄ $\in \{a, \perp\}$ is either the a value sent by a honest ID_C , or \perp if ID_C is corrupt or if the corrupt S impersonates ID_C .
- INIT₅ $\in \{\text{ID}_C^*, \perp\}$ is either the client name, delivered to the HSM, or set to \perp if the server is fully corrupt;
- INIT₆ $\in \{\text{kid}, \perp\}$ is the number of initializations started by ID_C^* , as seen by the HSM, or set to \perp if the server is fully corrupt;
- INIT₇ $\in \{a^*, \perp\}$ indicates the a value, delivered to the HSM, or set to \perp if the server is fully corrupt;
- INIT₈ $\in \{\text{sk}_{\text{ENC}}, \perp\}$ is the encryption secret key, generated by the HSM for the current sub-session, or set to \perp if the server is fully corrupt;
- INIT₉ $\in \{\text{sk}_C, \perp\}$ is the client's signature secret key either decrypted from c or simulated for an honest client. It is set to \perp if it cannot be extracted from a corrupted initialization;

Simulator $\text{SIM}^{\text{OPRF-PPKR}}$, part 1

$\text{SIM}^{\text{OPRF-PPKR}}$ stores a list $\text{kid}[\cdot]$. Initially $\text{SIM}^{\text{OPRF-PPKR}}$ executes $(\text{pk}, \text{sk}) \xleftarrow{\$} \text{Sig.KeyGen}(1^\lambda)$ and stores $(\text{PK}, \text{pk}, \text{sk})$. Whenever it outputs a message towards some ID_C or from $\mathcal{F}_{\text{HSM}}^{\text{OPRF-PPKR}}$ to a corrupt S , it attests the message with sk .

Wait for X means that $\text{SIM}^{\text{OPRF-PPKR}}$ does not proceed to the next instruction before receiving X and meanwhile gives back activation to \mathcal{A} . Once it receives X , it first proceeds with the instructions on input X and then continues at the instruction, where it waited for X . If a record cannot be retrieved, the query is ignored. For brevity we omit session identifier sid from all inputs, outputs, and records. (**G**₁)

On $(\text{INIT}_C, \text{ssid}, \text{ID}_C)$ from $\mathcal{F}_{\text{PPKR}}$: // honest ID_C

I.1 If $\text{kid}[\text{ID}_C]$ is undefined, set $\text{kid}[\text{ID}_C] \leftarrow 0$. Otherwise, set $\text{kid}[\text{ID}_C] \leftarrow \text{kid}[\text{ID}_C] + 1$. (**G**₃) // new OPRF key for each init

I.2 Give input $(\text{EVAL}, \text{sid}, \text{ID}_C \parallel \text{kid}[\text{ID}_C], \text{ssid}, S, 0)$ to $\mathcal{F}_{\text{OPRF}}$ from ID_C . This triggers Step 5 of SIM^{OPRF} , which outputs $(\text{sid}, \text{ID}_C, \text{ssid}, a)$. (**G**₃)
Record $(\text{INIT}, \text{ID}_C, \text{kid}[\text{ID}_C], \text{ssid}, a, \perp, \perp, \perp, \perp, \perp, \perp)$ and send $(\text{ssid}, a, \text{ID}_C)$ as message from ID_C to S . (**G**₁) // Output of $\mathcal{F}_{\text{OPRF}}$ (eval. of 0) ignored

On $(\text{INIT}_S, \text{ssid}, \text{ID}_C)$ from $\mathcal{F}_{\text{PPKR}}$: // honest S

I.3 If S honest, wait for $(\text{ssid}, a, \text{ID}_C)$ from ID_C to S . (**G**₁) // See $(\text{ssid}, a, \text{ID}_C)$ interface

On $(\text{Init}, \text{ssid}, a, \text{ID}_C)$ from \mathcal{A} to S on behalf of ID_C : // any ID_C , honest S

Ia.1 If S is honest, wait for $(\text{INIT}_S, \text{ssid}, \text{ID}_C)$ from $\mathcal{F}_{\text{PPKR}}$. (**G**₁)

Ia.2 If ID_C is corrupt:

- If $\text{kid}[\text{ID}_C]$ is undefined, set $\text{kid}[\text{ID}_C] \leftarrow 0$, otherwise set $\text{kid}[\text{ID}_C] \leftarrow \text{kid}[\text{ID}_C] + 1$. Set $\text{kid} \leftarrow \text{kid}[\text{ID}_C]$. (**G**₃) // Ensure new oprf key. If ID_C is honest this is ensured in INIT_C
- Record $(\text{INIT}, \perp, \perp, \text{ssid}, \perp, \perp, \perp, \perp, \perp, \perp, \perp, \perp)$ (**G**₁)

Ia.3 If ID_C is honest, retrieve $(\text{INIT}, \text{ID}_C, [\text{kid}], \text{ssid}, *, \perp, \perp, \perp, \perp, \perp, \perp)$. (**G**₃)

Ia.4 Give input $(\text{INIT}, \text{sid}, \text{ID}_C \parallel \text{kid})$ to the simulated $\mathcal{F}_{\text{OPRF}}$ from the simulated S . This triggers Step 2 of SIM^{OPRF} . (**G**₃) // Create new OPRF key

Ia.5 Give input $(\text{SNDRCOMPLETE}, \text{sid}, \text{ID}_C \parallel \text{kid}, \text{ssid})$ to the simulated $\mathcal{F}_{\text{OPRF}}$ from the simulated S . This triggers Step 6 of SIM^{OPRF} , which outputs $(\text{sid}, \text{ssid}, b)$. (**G**₃)

Ia.6 Compute $(\text{pk}_{\text{Enc}}, \text{sk}_{\text{Enc}}) \xleftarrow{\$} \text{PKE.Gen}(1^\lambda)$ and overwrite the entries 6-11 in $(\text{INIT}, *, *, \text{ssid}, *, *, *, *, *, *, *)$ with $\text{ID}_C, \text{kid}, a, \text{sk}_{\text{Enc}}, \perp, b$. Send $(\text{ssid}, b, \text{ID}_C, \text{pk}_{\text{Enc}})$ as HSM-signed message from S to ID_C . (**G**₁) // Can later access OPRF keys via F records in SIM^{OPRF}

On $(\text{Init}, \text{ssid}, a, \text{ID}_C)$ from \mathcal{A} to S on behalf of corrupt S : // any ID_C

Ia.7 If no record $(\text{INIT}, *, [\text{kid}], \text{ssid}, *, *, *, *, *, *)$ exists, create $(\text{INIT}, \perp, \perp, \text{ssid}, \perp, \perp, \perp, \perp, \perp, \perp, \perp, \perp)$. (**G**₁)
Additionally, if $\text{kid}[\text{ID}_C]$ is undefined set $\text{kid}[\text{ID}_C] \leftarrow 0$ and otherwise $\text{kid}[\text{ID}_C] \leftarrow \text{kid}[\text{ID}_C] + 1$ and set $\text{kid} \leftarrow \text{kid}[\text{ID}_C]$. (**G**₃) // ID_C corrupt and S just forwards or S impersonates ID_C

Ia.8 Execute Steps **Ia.4**-**Ia.6**, except that the output is returned to S instead of sending it from S to ID_C . (**G**₃)

On $(\text{ssid}, b, \text{ID}_C^*, \text{pk}_{\text{Enc}})$ from \mathcal{A} to ID_C : // honest ID_C , any S

Ib.1 Retrieve $(\text{INIT}, \text{ID}_C, [\text{kid}], \text{ssid}, *, *, *, *, *, *)$. (**G**₁)

Ib.2 If S is not honest, give input $(\text{INIT}_S, \text{ssid}, \text{ID}_C)$ to $\mathcal{F}_{\text{PPKR}}$. On response $(\text{INIT}_S, \text{ssid}, \text{ID}_C)$ from $\mathcal{F}_{\text{PPKR}}$ or if $\mathcal{F}_{\text{PPKR}}$ ignores the input, continue below. (**G**₁₀)
// Ensures that we can produce output for ID_C

Ib.3 If S is not fully corrupt and $(\text{ssid}, b, \text{ID}_C^*, \text{pk}_{\text{Enc}})$ was never output by $\text{SIM}^{\text{OPRF-PPKR}}$ on behalf of S or $\mathcal{F}_{\text{HSM}}^{\text{OPRF-PPKR}}$, or if S is fully corrupt and the attestation signature does not verify, send $(\text{COMPLETEINIT}_C, \text{ssid}, 0)$ to $\mathcal{F}_{\text{PPKR}}$. (**G**₁₀)

Ib.4 If $\text{ID}_C \neq \text{ID}_C^*$, send $(\text{COMPLETEINIT}_C, \text{ssid}, 0)$ to $\mathcal{F}_{\text{PPKR}}$. (**G**₁₀) // ID_C 's a rerouted

Ib.5 Compute $(\text{sk}_C, \text{pk}_C) \leftarrow \text{Sig.Gen}(1^\lambda)$. (**G**₁)

Ib.6 If S is fully corrupt:

- (a) Send $(\text{COMPLETEINIT}_S, \text{ssid}, 1)$ to $\mathcal{F}_{\text{PPKR}}$. (**G**₁₀)
- (b) For each record $(H_2, [\text{pw}] \parallel \text{ID}_C, *, [y])$ marked CONSISTENT (**G**₁₄) send $(\text{OFFLINEATTACK}, \text{ID}_C, \text{pw}, i + 1)$ to $\mathcal{F}_{\text{PPKR}}$, where $i \in \mathbb{N}$ is the largest number such that a record $(\text{LEAKED}, \text{ID}_C, *, *, *, i)$ exists. (**G**₁₀)
- (c) If for any query the output is $K \neq \text{FAIL}$ (**G**₁₀), compute $c \xleftarrow{\$} \text{AE.Enc}(y, (K, \text{sk}_C))$. Else set $c \xleftarrow{\$} \text{SIM}_{\text{EQV}}(\lambda + |\text{sk}_C|)$. (**G**₆)
- (d) Record $(\text{FILE}, \text{ID}_C, \text{kid}, \text{pk}_C, \text{sk}_C, c)$ and $(\text{LEAKED}, \text{ID}_C, \text{kid}, \text{sk}_C, c, i + 1)$, where $i \in \mathbb{N}$ is the largest number such that a record $(\text{LEAKED}, \text{ID}_C, *, *, *, i)$ exists or $i = 0$ if no such record exists. (**G**₁)
- (e) Compute $C \xleftarrow{\$} \text{PKE.Enc}(\text{pk}_{\text{Enc}}, (\text{ssid}, \text{pk}_C, c))$. (**G**₁)
- (f) Overwrite the three last entries of the INIT record retrieved above with sk_C, b, C . (**G**₁)

Ib.7 If S is not fully corrupt:

- (a) Compute $c \xleftarrow{\$} \text{SIM}_{\text{EQV}}(\lambda + |\text{sk}_C|)$. (**G**₆)
- (b) Compute $C \leftarrow \text{PKE.Enc}(\text{sk}_{\text{Enc}}, \perp)$. (**G**₁₃)
- (c) overwrite the tenth entry of the record retrieved above with sk_C and the last entry with C . (**G**₁)

Ib.8 Send (ssid, C) to S (**G**₁)
and $(\text{COMPLETEINIT}_C, \text{ssid}, 1)$ to $\mathcal{F}_{\text{PPKR}}$. (**G**₁₀). Store $(C, \text{ssid}, \text{pk}_C, c)$ (**G**₁₃).

