

# Single-Message Credential-Hiding Login

Kevin Lewi<sup>\*</sup>, Payman Mohassel<sup>†</sup>, and Arnab Roy<sup>‡</sup>

<sup>\*</sup>Novi Research

<sup>†</sup>Facebook

<sup>‡</sup>Fujitsu Laboratories of America

## Abstract

The typical login protocol for authenticating a user to a web service involves the client sending a password over a TLS-secured channel to the service, occasionally deployed with the password being prehashed. This widely-deployed paradigm, while simple in nature, is prone to both inadvertent logging and eavesdropping attacks, and has repeatedly led to the exposure of passwords in plaintext.

Partly to address this problem, symmetric and asymmetric PAKE protocols were developed to ensure that the messages exchanged during an authentication protocol reveal nothing about the passwords. However, these protocols inherently require at least two messages to be sent out: one from each party. This limitation hinders wider adoption, as the most common login flow consists of a single message from client to the login server. The ideal solution would retain the password privacy properties of asymmetric PAKEs while allowing the protocol to be a drop-in replacement into legacy password-over-TLS deployments.

With these requirements in mind, we introduce the notion of *credential-hiding login*, which enables a client to authenticate itself by sending a *single message* to the server, while ensuring the correct verification of credentials and maintaining credential privacy in the same strong sense as guaranteed by asymmetric PAKEs. We initiate a formal study of this primitive in the Universal Composability framework, design and implement a practical password-based protocol using identity-based encryption, and report on its performance. We also construct a variant of credential-hiding login for fuzzy secrets (e.g. biometrics), proven secure based on the Learning With Errors (LWE) assumption.

## 1 Introduction

In a traditional password-based account login protocol over a secure TLS channel, a client submits their username and password combination to a login server, which checks the password against a hashed record of the password established upon account registration. The database containing these hashed records is typically secured using a mixture of techniques including the use of a slow password hashing algorithm along with a salt [PM99, MKR17, Per09, BDK16], and encryption on top of the hashed records under keys stored in a separate database [Muf15, ECS<sup>+</sup>15, BLMR13, AMMR18]. These methods were designed primarily to protect passwords in the event of a server compromise.

Nevertheless, there have been a recurring series of incidents from major technology companies that have self-reportedly stored passwords in plaintext, including GitHub [Git18] and Twitter [Twi18] in 2018, and Facebook [Fac19], Google [Goo19], Robinhood [Rob19], and Coinbase [Coi19]

in 2019. In each of these incidents, user passwords underwent hashing server side before being stored, and yet were still persisted in plaintext in some form, frequently through inadvertent logging during login time.

To avoid the logging of plaintext passwords when users log in, servers can employ client-side cryptographic solutions. A simple approach, which we will refer to as “Hash-then-Encrypt” (HtE), requires that clients first hash their passwords, and then encrypt the resulting hash under a single public encryption key for which the server has access to the corresponding decryption key, before sending the resulting payload to the server. The central issue with this approach, which we refer to as “HtE-over-TLS”, is that confidentiality of any logged transcript is lost if the login server’s decryption key (which needs to constantly persist on the login server) is ever exposed since the same decryption key exposes all hashed passwords to an adversary who can run an offline dictionary attack on any previously logged transcript.<sup>1</sup> These privacy implications are even more pronounced when we consider legitimate but unsuccessful login attempts wherein registered users mistakenly use their credential from one service to login to another, running the risk of having their hashed password decrypted and logged by the wrong server.

**Password authenticated key exchange.** A promising approach to addressing the confidentiality of passwords during login involves using an asymmetric password authenticated key exchange (aPAKE) protocol. In an aPAKE, a client who holds *only* a password (as opposed to an authentic public key from the server) can establish a secure channel with a server who holds a one-way function of the password. When applied to the account login scenario, the client and server can exchange messages which allow the server to learn whether or not a client has the correct login password without revealing the contents of the password attempt to the server. Using an aPAKE for password-based login is essentially the state-of-the-art solution for mitigating the security issues with persisting the password in a database and handling the credentials during user login.

**Crossing the application and network layer boundary.** One approach for using aPAKEs for user login, which we refer to as the “aPAKE-into-TLS” approach, involves using aPAKE to bootstrap the TLS session as described in [SKFB19]. Taking the aPAKE-into-TLS approach for user login means that we can benefit from the improved security for passwords without introducing any extra rounds of communication between client and server. However, the main disadvantage is that it requires the application server to tightly integrate its password handling logic with its TLS session setup and management logic. Typically, the servers which handle TLS channel setup and maintenance sit at the edge of the service’s network boundary, and have been heavily optimized to reduce latency as much as possible. In other words, this kind of integration requires a merging of functionalities between the network layer and the application layer, which is not always possible or desirable.

In order to maintain the separation between network logic and application logic, we can also consider the “aPAKE-over-TLS” approach, in which the TLS session is first established, *and then* the aPAKE exchange is executed between client and server, purely within the application layer. However, we note that when comparing this solution to the HtE-over-TLS approach, the aPAKE protocol itself introduces an extra round of interaction between the client and server, whereas

---

<sup>1</sup>This kind of attack holds for any efficiently computable deterministic function applied to the user password before it is sent to the server.

HtE-over-TLS only requires the client to send a *single message* to the server based on some public parameters.

**Minimizing round complexity for login.** The introduction of the extra message incurred by an aPAKE-over-TLS approach, while acceptable in some scenarios, introduces an overall latency and reliability hit that is exacerbated by low-connectivity environments, and from a purely infrastructural perspective, can pose a challenge when attempting to integrate with systems that are built around the assumption that password-based login can conclude after a single message from the client. In general, the extra round trip required for migration from a HtE-over-TLS login form to an aPAKE-based login form is an undesirable side effect of the security benefits of using an aPAKE.

While an extra round trip seems like a fair price to pay for the added security benefits provided by aPAKEs compared with using HtE (both over TLS), it is often the case, especially in countries with slow networks, that this extra round trip results in seconds of delay during the login protocol. Based on a study on mobile performance optimization [Eve13], just two seconds of latency increase is enough to double the abandonment rate of users, and every additional 100 milliseconds of delay can cause a noticeable impact on revenue for large companies with user populations in low-connectivity regions.

A natural question we ask is if the additional round complexity of PAKEs is truly necessary in order to capture the desired security properties for account login. Typically, in a user login setting between a client and server on the web, a secure channel has already been established through the TLS certificate of the website and the public key infrastructure that it relies on. This means that the protections to address adversaries who aim to impersonate the server to the client do not need to be incorporated into the login protocol.

Given the distinctions between the login-within-TLS setting and PAKE, we introduce and formalize the notion of credential-hiding login, in which our aim is to formally define and capture a primitive which can act as a drop-in improvement in credential privacy for the HtE-over-TLS approach that most companies employ today, without resorting to altering infrastructure to support multiple rounds and the increase to latency that it causes.

## 1.1 Revisiting the Login Problem

Note that in the single-message setting, there are inherent limitations to the amount of privacy that can be guaranteed for the client’s login message to the server. The first issue is that the login server is able to learn information about the credential encoded in this message by simply testing it against any existing registered password files in its database. This is inherent to the functionality of the login scheme, and so this leakage cannot be avoided. The second issue is that the login server can essentially execute a brute-force dictionary attack by running the registration procedure to obtain password files for its password guesses, allowing it to again learn information about the password encoded in the message, even if the password has not been registered previously. Given this attack, it seems that we cannot do any better than HtE in terms of providing confidentiality (since HtE already forces an adversarial server to perform an offline dictionary attack to learn information about the encoded password).

In practice, however, there are usually significant distinctions between the flows which handle user registration and the flows which allow users to log in. For one, the login attack surface

is usually much larger, in terms of frequency (registrations often happen at a frequency orders of magnitude lower than login) and the overall complexity of different mechanisms necessary for login. Furthermore, it is usually the case that the servers that handle registration can be separated from the servers that handle login and their secrets be protected using hardware solutions due to the less stringent latency requirements.

**Separating registration from login.** In order to provide a more meaningful and applicable approach to defining privacy for login credentials, we capture this intuition formally by allowing the registration procedure to accept secret parameters that the login procedure is unable to access. This means that the registration server can embed secrets into the password file which the login server cannot generate on its own, thereby avoiding the impossibility result stated above. Of course, we still require that the login server be able to determine whether or not a password attempt should succeed *without* access to these secret parameters, and without any interaction with the registration server.

**Going beyond passwords.** In addition, we apply this framework to handle authentication schemes that are more complicated than checking if a login password is equal to the password registered to an account. Although many prior works have investigated the feasibility of biometric authentication (in which a distance metric is computed between two vectors to determine whether or not they are “close” enough), these works present schemes which do not satisfy the security properties we will lay out for credential-hiding login.

As a result, we believe that if these biometric authentication schemes were deployed in practice, they would in fact provide less security (in terms of protecting the credential itself) than a traditional password-based login scheme, while arguably being an even more important and sensitive type of credential to protect than user passwords.

Our goal with credential-hiding login is to ensure that we can design biometric (and other types of credentials) authentication schemes which piggyback on the benefits of decades of research work that have led to the state-of-the-art best practices for protecting passwords, instead of rediscovering the same pitfalls of insecure credential treatment from the password security literature.

Our central focus in this work is to study a primitive which represents the highest level of security we can achieve for login while also taking into account low-latency practical settings in which multiple rounds of interaction are a nonstarter. To do so, we take into account the distinction made between the registration server and the login server, and the separation between the server validating credentials versus the one performing a key exchange (unlike the aPAKE-into-TLS approach), and define a new primitive, called credential-hiding login (CHL), which realizes the security and efficiency properties we look for in an account login protocol. In Table 1, we illustrate the benefits of using CHL-over-TLS compared with the other approaches discussed previously in this section.

## 1.2 Credential-Hiding Login Schemes

A single-message credential-hiding login scheme consists of a setup, registration, client login, and server login procedure. The setup algorithm is run to produce public parameters that are issued to clients, typically after already establishing a secure channel between client and server. The registration algorithm is executed between the client and a registration server, and in the simplest case, involves the client sending their credential to the registration server in order for the server to

Login Protocol	Min # of Messages	OSI Separation	Security
HtE-over-TLS	3	✓	Weak
aPAKE-into-TLS	3	✗	Strong
aPAKE-over-TLS	4	✓	Strong
CHL-over-TLS	3	✓	Strong

Table 1: A comparison of the various approaches discussed for handling password-based login. Note that OSI Separation refers to whether or not the protocol respects the separation of application logic from network logic according to the Open Systems Interconnection Model. Our work is described by the last row of the table.

compute the password file. Then, when a client is asked to log in to their account, they run the client login algorithm to produce a credential-embedded message that they send to a login server (credential verifier). The login server then executes the server login algorithm to verify this message and determine whether or not the login attempt is successful.

We formalize the notion of a single-message credential-hiding login scheme in Section 2, along with a security definition under the Universal Composability (UC) framework. Informally, we highlight the two main security properties we capture in our UC definition of security:

- **Credential privacy**, which ensures that a client’s message representing the credentials are independent of the credentials, so long as the adversary does not have access to the client’s password file or if the client has not registered this particular credential with the registration server.
- **Credential verifier security**, which captures the notion that an adversary who has access to the server-side storage of the user’s credentials established during registration would still need to execute brute-force dictionary attacks to recover the underlying credentials, even if all secrets in the scheme have been leaked.

We note that these properties alone are not enough to satisfy the formal UC definition we present in Section 2, but merely serve as a more intuitive framing of the requirements that a credential-hiding login scheme aims to satisfy.

Satisfying credential privacy implies that the client login algorithm produces messages which do not encode information about the password from the login server’s perspective. This stands in contrast to the HtE approach, where the login server uses a master decryption key to effectively remove all entropy from the client login’s message. A credential-hiding login scheme ensures that the login server is unable to remove this entropy without access to the client’s password file.

Our constructions leverage hash functions that we model as random oracles to satisfy credential verifier security. Although not ideal, such a non-black box assumption is required to detect offline guesses from the adversary and has been traditionally used in UC aPAKEs [GMR06, JKX18, JR18]. We can instantiate these hash functions with computationally-expensive and memory-hard hash functions such as scrypt [PJ16], Argon2 [BDK16] or balloon hashing [BCS16] in order to increase the difficulty of executing an offline dictionary attack on credentials which come from a low-entropy distribution.

**Defining CHL under universal composability.** We model our UC definition for credential-hiding login along similar lines as aPAKEs, but with several significant distinctions. For one, aPAKEs aim to establish a shared secret, whereas CHL simply aims for login success or failure. Also for aPAKEs, a client only needs to assume that the public parameters come from a trusted source, whereas in CHL, the public parameters come from the registration server which also holds a secret corresponding to this parameter. While we do not allow the registration server to generate these parameters maliciously, we do allow the adversary the choice to obtain the secret upfront. This distinction in our setting seems intrinsic to our design, as we elaborate in more depth in Section 2.3. Furthermore, we do not model pre-computation attacks for CHL, as is done by strong aPAKEs such as OPAQUE [JKX18]. And finally, our definition of CHLs accommodates more general credential verification predicates, rather than just equality checks for passwords.

To reiterate, these properties are not new in the password-based authentication literature—indeed, many works have defined similar notions for authentication protocols which capture a subset or superset of the properties we study here. However, our aim is to focus specifically on the most important properties which are desirable in a *login protocol* (as opposed to key exchange) and provide a framework that enables future works to explore extensions to more general credentials beyond passwords and equality checks.

**Authenticated login and composition with TLS.** In Section 2.4, we define a simpler and more general abstraction of the CHL ideal functionality, which we refer to as “authenticated login”. The authenticated login definition abstracts away passwords and credentials analogously to standard key exchange versus password-based key exchange. We show that the UC notion of secure channels can also be obtained in the authenticated login UC hybrid model.

However, the typical deployment of authenticated login in practice that we envision within TLS *does not require such a bootstrapping to secure channels*. We instead focus on leveraging authenticated login to establish the authenticity of the user and the continuation of the subsequent interactions based on the *already established* server-authenticated TLS channel. Recall that this corresponds to the CHL-over-TLS description given earlier. Although this suffices for the primary applications that we target (namely, a login form over the web), we note that this model provides weaker guarantees than a standard secure channel, as it does impose on the user application to not send any sensitive content which assumes mutual trust before the authenticated login flow is completed.

### 1.3 Our Contributions

In this work, we initiate the study of credential-hiding login with a concrete formulation of security in the UC framework, formalized in Section 2 along with the definition of authenticated login. Additionally, we present:

- A construction  $\Pi_{\text{pwd}}$  from anonymous identity-based encryption that handles password credentials,
- An implementation and concrete instantiation of  $\Pi_{\text{pwd}}$  along with benchmarks, and
- A generalized scheme  $\Pi_{\text{fz}}$  which achieves credential-hiding login for a set of distance functions including Hamming distance based on the Learning With Errors (LWE) assumption.

**Credential-hiding login for passwords.** For  $\Pi_{\text{pwd}}$ , the password verifier (i.e. password file) for a password  $y$  consists of two components: an IBE secret key for the identity  $y$ , along with a one-way function of a random oracle evaluation of  $y$ . The key intuition here is that we use the password verifiers to hold IBE secret keys, and the client login algorithm produces IBE ciphertexts where the identity is set to the password  $y$  and the payload being encrypted is a random oracle output of  $y$ . Then, when the server wants to verify the credential, it can retrieve the password verifier containing the IBE secret key and attempt to decrypt the ciphertext. Successful decryption occurs only if the password tied to the password verifier matches the password embedded in the ciphertext, as desired. Credential privacy follows from the anonymity and IND-CCA security of the IBE scheme; in particular, we need the property that the ciphertexts completely hide both the identity and the payload when the corresponding secret key which can decrypt the ciphertext is not present. This allows us to prove that credential login transcripts without a corresponding password file truly hide the credential from the login server. Furthermore, the use of a random oracle output for the scheme ensures that even if the registration key is leaked along with the entire password verifier database, an attacker would still need to run a brute-force attack on the hash function we model after a random oracle in order to recover any plaintext passwords. A formal proof of these properties, along with a complete description of  $\Pi_{\text{pwd}}$  can be found in Section 3.

We also provide a concrete instantiation of  $\Pi_{\text{pwd}}$  in Section 4, along with performance benchmarks which show that incorporating  $\Pi_{\text{pwd}}$  into an existing login protocol is indeed realistic. We emphasize that this protocol offers better security when compared to the traditional hash-then-encrypt approach, while also avoiding multiple-message flows that are inherent with aPAKEs. Although our implementation serves as a stepping stone towards realizing credential-hiding login in practice, it still remains primarily as a proof-of-concept, and we hope that it encourages further improvements to the efficiency of credential-hiding login schemes, as well as more production-ready integrations into existing certificate-based server authentication protocols.

**Going beyond password credentials.** In Section 5, we present our generalized  $\Pi_{\text{fz}}$  construction, which leverages obfuscation primitives while still being provably secure from “standard” assumptions (albeit with the use of a random oracle). Although this construction is significantly less practical than  $\Pi_{\text{pwd}}$ , it establishes the feasibility of achieving credential-hiding login for more complex credentials than passwords, and that further research work in this area could directly benefit the field of biometric (and other) forms of authentication without having to reinvent and redefine the security principles that the community has established for password-based authentication.

Recently, Erwig et al. [EHOR20] showed how to construct a fuzzy aPAKE from error correcting codes and oblivious transfer. However, for an  $n$ -bit credential, their construction only handles  $O(\log n)$  errors, essentially by enumerating through all possible credentials that are close to the original registered credential. Our construction  $\Pi_{\text{fz}}$  obtains its error rate from the underlying fuzzy extractor, which is not bound by this brute-force technique and can handle a linear fraction of errors while also being asymptotically more efficient.

We conclude in Section 6 with a list of interesting directions for future work.

## 1.4 Related Work

In the following, we review some of the related areas of research which could be applied to the login scenario.

**Password authenticated key exchange.** As discussed previously, password-based authenticated key exchange is the most promising prior work for addressing confidentiality of passwords in a login protocol. Symmetric PAKE where both parties hold a low entropy password and want to establish a secure channel was first studied by Bellare and Merritt [BM93a], and later formalized by Bellare et al. [BPR00] using the game-based indistinguishability approach. Canetti et al. [CHK<sup>+</sup>05] provided the first formalization of PAKE in the UC framework [Can01]. The asymmetric variant of PAKE where one party (typically the server) only stores a one-way function of the password to protect against server compromise was first introduced by Bellare and Merritt [BM93b] and later formalized and studied in the simulation-based paradigm [BMP00, Mac01, MPS00]. Several follow-up works study aPAKE in the UC framework [GMR06, JR18], while the recent work of [JKX18] enhances aPAKE by providing protection against pre-computation attacks.

A number of PAKE constructions [AP05, HR10, HS14, HL19, PW17, JKX18] are also being discussed and reviewed by the Crypto Forum Research Group (CFRG) [Gro20] for standardization, with OPAQUE being the winner for aPAKEs and CPace the winner for (balanced) PAKEs.

Kiefer and Manulis [KM14] developed the notion of ZKPPC (zero-knowledge password policy checks) to enable the registration server to check if the client password satisfies a certain policy, while being oblivious of the raw password itself. Although their setting shares some similarities with applications of credential-hiding login in the use of a server-authenticated secure channel, we leave the construction of an analogous ZKPPC for credential-hiding login as an open problem.

**Private biometric authentication.** Protecting privacy of biometric credentials is even more critical given that they are impossible to change. A line of work building on information reconciliation [BBR88] and fuzzy extractors [DORS08, BDK<sup>+</sup>05, Boy04, DS05, CFP<sup>+</sup>16a, WZ17] propose potential solutions to this problem but they are limited in the biometric distributions and distances they support, do not provide password confidentiality during the login process, and do not provide protection against offline dictionary attacks (even in the presence of eavesdroppers).

In the multi-round setting, a line of work based on secure two-party computation [OPJM10, EHKM11, HEKM11, BCF<sup>+</sup>14], and the recent generalization of PAKE to the fuzzy setting [DHP<sup>+</sup>18] move closer to the credential privacy guarantees we aim for but are not secure in the face of server compromise (where the server stores the plaintext biometrics), and require three or more rounds of interaction between the client and the server.

