

LATKE: A Framework for Constructing Identity-Binding PAKEs

Jonathan Katz[Ⓛ]
Google and University of Maryland
jkatz2@gmail.com

Michael Rosenberg[Ⓛ]
University of Maryland
micro@cs.umd.edu

Version 1.0, July 6, 2024

Abstract

Motivated by applications to the internet of things (IoT), Cremers, Naor, Paz, and Ronen (CRYPTO '22) recently considered a setting in which multiple parties share a common password and want to be able to pairwise authenticate. They observed that using standard password-authenticated key exchange (PAKE) protocols in this setting allows for *catastrophic impersonation* attacks whereby compromise of a single party allows an attacker to impersonate any party to any other. To address this, they proposed the notion of *identity-binding PAKE* (iPAKE) and showed constructions of iPAKE protocol CHIP.

We present LATKE, a framework for iPAKE that allows us to construct protocols with features beyond what CHIP achieves. In particular, we can instantiate the components of our framework to yield an iPAKE protocol with post-quantum security and *identity concealment*, where one party hides its identity until it has authenticated the other. This is the first iPAKE protocol with either property.

To demonstrate the concrete efficiency of our framework, we implement various instantiations and compare the resulting protocols to CHIP when run on commodity hardware. The performance of our schemes is very close to that of CHIP, while offering stronger security properties.

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

1 Introduction

Short passwords are one of the most common methods of authentication today. Passwords are used for logging into websites, performing secure file transfers, connecting to WiFi, connecting to mesh networks, storing encrypted backups, and logging into remote servers. Dedicated protocols are necessary to handle authentication and key exchange in the *password-only* setting where no public-key infrastructure exists and the authenticating parties share only a low-entropy

password. This is precisely the setting *password authenticated key-exchange* (PAKE) protocols address.

Since their introduction by Bellare and Merritt [BM92], PAKE protocols have been extended to offer stronger security (see further discussion below) and have been deployed in several practical scenarios. In the context of the internet of things (IoT), PAKE protocols have been deployed for authentication between devices [Thr15, All22a], and the WiFi specification includes the SAE PAKE [Har08] for mutual authentication of stations in a network [All22b].¹ For IoT and mesh networks in which multiple entities share a common password, *identity-binding PAKE* (iPAKE)—recently introduced by Cremers et al. [CNPR22]—offers a strong form of security. This form of PAKE allows each party to bind their password to their identity, such that, if an adversary compromises a party, then it can impersonate only that party, and no other parties. In particular, upon compromising a party, the adversary does not learn the shared password, and can only recover the password via a brute-force attack. Besides introducing the notion, Cremers et al. also proposed the iPAKE protocol CHIP, and a precomputation-resistant variant, CRISP.

We present LATKE, a generic framework for constructing iPAKE protocols. Our framework allows us to construct iPAKE protocols that have features beyond what CHIP achieves (and does not compare to CRISP, as we do not include precomputation resistance). Specifically, we can instantiate our framework to yield an iPAKE protocol based on post-quantum assumptions that also offers identity concealment; this is the first iPAKE protocol with either property. This comes at the cost of a concretely small overhead in runtime and communication as compared to CHIP, and a slight weakening of the security model, that we argue is natural in the context of iPAKE.

In what follows, we review prior work on PAKE and its generalizations (including iPAKE), discuss limitations of existing iPAKE protocols, and then describe our contributions in greater detail.

1.1 Background

PAKE. Bellare and Merritt [BM92] proposed using *password authenticated key-exchange* (PAKE) protocols for allowing two parties to authenticate and generate a cryptographically strong shared key based on a low-entropy shared password. PAKE protocols ensure, in particular, that 1) a passive adversary learns nothing about the password; 2) an active adversary cannot learn the password any faster than by doing an online brute-force attack whereby it guesses a single password per protocol execution; 3) barring a correct password guess in an active attack, the session key shared by the parties is indistinguishable from uniform for the adversary. Note that, unlike most other cryptographic constructions, an active attacker can succeed in impersonating one of the parties with non-negligible probability by guessing the password using sufficiently many online attempts. This is as good a guarantee as possible, given the constraint

¹See doi.org/10.1109/IEEESTD.2021.9363693.

that honest users must be able to establish a shared key using only a low-entropy password.

aPAKE. PAKE protocols require both parties to have a copy of the password when executing the protocol. This is fine for a human user who memorizes a password and then types it to authenticate, but is less desirable for the server authenticating that user to store the password since, in particular, it would imply that if the server is compromised then an attacker would immediately learn the user’s password. *Augmented password authenticated key-exchange* (aPAKE) protocols—proposed by Jablon [Jab97] and later formalized by Gentry, MacKenzie, and Ramzan [GMR06]—extends PAKE to afford more security against server compromise. Such protocols allow the server to store a (per-user) *password file* that is computed as some hard-to-invert function of the password, and ensure that if an adversary compromises the server then they must perform an offline brute-force attack in order to learn the password. (Note that the possibility of such an attack is inherent.) Many protocols rely on aPAKE, including iCloud key recovery [App], WhatsApp message history recovery [DFG⁺23], and 1Password [Fil18].