 Figure 14: Simulator $\text{SIM}^{\text{OPRF-PPKR}}$ for $\pi^{\text{OPRF-PPKR}}$, part 1.

$\text{INIT}_{10} = b$ is either the b value computed by the honest HSM or the value received by ID_C if S is fully corrupt;
 $\text{INIT}_{11} \in \{C, \perp\}$ is the ciphertext C computed by ID_C or \perp if ID_C is corrupt.

Similarly, the simulator works with the following session state records

$$\langle \text{REC}, \text{ID}_C, \text{kid}, \text{ssid}, a, \text{match}, \text{ID}_C^*, \text{kid}^*, a^*, \text{pk}_C, b, c \rangle$$

in the recovery. Again, all values are initialized to \perp and potentially updated through a recovery. At the end of the recovery by ID_C , the record has the following semantics.

- $\text{REC}_1 \in \{\text{ID}_C, \perp\}$ is either the ID_C of an honest client running recovery or it is set to \perp if the client is corrupt or if the corrupt S impersonates ID_C ;
- $\text{REC}_2 \in \{\text{kid}, \perp\}$ is either the number of initializations started by ID_C when ID_C starts this recovery, as seen by the HSM, or it is set to \perp if the client is corrupt or if the corrupt S impersonates ID_C ;
- $\text{REC}_3 = \text{ssid}$ indicates the sub-session id used by ID_C ;
- $\text{REC}_4 \in \{a, \perp\}$ is either the a' value sent by a honest ID_C , or \perp if ID_C is corrupt or if the corrupt S impersonates ID_C .
- $\text{REC}_5 \in \{0, 1\}$ indicates whether the password in the recovery is correct
- $\text{REC}_6 \in \{\text{ID}_C^*, \perp\}$ is either the client name, delivered to the HSM, or set to \perp if the server is fully corrupt;
- $\text{REC}_7 \in \{\text{kid}^*, \perp\}$ is the number of initializations started by ID_C^* when S receives the first message in this recovery, as seen by the HSM, or set to \perp if the server is fully corrupt;
- $\text{REC}_8 \in \{a^*, \perp\}$ indicates the a' value, delivered to the HSM, or set to \perp if the server is fully corrupt;
- $\text{REC}_9 \in \{\text{pk}_C, \perp\}$ is the signature public key used by the HSM in this recovery, or set to \perp if the server is fully corrupt;
- $\text{REC}_{10} = b'$ is either the b value computed by the honest HSM or the value received by ID_C if S is fully corrupt;
- $\text{REC}_{11} \in \{c\}$ is the ciphertext c sent by the honest HSM or \perp if S is fully corrupt.

We construct a sequence of hybrid games \mathbf{G}_0 to \mathbf{G}_{14} where we gradually change the real-world execution of the protocol $\pi^{\text{OPRF-PPKR}}$ (interacting with the hybrid functionality $\mathcal{F}_{\text{HSM}}^{\text{OPRF-PPKR}}$) to reach the ideal-world execution, where the environment interacts with the simulator from Figures 14 to 17 and the ideal functionality $\mathcal{F}_{\text{PPKR}}$. We write $\text{Pr}[\mathbf{G}_i]$ to denote the probability that the environment outputs 1 in the hybrid game \mathbf{G}_i .

Game \mathbf{G}_0 : Real world. This is the real world.

Game \mathbf{G}_1 : Create simulator. In this game we create two new entities called the ideal functionality \mathcal{F} and the simulator SIM . Initially, \mathcal{F} just forwards the input of the dummy parties to SIM and outputs what SIM instructs it to output. In particular, \mathcal{F} has interfaces $(\text{INIT}_C, \text{ssid}, \text{ID}_C, \text{pw})$, $(\text{INIT}_S, \text{ssid}, \text{ID}_C)$, $(\text{REC}_C, \text{ssid}, \text{ID}_C, \text{pw}')$, and $(\text{REC}_S, \text{ssid}, \text{ID}_C)$ that just forward the input to SIM . Additionally, \mathcal{F} has the corruption and attack interfaces LEARFILE , FULLYCORRUPT , MALICIOUSINIT , MALICIOUSREC , and OFFLINEATTACK with the exact same code as in $\mathcal{F}_{\text{PPKR}}$. Note, however, that SIM does not use these interfaces yet.

The simulator executes the code of all honest parties of the protocol internally on the input that it is provided by \mathcal{F} and it internally runs the code of the hybrid functionality $\mathcal{F}_{\text{HSM}}^{\text{OPRF-PPKR}}$. Additionally, it creates records exactly as $\text{SIM}^{\text{OPRF-PPKR}}$ does, but at this point never uses the values from any of these records. Note that these are just syntactical changes and the protocol is still executed as in the real world. We have

$$\Pr[\mathbf{G}_1] = \Pr[\mathbf{G}_0].$$

Game \mathbf{G}_2 : Output FAIL on tampered message from HSM.

In this game, we change how the simulator reacts when an honest client receives a message from the HSM that was tampered with. Whenever an honest client receives a message (INITRES , ssid , Succ) while S is not fully corrupt, where SIM never produced this message on behalf of S or $\mathcal{F}_{\text{HSM}}^{\text{OPRF-PPKR}}$, then SIM makes the client output FAIL (see [Ib.3](#)). Similarly, when an honest client receives a message (RECREs , ssid , out) that was never produced by SIM on behalf of the HSM, then SIM makes the client output FAIL (see [Rb.5](#)).

This game is identical to \mathbf{G}_1 unless \mathcal{A} delivers a message to ID_C that was never output by SIM but the attestation signature is still valid. If an environment \mathcal{Z}^* exists that can distinguish between \mathbf{G}_2 and \mathbf{G}_1 , we can construct an adversary \mathcal{B}_1 against the sEUF-CMA security of the signature scheme Sig used for attestation. \mathcal{B}_1 internally runs the whole experiment including \mathcal{Z}^* and outputs the above described message with the signature to its challenger. We get

$$|\Pr[\mathbf{G}_2] - \Pr[\mathbf{G}_1]| \leq \text{Adv}_{\text{Sig}, \mathcal{B}_1}^{\text{sEUF-CMA}}(\lambda).$$

Game \mathbf{G}_3 : Simulate the OPRF execution. In this game, we replace the 2HashDH protocol with its ideal execution, namely functionality $\mathcal{F}_{\text{OPRF}}^\ell$ of Figure 7 and simulator $\text{SIM}_{\text{OPRF}}(\text{sid}, H_1, H_2, N)$ of Figure 8. We consequently let the simulated parties use the interfaces of $\mathcal{F}_{\text{OPRF}}^\ell$, e.g., instead of computing $a \leftarrow H_1(\text{pw} \parallel \text{ID}_C)^r$, a party sends (EVAL , sid , $\text{ID}_C \parallel \text{kid}$, ssid , S , $\text{pw} \parallel \text{ID}_C$) to $\mathcal{F}_{\text{OPRF}}^\ell$, where for the first initialization by ID_C we have $\text{kid} = 0$ and with each subsequent initialization by ID_C it is incremented by 1. Whenever a party is supposed to obtain a PRF value, the simulator calls the RcvCOMPLETE interface of $\mathcal{F}_{\text{OPRF}}^\ell$. Since in this ideal OPRF execution, OPRF keys are chosen by $\text{SIM}_{\text{OPRF}}(\text{sid}, H_1, H_2, N)$ (see Step 2 of Figure 8), the simulator takes these keys whenever it has to store a file. Finally, the 2HashDH protocol messages a, b sent by honest parties are now simulated by $\text{SIM}_{\text{OPRF}}(\text{sid}, H_1, H_2, N)$. The formal changes can be read from Figures 14 to 17, marked with [\(G3\)](#).

This and the previous game are indistinguishable up to the advantage of distinguishing the real 2HashDH execution from the ideal execution of $\mathcal{F}_{\text{OPRF}}^\ell$ and $\text{SIM}_{\text{OPRF}}(\text{sid}, H_1, H_2, N)$, which by Theorem 4 (originally Theorem 2 of [DFG⁺23a]) is bounded by

$$|\Pr[\mathbf{G}_3] - \Pr[\mathbf{G}_2]| \leq q_{\text{INIT}} \text{Adv}_{\mathcal{B}_2, \mathbb{G}}^{(q_E + q_H, q_E) - \text{OMDH}}(\lambda) + (q_E + q_H)^2 / q,$$

where q_{INIT} is the number of INIT inputs, q_E the number of overall INIT and REC inputs, q_H the number of H_1 queries, q is the order of \mathbb{G} and $\text{Adv}_{\mathcal{B}_2, \mathbb{G}}^{(x, y) - \text{OMDH}}(\lambda)$ denotes the advantage of breaking the one-more Diffie-Hellman problem in \mathbb{G} (see Definition 4).

Game \mathbf{G}_4 : Abort upon ambiguous ciphertexts. We let the simulator abort if it obtains an adversarially-generated AE ciphertext c such that for two values k, k' from records $\langle F, S, [\text{kid}], k, * \rangle, \langle F, S, [\text{kid}'], k', * \rangle$ stored by the OPRF simulator SIM_{OPRF} , c successfully decrypts under k and k' , i.e., the output of AE.Dec is not \perp (cf. [IC.2\(c\)](#), [IC.7\(c\)\(ii\)](#), and [Rb.3\(a\)](#)).