## 2 Definitions

For a credential domain  $\mathcal{X}$  and verifier domain  $\mathcal{Y}$ , a credential-hiding login scheme  $\Pi = (\text{Setup}, \text{Register}, \text{ClientLogin}, \text{ServerLogin})$  for a circuit  $\mathcal{C} : \mathcal{X} \times \mathcal{X} \rightarrow \{0, 1\}$  consists of four algorithms defined as follows:

- $\text{Setup}(1^\lambda) \rightarrow (\text{pp}, \text{sk})$ . The **Setup** algorithm outputs public parameters  $\text{pp}$  and a secret key  $\text{sk}$ .
- $\text{Register}(\text{sk}, \text{uid}, \text{y}) \rightarrow \gamma$ . The **Register** algorithm takes as input the secret key  $\text{sk}$ , a unique identifier  $\text{uid}$ , and a credential  $\text{y} \in \mathcal{X}$ , and outputs a verifier  $\gamma \in \mathcal{Y}$ .
- $\text{ClientLogin}(\text{pp}, \text{ssid}, \text{tok}, \text{uid}, \text{x}) \rightarrow \alpha$ . The client login algorithm **ClientLogin** takes as input a set of public parameters  $\text{pp}$ , a subsession id  $\text{ssid}$ , a token  $\text{tok}$ , an identifier  $\text{uid}$ , an input  $\text{x} \in \mathcal{X}$ , and outputs a message  $\alpha$ .



- $\text{ServerLogin}(\text{pp}, \text{ssid}, \gamma, \alpha) \rightarrow \{\text{tok}, \perp\}$ . The server login algorithm  $\text{ServerLogin}$  takes in the public parameters  $\text{pp}$ , a subsession id  $\text{ssid}$ , a verifier  $\gamma$ , the client message  $\alpha$ , and outputs either a token  $\text{tok}$  or  $\perp$ .

**Terminology.** We use  $\text{uid}$  to represent a user identifier,  $\text{sid}$  to represent a session id,  $\text{ssid}$  to represent a subsession id, and  $\text{tok}$  as an optional string that can be input to  $\text{ClientLogin}$  by the client. To draw an analogy to a typical real-world setting of a login form between a user and a web service, we can think of  $\text{uid}$  as the unique username used for login by the client,  $\text{sid}$  a unique identifier to represent the name of the web service,  $\text{ssid}$  as the session identifier produced by the server-established secure channel, and  $\text{tok}$  an optional informational string (e.g. session cookie) that the client can supply to the server on a successful login attempt. We adopt the convention within the UC framework that the  $\text{sid}$  parameter does not appear in the scheme definition, but is present in our real world experiment.

We use  $\mathcal{X}$  to denote the credential domain,  $\mathcal{I}_1$  to denote the domain of service identifiers (corresponding to  $\text{sid}$ ),  $\mathcal{I}_2$  to denote the domain of session identifiers (corresponding to  $\text{ssid}$ ).

**Correctness.** We say that  $\Pi$  is *correct* if for any two credentials  $x, y \in \mathcal{X}$ , identifiers  $\text{uid} \in \mathcal{U}$ ,  $\text{ssid} \in \mathcal{I}_2$  and token  $\text{tok} \in \{0, 1\}^\lambda$ , for  $(\text{pp}, \text{sk}) \leftarrow \text{Setup}(1^\lambda)$ , if  $\mathcal{C}(x, y) = 1$ , then  $\text{ServerLogin}(\text{pp}, \text{ssid}, \text{Register}(\text{sk}, \text{uid}, y), \text{ClientLogin}(\text{pp}, \text{ssid}, \text{tok}, \text{uid}, x)) = \text{tok}$  with overwhelming probability.

**Security within the UC framework.** In the following subsections, we define security within the UC framework in the context of a “real world” and an “ideal world”. Intuitively, security will hold if for all environments  $\mathcal{E}$  and adversaries  $\mathcal{A}$ , there exists an efficient simulator  $\mathcal{A}^*$  for which the interactions in the real and ideal worlds are indistinguishable to  $\mathcal{E}$ .

## 2.1 Real World

We describe the “real” experiment between an adversary  $\mathcal{A}$  and challenger  $\mathcal{B}$  as follows. In the real experiment  $\text{Expt}_{\text{real}}(\mathcal{A})$ , the adversary has access to a series of oracles that  $\mathcal{B}$  responds to. Let  $\mathcal{H} : \{0, 1\}^* \rightarrow \{0, 1\}^*$  be a hash function which we will model after a random oracle.  $\mathcal{B}$  maintains a collection of databases (initialized to be empty) as follows:

- A database of registration keys with setup types  $\mathcal{D}_{\text{key}} \leftarrow \{\}$  which is used to keep track of setup keypairs  $(\text{pp}, \text{sk})$  along with their setup type, being either  $\text{HIDEREGKEY}$  or  $\text{REVEALREGKEY}$ ,
- A database of registered credentials  $\mathcal{D}_{\text{Reg}} \leftarrow \{\}$  which is used to keep track of the credentials produced by queries to a registration oracle, and
- A map of credential verifier query restriction statuses  $\mathcal{D}_{\text{status}} : \{\} \rightarrow \{\text{REGISTERED}, \text{RESTRICTED}, \text{STOLEN}\}$  which is kept in tandem with  $\mathcal{D}_{\text{Reg}}$  and used to keep track of the status of the credential verifiers.

The challenger responds to oracle queries made by the adversary  $\mathcal{A}$  as follows. The reader will notice some artificial restrictions we place on the protocol; these are necessary to ensure that security is not trivially broken or our construction can provably ensure security. We explain these restrictions, after defining these oracles:

- **Setup Oracle:** On an input  $(\text{SETUP}, \text{sid}, \text{SetupType})$  from  $S$ , with  $\text{SetupType} \in \{\text{HIDEREGKEY}, \text{REVEALREGKEY}\}$ , if there already exists an entry of the form  $(\text{sid}, \star, \star, \star) \in \mathcal{D}_{\text{key}}$ , then  $\mathcal{B}$  outputs  $\perp$ . Otherwise,  $\mathcal{B}$  computes  $(\text{pp}, \text{sk}) \leftarrow \text{Setup}(1^\lambda)$ . If  $\text{SetupType} = \text{HIDEREGKEY}$ , then  $\mathcal{B}$  only returns  $\text{pp}$  to  $\mathcal{A}$ ; otherwise, it returns  $(\text{pp}, \text{sk})$  to  $\mathcal{A}$ . Finally,  $\mathcal{B}$  adds the entry  $(\text{sid}, \text{pp}, \text{sk}, \text{SetupType})$  to  $\mathcal{D}_{\text{key}}$ .
- **Registration Oracle:** On an input  $(\text{STORECREDENTIAL}, \text{sid}, \text{uid}, \text{y})$  from  $S$ ,  $\mathcal{B}$  computes  $\gamma = \text{Register}(\text{sk}, \text{uid}, \text{y})$ , adds the entry  $(\text{sid}, \text{uid}, \text{y}, \gamma)$  to  $\mathcal{D}_{\text{Reg}}$ , and sets  $\mathcal{D}_{\text{status}}[\text{sid}, \text{uid}] = \text{REGISTERED}$ .
- **Steal Credential Verifier Oracle:** On an input  $(\text{sid}, \text{uid})$  from  $\mathcal{A}$ , if for some verifier  $\gamma$  there is an entry of the form  $(\text{sid}, \text{uid}, \star, \gamma)$  in  $\mathcal{D}_{\text{Reg}}$  and  $\mathcal{D}_{\text{status}}[\text{sid}, \text{uid}] = \text{REGISTERED}$ , then  $\mathcal{B}$  returns  $\gamma$  to  $\mathcal{A}$  and sets  $\mathcal{D}_{\text{status}}[\text{sid}, \text{uid}] = \text{STOLEN}$ . If no such entry in  $\mathcal{D}_{\text{Reg}}$  or  $\mathcal{D}_{\text{status}}$  exists, then  $\mathcal{B}$  returns  $\perp$  to  $\mathcal{A}$ .
- **Client Login Oracle:** On an input  $(\text{CLIENTLOGIN}, \text{sid}, \text{ssid}, \text{tok}, \text{uid}, \text{x})$  from  $U$ , if there is no entry of the form  $(\text{sid}, \star, \star, \text{HIDEREGKEY})$  in  $\mathcal{D}_{\text{key}}$ , then  $\mathcal{B}$  returns  $\perp$ . Similarly, if  $\mathcal{D}_{\text{status}}[\text{sid}, \text{uid}] = \text{STOLEN}$  and there exists an entry  $(\text{sid}, \text{uid}, \text{y}, \star)$  in  $\mathcal{D}_{\text{Reg}}$  such that  $\mathcal{C}(\text{x}, \text{y}) = 1$ , then  $\mathcal{B}$  also returns  $\perp$ . Otherwise, let  $(\text{pp}, \text{sk})$  be the keypair for which  $(\text{sid}, \text{pp}, \text{sk}, \text{HIDEREGKEY})$  exists in  $\mathcal{D}_{\text{key}}$ .  $\mathcal{B}$  sets  $\alpha \leftarrow \text{ClientLogin}(\text{pp}, \text{ssid}, \text{tok}, \text{uid}, \text{x})$  and returns  $(\text{sid}, \text{ssid}, \text{uid}, \alpha)$  to  $\mathcal{A}$ . Finally, if either there is no entry of the form  $(\text{sid}, \text{uid}, \star, \star)$  in  $\mathcal{D}_{\text{Reg}}$ , or there exists an entry of the form  $(\text{sid}, \text{uid}, \text{y}, \star)$  in  $\mathcal{D}_{\text{Reg}}$  for which  $\mathcal{C}(\text{x}, \text{y}) = 1$  and  $\mathcal{D}_{\text{status}}[\text{sid}, \text{uid}] = \text{REGISTERED}$ , then  $\mathcal{B}$  sets  $\mathcal{D}_{\text{status}}[\text{sid}, \text{uid}] = \text{RESTRICTED}$ .
- **Server Login Oracle:** On an input  $(\text{SERVERLOGIN}, \text{sid}, \text{ssid}, \text{uid})$  from  $S$ , wait and receive  $(\text{sid}, \text{ssid}, \text{uid}, \alpha)$  from  $\mathcal{A}$ . Then  $\mathcal{B}$  performs a lookup in  $\mathcal{D}_{\text{Reg}}$  to find an entry of the form  $(\text{sid}, \text{uid}, \star, \gamma)$  for some verifier  $\gamma$ , and another lookup in  $\mathcal{D}_{\text{key}}$  to find an entry of the form  $(\text{sid}, \text{pp}, \star)$ . If no such entry exists in either lookup,  $\mathcal{B}$  sets  $\text{tok} \leftarrow \perp$ . Otherwise,  $\mathcal{B}$  sets  $\text{tok} \leftarrow \text{ServerLogin}(\text{pp}, \text{ssid}, \gamma, \alpha)$ . If  $\text{tok} \neq \perp$ , then  $\mathcal{B}$  sets  $\text{flag} \leftarrow \text{SUCCESS}$  (otherwise,  $\text{flag} \leftarrow \text{FAILURE}$ ). Finally,  $\mathcal{B}$  returns the tuple  $(\text{sid}, \text{ssid}, \text{uid}, \text{tok}, \text{flag})$  to  $S$ .
- **Random Oracle:** On an input  $x \in \mathcal{X}$  from  $\mathcal{A}$ ,  $\mathcal{B}$  responds with  $\mathcal{H}(x)$ .

**Enforcing restrictions.** We provide an intuitive justification for the restrictions in each of the oracle query responses above. First, consider an adversary which obtains the setup secret at the beginning of the experiment. To such an adversary, the client login messages can no longer retain semantic security. Hence, we disallow client login messages to be produced in this case, ensured by referencing the  $\mathcal{D}_{\text{key}}$  database. However, even in this case, we still require that the adversary cannot successfully login without providing the correct credentials; this is ensured by keeping the server login oracle oblivious of whether the setup secret was stolen.

Also, consider an adversary which obtains a credential verifier corresponding to  $\text{uid}$ . To such an adversary, a correctly formed client login message from  $\text{uid}$  will no longer be semantically secure. Hence, we disallow the Client Login oracle to produce a message based on a correct credential in this case. For the same reason, we also restrict stealing of the credential verifier for a  $\text{uid}$  for which the adversary obtained a client login message based on correct credentials. All of this extra bookkeeping is ensured by the  $\mathcal{D}_{\text{status}}$  database. On the other hand, we still require that client messages with incorrect credentials must be simulatable, regardless of whether the credential verifier was stolen.

## 2.2 Ideal World

For an adversary  $\mathcal{A}$  and simulator  $\mathcal{A}^*$ , we describe the experiment  $\text{Expt}_{\text{ideal}}(\mathcal{A}, \mathcal{A}^*)$  with respect to the functionality  $\mathcal{F}_{\text{CHL}}$  which interacts with both the environment  $\mathcal{E}$  (which triggers the parties U and S) and simulator  $\mathcal{A}^*$  in the ideal world. The signature for each of the functions is described below:

- The queries of type SETUP, CLIENTLOGIN, and SERVERLOGIN take input from  $\mathcal{E}$  and output to  $\mathcal{A}^*$ .
- The queries of type STORECREDENTIAL take input from  $\mathcal{E}$  and outputs to  $\mathcal{E}$ .
- The queries of type STEALCREDENTIAL, OFFLINETESTCRED, TESTCRED, and LOGIN take input from  $\mathcal{A}^*$  and output to  $\mathcal{A}^*$ .

For each of the oracle query types made by either U or S in  $\text{Expt}_{\text{real}}$ , there is a corresponding function defined in  $\mathcal{F}_{\text{CHL}}$  which takes input from  $\mathcal{E}$  in  $\text{Expt}_{\text{ideal}}$ : SETUP invoked by Setup, STORECREDENTIAL invoked by Registration, STEALCREDENTIAL invoked by Steal Credential Verifier, CLIENTLOGIN invoked by Client Login, SERVERLOGIN invoked by Server Login. The Random Oracle is an exception to this case, since it has no interaction with  $\mathcal{F}_{\text{CHL}}$  and is simulated entirely by the simulator  $\mathcal{A}^*$ . For these functions for which  $\mathcal{F}_{\text{CHL}}$  returns its response to  $\mathcal{A}^*$ , it is then up to the design of  $\mathcal{A}^*$  to return the corresponding oracle query response back to  $\mathcal{E}$ .

The ideal functionality  $\mathcal{F}_{\text{CHL}}$  keeps an internal state consisting of records of the following types:

- **File:**  $\langle \text{FILE}, \text{sid}, \text{uid}, \text{x} \rangle$  which is marked as being UNCOMPROMISED or COMPROMISED.
- **Offline:**  $\langle \text{OFFLINE}, \text{sid}, \text{x} \rangle$ , for each offline credential guess.
- **Session:**  $\langle \text{sid}, \text{ssid}, \text{uid}, \text{x}, \text{type}, \text{tok} \rangle$ , where  $\text{type} \in \{\text{CORRECT}, \text{INCORRECT}, \text{INVALID}\}$ , and for which the record is marked with a **state** that is either: FRESH, COMPROMISED, INTERRUPTED, or COMPLETED.
- **Server:**  $\langle \text{SERVER}, \text{sid}, \text{ssid}, \text{uid}, \text{y} \rangle$  that is marked as either FRESH or COMPLETED.

The complete definition of  $\mathcal{F}_{\text{CHL}}$  is encapsulated in Figure 1. We use U to represent the user and S to represent the server, as the participating parties controlled by the environment in the protocol, and  $\mathcal{A}^*$  to represent the simulator in the ideal world.

## 2.3 Comparison to UC aPAKE Functionality

Although the  $\mathcal{F}_{\text{CHL}}$  functionality shares many commonalities with UC aPAKEs, it also has significant departures, as we discuss below. We recall the UC aPAKE functionality definition in Appendix D. In the following, we describe the relevant features of aPAKEs in-line, and we encourage the reader to refer to the definitions also presented in [GMR06, JKX18] as a reference, for completeness.

### Setup and Registration

- On (SETUP, sid, SetupType) from S, if this is the first SETUP message for sid, check if SetupType  $\in$  {HIDEREGKEY, REVEALREGKEY}, and then forward (Setup, sid, SetupType) to  $\mathcal{A}^*$ .
- On (STORECREDFILE, sid, uid, y) from S, if this is the first STORECREDFILE message for (sid, uid), record  $\langle$ FILE, sid, uid, y $\rangle$  and mark it UNCOMPROMISED.

### Stealing Credential Data

- On (STEALCREDFILE, sid, uid) from  $\mathcal{A}^*$ , if there is no record  $\langle$ FILE, sid, uid, y $\rangle$ , return “NO CRED FILE” to  $\mathcal{A}^*$ . Otherwise, if the record is marked UNCOMPROMISED, mark it COMPROMISED.
  - If there is a record  $\langle$ OFFLINE, sid, x $\rangle$ , send x to  $\mathcal{A}^*$ .
  - Else, return “CRED FILE STOLEN” to  $\mathcal{A}^*$ .
- On (OFFLINETESTCRED, sid, uid, x) from  $\mathcal{A}^*$ , do:
  - If there is a record  $\langle$ FILE, sid, uid, y $\rangle$  marked COMPROMISED, do: if  $\mathcal{C}(x, y) = 1$ , return “CORRECT GUESS” to  $\mathcal{A}^*$ ; else return “INCORRECT GUESS”.
  - Else if  $\mathcal{C}(x, y) = 1$ , then record  $\langle$ OFFLINE, sid, x $\rangle$ .

### Credential Login

- On (CLIENTLOGIN, sid, ssid, uid, x, tok) from U, if this the first CLIENTLOGIN message for (sid, ssid, uid), try to retrieve  $\langle$ FILE, sid, uid, y $\rangle$  and do:
  - if the retrieval did not work, set type = INVALID.
  - if  $\mathcal{C}(x, y) = 1$ , set type = CORRECT.
  - else set type = INCORRECT.Record  $\langle$ sid, ssid, uid, x, type, tok $\rangle$  and mark it FRESH. Send (sid, ssid, uid, type) to  $\mathcal{A}^*$ .
- On (SERVERLOGIN, sid, ssid, uid) from S, if this is the first SERVERLOGIN message for (sid, ssid, uid), retrieve  $\langle$ FILE, sid, uid, y $\rangle$  and send (sid, ssid, uid) to  $\mathcal{A}^*$ . Record  $\langle$ SERVER, sid, ssid, uid, y $\rangle$  and mark it FRESH.

### Active Session Attacks

- On (TESTCRED, sid, ssid, uid, x) from  $\mathcal{A}^*$ , if there is a record  $\langle$ SERVER, sid, ssid, uid, y $\rangle$  marked FRESH, do: If  $\mathcal{C}(x, y) = 1$ , return “CORRECT GUESS” to  $\mathcal{A}^*$  and set state = COMPROMISED; else return “INCORRECT GUESS” and set state = INTERRUPTED. Mark  $\langle$ sid, ssid, uid, x, type, tok $\rangle$  with state.

### Login and Authentication

- On (LOGIN, sid, ssid, uid, tok', flag) from  $\mathcal{A}^*$ , if there is a record  $\langle$ SERVER, sid, ssid, uid, y $\rangle$  marked FRESH, then mark it COMPLETED:
  - If there is a record  $\langle$ sid, ssid, uid, x, type, tok $\rangle$  marked FRESH, then mark it COMPLETED. Set flag = SUCCESS if type = CORRECT and flag = FAILURE, otherwise.
  - If there is a record  $\langle$ sid, ssid, uid, x, type, tok $\rangle$  marked INTERRUPTED, then mark it COMPLETED and set (tok, flag) = ( $\perp$ , FAILURE).
  - If there is a record  $\langle$ sid, ssid, uid, x, type, tok $\rangle$  marked COMPROMISED, then mark it COMPLETED and set tok = tok'.

Output (sid, ssid, uid, tok, flag) to S.

Figure 1: The ideal functionality  $\mathcal{F}_{\text{CHL}}$ .

**Generalization to credentials.** The  $\mathcal{F}_{\text{CHL}}$  functionality allows more general forms of authorization than just password equality verification.