As just noted, aPAKE protocols ensure that an adversary who obtains a server’s password file must perform a brute-force attack in order to recover the corresponding password. But they allow for the possibility that the adversary can use precomputation to carry out the bulk of the work *before* learning the password file (and then recovering the password immediately upon corruption of the server). *Strong aPAKE* (saPAKE) protocols [JKX18] prevent this, and force the attacker to do work linear in the size of the password dictionary *after* the attacker learns the password file. Jarecki et al. [JKX18] show a generic approach for compiling any aPAKE protocol into an saPAKE protocol using an oblivious pseudorandom function (OPRF).

dPAKE. Although (s)aPAKE protocols offer protection against compromise of one of the parties, they still require the other party to use a plaintext password; thus, in scenarios where both parties to the protocol are machines² storing the password, compromise of one of the parties reveals the password immediately. *Doubly augmented password authenticated key-exchange* (dPAKE) protocols [Ham15] allow both parties to store a transformed version of the password, so that if *either* party is compromised the adversary cannot do better than impersonating the compromised party or performing a brute-force search for the password. By analogy with aPAKE, strong dPAKE (sdPAKE) makes precomputation ineffective. As with aPAKE, it is possible to generically construct sdPAKE from dPAKE using an OPRF. At the time of writing, (s)dPAKE has not received any formal treatment. But it is a useful conceptual extension of aPAKE, and has been proposed for use in WiFi stations [Tho22b].

²One may question why passwords—rather than high-entropy cryptographic keys—are used in this setting at all. First, as discussed further below, there may be more than two entities sharing the same secret information, one of which might be a human user. Moreover, a human user might be required to enter the secret information into the devices upon initialization. In any case, real-world deployments of IoT/mesh networks often use passwords.

iPAKE. Even (s)dPAKE is insufficient for some use cases, since it assumes that there are only two roles: client and server. In some applications, specifically IoT/mesh networks, there are multiple parties—having multiple roles—who share a common password and want to potentially authenticate each other (pairwise). Motivated by exactly this scenario, Cremers et al. [CNPR22] introduced the notion of *identity-binding password authenticated key exchange* (iPAKE), a strengthening of dPAKE in which each party’s password file is bound to an arbitrary identity string. Similar to dPAKE, an attacker who compromises a party can impersonate that party, but it cannot impersonate any other party without performing an offline brute-force attack to determine the password. *Strong* iPAKE (siPAKE) is also defined in the natural way.

Cremers et al. [CNPR22] define (s)iPAKE in the Universal Composability (UC) framework [Can01]. They also present CHIP, an efficient iPAKE protocol constructed generically from any PAKE protocol and any identity-based key-exchange (IBKE) protocol with certain properties (explained in the next section). Finally, they also present CRISP, an efficient siPAKE protocol that they prove secure in the generic-group model.

1.2 Limitations of Existing dPAKEs and iPAKEs

As noted above, CHIP is an efficient iPAKE protocol that can be based on any PAKE protocol and any IBKE protocol with certain properties. Specifically, the IBKE protocol is required to have key-compromise impersonation resistance (KCIR), a standard property, but also to be *msk-independent*, i.e., the message flow of the protocol must be statistically independent of the main secret key *msk*. While there are a handful of KCIR, *msk-independent* IBKE protocols [Oka88, Gün90, FG10, Shi03, Wan13, CC07b, CC07a], they are insufficient to give CHIP some desirable properties.

Quantum resistance. All the aforementioned IBKEs rely on the commutativity of certain group operations, and depend on the Diffie-Hellman assumption or a generalization thereof over bilinear pairings for security. Hence, there is no clear way to build iPAKE from *post-quantum assumptions*, i.e., cryptographic assumptions that plausibly hold for quantum computers.

Beyond post-quantum iPAKE, there is also no clear way to construct an *efficient* post-quantum dPAKE. The OPAQUE aPAKE compiler (not to be confused with the same paper’s saPAKE compiler) [JKX18]—doubly augmented to be a dPAKE [Ham15]—can be instantiated with entirely post-quantum primitives, namely a post-quantum authenticated key exchange (AKE) protocol, and a post-quantum oblivious pseudorandom function (OPRF). However, the only options for post-quantum OPRF at the time of writing are either broken [BKW20, BKM⁺21], or require two to four orders of magnitude greater computation time or communication overhead than Diffie-Hellman-based OPRFs [FOO23, Dod23, BDFH24], or have unknown concrete runtime [Bas22, Bas23]. Thus, practical use is limited.