This and the previous game are indistinguishable by the random-key robustness of AE. The reduction is straightforward, running on two challenge keys k, k' : it randomly chooses two occasions where $\mathcal{F}_{\text{OPRF}}$ samples $F_{\text{sid}, S, \text{kid}}(x) \xleftarrow{\$} \{0, 1\}^\ell$ and instead sets $F_{\text{sid}, S, \text{kid}}(x) \leftarrow k$ for the first one and $\leftarrow k'$ for the second. Note that this does not change the distribution as k, k' are uniformly random. Ciphertexts c passing the winning condition of the random-key robustness game can be detected by trial-decrypting with k and k' . We hence have

$$|\Pr[\mathbf{G}_4] - \Pr[\mathbf{G}_3]| \leq q_E^2 \text{Adv}_{\mathcal{B}_3, \text{AE}}^{\text{rkr}}(\lambda).$$

Game \mathbf{G}_5 : Honest ID_C fails upon malicious AE ciphertext.

In this game, we modify how SIM proceeds when receiving an adversarial message or honestly delivered message (ssid , b' , c , ID_C) from \mathcal{A} to some honest ID_C . SIM checks whether all of the following conditions hold:

- c was not produced by SIM ,
- There is no value $F_{\text{sid}, S, *}(pw' \parallel \text{ID}_C)$ defined in $\mathcal{F}_{\text{OPRF}}$ that successfully decrypts c ,
- There is a value $F_{\text{sid}, S, [\text{kid}]}(pw' \parallel \text{ID}_C)$ defined in $\mathcal{F}_{\text{OPRF}}$ that successfully decrypts c , but in this recovery a malicious S did not use the OPRF key with kid .

Formally, SIM executes these checks via the code in [Rb.3\(a\)](#) and [Rb.3\(b\)](#). If all conditions hold, then SIM lets ID_C immediately output (RECREs , ssid , FAIL) instead of following the protocol as in \mathbf{G}_4 . If any condition does not hold, it proceeds as in \mathbf{G}_4 . Note that \mathbf{G}_5 only differs if ID_C outputs (RECREs , ssid , Succ) in \mathbf{G}_4 and (RECREs , ssid , FAIL) in \mathbf{G}_5 , which we denote by the event E_{AE} . We now construct an adversary \mathcal{B}_4 that breaks the INT-CTXT-security of AE given any environment \mathcal{Z} that causes E_{AE} .

\mathcal{B}_4 acts like SIM and chooses $i^* \xleftarrow{\$} \{1, \dots, q_{\text{INIT}}\}$. Let ID_C^* denote the ID_C that executes the i^* -th initialization. If ID_C^* is corrupt, \mathcal{B}_4 aborts. Otherwise, it does not compute c by encrypting (K, sk_C) under ρ^* but instead submits (K, sk_C) to its encryption oracle. Note that this implicitly programs ρ^* to the key k^* chosen by the INT-CTXT experiment. Since k^* is chosen uniformly at random and $\mathcal{F}_{\text{OPRF}}$ chooses its outputs uniformly at random, this means that there is no change in the distribution of c .

In any subsequent recovery by ID_C^* before ID_C^* executes another initialization, \mathcal{B}_4 checks the conditions listed above, and if they hold, it outputs the c' received in that recovery to the INT-CTXT experiment. Further, if the file from the

Simulator $\text{SIM}^{\text{OPRF-PPKR}}$, part 2	
On (ssid, C) from \mathcal{A} to S on behalf of ID_C : // any ID_C , honest S	
IC.1	Retrieve $\langle \text{INIT}, *, *, \text{ssid}, *, *, [\text{kid}], *, [\text{sk}_{\text{Enc}}], [\text{sk}_C], *, [C'] \rangle$. (G ₁)
IC.2	If ID_C is corrupt: (G ₈) <ul style="list-style-type: none"> (a) Execute HSM code, (G₁) <ul style="list-style-type: none"> where the entry k_{OPRF} of the stored FILE record is set to \perp, (G₃) up to determining out and denote the decryption by $(\text{ssid}', \text{pk}_C, c)$. (G₁) If $out = \text{FAIL}$, send $(\text{COMPLETEINIT}, \text{ssid}, 0)$ to $\mathcal{F}_{\text{PPKR}}$ and otherwise continue. (G₁₁) // SIM_{OPRF} has the OPRF key (b) Find a record $\langle H_2, [\text{pw}] \parallel *, *, [y] \rangle$ that is marked CONSISTENT such that $\perp \neq (K, \text{sk}_C) \leftarrow \text{AE.Dec}(y, c)$. Otherwise set $\text{pw} \leftarrow \perp$. (G₈) (c) If more than one consistent record is found, abort the simulation. (G₄) (d) Send input $(\text{INIT}, \text{ssid}, \text{pw})$ to $\mathcal{F}_{\text{PPKR}}$ on behalf of ID_C. (G₈) <ul style="list-style-type: none"> On response $(\text{INIT}, \text{ssid}, \text{ID}_C)$ from $\mathcal{F}_{\text{PPKR}}$ send message $(\text{COMPLETEINIT}, \text{ssid}, 1)$ to $\mathcal{F}_{\text{PPKR}}$. (G₁₁).
IC.3	If ID_C is honest, retrieve $(C, \text{ssid}, [\text{pk}_C], [c])$. (G ₁₃) <ul style="list-style-type: none"> Store $(\text{FILE}, \text{ID}_C, \text{pk}_C, c, \perp, 10)$ overwriting any $(\text{FILE}, \text{ID}_C, *, *, *, *)$ (G₁), and send message $(\text{COMPLETEINIT}, \text{ssid}, 1)$ to $\mathcal{F}_{\text{PPKR}}$. (G₁₀)
On (ssid, C) from \mathcal{A} to $\mathcal{F}_{\text{HSM}}^{\text{OPRF-PPKR}}$ on behalf of corrupt S : // any ID_C , corrupt S	
IC.4	Retrieve $\langle \text{INIT}, [\text{ID}_C], *, *, \text{ssid}, [a], [\text{ID}_C^*], [\text{kid}], [a^*], [\text{sk}_{\text{Enc}}], [\text{sk}_C], *, [C'] \rangle$. (G ₁)
IC.5	If $C' = \perp$, give input $(\text{INIT}, \text{ssid}, \text{ID}_C)$ to $\mathcal{F}_{\text{PPKR}}$. On response $(\text{INIT}, \text{ssid}, \text{ID}_C)$ from $\mathcal{F}_{\text{PPKR}}$ continue below. (G ₁₀) // $C' \neq \perp$ implies that b was delivered to ID_C , where the simulator gave input INIT to $\mathcal{F}_{\text{PPKR}}$
IC.6	If $C = C'$, retrieve record $\langle C, [\text{ssid}], [\text{pk}_C], [c] \rangle$ (G ₁₃). Give input $(\text{COMPLETEINIT}, \text{ssid}, 1)$ to $\mathcal{F}_{\text{PPKR}}$ (G ₁₀). <ul style="list-style-type: none"> record $(\text{FILE}, \text{ID}_C^*, \text{kid}, \text{pk}_C, \text{sk}_C, c)$ overwriting any $(\text{FILE}, \text{ID}_C^*, *, *, *, *)$, and send $(\text{INITRES}, \text{ssid}, \text{Succ})$ to S as output of the HSM. (G₁) // honest delivery of C, which also implies honest delivery of a and $\text{ID}_C = \text{ID}_C^*$ as otherwise $C' = \perp$
IC.7	If $C \neq C'$: <ul style="list-style-type: none"> (a) Search for a record $\langle C, [\text{ssid}'], [\text{pk}_C], [c] \rangle$. (G₁₃) <ul style="list-style-type: none"> If $\text{ssid}' \neq \text{ssid}$, send $(\text{INITRES}, \text{ssid}, \text{FAIL})$ to S. Otherwise, record $(\text{FILE}, \text{ID}_C^*, \text{kid}, \text{pk}_C, \text{sk}_C, c)$ overwriting any $(\text{FILE}, \text{ID}_C^*, *, *, *, *)$. (G₁) (b) If no such record $\langle C, *, *, * \rangle$ exists (G₁₃), execute the HSM code on input (ssid, C), (G₁) <ul style="list-style-type: none"> where the entry k_{OPRF} of the stored FILE record is set to \perp, (G₃) up to determining out and denote the result of the decryption by $(\text{ssid}', \text{pk}_C, c)$. If $out = \text{FAIL}$, send $(\text{INITRES}, \text{ssid}, \text{FAIL})$ to S. If $out \neq \text{FAIL}$, record $(\text{FILE}, \text{ID}_C^*, \text{kid}, \text{pk}_C, \text{sk}_C, c)$ overwriting any $(\text{FILE}, \text{ID}_C^*, *, *, *, *)$. (G₁) (c) Search a record $\langle H_2, [\text{pw}] \parallel *, *, [y] \rangle$ that is marked CONSISTENT such that $\perp \neq (K, \text{sk}_C) \leftarrow \text{AE.Dec}(y, c)$. Overwrite the third to last entry of the REC record retrieved above with sk_C. (G₈) <ul style="list-style-type: none"> (i) If no such record exists, send $(\text{MALICIOUSINIT}, \text{ID}_C^*, \perp, \perp)$ to $\mathcal{F}_{\text{PPKR}}$. (G₈) <ul style="list-style-type: none"> // \mathcal{Z} has not yet computed the PRF value for a (ii) If more than one consistent record is found, abort the simulation. (G₄) (iii) In all other cases, send $(\text{MALICIOUSINIT}, \text{ssid}, \text{ID}_C^*, \text{pw}, K)$ to $\mathcal{F}_{\text{PPKR}}$. (G₈) <ul style="list-style-type: none"> // \mathcal{Z} can decrypt c (d) Send $(\text{INITRES}, \text{ssid}, \text{Succ})$ to S as output of the HSM. (G₁)
On random oracle query $H_1(\text{pw} \parallel \text{ID}_C)$ from \mathcal{A} :	
H1.1	Query $H_1(\text{pw} \parallel \text{ID}_C)$ to $\text{SIM}_{\text{OPRF}} \rightarrow$ Step 4 of SIM_{OPRF} . Return output. (G ₃)
On random oracle query $H_2(\text{pw} \parallel \text{ID}_C, x)$ from \mathcal{A} :	
H2.1	If a record $\langle H_2, \text{pw} \parallel \text{ID}_C, x, y \rangle$ exists, output y . (G ₁)
H2.2	If there exists a record $\langle F, S, \text{ID}_C \parallel [\text{kid}], [k_{\text{OPRF}}], * \rangle$ in SIM_{OPRF} such that $x = H_1(\text{pw} \parallel \text{ID}_C)^{k_{\text{OPRF}}}$: (G ₁₄) <ul style="list-style-type: none"> (a) Retrieve $(\text{LEAKED}, \text{ID}_C, \text{kid}, [\text{sk}_C], [c], [i])$. If multiple exist, choose smallest i. (G₁₀) (b) Send $(\text{OFFLINEATTACK}, \text{ID}_C, \text{pw}, i)$ to $\mathcal{F}_{\text{PPKR}}$. If it outputs $K \neq \text{FAIL}$, (G₁₀) run $y \stackrel{\perp}{\leftarrow} \text{SIM}_{\text{EQV}}(c, (K, \text{sk}_C))$, record $\langle H_2, \text{pw} \parallel \text{ID}_C, x, y \rangle$ and output y. (G₆) Otherwise continue.
H2.3	Query $H_2(\text{pw} \parallel \text{ID}_C, x)$ to SIM_{OPRF} , which triggers Step 8 of SIM_{OPRF} . Let y denote its output. Record $\langle H_2, \text{pw} \parallel \text{ID}_C, x, y \rangle$. If y was produced by $\mathcal{F}_{\text{OPRF}}$ mark the record as CONSISTENT. Output y . (G ₃)
On $(\text{REC}, \text{ssid}, \text{ID}_C, \text{match})$ from $\mathcal{F}_{\text{PPKR}}$ // Any ID_C	
R.1	If no record $(\text{FILE}, \text{ID}_C, [\text{kid}], *, *, *)$ exists, set $\text{kid} \leftarrow \perp$. (G ₃)
R.2	Give input $(\text{EVAL}, \text{sid}, \text{ID}_C \parallel \text{kid}, \text{ssid}, S, 0)$ to the simulated $\mathcal{F}_{\text{OPRF}}$ from the simulated ID_C . This triggers Step 5 of SIM_{OPRF} , which outputs $(\text{sid}, \text{ID}_C, \text{ssid}, a')$. (G ₃)
R.3	If a record $\langle \text{REC}, \perp, \perp, \text{ssid}, \perp, \perp, *, *, *, *, * \rangle$ exists, overwrite the second, third, fifth, and sixth entry with $(\text{ID}_C, \text{kid}, a', \text{match})$. (G ₁) // This record exists if RecS was executed previously
R.4	Otherwise, record $\langle \text{REC}, \text{ID}_C, \text{kid}, \text{ssid}, a', \text{match}, \perp, \perp, \perp, \perp, \perp \rangle$. (G ₁)
R.5	Send $(\text{Rec}, \text{ssid}, a', \text{ID}_C)$ as message from ID_C to S . (G ₁)
On $(\text{RecS}, \text{ssid}, \text{ID}_C, \text{match})$ from $\mathcal{F}_{\text{PPKR}}$: // honest, corrupt, or fully corrupt S	
R.6	If a record $\langle \text{REC}, *, *, \text{ssid}, *, \perp, *, *, *, *, * \rangle$ exists, overwrite \perp with match . (G ₁) // This record exists if RecC was executed previously
R.7	Otherwise, record $\langle \text{REC}, \perp, \perp, \text{ssid}, \perp, \text{match}, \perp, \perp, \perp, \perp, \perp \rangle$. (G ₁)
R.8	Wait for $(\text{Rec}, \text{ssid}, a', \text{ID}_C)$. (G ₁) // See the $(\text{Rec}, \text{ssid}, a', \text{ID}_C)$ interface below.