**Setup and registration.** Unlike aPAKEs,  $\mathcal{F}_{\text{CHL}}$  allows for two levels of server secrets: one is a universal setup secret which applies to all parties and sessions within the scope of the `sid`, and the other is a set of credential verifiers (password files), each corresponding to a different user identifier `uid`. In practice, the setup secret is only kept with a registration server and the verifiers are kept with the login server.

In aPAKEs, the server can generate a credential file itself, given the user credential, but in  $\mathcal{F}_{\text{CHL}}$ , only the registration server can generate the verifier file—the login server cannot do this on its own. On the other hand, the login process does not involve the registration server.

As in aPAKEs, we consider security properties even when all the server secrets are revealed to the adversary, including the setup secret. However, our constructions require these to be handled differently from when only the login server secrets are stolen. In particular, the setup secret revelation can only be performed at the beginning, whereas the verifier files can be stolen adaptively at any point.

**Stealing credential data.** This functionality remains exactly same as in aPAKEs, since we model the same behavior that an adversary can perform in order to obtain the credential file for registered users.

**Credential login.** The `CLIENTLOGIN` query also takes a `tok` field in addition to the credential. The reason for this extra field is fairly technical, but we provide an intuitive perspective below. The token enables the environment to check whether there is a match between the client session for a given `ssid` and the server session for the same `ssid`. In aPAKEs, this is naturally observable by the environment, since in a successful session, the agreed keys would come out to be the same in both the matching sessions. The key is also independently random from all other components of the protocol in the functionality, and so the simulator does not have the ability to obtain it, in case the honest messages were not tampered with. However, in the case of  $\mathcal{F}_{\text{CHL}}$ , there is no key agreed on; only the output representing login success or failure. This could allow for simulator strategies for insecure protocols, where the simulator can enforce a success without even extracting a credential and calling the `TESTCRED` functionality, even when the adversary actively modifies messages. By allowing a token to come from the environment, and requiring a simulator to work for all possible environments, we prevent these kinds of simulator strategies that could work for inherently insecure protocols.

Although the functionality does not transmit the credential to the simulator, it does indicate three types: `CORRECT` to signify that the proper credential is being used; `INCORRECT` to indicate that an incorrect credential is being used; and `INVALID` to indicate that there is no registered credential. Even though this is more information than aPAKEs allow, it still does not harm credential privacy. We need this extra information to allow our constructions to have a simulation strategy for each type. Note also that, unlike aPAKEs, we explicitly handle the case where a client login is initiated, even if there is no corresponding registration. This is to address credential privacy in the practical situations where a client inadvertently logs in to this server mistaking it for some other server.

The Server Login portion remains the same as aPAKEs.

**Active session attacks.** The TESTCRED functionality is essentially same as in aPAKEs, except for the generalization to circuits, rather than just equality checks. In aPAKEs, there is also an IMPERSONATE functionality to allow the simulator to compromise a client session if the corresponding credential file was stolen. We do not require this functionality in  $\mathcal{F}_{\text{CHL}}$  because the client session does not output anything to the environment.

**Login and authentication.** This functionality demonstrates the difference in objectives between aPAKEs and  $\mathcal{F}_{\text{CHL}}$ . While in aPAKEs, the objective is to output a mutually agreed upon random key, here the login part just outputs the success or failure of the protocol. As discussed before, the token field is also output to the environment. In aPAKEs, there is a TESTABORT functionality to allow aborting if certain verifications fail. This functionality is redundant in  $\mathcal{F}_{\text{CHL}}$ , since login failure can already account for such failed verifications.

**Corruptions.** While in aPAKEs, the only public parameters come from a trusted third party, in  $\mathcal{F}_{\text{CHL}}$ , the public parameters come from the registration server which also holds a secret corresponding to this parameter. While we do not allow the registration server to generate these parameters maliciously, we do allow the adversary the choice to obtain the secret at the very beginning.

This distinction has some important consequences. Although trusting the registration server is not ideal, we argue that it is still limited in scope. Firstly, the public parameter generation is only done once in the scope of an `sid`. Secondly, even if the registration server secret is revealed, an attacker still has to do an offline attack, not only to recover a credential, but also to successfully login. Thirdly, the registration server does not get involved at all at the far more frequent login sessions - it's only used to register credentials and transmit a verifier file to the login server.

We also observe that this level of trust seems to be intrinsic to our setting. If the registration server held no setup secret, then that would allow the login server to be able to generate verifier from the public parameters themselves using its credential guesses. This would invalidate semantic security for client messages containing inadvertent credentials, an important practical scenario that we did aim to address. On the other hand, our trust assumption for the login server remains the same as in aPAKEs. We allow the adversary to adaptively steal credential verifiers from the login server.

## 2.4 Authenticated Login

The primary objective of two parties participating in a key exchange protocol is to establish a secure channel. It was shown in [CK02] that the UC notion of secure channel is realizable using a UC key exchange protocol as a subroutine, or more technically in the UC key exchange hybrid model.

To relate our notion of CHL to these existing notions, we define a simpler abstraction of the CHL ideal functionality, called “authenticated login”, which abstracts away passwords/credentials in the same sense as standard key exchange is to password-based key exchange. We show that the UC notion of secure channels can also be obtained in the authenticated login hybrid model.

However, the typical use case we consider within TLS does not require such a bootstrapping to secure channels. Instead, we use the CHL protocol to establish the authenticity of the user and continue the subsequent interactions based on the already established server-authenticated TLS channel, which corresponds to the CHL-over-TLS approach described in Section 1. Although this

suffices for our applications, we note that this is still weaker than a standard secure channel as it does impose on the user application to not send any sensitive content which assumes mutual trust before the CHL flow is completed.

Note that the compilation of an authenticated login scheme into a secure channel requires stronger security properties from the wrapper protocol than those needed for the corresponding wrapper protocol needed to wrap key exchange into a secure channel. While a key exchange protocol ensures that the exchanged key has strong entropy and computational independence properties, an authenticated login protocol instead receives the analogue of this key (the token  $\text{tok}$ ) from the environment, and hence does not by itself enjoy such entropy properties. In the presence of stronger wrapper protocols which are already practically deployed, such as TLS with server-side PKI authentication, one can indeed relax such requirements on the core credential verification system itself. Weakening the requirements enables the construction of efficient schemes which are also amenable to modular deployment in common scenarios existing today.

The formal descriptions and proofs are given in Appendix C. The description of the authenticated login functionality  $\mathcal{F}_{\text{AL}}$  is given in Figure 2. We provide a simple PKI-based protocol in Figure 3 and prove that it securely realizes the  $\mathcal{F}_{\text{AL}}$  functionality. We then show that the UC secure channel functionality  $\mathcal{F}_{\text{SC}}$ , described in Figure 4 can be securely realized in the  $\mathcal{F}_{\text{AL}}$ -hybrid model, by a protocol described in Figure 5.

### 3 Construction from Identity-Based Encryption

In this section, we describe a construction of Credential-Hiding Login (CHL) using an anonymous identity-based encryption (IBE) scheme. We will implicitly use the IND-anon-ID-CCA property of the IBE scheme to satisfy our notion of credential privacy for a CHL scheme. Then, coupled with the use of random oracles, we show that this instantiation satisfies UC security according to the definitions from Section 2.

**Identity-based encryption.** We first recall the definitions of an IBE scheme  $\text{IBE} = (\text{IBE.Setup}, \text{IBE.KeyGen}, \text{IBE.Encrypt}, \text{IBE.Decrypt})$ , consisting of four algorithms as follows. The setup algorithm  $\text{IBE.Setup}(1^\lambda) \rightarrow (\text{pp}, \text{msk})$  outputs a pair of public parameters  $\text{pp}$  and a master secret key  $\text{msk}$ . The key generation algorithm  $\text{IBE.KeyGen}(\text{msk}, \text{id}) \rightarrow \text{sk}_{\text{id}}$  takes the master secret key  $\text{msk}$  and an identity  $\text{id} \in \{0, 1\}^*$  and outputs a secret key for the identity  $\text{sk}_{\text{id}}$ . The encrypt algorithm  $\text{IBE.Encrypt}(\text{pp}, \text{id}, \text{m}) \rightarrow \text{c}$  takes as input the public parameters  $\text{pp}$ , an identity  $\text{id} \in \{0, 1\}^*$ , a message  $\text{m} \in \{0, 1\}^*$ , and outputs a ciphertext  $\text{c} \in \{0, 1\}^*$ . The decrypt algorithm  $\text{IBE.Decrypt}(\text{sk}_{\text{id}}, \text{c}) \rightarrow \{\text{m}, \perp\}$  takes as input a secret key  $\text{sk}_{\text{id}}$  and a ciphertext  $\text{c}$ , and outputs either  $\text{m}$  if decryption was successful, or  $\perp$  otherwise. The correctness guarantee of an IBE scheme is that for any message  $\text{m} \in \{0, 1\}^*$  and identity  $\text{id} \in \{0, 1\}^*$ , with  $(\text{pp}, \text{msk}) \leftarrow \text{IBE.Setup}(1^\lambda)$  and  $\text{sk}_{\text{id}} \leftarrow \text{IBE.KeyGen}(\text{msk}, \text{id})$ ,  $\text{IBE.Decrypt}(\text{sk}_{\text{id}}, \text{IBE.Encrypt}(\text{pp}, \text{id}, \text{m})) = \text{m}$  with overwhelming probability.

We define the experiment  $\text{Expt}_b^{\text{anon-cca}}(\mathcal{A})$  as follows:

**Definition 1** ( $\text{Expt}_b^{\text{anon-cca}}(\mathcal{A})$ ). For an adversary  $\mathcal{A}$  and bit  $b \in \{0, 1\}$ , we define  $\text{Expt}_b^{\text{anon-cca}}(\mathcal{A})$  as follows. First, the challenger computes  $(\text{pp}, \text{msk}) \leftarrow \text{IBE.Setup}(1^\lambda)$  and sends  $\text{pp}$  to  $\mathcal{A}$ . The adversary has access to a key generation oracle which, on input an identity  $\text{id}$ , the oracle returns  $\text{sk}_{\text{id}} \leftarrow \text{IBE.KeyGen}(\text{msk}, \text{id})$ . The adversary also has access to a challenge oracle, which on input

two pairs of inputs  $(id_1, m_1), (id_2, m_2)$ , the challenger returns  $\text{IBE.Encrypt}(\text{pp}, id_b, m_b)$  to  $\mathcal{A}$ . The adversary also has access to a decryption oracle, which on input a secret key  $\text{sk}_{id}$  and ciphertext  $c$ , returns  $\text{IBE.Decrypt}(\text{sk}_{id}, c)$ . If at any point in time, there is a key generation oracle query made for an identity that was also submitted to the challenge oracle, or if the decryption oracle is queried on  $(\text{sk}_{id_1}, m_1)$  or  $(\text{sk}_{id_2}, m_2)$ , then the experiment outputs  $\perp$ . At the end of the experiment,  $\mathcal{A}$  outputs a bit, which the experiment also outputs.

We say that an IBE scheme satisfies IND-anon-ID-CCA security if the quantity

$$|\Pr[\text{Expt}_b^{\text{anon-cca}}(\mathcal{A}) = 1] - \Pr[\text{Expt}_1^{\text{anon-cca}}(\mathcal{A}) = 1]|$$

is negligible.

**The construction  $\Pi_{\text{pwd}}$ .** Let  $\mathcal{X}$  be a domain of credentials,  $\mathcal{U}$  be a domain of user identifiers, and  $\mathcal{H} : \mathcal{U} \times \mathcal{X} \rightarrow (\{0, 1\}^\lambda, \{0, 1\}^\lambda)$  and  $f : \{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda$  be hash functions which we will model after random oracles. We describe the construction  $\Pi_{\text{pwd}}$  as follows:

- $\text{Setup}(1^\lambda) \rightarrow (\text{pp}, \text{sk})$ . The Setup algorithm outputs  $(\text{pp}, \text{sk}) \leftarrow \text{IBE.Setup}(1^\lambda)$ .
- $\text{Register}(\text{sk}, \text{uid}, y) \rightarrow \gamma$ . The Register algorithm takes as input a credential  $y \in \mathcal{X}$ , sets  $(s, t) = \mathcal{H}(\text{uid}, y)$ , and outputs  $(\text{IBE.KeyGen}(\text{sk}, s), f(t))$ .
- $\text{ClientLogin}(\text{pp}, \text{ssid}, \text{tok}, \text{uid}, x) \rightarrow \alpha$ . The client login algorithm takes as input a set of public parameters  $\text{pp}$ , ssid  $\text{ssid}$ , token  $\text{tok}$ , user id  $\text{uid} \in \mathcal{U}$ , a credential  $x \in \mathcal{X}$ , computes  $(s, t) = \mathcal{H}(\text{uid}, x)$ , and outputs a message

$$\alpha := \text{IBE.Encrypt}(\text{pp}, \text{id} = s, (t, \text{ssid}, \text{tok})).$$

- $\text{ServerLogin}(\text{pp}, \text{ssid}, \text{uid}, \gamma, \alpha) \rightarrow \{\text{tok}, \perp\}$ . The server login algorithm takes in the public parameters  $\text{pp}$ , ssid  $\text{ssid}$ , user id  $\text{uid} \in \mathcal{U}$ , a credential verifier  $\gamma = (\gamma_1, \gamma_2)$ , the client message  $\alpha$ , and computes  $(t', \text{ssid}', \text{tok}') \leftarrow \text{IBE.Decrypt}(\gamma_1, \alpha)$ . If  $t' = \perp$  or  $f(t') \neq \gamma_2$  or  $\text{ssid} \neq \text{ssid}'$ , then return  $\perp$ . Otherwise, return  $\text{tok}'$ .

**Correctness.** For any identifier  $\text{uid} \in \mathcal{U}$  and any two credentials  $x, y \in \mathcal{X}$ , for  $(\text{pp}, \text{sk}) \leftarrow \text{IBE.Setup}(1^\lambda)$ , for  $(s_x, t_x) = \mathcal{H}(\text{uid}, x)$  and  $(s_y, t_y) = \mathcal{H}(\text{uid}, y)$ , let the verifier be set as  $(\gamma_1, \gamma_2) = (\text{IBE.KeyGen}(\text{sk}, s_y), f(t_y)) \leftarrow \text{Register}(\text{sk}, \text{uid}, y)$ . Let  $\alpha = \text{IBE.Encrypt}(\text{pp}, \text{id} = s_x, (t_x, \text{ssid}, \text{tok}))$ . If  $x = y$ , then by the correctness of the IBE scheme, we have that  $\text{ServerLogin}(\text{pp}, (\gamma_1, \gamma_2), \alpha) = \text{tok}$ .

### 3.1 Description of the Simulator

The simulator  $\mathcal{A}^*$  will maintain the following databases across oracle queries:

- A database of registration keys with setup types  $\mathcal{D}_{\text{key}} \leftarrow \{\}$  which is used to keep track of setup keypairs  $(\text{pp}, \text{sk})$  along with their setup type, being either `HIDEREGKEY` or `REVEALREGKEY`,
- A map of credential file statuses  $\mathcal{D}_{\text{status}} : \mathcal{I}_1 \times \mathcal{U} \rightarrow \{\text{REGISTERED}, \text{RESTRICTED}, \text{STOLEN}\}$ ,
- A map  $\mathcal{D}_{\text{Reg}} : \mathcal{I}_1 \times \mathcal{U} \rightarrow \mathcal{X}$  of known registered passwords collected by the simulator,



- A database of client login requests  $\mathcal{D}_{\text{sent}} \leftarrow \{\}$  which is used to keep track of entries submitted to the Client Login oracle, and
- A database of random oracle queries  $\mathcal{D}_{\text{RO}} : \mathcal{I}_1 \times \mathcal{U} \times \mathcal{X} \rightarrow (\{0, 1\}^\lambda \times \{0, 1\}^\lambda)$ .

For each oracle query made by an adversary  $\mathcal{A}$ , we show how the simulator  $\mathcal{A}^*$  responds by taking advantage of the interface provided by the ideal functionality  $\mathcal{F}_{\text{CHL}}$ .

- **Setup Oracle:** On input  $(\text{sid}, \text{SetupType})$  from  $\mathcal{F}_{\text{CHL}}$ , the simulator samples  $(\text{pp}, \text{sk}) \leftarrow \Pi_{\text{pwd}}.\text{Setup}(1^\lambda)$ . If  $\text{SetupType} = \text{HIDEREGKEY}$ , then send  $\text{pp}$  to  $\mathcal{A}$  and retain  $\text{sk}$ . If  $\text{SetupType} = \text{REVEALREGKEY}$ , then send  $(\text{pp}, \text{sk})$  to  $\mathcal{A}$ . In either case, add the entry  $(\text{sid}, \text{pp}, \text{sk}, \text{SetupType})$  to  $\mathcal{D}_{\text{key}}$ .
- **Steal Credential Verifier Oracle:** On input  $(\text{sid}, \text{uid})$  from  $\mathcal{A}$ , if  $\mathcal{D}_{\text{status}}[\text{sid}, \text{uid}] = \text{RESTRICTED}$ , the simulator returns  $\perp$  to  $\mathcal{A}$ . Otherwise, the simulator  $\mathcal{A}^*$  first calls  $\mathcal{F}_{\text{CHL}}$  with the query  $(\text{STEALCREDFILE}, \text{sid}, \text{uid})$ . If it receives “NO CRED FILE” as a response, then it also sends  $\perp$  to  $\mathcal{A}$ . Otherwise, there are two cases to consider:
  - **Case 1:**  $\mathcal{F}_{\text{CHL}}$  returns “CRED FILE STOLEN”. Then, if there is an entry  $(\text{sid}, \text{uid})$  in  $\mathcal{D}_{\text{Reg}}$ , then set  $\text{x} \leftarrow \mathcal{D}_{\text{Reg}}[\text{sid}, \text{uid}]$  and set  $(\text{s}, \text{t}) \leftarrow \mathcal{D}_{\text{RO}}[\text{sid}, \text{uid}, \text{x}]$ . Otherwise, the simulator samples  $(\text{s}, \text{t}) \leftarrow_{\text{R}} \{0, 1\}^\lambda \times \{0, 1\}^\lambda$  uniformly at random and sets  $\mathcal{D}_{\text{RO}}[\text{sid}, \text{uid}, \star] \leftarrow (\text{s}, \text{t})$ .
  - **Case 2:**  $\mathcal{F}_{\text{CHL}}$  returns a credential  $\text{x}$ . Then, the simulator sets  $(\text{s}, \text{t}) \leftarrow \mathcal{D}_{\text{RO}}[\text{sid}, \text{uid}, \text{x}]$ .<sup>2</sup>

In either case,  $\mathcal{A}^*$  returns  $\gamma \leftarrow (\text{IBE.KeyGen}(\text{sk}, \text{s}), \text{f}(\text{t}))$  to the adversary.