Figure 15: Simulator $\text{SIM}^{\text{OPRF-PPKR}}$ for $\pi^{\text{OPRF-PPKR}}$, part 2.

i^* -th initialization was leaked, in any subsequent recovery by ID_C^* , \mathcal{B}_4 checks the same conditions and outputs the c' received in that recovery to the INT-CTXT experiment if they hold.

If such a recovery by ID_C^* is the first that causes the event E_{AE} , which happens with probability $1/q_{\text{INIT}}$, \mathcal{B}_4 wins the INT-CTXT experiment due to the following. Let c^* be the ciphertext that causes E_{AE} , which means that in G_4 , ID_C^* would successfully decrypt c^* using ρ^* . Since \mathcal{B}_4 implicitly set $\rho^* = k^*$, c^* successfully decrypts under k^* and constitutes a forgery in the INT-CTXT experiment. Therefore, we have

$$|\Pr[\text{G}_5] - \Pr[\text{G}_4]| \leq q_{\text{INIT}} \text{Adv}_{\mathcal{B}_4, \text{AE}}^{\text{INT-CTXT}}(\lambda).$$

Simulator $\text{SIM}^{\text{OPRF-PPKR}}$, part 3.	
On $(\text{Rec}, \text{ssid}, a', \text{ID}_C)$ from \mathcal{A} to \mathcal{S} on behalf of ID_C : // any ID_C , any (if jumping here from the $(\text{Rec}, \text{ssid}, a', \text{ID}_C)$ from \mathcal{A} to $\mathcal{F}_{\text{HSM}}^{\text{OPRF-PPKR}}$ interface) \mathcal{S}	
Ra.1	If \mathcal{S} is honest, wait for $(\text{RECS}, \text{ssid}, \text{ID}_C, \text{match})$ message from $\mathcal{F}_{\text{PPKR}}$. (\mathbf{G}_1) // See $(\text{RECS}, \text{ssid}, \text{ID}_C, \text{match})$ interface.
Ra.2	Retrieve $(\text{FILE}, \text{ID}_C, [\text{kid}], [\text{pk}_C], *, [c'])$. If no such record exists and \mathcal{S} is honest, send $(\text{COMPLETEREC}, \text{ssid}, 0)$ to $\mathcal{F}_{\text{PPKR}}$. (\mathbf{G}_{12})
Ra.3	Execute HSM code on input $(\text{Rec}, \text{ssid}, a', \text{ID}_C)$ up to determining the output (\mathbf{G}_1) and with the computation of b' substituted as follows. Give input $(\text{SNDRCOMPLETE}, \text{sid}, \text{ID}_C \parallel \text{kid}, \text{ssid})$ to the simulated $\mathcal{F}_{\text{OPRF}}$ from the simulated \mathcal{S} . This triggers Step 6 of STM_{OPRF} , which outputs $(\text{sid}, \text{ssid}, b')$. (\mathbf{G}_3)
Ra.4	If the HSM output is $\text{ssid}, \text{DELREC}$, record $(\text{DELREC}, \text{ssid})$ and send $(\text{ssid}, \text{DELREC})$ as HSM-signed message from \mathcal{S} to ID_C . If the output is $\text{ssid}, b', c, \text{ID}_C$, overwrite the six last entries of $(\text{REC}, *, *, \text{ssid}, *, *, *, *, *, *, *, *)$ with $\text{ID}_C, \text{kid}, a', \text{pk}_C, b', c'$ and send $(\text{ssid}, b', c', \text{ID}_C)$ as HSM-signed message to ID_C . (\mathbf{G}_1)
On $(\text{Rec}, \text{ssid}, a', \text{ID}_C)$ from \mathcal{A} to $\mathcal{F}_{\text{HSM}}^{\text{OPRF-PPKR}}$ on behalf of corrupt \mathcal{S} :	
Ra.5	If no record $(\text{REC}, *, *, \text{ssid}, *, *, *, *, *, *, *, *)$ exists, create $(\text{REC}, \perp, \perp, \text{ssid}, \perp, \perp, \perp, \perp, \perp, \perp, \perp, \perp)$. (\mathbf{G}_1)
Ra.6	Execute Steps Ra.2–Ra.4 except that the output is returned to \mathcal{S} instead of sending it from \mathcal{S} to ID_C . (\mathbf{G}_3)
On $(\text{ssid}, \text{DELREC})$ from \mathcal{A} to ID_C on behalf of \mathcal{S} : // honest ID_C , any \mathcal{S}	
Del.1	If no record $(\text{REC}, \text{ID}_C, *, \text{ssid}, *, *, *, *, *, *, *, *)$ exists or if there already was a message $(\text{ssid}, b', c, \text{ID}_C^*)$ to ID_C , ignore. (\mathbf{G}_1)
Del.2	If \mathcal{S} is not honest, send $(\text{RECS}, \text{ssid}, \text{ID}_C, 0, \text{DELREC}, 0)$ to $\mathcal{F}_{\text{PPKR}}$. On response $(\text{RECS}, \text{ssid}, \text{ID}_C, \text{match})$ from $\mathcal{F}_{\text{PPKR}}$ continue below. (\mathbf{G}_{10})
Del.3	If HSM signature verifies, send $(\text{COMPLETEREC}, \text{ssid}, 1)$ to $\mathcal{F}_{\text{PPKR}}$ and otherwise $(\text{COMPLETEREC}, \text{ssid}, \text{ID}_C, 0)$. (\mathbf{G}_{10})
On $(\text{ssid}, b', c, \text{ID}_C^*)$ from \mathcal{A} to ID_C on behalf of \mathcal{S} : // honest ID_C , any \mathcal{S}	
Rb.1	Retrieve $(\text{REC}, \text{ID}_C, *, \text{ssid}, [a'], [\text{match}], *, [\text{kid}], *, *, *, *)$. (\mathbf{G}_1) If no such record exists or if there already was a message $(\text{ssid}, \text{DELREC})$ to ID_C , ignore. (\mathbf{G}_1) // ID_C did not start recovery or record deleted.
Rb.2	If \mathcal{S} is corrupt, set $\text{pw} \leftarrow \perp, K \leftarrow \perp, i \leftarrow \perp$. (\mathbf{G}_{12})
Rb.3	If \mathcal{S} is fully corrupt: (a) If record $(H_2, [\text{pw}] \parallel \text{ID}_C, [u], [y])$ marked CONSISTENT exists (\mathbf{G}_{14}) and records $(\text{ID}_C \parallel \text{kid}, \text{ssid}, \text{ID}_C, [r'])$ and $(H_1, \text{pw} \parallel \text{ID}_C, [r])$ exist in SIM_{OPRF} s.t. $\perp \neq (K, \text{sk}_C) \leftarrow \text{AE.Dec}(y, c)$, and $b'^{1/r'} = u^{1/r}$, (\mathbf{G}_5) set $i \leftarrow 0$. (\mathbf{G}_{10}) If more than one consistent record is found, abort the simulation. (\mathbf{G}_4) // Adversarial OPRF key (b) Otherwise, if records $(\text{LEAKED}, \text{ID}_C, [\text{kid}'], [\text{sk}_C], c, [i])$ and $(\text{INIT}, \text{ID}_C, \text{kid}', *, [a], *, *, *, *, [b], *)$ exist and a record $(F, S, *, [k], *)$ exists in SIM_{OPRF} such that $a^k = b$ and $a^{k'} = b'$, (\mathbf{G}_5) set $\text{pw} \leftarrow \perp, K \leftarrow \perp$. (\mathbf{G}_{10}) // Impersonation with old file (c) Otherwise, set $\text{pw} \leftarrow 0, K \leftarrow \text{FAIL}, i \leftarrow 0, \text{sk}_C \leftarrow \perp$. (\mathbf{G}_{10})
Rb.4	If \mathcal{S} is not honest, give input $(\text{RECS}, \text{ssid}, \text{ID}_C^*, \text{pw}, K, i)$ to $\mathcal{F}_{\text{PPKR}}$. On response $(\text{RECS}, \text{ssid}, \text{ID}_C^*, \text{match})$ overwrite the sixth entry in the record retrieved in Step 1 with match . (\mathbf{G}_{10}) // Ensures that we can produce output for ID_C
Rb.5	If \mathcal{S} is not fully corrupt and $(\text{ssid}, b', c, \text{ID}_C^*)$ was never output by $\text{SIM}^{\text{OPRF-PPKR}}$ on behalf of \mathcal{S} or $\mathcal{F}_{\text{HSM}}^{\text{OPRF-PPKR}}$, or if \mathcal{S} is fully corrupt and the attestation signature does not verify, send $(\text{COMPLETEINITC}, \text{ssid}, 0)$ to $\mathcal{F}_{\text{PPKR}}$. (\mathbf{G}_{10})
Rb.6	If $\text{ID}_C \neq \text{ID}_C^*$, send $(\text{COMPLETEREC}, \text{ssid}, 0)$ to $\mathcal{F}_{\text{PPKR}}$. (\mathbf{G}_{10}) // ID_C^* 's a rerouted
Rb.7	If \mathcal{S} is not fully corrupt, retrieve $(\text{FILE}, \text{ID}_C, \text{kid}, *, [\text{sk}_C], *)$. (\mathbf{G}_{14})
Rb.8	If $\text{match} = 1$ and $\text{sk}_C \neq \perp$, compute $\sigma \leftarrow \text{Sig.Sign}(\text{sk}_C, (a', \text{ID}_C, \text{ssid}, b', c))$, send (ssid, σ) as message from ID_C to \mathcal{S} (\mathbf{G}_{14}) and send $(\text{COMPLETEREC}, \text{ssid}, 1)$ to $\mathcal{F}_{\text{PPKR}}$. (\mathbf{G}_{10})
Rb.9	If $\text{match} = 0$ or $\text{sk}_C = \perp$, (\mathbf{G}_{14}) send $(\text{COMPLETEREC}, \text{ssid}, 0)$ to $\mathcal{F}_{\text{PPKR}}$. (\mathbf{G}_{10})