- **Client Login Oracle:** Let  $(\text{CLIENTLOGIN}, \text{sid}, \text{ssid}, \text{uid}, \text{x}^*, \text{tok}^*, \text{type})$  be the query sent to  $\mathcal{F}_{\text{CHL}}$ . On input  $(\text{sid}, \text{ssid}, \text{uid}, \text{type})$  from  $\mathcal{F}_{\text{CHL}}$ , if there is an entry of the form  $(\text{sid}, \star, \star, \text{REVEALREGKEY})$  in  $\mathcal{D}_{\text{key}}$ , or if there is no entry of the form  $(\text{sid}, \star, \star, \star)$  in  $\mathcal{D}_{\text{key}}$ , or  $\mathcal{D}_{\text{status}}[\text{sid}, \text{uid}] = \text{“CRED FILE STOLEN”}$  and  $\text{type} = \text{CORRECT}$ , then  $\mathcal{A}^*$  returns  $\perp$  to  $\mathcal{A}$ . Otherwise, let  $\text{pp}$  be the unique parameter for which  $(\text{sid}, \text{pp}, \star, \star)$  exists in  $\mathcal{D}_{\text{key}}$ . Then,  $\mathcal{A}^*$  randomly samples  $\text{id} \leftarrow_{\text{R}} \{0, 1\}^\lambda$  and  $\text{m} \leftarrow_{\text{R}} \{0, 1\}^\lambda$ , and sends  $\alpha \leftarrow \text{IBE.Encrypt}(\text{pp}, \text{id}, \text{m})$  to  $\mathcal{A}$ . Finally, if  $\text{type} \in \{\text{INVALID}, \text{CORRECT}\}$ , then  $\mathcal{A}^*$  sets  $\mathcal{D}_{\text{status}}[\text{sid}, \text{uid}] \leftarrow \text{RESTRICTED}$  and adds the entry  $(\text{sid}, \text{ssid}, \text{uid}, \alpha)$  to  $\mathcal{D}_{\text{sent}}$ .
- **Server Login Oracle:** On input  $(\text{sid}, \text{ssid}, \text{uid}, \alpha)$  from  $\mathcal{A}$ , let  $(\alpha_1, \alpha_2) = \alpha$ . There are two cases to consider:
  - **Case 1:** There is already an entry of the form  $(\star, \star, \star, \alpha)$  in  $\mathcal{D}_{\text{sent}}$ . Let  $(\text{sid}', \text{ssid}', \text{uid}', \alpha)$  be this entry. If the two tuples  $(\text{sid}, \text{ssid}, \text{uid})$  and  $(\text{sid}', \text{ssid}', \text{uid}')$  are equal in all positions, then set  $\text{tok} \leftarrow \perp$  and  $\text{flag} \leftarrow \perp$ . Otherwise, set  $\text{tok} \leftarrow \perp$ ,  $\text{flag} \leftarrow \text{FAILURE}$ , and call  $(\text{TESTCRED}, \text{sid}, \text{ssid}, \text{uid}, \perp)$  from  $\mathcal{F}_{\text{CHL}}$ .
  - **Case 2:** There is no entry of the form  $(\star, \star, \star, \alpha)$  in  $\mathcal{D}_{\text{sent}}$ . There are two subcases to consider:
    - \* **Case 2a:** There is an entry of the form  $(\text{sid}, \text{uid}, \star)$  in  $\mathcal{D}_{\text{RO}}$ . Let  $(\text{s}, \text{t}) \leftarrow \mathcal{D}_{\text{RO}}[\text{sid}, \text{uid}, \star]$ , and compute  $(\text{v}, \text{ssid}', \text{tok}) \leftarrow \text{IBE.Decrypt}(\text{IBE.KeyGen}(\text{sk}, \text{s}), \alpha)$ . If  $\text{f}(\text{v}) = \text{f}(\text{t})$ , then the simulator aborts.<sup>3</sup> Otherwise, set  $\text{tok} \leftarrow \perp$  and  $\text{flag} \leftarrow \text{FAILURE}$ , and call  $\mathcal{F}_{\text{CHL}}$

<sup>2</sup>Note that at this point, the simulator must have already called  $\mathcal{F}_{\text{CHL}}$  with an `OFFLINETESTCRED` query with the correct password.

<sup>3</sup>We will show in the security proof that the probability of this event happening is negligible.

with  $(\text{TESTCRED}, \text{sid}, \text{ssid}, \text{uid}, \perp)$ .

- \* **Case 2b:** There is no entry of the form  $(\text{sid}, \text{uid}, \star)$  in  $\mathcal{D}_{\text{RO}}$ . Then, perform a linear search for every entry in  $\mathcal{D}_{\text{RO}}$  of the form  $(\text{sid}, \text{uid}, x) \rightarrow (s, t)$ , by checking if there exists a tuple  $(t', \text{ssid}', \text{tok})$  such that  $\text{IBE.Decrypt}(\text{IBE.KeyGen}(\text{sk}, s), \alpha) = (t', \text{ssid}', \text{tok})$ ,  $f(t') = f(t)$ , and  $\text{ssid} = \text{ssid}'$ . If so, then call  $\mathcal{F}_{\text{CHL}}$  with  $(\text{TESTCRED}, \text{sid}, \text{ssid}, \text{uid}, x)$ . If  $\mathcal{F}_{\text{CHL}}$  returns “CORRECT GUESS”, then set  $\text{flag} \leftarrow \text{SUCCESS}$  and set  $\mathcal{D}_{\text{Reg}}[\text{sid}, \text{uid}] \leftarrow x$ ; otherwise, set  $\text{tok} \leftarrow \perp$  and  $\text{flag} \leftarrow \text{FAILURE}$ . If the linear search fails, then set  $\text{tok} \leftarrow \perp$ ,  $\text{flag} \leftarrow \text{FAILURE}$ , and call  $\mathcal{F}_{\text{CHL}}$  with  $(\text{TESTCRED}, \text{sid}, \text{ssid}, \text{uid}, \perp)$ .

After all cases, call  $\mathcal{F}_{\text{CHL}}$  with  $(\text{LOGIN}, \text{sid}, \text{ssid}, \text{uid}, \text{tok}, \text{flag})$ .

- **Random Oracle:** On input  $(\text{sid}, \text{uid}, x)$  from  $\mathcal{A}$ , check if this entry is in  $\mathcal{D}_{\text{RO}}$ , returning  $(s, t) \leftarrow \mathcal{D}_{\text{RO}}[\text{sid}, \text{uid}, x]$  if so. Otherwise, call  $\mathcal{F}_{\text{CHL}}$  with the query  $(\text{OFFLINETESTCRED}, \text{sid}, \text{uid}, x)$ . There are two cases to consider:
  - $\mathcal{F}_{\text{CHL}}$  return “CORRECT GUESS”. Then first check if there exists a record  $(\text{sid}, \text{uid}, \star) \rightarrow (s, t)$  in  $\mathcal{D}_{\text{RO}}$ , replacing it with  $\mathcal{D}_{\text{RO}}[\text{sid}, \text{uid}, x] \leftarrow (s, t)$  if possible. Otherwise,  $\mathcal{A}^*$  samples  $(s, t) \leftarrow_{\text{R}} \{0, 1\}^\lambda \times \{0, 1\}^\lambda$  uniformly at random, sets  $\mathcal{D}_{\text{RO}}[\text{sid}, \text{uid}, x] \leftarrow (s, t)$  and sets  $\mathcal{D}_{\text{Reg}}[\text{sid}, \text{uid}] \leftarrow x$ .
  - $\mathcal{F}_{\text{CHL}}$  returns “INCORRECT GUESS”. Then  $\mathcal{A}^*$  samples  $(s, t) \leftarrow_{\text{R}} \{0, 1\}^\lambda \times \{0, 1\}^\lambda$  uniformly at random and sets  $\mathcal{D}_{\text{RO}}[\text{sid}, \text{uid}, x] \leftarrow (s, t)$ .

After both cases,  $\mathcal{A}^*$  finishes by returning  $(s, t)$  to  $\mathcal{A}$ .

### 3.2 Security

We show that our scheme is secure in the UC model. While we defer the full proof to Appendix A, here we provide an intuitive proof sketch.

**Theorem 1.** *The construction  $\Pi_{\text{pwd}}$  is a secure credential-hiding login protocol, assuming IBE is an IND-anon-ID-CCA IBE scheme, and  $\mathcal{H}$  and  $f$  are random oracles.*

**Proof Sketch.** We show that a protocol adversary cannot distinguish between the real world challenger and the simulator that just uses the  $\mathcal{F}_{\text{CHL}}$  functionality. This is done by transitioning through a series of hybrid games (denoted  $\text{Game}_i$  for  $i \in [1, 4]$ ), beginning with the simulator  $\mathcal{A}^*$  in the ideal world, and ending with the real world challenger.

While  $\mathcal{A}^*$  does not see any passwords in the clear, the challenger for  $\text{Game}_1$  obtains them from the environment. This enables the  $\text{Game}_1$  challenger to resolve the case where the simulator aborts in response to a Server Login Oracle query. Indistinguishability then follows by observing that the probability of the abort event is negligible, assuming  $f$  is a random oracle. The random oracle simulation used here is critical, since if  $f$  were only collision-resistant and one-way, it is not clear how to provide a consistent simulation.

The challenger for  $\text{Game}_2$  directly uses the password information, which the  $\text{Game}_1$  challenger was still performing a search on. Indistinguishability follows by using the random oracle properties of  $\mathcal{H}$  and the collision resistance property of  $f$ .

While the  $\text{Game}_2$  challenger was encrypting random elements for the Client Login code, the  $\text{Game}_3$  challenger instead encrypts using the real password. Indistinguishability follows by using

the IND-anon-ID-CCA property of the IBE scheme. Anonymity is needed because the id also depends on the password. CCA security is needed because in the Server Login code, we may need to perform decryption on the same id for which the encryption oracle may have responded in the Client Login code.

Finally, in  $\text{Game}_4$  we change the challenger code which was keeping track of previously sent encryptions, to instead just behave as the real world challenger. Indistinguishability follows by the random oracle assumption on  $\mathcal{H}$  and the collision resistance property of  $f$ . We then observe that the  $\text{Game}_4$  challenger mirrors real world challenger, which concludes the proof.

## 4 Implementation of IBE Construction

In order to assess the practicality of  $\Pi_{\text{pwd}}$ , we give a concrete instantiation of the credential-hiding login scheme along with performance benchmarks. In the following, we describe this direct construction using a concrete CCA-secure anonymous IBE scheme, aimed for practical applications which intend to use credential-hiding login for the specific case of equality checks (passwords).

**The  $\Pi_{\text{BF}}$  construction.** More specifically, we can instantiate  $\Pi_{\text{pwd}}$  with the Boneh-Franklin anonymous IBE [BF01], where security is based on the Decisional Bilinear Diffie Hellman assumption. Let  $p \in \Omega(\text{poly}(\lambda))$  be the order of the groups  $\mathbb{G}_1$ ,  $\mathbb{G}_2$ , and  $\mathbb{G}_T$ , and let  $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$  represent the pairing operation for a target group  $\mathbb{G}_T$ . Recall that we use  $\mathcal{X}$  to represent the credential domain. Let  $\mathcal{H}_1 : \mathcal{U} \times \mathcal{X} \rightarrow (\mathbb{G}_1, \mathbb{G}_T)$ ,  $\mathcal{H}_2 : \mathbb{G}_T \rightarrow \{0, 1\}^\lambda$ ,  $\mathcal{H}_3 : \{0, 1\}^\lambda \times \mathcal{X} \rightarrow \mathbb{Z}_p^*$ ,  $\mathcal{H}_4 : \{0, 1\}^\lambda \rightarrow (\{0, 1\}^\lambda \times \mathcal{I}_2 \times \{0, 1\}^\lambda)$ , and  $\mathcal{H}_5 : \{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda$  be hash functions modeled as random oracles. We describe the construction  $\Pi_{\text{BF}}$  as follows:

- $\text{Setup}(1^\lambda) \rightarrow (\text{pp}, \text{sk})$ . The **Setup** algorithm samples a random group element  $\mathbf{g} \in \mathbb{G}_2$ , a random secret scalar  $k \in \mathbb{Z}_p^*$ , and computes  $\mathbf{h} = \mathbf{g}^k$ . The public parameters are  $\text{pp} = (\mathbf{g}, \mathbf{h})$  and the secret key is  $\text{sk} = k$ .
- $\text{Register}(\text{sk}, \text{uid}, \mathbf{y}) \rightarrow \mathbf{v}$ . The **Register** algorithm takes as input a credential  $\mathbf{y} \in \mathcal{X}$ , sets  $(\mathbf{s}, \mathbf{t}) = \mathcal{H}_1(\text{uid}, \mathbf{y})$ , and outputs  $\gamma = (\mathbf{s}^k, \mathcal{H}_5(\mathbf{t}))$ .
- $\text{ClientLogin}(\text{pp}, \text{ssid}, \text{tok}, \text{uid}, \mathbf{x}) \rightarrow \alpha$ . The **client login** algorithm takes as input a set of public parameters  $\text{pp}$ , a subsession id  $\text{ssid} \in \mathcal{I}_2$ , a token  $\text{tok} \in \{0, 1\}^\lambda$ , a user identifier  $\text{uid} \in \mathcal{U}$ , and credential  $\mathbf{x} \in \mathcal{X}$ . It computes  $(\mathbf{s}, \mathbf{t}) \leftarrow \mathcal{H}_1(\text{uid}, \mathbf{x})$ , selects a random  $\sigma \leftarrow_{\mathbb{R}} \{0, 1\}^\lambda$  and sets  $r \leftarrow \mathcal{H}_3(\sigma, \mathbf{t})$ , and then outputs the message

$$\alpha := (\mathbf{g}^r, \sigma \oplus \mathcal{H}_2(e(\mathbf{s}, \mathbf{h}^r)), (\mathbf{t}, \text{ssid}, \text{tok}) \oplus \mathcal{H}_4(\sigma)).$$

- $\text{ServerLogin}(\text{pp}, \text{ssid}, \gamma, \alpha) \rightarrow \text{tok}$ . The **server login** algorithm takes in the public parameters  $\text{pp}$ , a subsession id  $\text{ssid}^*$ , a credential verifier  $(\gamma_1, \gamma_2) := \gamma$ , and the client message  $(\alpha_1, \alpha_2, \alpha_3) := \alpha$ . It computes  $\sigma \leftarrow \alpha_2 \oplus \mathcal{H}_2(e(\alpha_1, \gamma_1))$ , and then sets  $(\mathbf{t}, \text{ssid}^*, \text{tok}) \leftarrow \alpha_3 \oplus \mathcal{H}_4(\sigma)$ . Next, it computes  $r \leftarrow \mathcal{H}_3(\sigma, \mathbf{t})$  and checks if  $\alpha_1 = \mathbf{g}^r$ ,  $\text{ssid}^* = \text{ssid}$ , and  $\gamma_2 = \mathcal{H}_5(\mathbf{t})$ . If any of these checks fail, then the output is  $\perp$ . Otherwise, it outputs  $\text{tok}$ .

Note that both the Client Login and Server Login steps only require a single pairing operation, which we have experimentally verified is on the same order of magnitude as doing an exponentiation.

Procedures	Hash Evaluations	Overall Runtime ( $\mu$ s)
Setup	0	668
Register	2	382
Client Login	4	3106
Server Login	4	2641

Table 2: Our experimental benchmarks measured in microseconds of wall time, averaged over 1000 iterations on a 2.4GHz Intel Skylake CPU with 16MiB L3 cache.

#### 4.1 Concrete Instantiation

Our implementation is written using Rust and is available open-source<sup>4</sup>. Overall, we target 128 bits of security for our selection of primitives. We use the MIRACL Rust library for our base pairing operations over the BLS12-381 curve. We model our SHA3 evaluations (in SHAKE256 mode) as random oracle evaluations in  $\Pi_{\text{BF}}$ . For producing random scalars from seeded randomness, we use the HKDF “Extract” and “Expand” functions to obtain pseudorandom outputs. We remark that in a production implementation of our credential-hiding login scheme (and more generally, any login protocol), the server must keep a list of credential files for each registered user, and typically uses the `uid` parameter as a key to locate the associated credential file.

Furthermore, the subsession id parameter (denoted by `ssid`) is a session identifier that can be optionally reused from an encapsulating TLS session if the goal is to establish a secure session. This parameter, along with the token `tok` used to signify a login success, can also be omitted (or simply set to the empty string) depending on the application.

**Password hashing functions.** We note that in a typical production implementation of a login protocol, the credentials stored within the server are protected against offline dictionary attacks in the event of a server compromise through the use of memory-hard password hashing functions such as `scrypt` and `PBKDF2`. The number of iterations for these password hashing functions is often selected to be as expensive as possible, without having a significant impact on user experience. Therefore, it is likely to be the case that the execution of the iterations of the password hashing function would dominate the execution time of the remaining components of  $\Pi_{\text{BF}}$ .

Nevertheless, we demonstrate in the following section that  $\Pi_{\text{BF}}$  is quite reasonably performant without factoring in the password hashing function, and when compared against other existing solutions.

**Evaluation.** Our benchmarks are presented in Table 2, highlighting the total number of hash evaluations and overall runtime for each procedure. We emphasize that these benchmarks are intended primarily as a means to gauge the overall practicality of  $\Pi_{\text{pwd}}$ . In addition to the performance, we note that each client login message consists of three 256-bit components, for a total of 96 bytes per login message<sup>5</sup>. The password file is a single group element and is hence only 32 bytes.

<sup>4</sup><https://github.com/pmohassel/ch1>

<sup>5</sup>This is assuming that `ssid` and `tok` are unused, which may often be the case for a normal password-based login mechanism over the web.

We believe that these metrics are quite comparable to existing password login based schemes today (both in terms of bandwidth per client login message and space occupied per password file).

## 5 Construction for Fuzzy Credentials

In this section, we present a construction of credential-hiding login for fuzzy credentials, along with a description of the simulator for our UC definition. Our construction relies on a multitude of primitives, which we review in the following definitions. Recall the definition of a CCA-secure IBE scheme from Section 3, though without the need for the anonymity property (unlike  $\Pi_{\text{pwd}}$ ). We then review the definitions for a CCA-secure public-key encryption scheme, a reusable  $t$ -fuzzy extractor, and a distributional virtual black-box (VBB) obfuscator.

### 5.1 Primitives

Let  $\text{PKE} = (\text{PKE.Setup}, \text{PKE.Encrypt}, \text{PKE.Decrypt})$  be a public-key encryption scheme. Security is captured with the following definition.

**Definition 2** ( $\text{Expt}_b^{\text{CCA}}(\mathcal{A})$ ). For a bit  $b$ , recall the experiment  $\text{Expt}_b^{\text{CCA}}$  being defined as follows. First, the challenger runs  $(\text{pp}, \text{sk}) \leftarrow \text{PKE.Setup}(1^\lambda)$  and sends  $\text{pp}$  to  $\mathcal{A}$ . For an integer  $Q$ , the adversary  $\mathcal{A}$  can (adaptively) make  $Q$  queries to a decryption oracle hosted by the challenger, by submitting ciphertexts  $c_i$  for each  $i \in [1, Q]$  and receiving  $\text{PKE.Decrypt}(\text{sk}, c_i)$  as a response from the challenger. At some point, the adversary submits two challenge messages  $m_0^*$  and  $m_1^*$ , and the challenger returns  $c^* \leftarrow \text{PKE.Encrypt}(\text{pp}, m_b^*)$  to  $\mathcal{A}$ . If there exists some  $i \in [1, Q]$  for which  $c^* = c_i$ , then the experiment halts and outputs  $\perp$ . The adversary outputs a bit  $b'$ , which is also the output of the experiment. We say that  $\text{PKE}$  is CCA-secure if for every efficient adversary  $\mathcal{A}$ , the quantity

$$\left| \Pr[\text{Expt}_0^{\text{CCA}}(\mathcal{A}) = 1] - \Pr[\text{Expt}_1^{\text{CCA}}(\mathcal{A}) = 1] \right|$$

is negligible in the security parameter  $\lambda$ .

Let  $\text{FuzzyExt} = (\text{Gen}, \text{Rep})$  be a reusable  $t$ -fuzzy extractor from [CFP<sup>+</sup>16b]. Correctness implies that, for a domain of inputs  $\{0, 1\}^\lambda$ , for every  $x, y \in \{0, 1\}^\lambda$ , if  $d(x, y) \leq t$ , then for  $(r, p) \leftarrow \text{Gen}(x)$ , we have that  $\text{Rep}(y, p) = r$ . Security is captured with the following definition.

**Definition 3** ( $\text{Expt}_b^{\text{FE}}(\mathcal{A})$ ). For a bit  $b$ , recall the experiment  $\text{Expt}_b^{\text{FE}}$  between a challenger and adversary  $\mathcal{A}$  being defined as follows. For a distribution  $W$  of inputs and integer  $N = \text{poly}(\lambda)$ , the challenger samples inputs  $w_1, \dots, w_N \leftarrow_{\text{R}} W$  and computes  $(r_i, p_i) \leftarrow \text{Gen}(w_i)$ . If  $b = 0$ , it sends  $(r_i, p_i)$  to  $\mathcal{A}$ ; otherwise if  $b = 1$ , it samples random  $u_1, \dots, u_N \leftarrow_{\text{R}} \{0, 1\}^\lambda$  and sends  $(u_i, p_i)$  to  $\mathcal{A}$ . The adversary outputs a bit, which is also the output of the experiment. We say that  $\text{FuzzyExt}$  is secure if for every efficient adversary  $\mathcal{A}$ , the quantity

$$\left| \Pr[\text{Expt}_0^{\text{FE}}(\mathcal{A}) = 1] - \Pr[\text{Expt}_1^{\text{FE}}(\mathcal{A}) = 1] \right|$$

is negligible in  $\lambda$ .

Let  $\text{Obf}$  represent a distributional VBB obfuscator as defined in [WZ17].  $\text{Obf}$  takes as input a program  $P : \mathcal{X} \rightarrow \{0, 1\}^\lambda \cup \{\perp\}$  and outputs an obfuscated program  $\hat{P}$ , such that for all inputs  $\alpha \in$

$\mathcal{X}$ ,  $P(\alpha) = \hat{P}(\alpha)$ . Security intuitively captures the property that having access to the description of the obfuscated program can be simulated by having only oracle access to the program itself, under the assumption that the program  $P$  is sampled from a sufficiently unpredictable distribution. In this work, we will focus on the obfuscation of multi-bit compute-and-compare (MBCC) programs. We write  $\mathbf{MBCC}[f, y, z]$  to denote the program that on input  $\alpha$ , outputs  $z$  if  $f(\alpha) = y$ , and outputs  $\perp$  otherwise.

**Definition 4** (Computationally unpredictable MBCC distribution). We say that a distribution  $D$  of MBCC program is unpredictable if the following holds. Denoting  $\mathbf{MBCC}[f, y, z]$  as the program sampled from  $D$  by a challenger, the probability that an adversary, given only  $f$  and  $z$ , can output  $y$ , is negligible in  $\lambda$ .

**Definition 5** (Distributional VBB obfuscator for MBCC programs). Let  $\mathcal{A}$  be an adversary, and let  $D$  be a computationally unpredictable MBCC distribution of programs. Then, there exists a simulator  $\text{Sim}$ , which has only oracle access to  $P(\cdot)$ , for which the quantity

$$\left| \Pr[\mathcal{A}(\text{Obf}(P)) = 1] - \Pr[\text{Sim}^{P(\cdot)}(1^\lambda) = 1] \right|$$

is negligible in  $\lambda$ .<sup>6</sup>

Wichs and Zirdelis [WZ17] and Goyal, Koppula and Waters [GKW17] show that a distributional VBB obfuscator for MBCC programs can be realized based on the Learning with Errors assumption.

## 5.2 The Construction $\Pi_{fz}$

For simplicity, we assume that all key domains will be in  $\{0, 1\}^\lambda$ , and let  $\mathcal{X}$  be the domain of the credentials. Let  $\text{PKE} = (\text{PKE.Setup}, \text{PKE.Encrypt}, \text{PKE.Decrypt})$  represent a CCA-secure public-key encryption scheme, and let  $\text{IBE} = (\text{IBE.Setup}, \text{IBE.KeyGen}, \text{IBE.Encrypt}, \text{IBE.Decrypt})$  represent a (public-key) CCA-secure identity-based encryption scheme. Let  $\text{Obf}$  represent an obfuscator for compute-and-compare programs which satisfies distributional virtual black-box security. Let  $\mathcal{H}$  be a hash function modeled after a random oracle. Let  $\text{FuzzyExt} = (\text{Gen}, \text{Rep})$  represent a  $t$ -reusable fuzzy extractor. We define the following function which represents an inner routine for the program we will obfuscate:

$$f_{\text{sk}_{\text{pke}}, \text{uid}, \sigma}(\alpha) := \mathcal{H}(\text{uid}, \text{Rep}(\text{PKE.Decrypt}(\text{sk}_{\text{pke}}, \alpha), \sigma)).$$

The construction  $\Pi_{fz}$  is described as follows:

- $\text{Setup}(1^\lambda) \rightarrow (\text{pp}, \text{sk})$ . The Setup algorithm computes  $(\text{pp}_{\text{pke}}, \text{sk}_{\text{pke}}) \leftarrow \text{PKE.Setup}(1^\lambda)$  and  $(\text{pp}_{\text{ibe}}, \text{sk}_{\text{ibe}}) \leftarrow \text{IBE.Setup}(1^\lambda)$  and outputs  $\text{pp} = (\text{pp}_{\text{pke}}, \text{pp}_{\text{ibe}})$  and  $\text{sk} = (\text{sk}_{\text{pke}}, \text{sk}_{\text{ibe}})$ .
- $\text{Register}(\text{sk}, \text{uid}, y) \rightarrow v$ . The Register algorithm takes as input a secret key  $\text{sk} = (\text{sk}_{\text{pke}}, \text{sk}_{\text{ibe}})$ , an identifier  $\text{uid}$ , a credential  $y$ , computes  $(r, \sigma) \leftarrow \text{Gen}(y)$ , computes  $\text{sk}_{\text{uid}} \leftarrow \text{IBE.KeyGen}(\text{sk}_{\text{ibe}}, \text{uid})$ , and outputs  $\gamma = \hat{P} \leftarrow \text{Obf}(1^\lambda, \mathbf{MBCC}[f_{\text{sk}_{\text{pke}}, \text{uid}, \sigma}, \mathcal{H}(\text{uid}, r), \text{sk}_{\text{uid}}])$ .
- $\text{ClientLogin}(\text{pp}, \text{ssid}, \text{tok}, \text{uid}, x) \rightarrow \alpha$ . The client login algorithm takes as input a set of public parameters  $\text{pp}$ , an identifier  $\text{uid}$ , a credential  $x$ , computes  $\alpha_1 \leftarrow \text{PKE.Encrypt}(\text{pp}_{\text{pke}}, x)$ ,  $\alpha_2 \leftarrow \text{IBE.Encrypt}(\text{pp}_{\text{ibe}}, (\text{ssid}, \text{tok}, \alpha_1), \text{uid})$ , and returns  $\alpha \leftarrow (\alpha_1, \alpha_2)$ .

---

<sup>6</sup>Note that we omit the inclusion of auxiliary information and the program parameters in this description for simplicity, but these are certainly accounted for in the security proof.

- $\text{ServerLogin}(\text{pp}, \text{ssid}, \gamma, \alpha) \rightarrow \text{tok}$ . The server login algorithm takes in the public parameters  $\text{pp}$ , a credential verifier  $\gamma = \hat{P}$ , the client message  $\alpha = (\alpha_1, \alpha_2)$ , and computes  $\text{sk}_{\text{uid}} \leftarrow \hat{P}(\alpha_1)$ , outputting  $\perp$  if  $\text{sk}_{\text{uid}} = \perp$ . Then, it computes  $(\text{ssid}', \text{tok}, \alpha'_1) \leftarrow \text{IBE.Decrypt}(\text{sk}_{\text{uid}}, \alpha_2)$ , outputting  $\perp$  if either decryption fails,  $\alpha_1 \neq \alpha'_1$ , or  $\text{ssid} \neq \text{ssid}'$ . Otherwise, it outputs  $\text{tok}$ .

Note that the security of all of the primitives used in this construction can be based on random oracles and the LWE assumption.

**Correctness.** For any identifier  $\text{uid}$  and any two credentials  $x, y \in \mathcal{X}$ , for  $(\text{pp}, \text{sk}) \leftarrow \text{PKE.Setup}(1^\lambda)$ , and  $\gamma \leftarrow \hat{P}$ , with  $(r, \sigma) \leftarrow \text{Gen}(y)$ , by the correctness of the obfuscator and fuzzy extractor, the  $\text{ServerLogin}$  algorithm outputs  $\hat{P}(\alpha) = \mathcal{H}(\text{uid}, \text{Rep}(\text{PKE.Decrypt}(\text{sk}, \text{PKE.Encrypt}(\text{pp}, x)), \sigma))$ , which is equal to  $\mathcal{H}(\text{uid}, r)$  if and only if  $d(x, y) \leq t$ . Hence, the obfuscator outputs 1 if and only if  $d(x, y) \leq t$ , which is the  $\text{ServerLogin}$  output as well.

### 5.3 Description of the Simulator

In the following description, we consider a simulator which additionally responds to Obfuscator Oracle queries made by  $\mathcal{A}$ . The adversary can submit queries to this obfuscator oracle for program identifiers for which it has received the corresponding obfuscations as computed by the Steal Credential Verifier Oracle. Consequently, the adversary does not receive any response from the Steal Credential Verifier Oracle other than a program identifier  $\text{pid}$ .

Note that this deviates from the adversary's interface described in the real world in Section 2. But in Appendix B, we show how, assuming our obfuscator satisfies the distributional VBB property, that the indistinguishability between the real world and ideal world with this extended adversary's interface implies the indistinguishability between the real world and ideal world *without* the addition of the obfuscator oracle. For the ease of exposition, we defer this argument to the conclusion of the proof of security, and begin with a description of the simulator with the obfuscator oracle as follows.

The simulator  $\mathcal{A}^*$  will maintain the following databases across oracle queries:

- A database of registration keys with setup types  $\mathcal{D}_{\text{key}} \leftarrow \{\}$  which is used to keep track of setup keypairs  $(\text{pp}, \text{sk})$  along with their setup type, being either `HIDEREGKEY` or `REVEALREGKEY`,
- A map of credential file statuses  $\mathcal{D}_{\text{status}} : \mathcal{I}_1 \times \mathcal{U} \rightarrow \{\text{REGISTERED}, \text{RESTRICTED}, \text{STOLEN}\}$ ,
- A database of client login requests  $\mathcal{D}_{\text{sent}} \leftarrow \{\}$  which is used to keep track of entries submitted to the Client Login oracle,
- A map of ciphertexts for PKE to their corresponding plaintexts  $\mathcal{D}_{\text{PKE}} : \{0, 1\}^* \rightarrow \{0, 1\}^*$  to keep track of encryption calls made for the Client Login oracle,
- A database of random oracle queries  $\mathcal{D}_{\text{RO}} : \mathcal{U} \times \{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda$ , and
- A map of program ids  $\text{pid}$  to program descriptions  $\mathcal{D}_{\text{obf}}$  that have been obfuscated.

For each oracle query made by an adversary  $\mathcal{A}$ , we show how the simulator  $\mathcal{A}^*$  responds by taking advantage of the interface provided by the ideal functionality  $\mathcal{F}_{\text{CHL}}$ .

- **Setup Oracle:** On input  $(\text{sid}, \text{SetupType})$  from  $\mathcal{F}_{\text{CHL}}$ , the simulator samples  $(\text{pp}, \text{sk}) \leftarrow \Pi_{\text{fz}}.\text{Setup}(1^\lambda)$ . If  $\text{SetupType} = \text{HIDEREKEY}$ , then send  $\text{pp}$  to  $\mathcal{A}$  and retain  $\text{sk}$ . If  $\text{SetupType} = \text{REVEALREGKEY}$ , then send  $(\text{pp}, \text{sk})$  to  $\mathcal{A}$ . In either case, add the entry  $(\text{sid}, \text{pp}, \text{sk}, \text{SetupType})$  to  $\mathcal{D}_{\text{key}}$ .
- **Steal Credential Verifier Oracle:** On input  $(\text{sid}, \text{uid})$  from  $\mathcal{A}$ , if  $\mathcal{D}_{\text{status}}[\text{sid}, \text{uid}] = \text{RESTRICTED}$ , the simulator returns  $\perp$  to  $\mathcal{A}$ . Otherwise, it computes  $\text{sk}_{\text{uid}} \leftarrow \text{IBE.KeyGen}(\text{sk}_{\text{ibe}}, \text{uid})$ , and then calls  $\mathcal{F}_{\text{CHL}}$  with the query  $(\text{STEALCREDFILE}, \text{sid}, \text{uid})$ . If  $\mathcal{A}^*$  receives “NO CRED FILE” as a response, then it also sends  $\perp$  to  $\mathcal{A}$ . Otherwise:
  - If it receives “CRED FILE STOLEN”, then Sim picks a random  $s \leftarrow_{\text{R}} \{0, 1\}^\lambda$  and sets  $\mathcal{D}_{\text{RO}}[\text{uid}, \perp] \leftarrow s$ .
  - If it receives a credential  $y$ , then Sim picks a random  $s \leftarrow_{\text{R}} \{0, 1\}^\lambda$  and sets  $\mathcal{D}_{\text{RO}}[\text{uid}, \perp] \leftarrow s$ .

In either case,  $\mathcal{A}^*$  responds by picking a unique identifier  $\text{pid}$ , adding the mapping  $\text{pid} \rightarrow \text{MBCC}[f_{\text{sk}_{\text{pke}}, \text{uid}, \sigma}, s, \text{sk}_{\text{uid}}]$  to  $\mathcal{D}_{\text{obf}}$ , and returning  $\text{pid}$  to  $\mathcal{A}$ .

- **Client Login Oracle:** Let  $(\text{CLIENTLOGIN}, \text{sid}, \text{ssid}, \text{uid}, \text{x}^*, \text{tok}^*, \text{type})$  be the query sent to  $\mathcal{F}_{\text{CHL}}$ . On input  $(\text{sid}, \text{ssid}, \text{uid}, \text{type})$  from  $\mathcal{F}_{\text{CHL}}$ , if there is an entry of the form  $(\text{sid}, \star, \star, \text{REVEALREGKEY})$  in  $\mathcal{D}_{\text{key}}$ , or if there is no entry of the form  $(\text{sid}, \star, \star, \star)$  in  $\mathcal{D}_{\text{key}}$ , or  $\mathcal{D}_{\text{status}}[\text{sid}, \text{uid}] = \text{“CRED FILE STOLEN”}$  and  $\text{type} = \text{CORRECT}$ , then  $\mathcal{A}^*$  returns  $\perp$  to  $\mathcal{A}$ . Otherwise, let  $\text{pp}$  be the unique parameter for which  $(\text{sid}, \text{pp}, \star, \star)$  exists in  $\mathcal{D}_{\text{key}}$ . Then,  $\mathcal{A}^*$  randomly samples a credential  $\text{x} \leftarrow \mathcal{X}$  and token  $\text{tok} \in \{0, 1\}^\lambda$  and returns  $\alpha = (\alpha_1, \alpha_2) \leftarrow \Pi_{\text{fz}}.\text{ClientLogin}(\text{pp}, \text{ssid}, \text{tok}, \text{uid}, \text{x})$  to  $\mathcal{A}$ . Finally, if  $\text{type} \in \{\text{INVALID}, \text{CORRECT}\}$ , then  $\mathcal{A}^*$  sets  $\mathcal{D}_{\text{status}}[\text{sid}, \text{uid}] \leftarrow \text{RESTRICTED}$ , adds the entry  $(\text{sid}, \text{ssid}, \text{uid}, \alpha)$  to  $\mathcal{D}_{\text{sent}}$ , and sets  $\mathcal{D}_{\text{PKE}}[\alpha_1] \leftarrow \text{x}$ .
- **Server Login Oracle:** On input  $(\text{sid}, \text{ssid}, \text{uid}, \alpha)$  from  $\mathcal{A}$ , let  $(\alpha_1, \alpha_2) = \alpha$ . There are two cases to consider:
  - **Case 1:** There is already an entry of the form  $(\star, \star, \star, \alpha)$  in  $\mathcal{D}_{\text{sent}}$ . Let  $(\text{sid}', \text{ssid}', \text{uid}', \alpha)$  be this entry. If the two tuples  $(\text{sid}, \text{ssid}, \text{uid})$  and  $(\text{sid}', \text{ssid}', \text{uid}')$  are equal in all positions, then set  $\text{tok} = \perp$  and  $\text{flag} = \perp$ . Otherwise, set  $\text{tok} = \perp$ ,  $\text{flag} = \text{FAILURE}$ , and call  $\mathcal{F}_{\text{CHL}}$  with the query  $(\text{TESTCRED}, \text{sid}, \text{ssid}, \text{uid}, \perp)$ .
  - **Case 2:** There is no entry of the form  $(\star, \star, \star, \alpha)$  in  $\mathcal{D}_{\text{sent}}$ . The simulator computes  $\text{sk}_{\text{uid}} \leftarrow \text{IBE.KeyGen}(\text{sk}_{\text{ibe}}, \text{uid})$  and  $(\text{ssid}', \text{tok}', \alpha'_1) \leftarrow \text{IBE.Decrypt}(\text{sk}_{\text{uid}}, \alpha_2)$  and checks if  $\text{ssid} = \text{ssid}'$  and  $\alpha_1 = \alpha'_1$ , and then computes  $\text{x} \leftarrow \text{PKE.Decrypt}(\text{sk}_{\text{pke}}, \alpha_1)$ . If all of these checks pass, then the simulator sets  $\text{tok} \leftarrow \text{tok}'$ ,  $\text{flag} = \text{SUCCESS}$ , and calls  $\mathcal{F}_{\text{CHL}}$  with  $(\text{TESTCRED}, \text{sid}, \text{ssid}, \text{uid}, \text{x})$ . Otherwise, if any of the above steps fail, then it sets  $\text{tok} = \perp$ ,  $\text{flag} = \text{FAILURE}$ , and calls  $\mathcal{F}_{\text{CHL}}$  with  $(\text{TESTCRED}, \text{sid}, \text{ssid}, \text{uid}, \perp)$ .

After all cases, call  $\mathcal{F}_{\text{CHL}}$  with  $(\text{LOGIN}, \text{sid}, \text{ssid}, \text{uid}, \text{tok}, \text{flag})$ .

- **Random Oracle:** On input  $(\text{uid}, r)$  from  $\mathcal{A}$ , the simulator returns  $\mathcal{D}_{\text{RO}}[\text{uid}, r]$  if it exists. Otherwise, it samples an  $s \leftarrow_{\text{R}} \{0, 1\}^\lambda$  uniformly at random, sets  $\mathcal{D}_{\text{RO}}[\text{uid}, r] \leftarrow s$ , and returns  $s$ .



- **Obfuscator Oracle:** On input a program identifier  $\text{pid}$  and input  $\alpha$  from  $\mathcal{A}$ , Sim first checks if  $\text{pid}$  is an index in  $\mathcal{D}_{\text{obf}}$ , returning  $\perp$  if not. Otherwise, let  $\text{MBCC}[f_{\text{sk}_{\text{pke}}, \text{uid}, \sigma}, s, \text{sk}_{\text{uid}}]$  be the program associated with  $\text{pid}$ . Sim checks if  $\alpha$  is set in  $\mathcal{D}_{\text{PKE}}$ , letting  $x \leftarrow \mathcal{D}_{\text{PKE}}[\alpha]$  if so. Otherwise, Sim computes  $x \leftarrow \text{PKE.Decrypt}(\text{sk}_{\text{PKE}}, \alpha)$  and calls  $\mathcal{F}_{\text{CHL}}$  with  $(\text{OFFLINETESTCRED}, \text{sid}, \text{uid}, x)$ . If the response is “CORRECT GUESS”, then:
    - Sim checks if there exists an entry of the form  $\mathcal{D}_{\text{RO}}[\text{uid}, \perp] = s$ , and if so, then Sim picks a random  $r \leftarrow_{\text{R}} \{0, 1\}^\lambda$  and sets  $\mathcal{D}_{\text{RO}}[\text{uid}, r] \leftarrow s$  (erasing the entry for  $\mathcal{D}_{\text{RO}}[\text{uid}, \perp]$ ).
    - Then, Sim returns  $\text{sk}_{\text{uid}}$  to  $\mathcal{A}$ .
- Otherwise, Sim returns  $\perp$  to  $\mathcal{A}$ .

We use the simulator defined above in order to prove the following theorem.

**Theorem 2.** *The protocol  $\Pi_{\text{fz}}$  securely realizes the  $\mathcal{F}_{\text{CHL}}$  functionality in the UC model, assuming FuzzyExt is a reusable fuzzy extractor, IBE satisfies IND-ID-CCA security, PKE satisfies CCA security,  $\mathcal{H}$  is a random oracle, and the input credential distribution  $\mathcal{X}$  has min-entropy at least  $\lambda$ .*

We note that all of the primitives described above can be achieved from the LWE assumption (except for the modeling of  $\mathcal{H}$  as a random oracle).