 Figure 16: Simulator $\text{SIM}^{\text{OPRF-PPKR}}$ for $\pi^{\text{OPRF-PPKR}}$, part 3.

Game \mathbf{G}_6 : Simulate ciphertexts of honest initializations without passwords. In this game, we change how SIM computes the AE ciphertext c in initializations of honest clients. Now, when SIM receives a message $(\text{ssid}, b, \text{ID}_C, \text{pk}_{\text{Enc}})$ and \mathcal{A} has not yet queried $(\text{pw} \parallel \text{ID}_C, b^{1/r'})$ to the random oracle H_2 , SIM does not compute $c \stackrel{\$}{\leftarrow} \text{AE.Enc}(\rho, (K, \text{sk}_C))$. Instead it uses the equivocability simulator of AE and runs $c \stackrel{\$}{\leftarrow} \text{SIM}_{\text{EQV}}(\lambda + |\text{sk}_C|)$. Note that this can only happen with non-negligible probability if \mathcal{S} is fully corrupt as otherwise \mathcal{A} would have to guess the OPRF key used by SIM_{OPRF} .

Then, SIM records $(\text{AE}, c, \text{pw} \parallel \text{ID}_C, b^{1/r}, K, \text{sk}_C)$. Recall that SIM still receives pw as input and can create these records. We emphasize that the creation of these records is only a temporary change that will be removed later in \mathbf{G}_{14} and any step introduced in this game that relies on these records will be changed later to work without them (cf. \mathbf{G}_{10} and \mathbf{G}_{14}).

We now use the AE records to ensure that recovery still works correctly. To that end, we change SIM such that when it receives a message $(\text{ssid}, b', c, \text{ID}_C)$ in a recovery of an honest ID_C , SIM retrieves the record $(\text{AE}, c, \text{pw}' \parallel \text{ID}_C, b'^{1/r'}, [K], [\text{sk}_C])$. If such a record exists, c was produced SIM_{EQV} , and SIM proceeds using K and sk_C from that record. If no such record exists, then c was computed by \mathcal{Z} , and SIM checks whether ID_C can decrypt it. For this, it again checks the conditions from \mathbf{G}_5 . If they hold, SIM let's ID_C output FAIL and otherwise decrypts c using $F_{\text{ssid}, \text{ID}_C}(\text{pw}' \parallel \text{ID}_C)$ to obtain K and sk_C . Note that the output behavior of ID_C here is the same as in \mathbf{G}_5 . Formally, all these changes that do not rely on the AE records are done in **Ib.6(c)** and **Ib.7(a)**.

To ensure that \mathcal{Z} cannot distinguish \mathbf{G}_6 from \mathbf{G}_5 , SIM additionally has to properly program the random oracle H_2 to account for password guessing. For each query $(\text{pw} \parallel \text{ID}_C, y)$ to H_2 , it retrieves the record $(\text{AE}, [c], \text{pw} \parallel \text{ID}_C, y, [K], [\text{sk}_C])$, then runs $\rho \leftarrow \text{SIM}_{\text{EQV}}(c, (K, \text{sk}_C))$, and sets $H_2(\text{pw}, y) \leftarrow \rho$ (**H2.2(b)**). That means, whenever \mathcal{A} queries H_2 on the input that would yield the key for c in \mathbf{G}_5 , SIM now equivocates c to the matching message and programs H_2 accordingly. If no such record exists, SIM forwards the query to the random oracle H_2 of SIM_{OPRF} as in \mathbf{G}_5 .

We claim that for any \mathcal{Z} that is able to distinguish \mathbf{G}_6 and \mathbf{G}_5 , we can construct an adversary \mathcal{B}_5 that wins the EQV experiment for AE. Since in \mathbf{G}_5 , ρ is chosen by $\mathcal{F}_{\text{OPRF}}$ uniformly at random and c is encrypted under ρ , the distribution of ρ and c in \mathbf{G}_5 is exactly the same as in the real game of the EQV experiment. In \mathbf{G}_6 , c is output by SIM_{EQV} and thus its distribution is exactly the same as in the ideal game of the EQV experiment. The distribution of ρ in \mathbf{G}_6 remains unchanged unless there is a query to H_2 such that SIM equivocates c . In that case ρ is distributed exactly the same as in the ideal world in the EQV experiment as it is computed by SIM_{EQV} . Thus, we have

$$|\Pr[\mathbf{G}_6] - \Pr[\mathbf{G}_5]| \leq \text{Adv}_{\mathcal{B}_5, \text{AE}}^{\text{EQV}}(\lambda).$$

Game \mathbf{G}_7 : Abort upon signature forgery. In this game, we let SIM abort upon observing a forged signature. Concretely, if SIM receives a message (ssid, σ^*) from \mathcal{A} to \mathcal{S} on behalf of ID_C then SIM aborts if all of the following conditions hold (cf. **G.1**):

- There is a record $(\text{FILE}, \text{ID}_C, [\text{pk}_C], *, *, *)$ such that $\text{Sig.Vfy}(\text{pk}_C, (a', \text{ID}_C, \text{ssid}, b', c), \sigma^*) = 1$,
- The signed c was not equivocated using SIM_{EQV} , i.e., SIM never computed $\rho \leftarrow \text{SIM}_{\text{EQV}}(c, (K, \text{sk}_C))$,
- σ^* was not computed by SIM on behalf of an honest ID_C .

Similarly, if SIM receives a message (ssid, σ^*) from \mathcal{A} to $\mathcal{F}_{\text{HSM}}^{\text{OPRF-PPKR}}$ on behalf of S then SIM aborts if all of the above conditions hold (cf. [σ.2](#)). We call this event E_{Sig} . If \mathcal{A} can provoke E_{Sig} to happen, then we can use \mathcal{A} to construct an adversary \mathcal{B}_6 against the sEUF-CMA security of Sig. \mathcal{B}_6 internally runs \mathcal{Z} and plays the role of SIM and \mathcal{F} in the execution of the protocol. \mathcal{B}_6 starts by randomly choosing $i^* \in \{1, \dots, q_{\text{INIT}}\}$, where q_{INIT} is the number of initializations. Let ID_C^* denote the ID_C that executes the i^* -th initialization. If ID_C^* is honest, then \mathcal{B}_6 does not generate a signing key pair $(\text{sk}_C, \text{pk}_C)$ but instead uses the public key pk^* that it gets from its challenger. Since it does not know the key sk^* , it stores \perp instead of sk_C in the AE record in the i^* -th initialization.

If ID_C^* is corrupt, then \mathcal{B}_6 aborts. Further, if \mathcal{A} makes an H_2 query such that SIM would equivocate c^* by computing $\rho \leftarrow \text{SIM}_{\text{EQV}}(c^*, (K, \text{sk}_C))$, where c^* is the AE ciphertext produced in the i^* -th initialization, then \mathcal{B}_6 also aborts. Now, whenever ID_C^* executes a recovery, \mathcal{B}_6 checks two conditions:

- ID_C^* receives a message $(\text{ssid}, b', c^*, \text{ID}_C^*)$, where c^* is the ciphertext it produced in the i^* -th initialization.
- The recovery is successful (in this game, \mathcal{B}_6 still gets the client input pw' and can check $\text{pw} = \text{pw}'$).