**Proof sketch.** To prove that  $\Pi_{\text{fz}}$  satisfies our UC formulation of a credential-hiding login protocol, we can proceed with a series of intermediate games (similar to the proof of security for  $\Pi_{\text{pwd}}$ ), with the notable exception that we ultimately rely on the distributional VBB obfuscation property to establish indistinguishability between the real-world protocol and the above simulator description which gives the adversary access to the Obfuscator Oracle.

In order to satisfy the requirements for applying distributional VBB obfuscation for a given function  $\text{MBCC}[f, y, z]$ , we claim that the value  $y$  is computationally unpredictable given access to the description of  $f$  and any auxiliary information required for evaluation. Recall that in our instantiation, we obfuscate the function  $\text{MBCC}[f_{\text{sk}_{\text{pke}}, \text{uid}, \sigma}, \mathcal{H}(\text{uid}, r), \text{sk}_{\text{uid}}]$ , where  $(r, \sigma) \leftarrow \text{Gen}(y)$ . The security of the fuzzy extractor states that if  $y$  comes from a distribution of sufficient min-entropy, then  $r$  is sufficiently unpredictable given only  $\sigma$ . Putting this together, we can then conclude that  $\mathcal{H}(\text{uid}, r)$  is indeed unpredictable to  $\mathcal{A}$ , which then allows us to replace the adversary’s access to  $\hat{P}$  with the oracle access to the Obfuscator Oracle in our simulator. We provide the full proof in Appendix B.

## 6 Conclusions

In this work, we introduced the notion of credential-hiding login, along with a formulation of security within the UC framework. Credential-hiding login can be seen as a special case of the more general primitive: authenticated login, which when used within an already-established server-authenticated channel, suffices for the user login scenario. We then presented two constructions of credential-hiding login:  $\Pi_{\text{pwd}}$  for handling password-based login that we also implemented as a proof-of-concept, and  $\Pi_{\text{fz}}$  for handling fuzzy credentials, as a generalization of the equality checks done by passwords to Hamming distance checks for approximate equality.

To conclude, we present the following directions for future study:

- Can we have CHL constructions resistant to precomputation attacks (as defined by the distinction between “strong aPAKEs” versus regular aPAKEs)?
- Can we further improve the performance of password-based credential-hiding login, and show how  $\Pi_{\text{pwd}}$  performs when instantiated within a production-grade deployment of a server-authenticated channel along with a real-world password login form?
- In our work, we have shown that the constructions  $\Pi_{\text{pwd}}$  and  $\Pi_{\text{fz}}$  satisfy the notion of credential-hiding login. Can a simple modification of these constructions be shown to meet the security notions for UC aPAKEs and UC fuzzy aPAKEs, respectively?
- More generally, can we show black box reductions or separations among CHL and aPAKE functionalities?

## References

- [AMMR18] Shashank Agrawal, Payman Mohassel, Pratyay Mukherjee, and Peter Rindal. Distributed symmetric-key encryption. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 1993–2010, 2018.
- [AP05] Michel Abdalla and David Pointcheval. Simple password-based encrypted key exchange protocols. In *Cryptographers’ track at the RSA conference*, pages 191–208. Springer, 2005.
- [BBR88] Charles H Bennett, Gilles Brassard, and Jean-Marc Robert. Privacy amplification by public discussion. *SIAM journal on Computing*, 17(2):210–229, 1988.
- [BCF<sup>+</sup>14] Julien Bringer, Herve Chabanne, Melanie Favre, Alain Patey, Thomas Schneider, and Michael Zohner. Gshade: faster privacy-preserving distance computation and biometric identification. In *Proceedings of the 2nd ACM workshop on Information hiding and multimedia security*, pages 187–198, 2014.
- [BCS16] Dan Boneh, Henry Corrigan-Gibbs, and Stuart E. Schechter. Balloon hashing: A memory-hard function providing provable protection against sequential attacks. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016*, volume 10031 of *Lecture Notes in Computer Science*, pages 220–248, 2016.
- [BDK<sup>+</sup>05] Xavier Boyen, Yevgeniy Dodis, Jonathan Katz, Rafail Ostrovsky, and Adam Smith. Secure remote authentication using biometric data. In *annual international conference on the theory and applications of cryptographic techniques*, pages 147–163. Springer, 2005.
- [BDK16] Alex Biryukov, Daniel Dinu, and Dmitry Khovratovich. Argon2: New generation of memory-hard functions for password hashing and other applications. In *IEEE European Symposium on Security and Privacy, EuroS&P 2016, Saarbrücken, Germany, March 21-24, 2016*, pages 292–302. IEEE, 2016.
- [BF01] Dan Boneh and Matt Franklin. Identity-based encryption from the weil pairing. In *Annual international cryptology conference*, pages 213–229. Springer, 2001.

- [BLMR13] Dan Boneh, Kevin Lewi, Hart Montgomery, and Ananth Raghunathan. Key homomorphic prfs and their applications. In *Annual Cryptology Conference*, pages 410–428. Springer, 2013.
- [BM93a] Steven M Bellovin and Michael Merritt. Augmented encrypted key exchange: a password-based protocol secure against dictionary attacks and password file compromise. In *CCS*, pages 244–250, 1993.
- [BM93b] Steven M Bellovin and Michael Merritt. Augmented encrypted key exchange: a password-based protocol secure against dictionary attacks and password file compromise. In *Proceedings of the 1st ACM Conference on Computer and Communications Security*, pages 244–250, 1993.
- [BMP00] Victor Boyko, Philip MacKenzie, and Sarvar Patel. Provably secure password-authenticated key exchange using diffie-hellman. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 156–171. Springer, 2000.
- [Boy04] Xavier Boyen. Reusable cryptographic fuzzy extractors. In *Proceedings of the 11th ACM conference on Computer and communications security*, pages 82–91, 2004.
- [BPR00] Mihir Bellare, David Pointcheval, and Phillip Rogaway. Authenticated key exchange secure against dictionary attacks. In *International conference on the theory and applications of cryptographic techniques*, pages 139–155. Springer, 2000.
- [Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *FOCS*, pages 136–145. IEEE, 2001.
- [CFP<sup>+</sup>16a] Ran Canetti, Benjamin Fuller, Omer Paneth, Leonid Reyzin, and Adam Smith. Reusable fuzzy extractors for low-entropy distributions. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 117–146. Springer, 2016.
- [CFP<sup>+</sup>16b] Ran Canetti, Benjamin Fuller, Omer Paneth, Leonid Reyzin, and Adam D. Smith. Reusable fuzzy extractors for low-entropy distributions. In *EUROCRYPT*, pages 117–146, 2016.
- [CHK<sup>+</sup>05] Ran Canetti, Shai Halevi, Jonathan Katz, Yehuda Lindell, and Philip D. MacKenzie. Universally composable password-based key exchange. In *EUROCRYPT*, pages 404–421, 2005.
- [CK02] Ran Canetti and Hugo Krawczyk. Universally composable notions of key exchange and secure channels. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 337–351. Springer, 2002.
- [Coi19] Post Mortem: A closer look at a password storage issue affecting 3,420 customers. <https://blog.coinbase.com/post-mortem-a-closer-look-at-a-password-storage-issue-affecting-3-420-customers-e23cfc8a0363>, August 16, 2019.

- [DHP<sup>+</sup>18] Pierre-Alain Dupont, Julia Hesse, David Pointcheval, Leonid Reyzin, and Sophia Yakoubov. Fuzzy password-authenticated key exchange. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 393–424. Springer, 2018.
- [DORS08] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM journal on computing*, 38(1):97–139, 2008.
- [DS05] Yevgeniy Dodis and Adam Smith. Correcting errors without leaking partial information. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 654–663, 2005.
- [ECS<sup>+</sup>15] Adam Everspaugh, Rahul Chaterjee, Samuel Scott, Ari Juels, and Thomas Ristenpart. The Pythia PRF Service. In *USENIX Security Symposium*, pages 547–562, 2015.
- [EHKM11] David Evans, Yan Huang, Jonathan Katz, and Lior Malka. Efficient privacy-preserving biometric identification. In *Proceedings of the 17th conference Network and Distributed System Security Symposium, NDSS*, volume 68, 2011.
- [EHOR20] Andreas Erwig, Julia Hesse, Maximilian Orlt, and Siavash Riahi. Fuzzy asymmetric password-authenticated key exchange. *IACR Cryptol. ePrint Arch.*, 2020:987, 2020.
- [Eve13] Tammy Everts. Rules for mobile performance optimization. *Communications of the ACM*, 56:52–59, 08 2013.
- [Fac19] Keeping Passwords Secure. <https://about.fb.com/news/2019/03/keeping-passwords-secure/>, March 21, 2019.
- [Git18] GitHub says bug exposed some plaintext passwords. <https://www.zdnet.com/article/github-says-bug-exposed-account-passwords/>, May 1, 2018.
- [GKW17] Rishab Goyal, Venkata Koppula, and Brent Waters. Lockable obfuscation. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 612–621. IEEE, 2017.
- [GMR06] Craig Gentry, Philip D. MacKenzie, and Zulfikar Ramzan. A method for making password-based key exchange resilient to server compromise. In *Advances in Cryptology - CRYPTO 2006, 26th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 2006, Proceedings*, pages 142–159, 2006.
- [Goo19] Notifying administrators about unhashed password storage. <https://cloud.google.com/blog/products/g-suite/notifying-administrators-about-unhashed-password-storage>, May 21, 2019.
- [Gro20] Crypto Forum Research Group. *CFRG PAKE Selection Process*, 2019 (accessed February, 2020).
- [HEKM11] Yan Huang, David Evans, Jonathan Katz, and Lior Malka. Faster secure two-party computation using garbled circuits. In *USENIX Security Symposium*, volume 201, pages 331–335, 2011.

- [HL19] Björn Haase and Benoît Labrique. Aucpace: Efficient verifier-based pake protocol tailored for the iiot. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pages 1–48, 2019.
- [HR10] Feng Hao and Peter Ryan. J-pake: authenticated key exchange without pki. In *Transactions on computational science XI*, pages 192–206. Springer, 2010.
- [HS14] Feng Hao and Siamak F. Shahandashti. The speke protocol revisited. *Cryptology ePrint Archive*, Report 2014/585, 2014. <https://eprint.iacr.org/2014/585>.
- [JKX18] Stanislaw Jarecki, Hugo Krawczyk, and Jiayu Xu. Opaque: an asymmetric pake protocol secure against pre-computation attacks. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 456–486. Springer, 2018.
- [JR18] Charanjit S Jutla and Arnab Roy. Smooth nizk arguments. In *Theory of Cryptography Conference*, pages 235–262. Springer, 2018.
- [KM14] Franziskus Kiefer and Mark Manulis. Zero-knowledge password policy checks and verifier-based pake. In *European Symposium on Research in Computer Security*, pages 295–312. Springer, 2014.
- [Mac01] Philip Mackenzie. More efficient password-authenticated key exchange. In *Cryptographers’ Track at the RSA Conference*, pages 361–377. Springer, 2001.
- [MKR17] Kathleen Moriarty, Burt Kaliski, and Andreas Rusch. Pkcs# 5: Password-based cryptography specification version 2.1. *Internet Engineering Task Force (IETF)*, 2017.
- [MPS00] Philip MacKenzie, Sarvar Patel, and Ram Swaminathan. Password-authenticated key exchange based on rsa. In *International conference on the theory and application of cryptology and information security*, pages 599–613. Springer, 2000.
- [Muf15] Alec Muffet. Facebook: Password hashing & authentication. *Presentation at Real World Crypto*, 2015.
- [OPJM10] Margarita Osadchy, Benny Pinkas, Ayman Jarrous, and Boaz Moskovich. Scifi-a system for secure face identification. In *2010 IEEE Symposium on Security and Privacy*, pages 239–254. IEEE, 2010.
- [Per09] Colin Percival. Stronger key derivation via sequential memory-hard functions, 2009.
- [PJ16] Colin Percival and Simon Josefsson. The scrypt Password-Based Key Derivation Function. RFC 7914, August 2016.
- [PM99] Niels Provos and David Mazieres. A future-adaptable password scheme. In *USENIX Annual Technical Conference, FREENIX Track*, pages 81–91, 1999.
- [PW17] David Pointcheval and Guilin Wang. Vtbpeke: verifier-based two-basis password exponential key exchange. In *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security*, pages 301–312, 2017.

- [Rob19] Robinhood stored passwords in plaintext, so change yours now. <https://techcrunch.com/2019/07/24/robinhood-stored-passwords-in-plaintext-so-change-yours-now/>, July 24, 2019.
- [SKFB19] Nick Sullivan, Dr. Hugo Krawczyk, Owen Friel, and Richard Barnes. Usage of OPAQUE with TLS 1.3. Internet-Draft draft-sullivan-tls-opaque-00, Internet Engineering Task Force, March 2019. Work in Progress.
- [Twi18] Twitter says bug exposed user plaintext passwords. <https://www.zdnet.com/article/twitter-says-bug-exposed-passwords-in-plaintext/>, May 3, 2018.
- [WZ17] Daniel Wichs and Giorgos Zirdelis. Obfuscating compute-and-compare programs under lwe. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 600–611. IEEE, 2017.

# Appendix

## A UC Security Proof for $\Pi_{\text{pwd}}$

In the following, we use  $\mathcal{A}^*$  to denote the simulator described in Section 3.1. For an adversary  $\mathcal{A}$ , we define a series of games, represented as  $\text{Game}_i(\mathcal{A})$ :

- **Game 0:**  $\text{Game}_0$  is identical to  $\text{Expt}_{\text{idéal}}(\mathcal{A}, \mathcal{A}^*)$ .
- **Game 1:**  $\text{Game}_1$  differs from  $\text{Game}_0$  in the following: when  $\mathcal{A}$  calls the Registration Oracle,  $\mathcal{D}_{\text{Reg}}$  is populated as in the real world, and then when responding to the Server Login Oracle query:
  - In the abort condition of Case 2a, rather than aborting, the simulator checks  $\mathcal{D}_{\text{Reg}}$  for the corresponding registered credential  $x$  and instead sets  $\text{flag} \leftarrow \text{SUCCESS}$  and calls  $\mathcal{F}_{\text{CHL}}$  with  $(\text{TESTCRED}, \text{sid}, \text{ssid}, \text{uid}, x)$ . All other steps remain unchanged.
  - In Case 2b, rather than performing a linear search over the entries of  $\mathcal{D}_{\text{RO}}$ , the simulator checks  $\mathcal{D}_{\text{Reg}}$  for the corresponding registered credential  $x$  and computes  $(s, t) \leftarrow \mathcal{H}(\text{sid}, \text{uid}, x)$ , and  $(v, \text{ssid}', \text{tok}) \leftarrow \text{IBE.Decrypt}(\text{IBE.KeyGen}(\text{sk}, \text{id} = s), \alpha)$ . The remainder steps remain the same.
- **Game 2:**  $\text{Game}_2$  differs from  $\text{Game}_1$  in responding to Client Login Oracle queries: instead of computing  $\alpha \leftarrow \text{IBE.Encrypt}(\text{pp}, \text{id}, m)$  for randomly sampled  $\text{id}, m \leftarrow_{\text{R}} \{0, 1\}^\lambda$ , it computes  $\alpha \leftarrow \text{IBE.Encrypt}(\text{pp}, s, t)$ , where  $(s, t) \leftarrow \mathcal{H}(\text{sid}, \text{uid}, x)$ . Note that this is identical to returning  $\alpha \leftarrow \text{ClientLogin}(\text{pp}, \text{ssid}, \text{tok}^*, \text{uid}, x^*)$ .
- **Game 3:**  $\text{Game}_3$  differs from  $\text{Game}_2$  in the Server Login Oracle query response: in Case 2, the simulator checks  $\mathcal{D}_{\text{Reg}}$  for the corresponding registered verifier  $\gamma$ , and simply returns the result of computing  $\text{ServerLogin}(\text{pp}, \text{ssid}, \gamma, \alpha)$ .

In the following lemmas, for two experiments  $E_1(\mathcal{A})$  and  $E_2(\mathcal{A})$ , we use the notation  $E_1(\mathcal{A}) \approx E_2(\mathcal{A})$  as shorthand to represent the notion that the quantity  $|\Pr[E_1(\mathcal{A}) = 1] - \Pr[E_2(\mathcal{A}) = 1]|$  is negligible in  $\lambda$ .

**Lemma 1.** *For an adversary  $\mathcal{A}$ ,  $\text{Game}_0(\mathcal{A}) \approx \text{Game}_1(\mathcal{A})$  when  $\mathcal{H}$  is modeled after a random oracle.*

*Proof.* To handle the change in Case 2a, we show that the probability of  $f(v) = f(t)$  is negligible given that  $f$  is a random oracle. So the indistinguishability of Games 0 and 1 for this specific case follows statistically. The presence of  $((\text{sid}, \text{uid}, *), (s, t))$  in  $\mathcal{D}_{\text{RO}}$  indicates that the password file was stolen. Now note that if we model  $f$  as a random oracle, then  $t$  and  $f(t)$  are independently random quantities. In case the adversary never does any successful online or offline attempt previously, it does not receive any more information about  $t$ , other than  $f(t)$  provided in the stolen password file. So there is a negligible chance of the event  $v = t$  occurring and hence also of  $f(v) = f(t)$ .

To handle the change in Case 2b, indistinguishability follows by again using the random oracle properties of  $\mathcal{H}$  and the collision resistance property of  $f$ . We need two key properties in order to ensure the consistency of the trial and error search over the entries of  $\mathcal{D}_{\text{RO}}$ :

Let  $(s, t) = \mathcal{H}(\text{sid}, \text{uid}, x)$  and  $(s', t') = \mathcal{H}(\text{sid}', \text{uid}', x')$  for any  $x \neq x'$  and identifiers  $\text{sid}, \text{sid}' \in \mathcal{I}_2$  and  $\text{uid}, \text{uid}' \in \mathcal{U}$ , and tokens  $\text{tok}, \text{tok}' \in \{0, 1\}^\lambda$ . Then, we argue that for any  $\alpha$ , the following holds with statistically overwhelming probability:

$$\begin{aligned} & \text{IBE.Decrypt}(\text{IBE.KeyGen}(\text{sk}, s), \alpha) = (t, \text{ssid}, \text{tok}) \\ \implies & \neg(\text{IBE.Decrypt}(\text{IBE.KeyGen}(\text{sk}, s'), \alpha) = (t', \text{ssid}, \text{tok}') \wedge f(t') = f(t)) \end{aligned}$$

This easily follows, as  $x \neq x'$  and  $\mathcal{H}$  is a random oracle, which makes  $t'$  independently random from  $s, s', t$ .

We also argue that if the random oracle  $\mathcal{H}$  is not called with  $(\text{sid}, \text{uid}, x)$ , then the following event is statistically negligible in probability:

$$\text{IBE.Decrypt}(\text{IBE.KeyGen}(\text{sk}, s), \alpha) = (t, \text{ssid}, \text{tok})$$

This is because  $t$  is independently random of the adversary's view if the random oracle is not called with the correct input. This concludes the proof.  $\square$

Note that if  $f$  were only a CROWF (Collision-Resistant OWF), the proof of Lemma 1 would not hold. Given a one-wayness challenge to invert  $u$ , we can implicitly set  $f(t) = u$  for a randomly chosen  $(\text{sid}, \text{uid}, x)$  registration. If the IBE decryption of an adversarially generated client login message yields a  $t'$ , such that  $f(t') = u$ , then this provides the one-wayness response. However, consider the case, where an offline test takes place with the correct password in between these two events. In this case, the simulator cannot provide a valid response without inverting the challenge itself. Hence, provability seems difficult just assuming CROWF.