If both conditions are satisfied, then \mathcal{B}_6 does not compute $\text{Sig}(\text{Sign}(\text{sk}_C, (a', \text{ID}_C, \text{ssid}, b', c^*)))$ but instead uses its $\text{Sign}(\cdot)$ oracle provided by its challenger to get a signature σ on the message $m = (a', \text{ID}_C, \text{ssid}, b', c^*)$. If any condition does not hold, \mathcal{B}_6 proceeds just like SIM in \mathbf{G}_6 . Now, \mathcal{B}_6 continues the simulation and observes all (ssid, σ) messages that it receives from \mathcal{A} . If there is a message such that E_{Sig} happens, then \mathcal{B}_6 outputs σ .

Note that the view of \mathcal{Z} in the reduction did not change with respect to \mathbf{G}_7 . The public key pk_C and the signatures σ are distributed exactly as in \mathbf{G}_6 . In the i^* -th initialization, \mathcal{B}_6 did not use sk_C , as it simulated $c^* \xleftarrow{\$} \text{SIM}_{\text{EQV}}(\lambda + |\text{sk}_C|)$ without sk_C as in \mathbf{G}_6 . \mathcal{B}_6 also did *not* have to equivocate c^* with $\rho \leftarrow \text{SIM}_{\text{EQV}}(c^*, (K, \text{sk}_C))$ (which it would not be able to do). If the first signature σ^* that caused the event E_{Sig} is valid under pk^* , which happens with probability at least $1/q_{\text{INIT}}$, \mathcal{B}_6 wins. Overall, we get

$$|\Pr[\mathbf{G}_7] - \Pr[\mathbf{G}_6]| \leq q_{\text{INIT}} \text{Adv}_{\mathcal{B}_6, \text{Sig}}^{\text{sEUF-CMA}}(\lambda).$$

Game \mathbf{G}_8 : Extract from malicious initialization. In this game, we change the behaviour of SIM whenever it receives a message (ssid, C) from a corrupt ID_C or the corrupt S queries (ssid, C) to $\mathcal{F}_{\text{HSM}}^{\text{OPRF-PPKR}}$. For this, we first modify \mathcal{F} in the following way. We change the interface $(\text{INITC}, \text{ssid}, \text{pw})$ such that it acts exactly as $\mathcal{F}_{\text{PPKR}}$ if and only if the ID_C that makes the query is corrupt. Further, we add an interface $(\text{COMPLETEINITS}, \text{ssid}, \text{ID}_C, \text{pw}, K)$ that executes the steps [CIS.3](#) and [CIS.4](#) of $\mathcal{F}_{\text{PPKR}}$. Again, the addition of this interface is only a temporary change and it is removed again in \mathbf{G}_{11} .

When receiving a message (ssid, C) from a corrupt party, SIM executes the protocol as in \mathbf{G}_7 and then searches for a record $\langle H_2, \text{pw} \parallel *, *, y \rangle$ such that c decrypts successfully

to some (K, sk_C) under y . Note that as of game \mathbf{G}_4 , there is at most one such y . If no such record exists, it sets $K \leftarrow \perp, \text{pw} \leftarrow \perp$.

If the message (ssid, C) came from a corrupt ID_C , SIM gives the input $(\text{INITC}, \text{ssid}, \text{pw})$ to $\mathcal{F}_{\text{PPKR}}$ on behalf of ID_C and sends $(\text{COMPLETEINITS}, \text{ssid}, \text{ID}_C, \text{pw}, K)$ to $\mathcal{F}_{\text{PPKR}}$. On the other hand, if (ssid, C) came from the corrupt S, SIM queries $(\text{MALICIOUSINIT}, \text{ssid}, \text{ID}_C, \text{pw}, K)$ to $\mathcal{F}_{\text{PPKR}}$, where ID_C is the value the corrupt S sent in the first message of the subsession ssid . Formally, the changes can be read from [IC.2](#) (Step [IC.2](#)(b) and the first part of Step [IC.2](#)(d)) and [IC.7](#)(c) (Steps [IC.7](#)(c)(i) and [IC.7](#)(c)(iii)).

The changes introduced in this game are only syntactical and do not affect any output of SIM. Essentially we only let \mathcal{F} store some records for initializations executed by a corrupt party, but \mathcal{F} never uses these records yet to produce an output for any party. Hence, we have

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

Game \mathbf{G}_9 : Extract from malicious Recoveries. We change the simulator whenever it receives message (ssid, σ) from a corrupt ID_C to S, or from a corrupt S to $\mathcal{F}_{\text{HSM}}^{\text{OPRF-PPKR}}$. As of game \mathbf{G}_7 , we know that if σ verifies under pk_C stored in the file of ID_C , either the adversary previously initialized (on behalf of corrupt ID_C or a corrupt S) that file record or it guessed the password of ID_C . As in the previous game, we let the simulator extract pw from that malicious initialization or password guess and submit $(\text{REC}, \text{ssid}, \text{pw})$ to $\mathcal{F}_{\text{PPKR}}$ on behalf of corrupt ID_C (see [σ.2](#)(a) and [σ.2](#)(b)), resp. $(\text{MALICIOUSREC}, \text{ID}_C^*, \text{pw}')$ for a corrupt S (see [σ.8](#)). Because the changes in the simulation again only affect the state of $\mathcal{F}_{\text{PPKR}}$ which is not yet influencing protocol outputs, the change is only syntactical, and we have

$$\Pr[\mathbf{G}_9] = \Pr[\mathbf{G}_8].$$

Game \mathbf{G}_{10} : Let \mathcal{F} produce the output of the client. In this game, we change the simulator such that the output for honest clients is generated by \mathcal{F} . To this end we have to introduce several changes to \mathcal{F} . We modify the interfaces $(\text{INITC}, \text{ssid}, \text{pw})$ and $(\text{REC}, \text{ssid}, \text{pw})$ to act exactly like $\mathcal{F}_{\text{PPKR}}$, except that they still forward all inputs to SIM. Next, we add the interfaces $(\text{INITS}, \text{ssid}, \text{ID}_C)$, $(\text{RECS}, \text{ssid}, \text{ID}_C)$, $(\text{COMPLETEINITC}, \text{ssid}, b_C)$, and $(\text{COMPLETEREC}, \text{ssid}, b_C)$ and let them act exactly like in $\mathcal{F}_{\text{PPKR}}$. Finally, we introduce the interfaces $(\text{COMPLETEINITS}, \text{ssid}, b_S)$ and $(\text{COMPLETERECS}, \text{ssid}, b_S)$, which both act as in $\mathcal{F}_{\text{PPKR}}$, except that they never give any output to S.

Furthermore, we change SIM to always use the interfaces COMPLETEINITC and COMPLETEREC of \mathcal{F} whenever it wants to produce output of some honest ID_C with b_C set appropriately ([Ib.3](#), [Ib.4](#), [Ib.8](#), [Del.3](#), [Rb.5](#), [Rb.6](#), [Rb.8](#), last step, and [Rb.9](#), last step). Note that these interfaces only produce output if both INITC and INITS , resp. REC and RECS , were queried previously and thus require SIM to give the inputs to \mathcal{F} on behalf of S if S is not honest ([Ib.2](#), [Del.2](#), [Rb.4](#)). Additionally, SIM has to ensure that \mathcal{F} internally creates FILE records that can be accessed during recoveries.

Simulator $\text{SIM}^{\text{OPRF-PPKR}}$, part 4.	
On (ssid, σ) from \mathcal{A} to S on behalf of ID_C : // any ID_C , honest S	
$\sigma.1$	Abort if all of the following conditions hold: (G7) <ul style="list-style-type: none"> • There is either a record $(\text{FILE}, \text{ID}_C, [\text{pk}_C], *, *, *)$ or SIM leaked pk_C to \mathcal{A} such that $\text{Sig.Vfy}(\text{pk}_C, (a', \text{ID}_C, \text{ssid}, b', c), \sigma^*) = 1$, (G7) • The signed c was not equivocated using SIMEQV, i.e., SIM never computed $\rho \leftarrow \text{SIMEQV}(c, (K, \text{sk}_C))$, (G7) • SIM never output σ on behalf of an honest ID_C. (G7)
$\sigma.2$	If ID_C is corrupt, do the following. <ol style="list-style-type: none"> Retrieve $(\text{REC}, *, *, \text{ssid}, *, *, \text{ID}_C, *, [a'], [\text{pk}_C], [b'], [c])$. If a record $(\text{DELREC}, \text{ssid})$ exists, ignore. (G1) Determine pw' as follows: <ul style="list-style-type: none"> • If $\text{Sig.Vfy}(\text{pk}_C, (a', \text{ID}_C, \text{ssid}, b', c), \sigma) = 0$, set $\text{pw}' = \perp$. (G9) • Otherwise, search for a record $(H_2, [\text{pw}'] \ *, *, [y])$ marked CONSISTENT such that $\perp \neq (K, \text{sk}_C) \leftarrow \text{AE.Dec}(y, c)$. If no such record exists, set $\text{pw}' = \perp$. (G9) Give input $(\text{REC}, \text{ssid}, \text{pw}')$ to $\mathcal{F}_{\text{PPKR}}$ on behalf of ID_C. On response $(\text{REC}, \text{ssid}, \text{ID}_C, \text{match})$ from $\mathcal{F}_{\text{PPKR}}$ continue below. (G9) Execute HSM code on input (ssid, σ) up to determining <i>out</i>. (G1) If <i>out</i> = FAIL, send message $(\text{COMPLETEREC}, \text{ssid}, 0)$ to $\mathcal{F}_{\text{PPKR}}$ (G12), and otherwise $(\text{COMPLETEREC}, \text{ssid}, 1)$ (G10)
$\sigma.3$	If ID_C is honest, execute HSM code on input (ssid, σ) and send message $(\text{COMPLETEREC}, \text{ssid}, 1)$ to $\mathcal{F}_{\text{PPKR}}$. (G12) // If ID_C is honest, we only simulate σ if <i>match</i> = 1 and do not need to check again here. Due to ID_C -authentication the adversary also cannot inject any messages
On (ssid, σ) from \mathcal{A} to $\mathcal{F}_{\text{HSM}}^{\text{OPRF-PPKR}}$ on behalf of corrupt S : // any ID_C	
$\sigma.4$	Abort if all of the following conditions hold: (G7) <ol style="list-style-type: none"> There is either a record $(\text{FILE}, \text{ID}_C, [\text{pk}_C], *, *, *)$ or SIM leaked pk_C to \mathcal{A} such that $\text{Sig.Vfy}(\text{pk}_C, (a', \text{ID}_C, \text{ssid}, b', c), \sigma^*) = 1$, (G7) The signed c was not equivocated using SIMEQV, i.e., SIM never computed $\rho \leftarrow \text{SIMEQV}(c, (K, \text{sk}_C))$, (G7) SIM never output σ on behalf of an honest ID_C. (G7)
$\sigma.5$	Retrieve $(\text{REC}, [\text{ID}_C], *, \text{ssid}, *, [\text{match}], [\text{ID}_C^*], *, [a'], [\text{pk}_C], [b'], [c])$. If a record $(\text{DELREC}, \text{ssid})$ exists, ignore. Execute HSM code on input (ssid, σ) up to determining <i>out</i> . (G1)
$\sigma.6$	If <i>match</i> = \perp , give input $(\text{REC}, \text{ssid}, \text{ID}_C^*, \perp, \perp, \perp)$ to $\mathcal{F}_{\text{PPKR}}$. On response $(\text{REC}, \text{ssid}, \text{ID}_C^*, \text{match})$ from $\mathcal{F}_{\text{PPKR}}$ continue below. (G10) // <i>match</i> = \perp indicates that REC input was not given to $\mathcal{F}_{\text{PPKR}}$ yet, which is necessary to decrease counter
$\sigma.7$	If <i>out</i> = FAIL, send $(\text{RECRES}, \text{ssid}, \text{FAIL})$ to S . (G1)
$\sigma.8$	If <i>out</i> = SUCC, search for a record $(H_2, [\text{pw}'] \ *, *, [y])$ marked CONSISTENT such that $\perp \neq (K, \text{sk}_C) \leftarrow \text{AE.Dec}(y, c)$, otherwise set $\text{pw}' \leftarrow \perp$. Then do: (G9) // need to reset <i>ctr</i> in $\mathcal{F}_{\text{PPKR}}$ <ol style="list-style-type: none"> both messages in this recovery come from corrupt S If $\text{ID}_C = \perp$, give input $(\text{MALICIOUSREC}, \text{ID}_C^*, \text{pw}')$ to $\mathcal{F}_{\text{PPKR}}$. (G9) Independent of the response from $\mathcal{F}_{\text{PPKR}}$, send $(\text{RECRES}, \text{ssid}, \text{SUCC})$ to S. (G1) // We give the same password as in the MALICIOUSINIT query, which ensures that the counter is reset a' message came from ID_C, σ from corrupt S If $\text{ID}_C \neq \perp$: <ul style="list-style-type: none"> • If ID_C is corrupt, input $(\text{REC}, \text{ssid}, \text{pw}')$ to $\mathcal{F}_{\text{PPKR}}$ (G9) • Give input $(\text{COMPLETEREC}, \text{ssid}, 1)$ to $\mathcal{F}_{\text{PPKR}}$ and send $(\text{RECRES}, \text{ssid}, \text{SUCC})$ to S. (G10)
On GETPK from anyone to $\mathcal{F}_{\text{HSM}}^{\text{OPRF-PPKR}}$:	
pk.1	Retrieve $(PK, [\text{pk}], *)$ and return pk . (G1)
On LEAKFILE from \mathcal{A} to $\mathcal{F}_{\text{HSM}}^{\text{OPRF-PPKR}}$ on behalf of S :	
LF.1	Send (LEAKFILE) to $\mathcal{F}_{\text{PPKR}}$ to obtain ID_C . (G10)
LF.2	$L \leftarrow \emptyset$. For each $(\text{ID}_C, \text{ctr})$ in ID_C : (G10) <ol style="list-style-type: none"> Retrieve $(\text{FILE}, \text{ID}_C, [\text{kid}], [\text{pk}_C], [\text{sk}_C], [c])$. (G10) Retrieve $(F, S, \text{ID}_C \ \text{kid}, [\text{kOPRF}], *)$ in OPRF simulator. (G10) If no record $(\text{LEAKED}, \text{ID}_C, *, *, *, *)$ exists, record $(\text{LEAKED}, \text{ID}_C, \text{kid}, \text{sk}_C, c, 1)$. Otherwise, record $(\text{LEAKED}, \text{ID}_C, \text{kid}, \text{sk}_C, c, i + 1)$ where $i \in \mathbb{N}$ is the biggest number such that a record $(\text{LEAKED}, \text{ID}_C, *, *, *, i)$ exists. Append $(\text{ID}_C, \text{pk}_C, c, \text{kOPRF}, \text{ctr})$ to L. (G10)
LF.3	Output L to \mathcal{A} . (G10)
On FULLYCORRUPT from \mathcal{A} to $\mathcal{F}_{\text{HSM}}^{\text{OPRF-PPKR}}$ on behalf of S : (G10)	
FC.1	Send $(\text{FULLYCORRUPT}, S)$ to $\mathcal{F}_{\text{PPKR}}$, which outputs ID_C . (G10)
FC.2	Execute Step 2 of LEAKFILE using ID_C . (G10)
FC.3	Retrieve $(PK, *, [\text{sk}])$ and output sk . (G10)

Figure 17: Simulator $\text{SIM}^{\text{OPRF-PPKR}}$ for $\pi^{\text{OPRF-PPKR}}$, part 4.

For this reason, whenever \mathcal{A} sends some (ssid, C) from an honest ID_C , SIM sends $(\text{COMPLETEINITS}, \text{ssid}, 1)$ to \mathcal{F} (**IC.3**, **IC.6**), once again first giving the INITS input to \mathcal{F} if S is not honest to ensure that the COMPLETEINITS query proceeds (**IC.5**). Lastly, SIM acts exactly like $\text{SIM}^{\text{OPRF-PPKR}}$ on

the queries LEAKFILE and $(\text{FULLYCORRUPT}, S)$ by \mathcal{A} (**LF.1-LF.3**, **FC.1-FC.3**).

A difficulty for SIM resulting from these changes is that for honest clients, K is now chosen by \mathcal{F} instead of SIM and is unknown to SIM . In particular, this affects the changes introduced in **G6** as SIM cannot use the records (AE, \dots) anymore in **G10** to store and obtain K . Instead, whenever SIM equivocates some ciphertext c during a query to H_2 , it has to extract K from \mathcal{F} . Note that \mathcal{A} can only make a query that requires SIM to equivocate if the corresponding file containing c was leaked or if S is fully corrupt as otherwise \mathcal{A} does not know the OPRF key used by SIM^{OPRF} . This means that in \mathcal{F} there is now a LEAKED record that SIM can address in the OFFLINEATTACK interface (**H2.2(a)** and the first part of **H2.2(b)**). This allows SIM to obtain K by using the password from the query by \mathcal{A} and properly equivocate c . SIM proceeds analogously if some honest ID_C receives a message $(\text{ssid}, b, \text{ID}_C^*, \text{pk}_{\text{Enc}})$ and \mathcal{A} has already queried $(\text{pw} \| \text{ID}_C, b^{1/r})$ to H_2 , except that it first has to send $(\text{COMPLETEINITS}, \text{ssid}, 1)$ to \mathcal{F} in order to let \mathcal{F} create the LEAKED record (**Ib.6(a)-Ib.6(c)**). Next, we change the handling of messages $(\text{ssid}, b', c, \text{ID}_C^*)$ in SIM . Here we need to appropriately choose the values pw^*, K^* , and i for the REC query if S is fully corrupt. For this, SIM again relies on similar checks as introduced in **G5**. If there is a record $F_{\text{ssid}, S, *}(pw' \| \text{ID}_C^*)$ that successfully decrypts c , SIM sets $\text{pw}^* \leftarrow \text{pw}'$ and $K^* \leftarrow K$, where K is the key obtained from decrypting c (**Rb.3(a)**). If no such record exists but S used the same OPRF key as in the initialization that produced c , SIM sets $\text{pw}^* \leftarrow \perp, K \leftarrow \perp$ and chooses i such that it indicates the LEAKED record that contains c (**Rb.3(b)**). Otherwise, it sets $\text{pw}^* \leftarrow 0$ and $K^* \leftarrow \text{FAIL}$, which ensures that ID_C outputs FAIL independent of the password (**Rb.3(c)**). Further, SIM uses the sk_C obtained from decrypting c , resp. the LEAKED record, to compute the signature σ .

Finally, we need to ensure that the counter for ID_C is updated in \mathcal{F} . Hence, whenever SIM gives the output $(\text{RECRES}, \text{ssid}, \text{SUCC})$ to S , it additionally sends $(\text{COMPLETEREC}, \text{ssid}, 1)$ to \mathcal{F} (last part of **$\sigma.2$ (c)**, **$\sigma.3$** , second step of **$\sigma.8$ (b)**).

Let us now argue why the changes introduced in **G10** are indistinguishable from **G9**. In **G9**, K was chosen uniformly at random by SIM , and K is chosen uniformly at random by \mathcal{F} in **G10**. Therefore the distribution of K obviously does not change. Moreover, the other significant change introduced in **G10** to the initialization phase is removing K from the records (AE, \dots) . However, as argued above, SIM is always able to extract K from \mathcal{F} whenever necessary. Thus, all outputs for any ID_C in the initialization phase remain unchanged.

In the recovery phase, the removal of K from the records (AE, \dots) has no effect, since in **G9** in the recovery phase the record was only used to obtain K when SIM needed to output it to ID_C in a successful recovery, where ID_C

received an equivocal c . However, this is not necessary in \mathbf{G}_{10} , as K is output by \mathcal{F} to ID_C .

In both phases, the outputs of S remain unchanged as S still gets its output from SIM and the COMPLETEINIT S and COMPLETEREC S interfaces introduced here do not produce output to S . Overall, we therefore have

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

Game \mathbf{G}_{11} : Let \mathcal{F} produce the output of honest servers in initialization. In this game, we change SIM such that it produces the output for honest servers by calling the appropriate interfaces of \mathcal{F} . We also change the COMPLETEINIT S interface of \mathcal{F} to provide output to S .