**Lemma 2.** *For an adversary  $\mathcal{A}$ ,  $\text{Game}_1(\mathcal{A}) \approx \text{Game}_2(\mathcal{A})$  based on the IND-anon-ID-CCA security of IBE.*

*Proof.* We show that Indistinguishability follows by IND-anon-ID-CCA property of the IBE scheme in Lemma 2. In the following proof, we say that  $\text{sid}$  is marked with a type  $\text{SetupType} \in \{\text{HIDEREGKEY}, \text{REVEALREGKEY}\}$  to represent that there exists an entry in  $\mathcal{D}_{\text{key}}$  of the form  $(\text{sid}, \star, \star, \text{SetupType})$ . We describe the simulator acting as an adversary for  $\text{Expt}_b^{\text{anon-cca}}$  as follows:

- **On receiving Setup from S.** Let the call parameter from S be  $(\text{sid}, \text{type})$ . The simulator calls  $\text{IBE.Setup}(1^\lambda)$  and receives  $\text{pp}$ . If  $\text{type} = \text{HIDEREGKEY}$ , then send  $\text{pp}$  to  $\mathcal{A}$ . Mark  $\text{sid}$  with  $\text{HIDEREGKEY}$ . If  $\text{type} = \text{REVEALREGKEY}$ , then the simulator ignores the IBE challenger and generates  $(\text{pp}, \text{sk}) \leftarrow \text{IBE.Setup}(1^\lambda)$  itself and sends  $(\text{pp}, \text{sk})$  to  $\mathcal{A}$ . Mark  $\text{sid}$  with  $\text{REVEALREGKEY}$ .
- **On receiving StealCredFile from  $\mathcal{A}$ .** This is same as  $\text{Game}_2$ , except it calls the  $\text{IBE.KeyGen}$  oracle provided by the challenger in order to construct its responses.
- **On receiving ClientLogin from U:** Let the call parameter from U be  $(\text{sid}, \text{ssid}, \text{uid}, \text{type})$ . If  $\text{sid}$  is marked with  $\text{REVEALREGKEY}$ , or if  $\mathcal{D}_{\text{status}}[\text{sid}, \text{uid}] = \text{STOLEN}$ , then the simulator sends  $\perp$  to the adversary. Otherwise, the simulator samples  $(\text{id}, \text{m})$  from random and computes  $(s, t) = \mathcal{H}(\text{sid}, \text{uid}, x')$  using the random oracle, and sends  $((\text{id}, \text{m}), (s, (t, \text{ssid}, \text{tok})))$  to the IBE Encryption oracle, obtaining  $\alpha$ . Note that this step is not performed if the registration key was revealed. It then sends  $(\text{sid}, \text{ssid}, \alpha)$  to  $\mathcal{A}$ . If  $\text{type} = \text{INVALID}$  or  $\text{type} = \text{CORRECT}$ , then set  $\mathcal{D}_{\text{status}}[\text{sid}, \text{uid}] \leftarrow \text{RESTRICTED}$ .



- **On receiving a Server Login Oracle query from S.** This is same as Game 2, except it uses the IBE Decryption oracle to respond to the query.

Now we argue that the admissibility conditions for the CCA security game are obeyed. Firstly, due to the conditions in the Server Login Oracle response, decryption is never called on messages responded with by the encryption queries. Secondly, due to the RESTRICTED conditions, the IBE KeyGen oracle is never called for  $(uid, x)$  tuples for which Client Login was simulated. Lastly, Client Login is never simulated if the registration secret was revealed.

Since  $\text{Expt}_0^{\text{anon-cca}}$  corresponds to  $\text{Game}_1$  and  $\text{Expt}_1^{\text{anon-cca}}$  corresponds to  $\text{Game}_2$ , the lemma follows.  $\square$

**Lemma 3.** *For an adversary  $\mathcal{A}$ ,  $\text{Game}_2(\mathcal{A}) \approx \text{Game}_3(\mathcal{A})$  when  $\mathcal{H}$  is modeled after a random oracle.*

*Proof.* The only difference between the two games is the way the following situation is handled:  $(sid', ssid', uid', \alpha')$  was sent by an honest client with associated type  $\text{type}$ . Let's divide the cases exhaustively as follows.

First consider the case  $(sid', uid', \text{type}) = (sid, uid, \text{CORRECT})$ . In this case, the decryption will be with the same id as the encryption. Now due to IBE correctness, login will succeed in Game 3 if the ssid's are same, and will fail if not. This behavior is same as in Game 2.

Second, where  $(sid', uid', \text{type}) \neq (sid, uid, \text{CORRECT})$ , we will have a different decryption id than was used for encryption. In this case login cannot succeed without  $f(t') = f(t)$  holding. This is either a collision for  $f$ , or we have  $t' = t$ . The probability of the former is bounded by the collision-resistance of  $f$  and the probability of the later is statistically negligible as the RHS is information theoretically independent of the LHS due to the random oracle. Indistinguishability thus holds by the collision resistance of  $f$  and the random oracle assumption on  $\mathcal{H}$ .  $\square$

Note that  $\text{Game}_3$  is identical to  $\text{Expt}_{\text{real}}(\mathcal{A})$ . Thus, putting together Lemmas 1, 2, and 3 concludes the proof of Theorem 1.

## B UC Security Proof for $\Pi_{fz}$

In the following proof, we will consider a slight extension of  $\text{Expt}_{\text{real}}$  and  $\text{Expt}_{\text{ideal}}$  to support an “obfuscator oracle”, in order to incorporate the use of the MBCC obfuscator — we will denote these experiments as  $\text{Expt}_{\text{real}}^*$  and  $\text{Expt}_{\text{ideal}}^*$ . Intuitively, this extension allows for the adversary to query an oracle that provides black-box access to the program associated with the credential verifier stored upon a call to the registration oracle in the real world. Thus, instead of receiving the credential verifier upon a successful call to the Steal Credential Verifier Oracle, the adversary must instead issue queries to the Obfuscator Oracle to receive information about the credential verifier.

For an input adversary  $\mathcal{A}$ , we define a series of games, represented as  $\text{Game}_i(\mathcal{A})$ :

- **Game 0:**  $\text{Game}_0$  is identical to  $\text{Expt}_{\text{ideal}}^*(\mathcal{A}, \text{Sim})$ .
- **Game 1:**  $\text{Game}_1$  differs from  $\text{Game}_0$  in responding to the Steal Credential Verifier Oracle query: after  $\text{Sim}$  picks  $s \leftarrow_{\mathcal{R}} \{0, 1\}^\lambda$ , it does not set any entries in  $\mathcal{D}_{\text{RO}}$ . Then, in response to the Obfuscator Oracle query, if the response is “CORRECT GUESS”,  $\text{Sim}$  picks a random  $r \leftarrow_{\mathcal{R}} \{0, 1\}^\lambda$  and sets  $s \leftarrow \mathcal{H}(\text{uid}, r)$ . Finally, on responses to the random oracle query,  $\text{Sim}$  simply returns  $\mathcal{H}(\text{uid}, r)$ . All other steps remain unchanged.

- **Game 2:**  $\text{Game}_2^{(i)}$  differs from  $\text{Game}_1$  in responding to the Steal Credential Verifier Oracle query: if it receives a credential  $y$  from  $\mathcal{F}_{\text{CHL}}$ , then  $\text{Sim}$  computes  $(r^*, \sigma) \leftarrow \text{Gen}(y)$  and sets  $s \leftarrow \mathcal{H}(\text{uid}, r^*)$ . Then, in response to the Obfuscator Oracle query, if the response is “CORRECT GUESS”, then  $\text{Sim}$  computes  $(r^*, \sigma) \leftarrow \text{Gen}(x)$  and sets  $s \leftarrow \mathcal{H}(\text{uid}, r^*)$ . All other steps remain unchanged.
- **Game 3:** Let  $Q$  be the total number of Client Login Oracle queries made by  $\mathcal{A}$ . For each  $i \in [Q]$ , we define  $\text{Game}_3^{(i)}$  to be the following:  $\text{Game}_3^{(i)}$  differs from  $\text{Game}_2$  in the Client Login Oracle’s construction of the value  $\alpha$ . Recall that in  $\text{Game}_2$ ,  $\alpha \leftarrow (\alpha_1, \alpha_2) = (\text{PKE.Encrypt}(\text{pp}_{\text{pke}}, x), \text{IBE.Encrypt}(\text{pp}_{\text{ibe}}, (\text{ssid}, \text{tok}, \alpha_1), \text{uid}))$ . In  $\text{Game}_3^{(i)}$ , we instead have, for the first  $i$  queries to the Client Login Oracle,  $\alpha_1^* \leftarrow \text{PKE.Encrypt}(\text{pp}_{\text{pke}}, x^*)$ , where  $x^*$  is the credential part of the query sent to  $\mathcal{F}_{\text{CHL}}$ , and the oracle returns  $\alpha = (\alpha_1^*, \alpha_2)$  and sets  $\mathcal{D}_{\text{PKE}}[\alpha_1^*] \leftarrow x^*$  (instead of setting  $\mathcal{D}_{\text{PKE}}[\alpha_1] \leftarrow x$ ). We define  $\text{Game}_3^{(Q)} = \text{Game}_3$ .
- **Game 4:** Let  $Q$  be the total number of Client Login Oracle queries made by  $\mathcal{A}$ . For each  $i \in [Q]$ , we define  $\text{Game}_4^{(i)}$  to be the following:  $\text{Game}_4^{(i)}$  differs from  $\text{Game}_3$  in the Client Login Oracle’s construction of  $\alpha$ : for the first  $i$  queries to the Client Login Oracle,  $\alpha^* \leftarrow (\alpha_1^*, \alpha_2^*) = (\text{PKE.Encrypt}(\text{pp}_{\text{pke}}, x^*), \text{IBE.Encrypt}(\text{pp}_{\text{ibe}}, (\text{ssid}, \text{tok}, \alpha_1^*), \text{uid}))$ . In other words, we have that  $\alpha^* \leftarrow \Pi_{\text{fz}}.\text{ClientLogin}(\text{pp}, \text{ssid}, \text{tok}, \text{uid}, x)$ . We define  $\text{Game}_4^{(Q)} = \text{Game}_4$ .

Note that, by construction,  $\text{Game}_4$  is identical to  $\text{Expt}_{\text{real}}^*(\mathcal{A})$ . In the following lemmas, for two experiments  $E_1(\mathcal{A})$  and  $E_2(\mathcal{A})$ , we use the notation  $E_1(\mathcal{A}) \approx E_2(\mathcal{A})$  as shorthand to represent that  $|\Pr[E_1(\mathcal{A}) = 1] - \Pr[E_2(\mathcal{A}) = 1]|$  is negligible in  $\lambda$ .

**Lemma 4.** *For an adversary  $\mathcal{A}$ ,  $\text{Game}_0(\mathcal{A}) \approx \text{Game}_1(\mathcal{A})$  when  $\mathcal{H}$  is modeled after a random oracle.*

*Proof.* Since the random oracle outputs of  $\mathcal{H}$  are distributed uniformly at random, the only way in which an adversary can distinguish  $\text{Game}_0$  from  $\text{Game}_1$  is if he happens to query the Random Oracle in between a Steal Credential Verifier Oracle and Obfuscator Oracle query on an input  $\text{uid}$  and  $r' \in \{0, 1\}^\lambda$  for which  $r'$  is also selected in the Obfuscator Oracle query response. However, this probability is negligible in  $\lambda$  given that the  $r$  input is selected uniformly at random in the Obfuscator Oracle query response.  $\square$

**Lemma 5.** *For an adversary  $\mathcal{A}$ ,  $\text{Game}_1(\mathcal{A}) \approx \text{Game}_2(\mathcal{A})$ , based on the security of the reusable fuzzy extractor.*

*Proof.* For an adversary  $\mathcal{A}$ , we construct a simulator  $\mathcal{B}$  which simulates  $\text{Game}_1$  and  $\text{Game}_2$  by interacting with the challenger of  $\text{Expt}_b^{\text{FE}}$  as follows. On each Steal Credential Verifier Oracle and Obfuscator Oracle query from  $\mathcal{A}$ ,  $\mathcal{B}$  simulates the query normally, except that when producing  $(r, \sigma)$ , obtains the corresponding pair from the challenger of  $\text{Expt}_b^{\text{FE}}$  instead. Note that  $\text{Expt}_0^{\text{FE}}$  corresponds to  $\text{Game}_1$ , and  $\text{Expt}_1^{\text{FE}}$  corresponds to  $\text{Game}_2$ . Hence, the claim follows.  $\square$

Recall by definition that  $\text{Game}_2 = \text{Game}_3^{(0)}$ .

**Lemma 6.** *For an adversary  $\mathcal{A}$ , with  $i \in [Q]$ ,  $\text{Game}_3^{(i-1)}(\mathcal{A}) \approx \text{Game}_3^{(i)}(\mathcal{A})$ , based on the CCA security of PKE.*

*Proof.* For a bit  $b$ , let  $\mathcal{B}$  be an adversary that participates in  $\text{Expt}_b^{\text{CCA}}$  against a challenger. The challenger begins by computing  $(\text{pp}_{\text{pke}}, \text{sk}_{\text{pke}}) \leftarrow \text{PKE.Setup}(1^\lambda)$  and returns  $\text{pp}_{\text{pke}}$  to  $\mathcal{B}$ . Note that for each Setup Oracle query, if  $\text{type} = \text{REVEALREGKEY}$  for the  $i^{\text{th}}$  Setup Oracle query, then  $\text{Game}_3^{(i-1)}$  and  $\text{Game}_3^{(i)}$  are identical, concluding the proof. If  $\text{type} = \text{HIDEREGKEY}$ ,  $\mathcal{B}$  returns  $\text{pp}_{\text{pke}}$  to  $\mathcal{A}$  as the response to the Setup Oracle query.

Then, for the  $i^{\text{th}}$  Client Login Oracle query,  $\mathcal{B}$  submits two challenge messages  $x$  and  $x^*$  to the challenger, who then returns a ciphertext  $c$ . For each Server Login Oracle query,  $\mathcal{B}$  makes a decryption oracle query to the challenger in order to obtain an output  $x'$  (in place of calling  $\text{PKE.Decrypt}$  by itself). Note that by construction,  $\text{Sim}$  only invokes  $\text{PKE.Decrypt}$  if the input to the Server Login oracle differs from every ciphertext produced in response to the Client Login Oracle query. Similarly, during a response to the Obfuscation Oracle,  $\text{Sim}$  only calls  $\text{PKE.Decrypt}$  on a ciphertext that was not produced by the Client Login Oracle query. Hence, in both types of queries, the admissibility condition of  $\text{Expt}_b^{\text{CCA}}$  is satisfied. Since  $\text{Expt}_0^{\text{CCA}}$  corresponds to  $\text{Game}_3^{(i-1)}$  and  $\text{Expt}_1^{\text{CCA}}$  corresponds to  $\text{Game}_3^{(i)}$ , the lemma follows.  $\square$

Recall by definition that  $\text{Game}_3^{(Q)} = \text{Game}_3 = \text{Game}_4^{(0)}$ .

**Lemma 7.** *For an adversary  $\mathcal{A}$ , with  $i \in [Q]$ ,  $\text{Game}_4^{(i-1)}(\mathcal{A}) \approx \text{Game}_4^{(i)}(\mathcal{A})$ , based on the ID-CCA security of IBE.*

*Proof.* For a bit  $b$ , let  $\mathcal{B}$  be an adversary that participates in  $\text{Expt}_b^{\text{IDCCA}}$  against a challenger. The challenger begins by computing  $(\text{pp}_{\text{ibe}}, \text{sk}_{\text{ibe}}) \leftarrow \text{IBE.Setup}(1^\lambda)$  and returns  $\text{pp}_{\text{ibe}}$  to  $\mathcal{B}$ . Note that for each Setup Oracle query, if  $\text{type} = \text{REVEALREGKEY}$  for the  $i^{\text{th}}$  Setup Oracle query, then  $\text{Game}_4^{(i-1)}$  and  $\text{Game}_4^{(i)}$  are identical, concluding the proof. If  $\text{type} = \text{HIDEREGKEY}$ ,  $\mathcal{B}$  returns  $\text{pp}_{\text{pke}}$  to  $\mathcal{A}$  as the response to the Setup Oracle query.

Then, for the  $i^{\text{th}}$  Client Login Oracle query,  $\mathcal{B}$  submits the challenge id  $\text{uid}^*$  and two challenge messages  $\text{tok}$  and  $\text{tok}^*$  to the challenger, who then returns a ciphertext  $c$ . For each Server Login Oracle query, after checking its preconditions,  $\mathcal{B}$  makes a decryption oracle query to the challenger in order to obtain an output  $\text{tok}'$  (in place of calling  $\text{IBE.KeyGen}$  followed by  $\text{IBE.Decrypt}$  by itself). In order to assert the admissibility of these queries, there are two properties to check:

- $\text{IBE.KeyGen}$  is called on the challenge id  $\text{uid}^*$ . Note that the only opportunity for  $\text{IBE.KeyGen}$  to be called is during a Steal Credential Verifier Oracle query. This must also happen *before* calling the Client Login Oracle (when the challenge messages are submitted to the challenger). In this case, note that  $\mathcal{D}_{\text{status}}[\text{sid}, \text{uid}^*]$  is set to  $\text{RESTRICTED}$ , and there are two cases to consider based on the value of  $\mathcal{C}(x, y)$  in the Client Login Oracle response. If  $\mathcal{C}(x, y) = 1$ , then the Client Login Oracle simulation aborts, and by construction  $\text{Game}_4^{(i-1)}$  and  $\text{Game}_4^{(i)}$  are identical. If  $\mathcal{C}(x, y) \neq 1$ , then in response to the Server Login Oracle,  $\mathcal{B}$  can simply set the appropriate parameters to  $\perp$  instead of submitting the decryption oracle query to the challenger. Hence, in either case, we have shown that the admissibility conditions for the IBE CCA game have not been violated.
- $\text{IBE.KeyGen}$  is not called on the challenge id  $\text{uid}^*$ . There are two subcases to consider based on the  $\alpha$  component of the query sent to the Server Login Oracle. If there is already an entry of the form  $(\star, \star, \star, \alpha)$  in  $\mathcal{D}_{\text{sent}}$ , then  $\mathcal{B}$  can complete the simulation by setting the appropriate parameters to  $\perp$ . If there is no such entry in  $\mathcal{D}_{\text{sent}}$ , then, for  $(\alpha_1, \alpha_2) = \alpha$ , we need only

consider cases where  $\alpha_2$  previously appeared in  $\mathcal{D}_{\text{sent}}$ . For each of these cases, the  $\alpha_1$  value must be such that when decrypting  $\alpha_2$ , the output of the decryption will not match either  $\text{ssid}$  or  $\alpha_1$ . Hence, it suffices for the simulator in these cases to set the appropriate parameters to  $\perp$ .

We have shown that the simulator obeys the admissibility conditions for  $\text{Expt}_b^{\text{IDCCA}}$ . Since  $\text{Expt}_0^{\text{IDCCA}}$  corresponds to  $\text{Game}_4^{(i-1)}$  and  $\text{Expt}_1^{\text{IDCCA}}$  corresponds to  $\text{Game}_4^{(i)}$ , the lemma follows.  $\square$

Now, by tying together the above lemmas, we have shown that  $\text{Expt}_{\text{ideal}}^*(\mathcal{A}, \text{Sim})$  is indistinguishable from  $\text{Expt}_{\text{real}}^*(\mathcal{A})$ , as formalized in the below lemma.

**Lemma 8.** *For any adversary  $\mathcal{A}$ , with  $\text{Sim}$  defined as the simulator from Section 5.3, we have that  $\text{Expt}_{\text{real}}^*(\mathcal{A}) \approx \text{Expt}_{\text{ideal}}^*(\mathcal{A}, \text{Sim})$ .*

*Proof.* Follows directly from applying Lemmas 4, 5, 6, and 7.  $\square$

## B.1 Applying the Obfuscator

Recall from Definition 5 that in order to apply a distributional VBB obfuscator to a program  $\text{MBCC}[f, y, z]$ , it must be the case that  $y$  is computationally unpredictable, even in the presence of  $f$  and  $z$ .

**Lemma 9.** *For every adversary  $\mathcal{A}$ , there exists a simulator  $\mathcal{A}^*$  such that  $\text{Expt}_{\text{real}}(\mathcal{A}) \approx \text{Expt}_{\text{real}}^*(\mathcal{A}^*)$ , based on the security of the distributional VBB obfuscator and the reusable fuzzy extractor and the assumption that  $\mathcal{X}$  has min-entropy at least  $\lambda$ .*

*Proof.* Defining  $f$ ,  $y$ , and  $z$  to be of the form in which each of the programs obfuscated  $\text{Expt}_{\text{ideal}}$  are of the form  $\text{MBCC}[f, y, z]$ , note that the only term that is not truly independent of  $y = \mathcal{H}(\text{uid}, z)$  is  $\sigma$  (contained in the description of  $f$ ).