Concretely, in [IC.2\(a\)](#), we still let SIM compute $(\text{ssid}', \text{pk}_C, c) \leftarrow \text{Dec}(\text{sk}_{\text{Enc}}, C)$ and check $\text{ssid} = \text{ssid}'$. However, if the check fails, SIM now sends $(\text{COMPLETEINIT}$ S, $\text{ssid}, 0)$ to \mathcal{F} to let the server output FAIL.

Further, we remove the $(\text{COMPLETEINIT}$ S, $\text{ssid}, \text{ID}_C, \text{pw}, K)$ interface from \mathcal{F} that we introduced in \mathbf{G}_8 . In \mathbf{G}_8 the interface was used by SIM to let \mathcal{F} create FILE records. However, as of \mathbf{G}_{10} , \mathcal{F} stores all the necessary records because of the added INITC and INIT S interfaces. Therefore, SIM can execute [IC.2\(d\)](#) exactly as SIM_{OPRF} , i.e., SIM sends $(\text{COMPLETEINIT}$ S, $\text{ssid}, 1)$ to \mathcal{F} instead of $(\text{COMPLETEINIT}$ S, $\text{ssid}, \text{ID}_C, \text{pw}, K)$.

We argue that the above changes do not alter the view of \mathcal{Z} : First, whenever in \mathbf{G}_{10} SIM would have given output $(\text{INITRES}, \text{ssid}, \text{ID}_C, \text{FAIL})$ directly to the honest S , we now use the $(\text{COMPLETEINIT}$ S, $\text{ssid}, 0)$ message to \mathcal{F} . The effect is the same and \mathcal{F} outputs $(\text{INITRES}, \text{ssid}, \text{ID}_C, \text{FAIL})$ to S . The same holds for the $(\text{COMPLETEINIT}$ S, $\text{ssid}, 1)$ messages and SUCC output. In addition, COMPLETEINIT S makes \mathcal{F} record FILE records that use $\text{ID}_C, \text{pw}, K$ from \mathcal{F} 's INIT records (instead of the values provided to COMPLETEINIT S). But these records were already created in \mathbf{G}_{10} , and therefore $\text{ID}_C, \text{pw}, K$ are the same in both games. We get

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

Game \mathbf{G}_{12} : Let \mathcal{F} produce the output of honest servers in recovery. In this game, we change SIM such that it produces the output for honest servers in recovery by calling the appropriate interfaces of \mathcal{F} . In \mathbf{G}_{11} , SIM only used the COMPLETEREC S interface when it wanted to produce the output SUCC for S . Hence, here we change SIM to also send the message $(\text{COMPLETEREC}$ S, $\text{ssid}, 0)$ to \mathcal{F} when it wants to produce the output FAIL for S ([Ra.2](#) and [σ.2\(c\)](#)). By setting $b_S = 0$, S always gets output FAIL from \mathcal{F} unless \mathcal{F} deletes the file for ID_C in this recovery, however in that case SIM ignores the message (ssid, σ) (see [σ.2\(a\)](#), [σ.5](#)). Therefore, whenever SIM sends $(\text{COMPLETEREC}$ S, $\text{ssid}, 0)$ to \mathcal{F} , this produces the output $(\text{RECRESSID}, \text{FAIL})$ to S .

Further, it is easy to verify that any $(\text{COMPLETEREC}$ S, $\text{ssid}, 1)$ query from SIM , where S is honest, produces the output $(\text{RECRES}, \text{ssid}, \text{SUCC})$ to S ([σ.2\(c\)](#), [σ.3](#)). In [σ.2](#), SIM is either able to extract the correct password due to \mathbf{G}_9 , which leads to COMPLETEREC S outputting SUCC , or if it cannot extract a password the file must have been created in a

malicious initialization and COMPLETEREC S then outputs SUCC as well. If ID_C is honest (cf. [σ.3](#)), the signature σ received in this interface is only valid if it was output by SIM , which only happens if ID_C used the correct password. Thus, COMPLETEREC S outputs SUCC as well and we have

$$\Pr[\mathbf{G}_{12}] = \Pr[\mathbf{G}_{11}].$$

Game \mathbf{G}_{13} : Simulate C during Init. In this game, we change how SIM computes the message (ssid, C) for honest ID_C as long as S is not fully corrupt. Whenever SIM receives a message $(\text{ssid}, b, \text{ID}_C, \text{pk}_{\text{Enc}})$ to an honest ID_C , the simulator does not compute $C \stackrel{\$}{\leftarrow} \text{PKE.Enc}(\text{pk}_{\text{Enc}}, (\text{ssid}, \text{pk}_C, c))$ but now computes $C \stackrel{\$}{\leftarrow} \text{PKE.Enc}(\text{pk}_{\text{Enc}}, \perp)$ (see [Ib.7\(b\)](#)). We also change SIM such that it does not decrypt the ciphertext C produced for an honest ID_C anymore. In [IC.3](#), when C was computed by SIM for an honest ID_C , then SIM retrieves the record containing pk_C and c that it stored when computing C . Similarly, in [IC.6](#), when C was computed by SIM , then SIM just retrieves the stored values pk_C and c . Additionally, we need to check if the corrupt S replays some C that was computed as $C \stackrel{\$}{\leftarrow} \text{PKE.Enc}(\text{pk}_{\text{Enc}}, \perp)$ and if so, retrieve the corresponding values from the record instead of decrypting C (cf. [IC.7\(a\)](#) and [IC.7\(b\)](#)). The only change in the view of \mathcal{Z} is that the distribution of the ciphertexts C . If \mathcal{Z} can detect this difference, we can construct an adversary \mathcal{B}_7 against the IND-CCA security of PKE as follows:

Let $q_{\text{INIT}} \in \mathbb{N}$ be the number of initializations. We construct a sequence of games $\mathbf{G}_{12}^{(0)}, \dots, \mathbf{G}_{12}^{(q_{\text{INIT}})}$, where in $\mathbf{G}_{12}^{(i)}$ the first i ciphertexts are computed as encryptions of \perp and the remaining ciphertexts are encrypted as in \mathbf{G}_{12} . We have $\mathbf{G}_{12} = \mathbf{G}_{12}^{(0)}$ and $\mathbf{G}_{13} = \mathbf{G}_{12}^{(q_{\text{INIT}})}$. Because \mathcal{Z} can distinguish \mathbf{G}_{12} from \mathbf{G}_{13} there must be an index $i^* \in [q_{\text{INIT}}]$ such that \mathcal{Z} has a non-negligible advantage in distinguishing $\mathbf{G}_{12}^{(i^*-1)}$ and $\mathbf{G}_{12}^{(i^*)}$. its challenger. Now, let ssid^* denote the ssid of the i^* -th initialization and ID_C^* denote the ID_C that executes that initialization. In that initialization the reduction \mathcal{B}_7 does not compute $(\text{sk}_{\text{Enc}}, \text{pk}_{\text{Enc}})$ itself in [Ia.6](#) but uses the pk^* provided by its challenger. Then, in [Ib.7\(b\)](#), \mathcal{B}_7 gives $m_0 := (\text{ssid}^*, \text{pk}_C, c)$ and $m_1 := \perp$ to the challenger and uses the returned C^* as ciphertext for the i^* -th initialization. In any subsequent recovery by ID_C^* in [IC.7\(b\)](#), when \mathcal{B}_7 receives a message (ssid, C) from \mathcal{A} to $\mathcal{F}_{\text{HSM}}^{\text{encPw}}$ on behalf of a corrupted server such that there is no record $(\text{INIT}, *, *, \text{ssid}^*, *, *, *, *, *, *, C)$ then \mathcal{B}_7 uses the challenger's decryption oracle to decrypt C . Similarly, in [IC.2\(a\)](#), if there is no record $(\text{INIT}, *, *, \text{ssid}^*, *, *, *, *, *, *, C)$ then SIM again uses the challenger's decryption oracle to decrypt C .

Then we have that if C^* encrypts m_0 the game is distributed as in $\mathbf{G}_{12}^{(i^*-1)}$ and if it encrypts m_1 then the game is distributed as in $\mathbf{G}_{12}^{(i^*)}$. We get

$$|\Pr[\mathbf{G}_{13}] - \Pr[\mathbf{G}_{12}]| \leq q_{\text{INIT}} \text{Adv}_{\text{PKE}, \mathcal{B}_7}^{\text{IND-CCA}}(\lambda).$$

Game \mathbf{G}_{14} : Remove password forwarding from \mathcal{F} . In this game, we do not give SIM any private input of the parties anymore. Up to \mathbf{G}_{13} , SIM still stored the records $(\text{AE}, c, \text{pw} \parallel$

$\text{ID}_C, b^{1/r}, \text{sk}_C$ and used them for three purposes: (1) to obtain the sk_C that is supposed to be contained in an equivocal c , (2) to check whether a client used the correct password in a recovery, and (3) to determine critical queries $(\text{pw} \parallel \text{ID}_C, b^{1/r})$ to H_2 that require SIM to equivocate c . As we now finally remove the AE records, we have to simulate these steps in another way.

Issue (1) can be solved trivially, and we instead use the FILE records, which also store c and sk_C (see Rb.7). For issue (2), we now rely on the bit *match* that SIM receives from \mathcal{F} and indicates whether the password used in the recovery is correct (see Rb.8 and Rb.9).

To solve issue (3), we can instead rely on the records $\langle H_2, \text{pw} \parallel \text{ID}_C, u, y \rangle$ and the CONSISTENT marking. An H_2 record if

marked CONSISTENT if and only if $u = H_1(\text{pw} \parallel \text{ID}_C)^k$ for some OPRF key k , where k can be either created by SIM_{OPRF} (cf. H2.2), i.e., when S is not fully corrupt, or adversarially chosen (cf. Ib.6(b)), i.e., when S is fully corrupt. Then, whenever SIM would query the OFFLINEATTACK interface in \mathbf{G}_{13} with the password pw obtained from the AE record, in this game we instead query the interface with all passwords pw from CONSISTENT H_2 records (Ib.6(b)).

With this change, we finally reach the point where $\text{SIM} = \text{SIM}^{\text{OPRF-PPKR}}$ and $\mathcal{F} = \tilde{\mathcal{F}}_{\text{PPKR}}$. We clearly added all interfaces of $\tilde{\mathcal{F}}_{\text{PPKR}}$ to \mathcal{F} . One can also verify that SIM is indeed $\text{SIM}^{\text{OPRF-PPKR}}$.