However, under the assumption that  $(\text{Gen}, \text{Rep})$  is a reusable fuzzy extractor, and that the inputs  $y$  are drawn from  $\mathcal{X}$  with sufficient min-entropy, since  $r$  is produced from  $\text{Gen}(y)$ , we can conclude that each  $r$  produced in this way also has entropy at least  $\lambda$ . This allows us to apply the distributional VBB obfuscator property as defined in Definition 5 on each obfuscation, which yields the existence of such a simulator  $\mathcal{A}^*$  that satisfies the claim.  $\square$

**Lemma 10.** *For every adversary  $\mathcal{A}$ , there exists a simulator  $\mathcal{A}^*$  and simulator  $\text{Sim}$  such that  $\text{Expt}_{\text{ideal}}(\mathcal{A}, \text{Sim}) \approx \text{Expt}_{\text{ideal}}^*(\mathcal{A}^*, \text{Sim})$ , based on the security of the distributional VBB obfuscator, the reusable fuzzy extractor, and the assumption that  $\mathcal{X}$  has min-entropy at least  $\lambda$ .*

*Proof.* This also follows similarly to the proof of Lemma B.1 except even more straightforwardly, since the inputs  $s$  to each obfuscated function are drawn independently and uniformly at random.  $\square$

Finally, by observing that for any adversary  $\mathcal{A}$ , we have that there exists a simulator  $\mathcal{A}^*$  and  $\text{Sim}$  for which  $\text{Expt}_{\text{real}}(\mathcal{A}) \approx \text{Expt}_{\text{real}}^*(\mathcal{A}^*)$  (from Lemma B.1),  $\text{Expt}_{\text{real}}^*(\mathcal{A}^*) \approx \text{Expt}_{\text{ideal}}^*(\mathcal{A}^*, \text{Sim})$  (from Lemma 8), and  $\text{Expt}_{\text{ideal}}^*(\mathcal{A}^*) \approx \text{Expt}_{\text{ideal}}(\mathcal{A}, \text{Sim})$ . Putting these together, we have that  $\text{Expt}_{\text{real}}(\mathcal{A}) \approx \text{Expt}_{\text{ideal}}(\mathcal{A}, \text{Sim})$ , which concludes the proof of Theorem 2.

<p><u>Session</u></p> <ul style="list-style-type: none"> <li>• On <math>(\text{CLIENTLOGIN}, \text{sid}, \text{ssid}, \text{uid}, \text{tok})</math> from <math>\mathcal{U}</math>, if this is the first <math>\text{CLIENTLOGIN}</math> message for <math>(\text{sid}, \text{ssid}, \text{uid})</math>, record <math>\langle \text{CLIENT}, \text{sid}, \text{ssid}, \text{uid}, \text{tok} \rangle</math> and mark it <math>\text{FRESH}</math>. send <math>(\text{CLIENTLOGIN}, \text{sid}, \text{ssid}, \text{uid})</math> to <math>\mathcal{A}^*</math>.</li> <li>• On <math>(\text{SERVERLOGIN}, \text{sid}, \text{ssid}, \text{uid})</math> from <math>\mathcal{S}</math>, if this is the first <math>\text{SERVERLOGIN}</math> message for <math>(\text{sid}, \text{ssid}, \text{uid})</math>, record <math>\langle \text{SERVER}, \text{sid}, \text{ssid}, \text{uid} \rangle</math> and mark it <math>\text{FRESH}</math>. Send <math>(\text{SERVERLOGIN}, \text{sid}, \text{ssid}, \text{uid})</math> to <math>\mathcal{A}^*</math>.</li> <li>• On <math>(\text{INTERRUPT}, \text{sid}, \text{ssid}, \text{uid})</math> from <math>\mathcal{A}^*</math>, if this is the first <math>\text{INTERRUPT}</math> message for <math>(\text{sid}, \text{ssid}, \text{uid})</math>, mark <math>\langle \text{SERVER}, \text{sid}, \text{ssid}, \text{uid} \rangle</math> as <math>\text{INTERRUPTED}</math>.</li> </ul> <p><u>Login and Authentication</u></p> <ul style="list-style-type: none"> <li>• On <math>(\text{LOGIN}, \text{sid}, \text{ssid}, \text{uid}, \text{tok}', \text{flag}')</math> from <math>\mathcal{A}^*</math>, if there is a record <math>\langle \text{SERVER}, \text{sid}, \text{ssid}, \text{uid} \rangle</math>, then mark it <math>\text{COMPLETED}</math> and do the following: <ul style="list-style-type: none"> <li>– If the client <math>\text{uid}</math> is corrupted, then set <math>(\text{tok}, \text{flag}) = (\text{tok}', \text{flag}')</math>.</li> <li>– Otherwise if there is a record <math>\langle \text{CLIENT}, \text{sid}, \text{ssid}, \text{uid}, \text{tok} \rangle</math> marked <math>\text{FRESH}</math>, then mark it <math>\text{COMPLETED}</math>. If the server record <math>\langle \text{SERVER}, \text{sid}, \text{ssid}, \text{uid} \rangle</math> is <math>\text{INTERRUPTED}</math>, then set <math>(\text{tok}, \text{flag}) = (\perp, \text{FAILURE})</math>. If the server record <math>\langle \text{SERVER}, \text{sid}, \text{ssid}, \text{uid} \rangle</math> is <math>\text{FRESH}</math>, then don't modify <math>\text{tok}</math> and set <math>\text{flag} = \text{SUCCESS}</math>.</li> </ul> </li> </ul> <p>Output <math>(\text{sid}, \text{ssid}, \text{uid}, \text{tok}, \text{flag})</math> to <math>\mathcal{S}</math>.</p>
---

Figure 2: The ideal functionality  $\mathcal{F}_{\text{AL}}$  for Authenticated Login.

## C Authenticated Login

We formally define the UC authenticated login functionality  $\mathcal{F}_{\text{AL}}$  in Figure 2. This is essentially a simplified version of  $\mathcal{F}_{\text{CHL}}$ , with the removal of password/credential verification and the credential stealing interfaces. As such, it is the analogue of key exchange in the same sense as CHL is analogous to aPAKE.

Authenticated Login carries some of the distinctions that CHL has, compared to aPAKE. For example, the environment passes a token to the client and expects the same token to be output by the server, when the protocol is correctly executed. Also, an  $\text{INTERRUPT}$  interface is needed for authenticated login, but not in key exchange. This is because, in key exchange, if the server simulation gets a modified message, the simulator does not call  $\text{NEWKEY}$ , effectively aborting the session. This is the same as in the real world, where there is no explicit abort. In authenticated login, if the server simulation gets a modified message, then it must signal a login failure to the environment; this is the purpose of the  $\text{INTERRUPT}$  functionality. The  $\text{INTERRUPT}$  interface is the equivalent of the  $\text{TESTCRED}$  interface in CHL, except there is no low entropy credential to test.

### C.1 PKI-based Authenticated Login Protocol

We describe a PKI-based protocol  $\Pi_{\text{AL}}$  in Figure 3, and show that it securely realizes  $\mathcal{F}_{\text{AL}}$ . This protocol is a straightforward application of a public-key encryption PKE and signature scheme  $\text{Sig}$  to ensure confidential and authenticated transport from the client to the server. Given that the client is required to have a signature certified by a certificate authority, this protocol clearly provides redundant security guarantees when integrated within an already-authenticated TLS channel. However,  $\Pi_{\text{AL}}$  illustrates the basic concepts around the authenticated login primitive itself.

<p><u>Setup</u>  Sample <math>(vk_{uid}, sk_{uid}) \leftarrow \text{Sig.Setup}(1^\lambda)</math> for all the client <math>uid</math>'s. Sample <math>(ek, dk) \leftarrow \text{PKE.Setup}(1^\lambda)</math> for the Server.</p> <p><u>Session</u></p> <ul style="list-style-type: none"> <li>• On <math>(\text{CLIENTLOGIN}, sid, ssid, uid, tok)</math> from <math>U</math>, if this the first <math>\text{CLIENTLOGIN}</math> message for <math>(sid, ssid, uid)</math>, then sample <math>enc \leftarrow \text{PKE.Encrypt}(ek, tok)</math> and <math>sig \leftarrow \text{Sig.Sign}(sk_{uid}, (sid, ssid, enc))</math>, and send <math>(sid, ssid, uid, enc, sig)</math> to <math>\mathcal{A}</math>.</li> <li>• On <math>(\text{SERVERLOGIN}, sid, ssid, uid)</math> from <math>S</math>, if this is the first <math>\text{SERVERLOGIN}</math> message for <math>(sid, ssid, uid)</math>, receive <math>(sid, ssid, uid, enc', sig')</math> from <math>\mathcal{A}</math>. If <math>\text{Sig.Verify}(vk_{uid}, (sid, ssid, sig'))</math> succeeds, then compute <math>tok' = \text{PKE.Decrypt}(dk, enc')</math> and set <math>flag' = \text{SUCCESS}</math>. Otherwise set <math>(tok', flag') = (\perp, \text{FAILURE})</math>. Output <math>(sid, ssid, uid, tok', flag')</math> to <math>S</math>.</li> </ul>
---

Figure 3: PKI-based Authenticated Login protocol  $\Pi_{AL}$ .

**Theorem 3.** *The protocol  $\Pi_{AL}$  realizes  $\mathcal{F}_{AL}$ , provided PKE satisfies IND-CPA security and Sig is existentially unforgeable.*

*Proof.* We first describe the simulator  $\mathcal{A}^*$  below:

- **Setup:** The simulator generates signature key pairs  $(vk_{uid}, sk_{uid})$  and encryption key pairs  $(ek, dk)$ .
- **Client Login:** The simulator receives  $(\text{CLIENTLOGIN}, sid, ssid, uid)$  from  $U$ . Sample  $enc \leftarrow \text{PKE.Encrypt}(ek, 0)$  and  $sig \leftarrow \text{Sig.Sign}(sk_{uid}, (sid, ssid, enc))$ . Send  $(sid, ssid, uid, enc, sig)$  to  $\mathcal{A}$ .
- **Server Login:** The simulator receives  $(\text{SERVERLOGIN}, sid, ssid, uid)$  from  $S$ . Then it waits to accept a message  $(sid, ssid, uid, enc', sig')$ . If this is same as the corresponding sent message, then just call  $(\text{LOGIN}, sid, ssid, uid, \perp, \perp)$ . If it is not the same message that was sent, then the simulator calls  $(\text{INTERRUPT}, sid, ssid, uid)$  first and then calls  $(\text{LOGIN}, sid, ssid, uid, \perp, \text{FAILURE})$ .

In order to prove that this simulation is indistinguishable to the adversary and the environment from the real protocol execution, note that the adversary only sees encryptions of 0, which are indistinguishable from the hybrid execution due to the semantic security of PKE. Also, the adversary cannot construct fake messages purporting to come from honest parties which pass signature verification, due to the existential unforgeability of the Sig scheme.  $\square$

## C.2 Secure Channel from Authenticated Login

In this section, we provide the UC secure channel functionality  $\mathcal{F}_{SC}$  in Figure 4, which is adapted from [CK02]. We then present in Figure 5 the protocol  $\Pi_{SC}$  in the  $\mathcal{F}_{AL}$ -hybrid model and show that it securely realizes  $\mathcal{F}_{SC}$ .

**Theorem 4.** *The protocol  $\Pi_{SC}$  realizes  $\mathcal{F}_{SC}$  in the  $\mathcal{F}_{AL}$ -hybrid model, provided that PKE satisfies IND-CCA security and MAC is unforgeable.*

*Proof.* We describe the simulator steps as follows:

- Upon receiving  $\mathcal{F}_{SC}(\text{NEWSESSION}, sid, U, S, \text{INITIATOR})$  from  $U$ , the simulator samples encryption and mac keys  $tok = (k_e, k_a)$  and calls  $\mathcal{F}_{AL}$  with  $(\text{CLIENTLOGIN}, sid, U, tok)$ .

$\mathcal{F}_{\text{SC}}$  proceeds as follows, running with parties  $U$  and  $S$  and  $\mathcal{A}^*$ :

- On  $(\text{NEWSESSION}, \text{sid}, U, S, \text{INITIATOR})$  from  $S$ , send  $(\text{NEWSESSION}, \text{sid}, U, S, \text{INITIATOR})$  to  $\mathcal{A}^*$ , and wait to receive a value  $(\text{NEWSESSION}, \text{sid}, S, U, \text{RESPONDER})$  from  $U$ . Once this value is received, set a boolean variable `active`, say that  $U$  and  $S$  are the partners of this session. Send  $(\text{NEWSESSION}, \text{sid}, S, U, \text{RESPONDER})$  to  $\mathcal{A}^*$ .
- On  $(\text{SEND}, \text{sid}, U, S, m)$  from  $U$ , and if `active` is set, send  $(\text{SEND}, \text{sid}, U, S, |m|, t)$  to the adversary, where  $t$  is a counter initialized to 0 and incremented by 1 at every send. Record  $(\text{sid}, U, S, m, t)$ .
- On  $(\text{DELIVER}, \text{sid}, U, S, m, t)$  from  $\mathcal{A}^*$ : if  $S$  is corrupted, then output  $(\text{RECEIVED}, \text{sid}, U, S, m, t)$ . Otherwise, lookup record  $(\text{sid}, U, S, m', t)$  and output  $(\text{RECEIVED}, \text{sid}, U, S, m', t)$ .
- The sends from  $S$  to  $U$  are handled analogously.

Figure 4: The ideal functionality  $\mathcal{F}_{\text{SC}}$  for Secure-Channel

- On  $\mathcal{F}_{\text{SC}}.(\text{NEWSESSION}, \text{sid}, U, S, \text{INITIATOR})$  from  $U$ , sample encryption and mac keys  $\text{tok} = (k_e, k_a)$  and call  $\mathcal{F}_{\text{AL}}$  with  $(\text{CLIENTLOGIN}, \text{sid}, U, \text{tok})$ .
- On  $\mathcal{F}_{\text{SC}}.(\text{NEWSESSION}, \text{sid}, S, U, \text{RESPONDER})$  from  $S$ , call  $\mathcal{F}_{\text{AL}}$  with  $(\text{SERVERLOGIN}, \text{sid}, U)$ .
- On  $(\text{sid}, \text{uid}, \text{tok}', \text{flag})$  from  $\mathcal{F}_{\text{AL}}. \text{LOGIN}$ : If `flag` = `SUCCESS`, then parse  $\text{tok}'$  as  $(k'_e, k'_a)$ .
- On  $\mathcal{F}_{\text{SC}}.(\text{SEND}, \text{sid}, U, S, m)$  from  $U$ , encrypt and mac  $(t, m)$  with counter  $t$  which is initialized to 0 and incremented by 1 for each send, using  $(k_e, k_a)$  and sends  $(\text{DELIVER}, \text{sid}, S, U, t, c)$  to  $\mathcal{A}^*$ .
- Upon receiving  $(\text{DELIVER}, \text{sid}, S, U, t, c')$  from  $\mathcal{A}^*$  use the corresponding key pair  $(k_e, k_a)$  for  $(\text{sid}, S, U)$ , checks authenticity and decrypts to  $m'$  and outputs  $\mathcal{F}_{\text{SC}}.(\text{RECEIVED}, \text{sid}, S, U, t, m')$ .
- The sends from  $S$  to  $U$  are handled analogously.

Figure 5: Protocol  $\Pi_{\text{SC}}$  to realize  $\mathcal{F}_{\text{SC}}$  in the  $\mathcal{F}_{\text{AL}}$ -hybrid model.

- Upon receiving  $\mathcal{F}_{\text{SC}}(\text{NEWSESSION}, \text{sid}, \text{S}, \text{U}, \text{RESPONDER})$  from  $\text{S}$ , the simulator calls  $\mathcal{F}_{\text{AL}}$  with  $(\text{SERVERLOGIN}, \text{sid}, \text{U})$ .
- Upon getting output  $(\text{sid}, \text{uid}, \text{tok}', \text{flag})$  from  $\mathcal{F}_{\text{AL}}.\text{LOGIN}$ , if  $\text{flag} = \text{SUCCESS}$ , then parse  $\text{tok}'$  as  $(k'_e, k'_a)$ .
- Upon receiving  $\mathcal{F}_{\text{SC}}(\text{SEND}, \text{sid}, \text{U}, \text{S}, t, |m|)$  from  $\text{U}$ , the simulator encrypts and macs  $(t, 0^{|m|})$  with counter  $t$  which is initialized to 0 and incremented by 1 for each send, using  $(k_e, k_a)$  and sends to  $\mathcal{A}$ .
- Upon receiving  $\mathcal{F}_{\text{SC}}(\text{sid}, \text{S}, \text{U}, t, c')$  from  $\mathcal{A}$ , the simulator checks if it sent  $c'$  for  $(\text{sid}, \text{S}, \text{U}, t)$ . If so, then it calls  $\mathcal{F}_{\text{SC}}$  with  $(\text{RECEIVED}, \text{sid}, \text{S}, \text{U}, t, \perp)$ . This has the effect of allowing  $\mathcal{F}_{\text{SC}}$  to output the recorded message and completely ignore  $c'$ . Otherwise, the simulator looks up for a recorded key pair  $(k_e, k_a)$  for  $(\text{sid}, \text{S}, \text{U})$ , checks authenticity and decrypts to  $m'$  and sends  $\mathcal{F}_{\text{SC}}(\text{RECEIVED}, \text{sid}, \text{S}, \text{U}, t, m')$ . Observe that the adversary can modify the message while passing authenticity only if it has corrupted the party or compromised the session in some way.

The sends from  $\text{S}$  to  $\text{U}$  are handled analogously.

Indistinguishability follows from the fact that the adversary only sees encryptions of 0, which are indistinguishable from the hybrid execution due to the CCA security of PKE. Also, the adversary is unable to construct fake messages purporting to come from honest parties which pass MAC verification, due to the unforgeability of the MAC scheme.  $\square$

## D UC aPAKE Definition

We reproduce the UC functionality for Asymmetric PAKEs from [JKX18], split between Figures 6 and 7. More details can be found in the original paper [GMR06].



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

#### Password Registration

- On  $(\text{STOREPWDFILE}, \text{sid}, U, \text{PW})$  from  $S$ , if this is the first  $\text{STOREPWDFILE}$  message, record  $\langle \text{FILE}, U, S, \text{PW} \rangle$  and mark it  $\text{UNCOMPROMISED}$ .

#### Stealing Password Data

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

#### Password Authentication

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

Figure 6: The ideal functionality for aPAKE (part 1 of 2).

### Active Session Attacks

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

### Key Generation and Authentication

- On (NEWKEY, sid, ssid, P, SK\*) from  $\mathcal{A}^*$  where  $|SK^*| = \ell$ , if there is a record  $\langle \text{ssid}, P, P', \text{PW}' \rangle$  not marked COMPLETED, do:
  - If the record is COMPROMISED, or P or P' is corrupted, set  $SK := SK^*$ .
  - Else if the record is FRESH, a  $(\text{sid}, \text{ssid}, SK')$  tuple was sent to P', and at that time there was a record  $(\text{ssid}, P', P)$  marked FRESH, set  $SK := SK'$ .
  - Else pick  $SK \leftarrow_{\mathcal{R}} \{0, 1\}^{\ell}$ .

Finally, mark  $\langle \text{ssid}, P, P', \text{PW}' \rangle$  COMPLETED and send  $(\text{sid}, \text{ssid}, SK)$  to P.

- On (TESTABORT, sid, ssid, P) from  $\mathcal{A}^*$ , if there is a record  $\langle \text{ssid}, P, P', \text{PW}' \rangle$  not marked COMPLETED, do:
  - If it is FRESH and there is a record  $\langle \text{ssid}, P', P, \text{PW}' \rangle$ , send SUCC to  $\mathcal{A}^*$ .
  - Else send FAIL to  $\mathcal{A}^*$  and (ABORT, sid, ssid) to P, and mark  $\langle \text{ssid}, P, P', \text{PW}' \rangle$  COMPLETED.

Figure 7: The ideal functionality for aPAKE (part 2 of 2).