Identity concealment. Another property not present in the listed IBKEs is

that of *identity concealment*, i.e., one of the parties remains anonymous until the other party has successfully authenticated itself, and both parties remain anonymous to a passive eavesdropper. This idea was introduced in the SIGMA-I/SIGMA-R protocols [Kra03] and is available in IKEv2 which is used in the widely deployed IPsec VPN protocol. If an iPAKE device responds to every incoming request with its identity string, an attacker can *wardrive*, i.e., move through a geographical area, attempt to connect to all the devices within range, and map the network’s topology [SPT13]. This may be problematic in the case of industrial IoT, e.g., leaking which building a specific sensor or controller resides in.

It is not clear if adding identity concealment to the listed IBKEs is possible, since none have explicit authentication, and adding explicit authentication in the form of signatures or MACs can easily break msk independence.

1.3 Our Contributions

In this work, we present low-assumption iPAKE (LATKE), a highly flexible framework for building iPAKE protocols. LATKE relies on an ordinary PAKE and an identity-based key-exchange (IBKE) protocol with key-compromise impersonation resistance (KCIR) and full forward secrecy (full FS). There are many IBKE protocols with the necessary properties, as opposed to the situation for the msk -independent IBKEs that CHIP requires. In particular, we show that it is possible to construct a LATKE-compatible IBKE using only a signature scheme and an authenticated key exchange (AKE) protocol.

We consider two variants: $\text{LATKE}^{\text{pre}}$ and $\text{LATKE}^{\text{post}}$. The former is intended to be used for *pre-specified peer* IBKEs, i.e., ones where the peers both know the identity of their intended partner in advance. The latter is intended to be used for *post-specified peer* IBKEs, i.e., ones where that is not the case. We will consider identity concealment in the post-specified peer setting, which is the more difficult of the two, since the parties to conceal from must also include a user’s intended partner.

We prove security of LATKE with respect to a slightly weaker version of the iPAKE functionality $\mathcal{F}_{\text{iPAKE}}$ given in [CNPR22]. We also argue why our relaxed functionality $\mathcal{F}'_{\text{iPAKE}}$ is a natural one. We prove that LATKE produces protocols that UC-realize $\mathcal{F}'_{\text{iPAKE}}$ under adaptive corruptions in the \mathcal{F}_{RO} -hybrid model with erasure. Since there exist post-quantum AKEs/IBKEs and post-quantum PAKEs, we conclude that LATKE can be instantiated as a post-quantum iPAKE.³ To the authors’ knowledge, this is the first description or implementation of a post-quantum or identity-concealing iPAKE. In addition, our construction also represents the most efficient post-quantum dPAKE, since the only known post-quantum dPAKE currently requires post-quantum OPRF, whose shortcomings were described above.

To demonstrate concrete efficiency, we instantiate LATKE and CHIP in various ways and benchmark the resulting protocols on a commodity WiFi router

³We mean here that it can be based on assumptions that are plausibly quantum-hard. The proof we give, however, only considers classical adversaries.

at the 128-bit security level. We find some pre-quantum instantiations of LATKE have computation cost within 5% of CHIP and with a communication overhead of 324B, and one post-quantum instantiation achieves computation cost within 3% of CHIP with a communication overhead of only 3kB.

2 Preliminaries

In this section we present the notation, definitions, and security models necessary for the construction of LATKE.

2.1 Notation

We write probabilistic algorithms as $\text{Alg}(x; r)$ where x denotes the input and r denotes the random coins. We write $y \leftarrow \text{Alg}(x)$ to denote sampling uniform r and setting $y := \text{Alg}(x; r)$, and $x \leftarrow_{\$} S$ to denote sampling a uniform value from a set S . For our security proofs we will consider probabilistic polynomial time (PPT) adversaries, and denote them with calligraphic letters \mathcal{A} . We denote the output x, x' of either side of an interactive protocol between parties A, B by $(x, x') \leftarrow (A \leftrightarrow B)$. For interactive protocols, a *protocol round* is the set of all messages that can be sent in parallel from any point in the protocol [Gon93].⁴ We use λ to denote the security parameter.

In our pseudocode for ideal functionalities, **retrieve** denotes retrieval of a record with a specific marking; if no such record is present, this returns “no record.” We write $\boxed{\text{IBKE}}$ to denote execution of an IBKE protocol, and $\boxed{\text{PAKE}_{\text{sid}, \text{ssid}}}$ to denote execution of a PAKE protocol with session and sub-session identifiers sid and ssid .

2.2 Universal Composability

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

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

⁴The number of messages is not the same as number of rounds. A protocol with 2 rounds and 4 messages can be converted into one with 3 rounds and 3 messages by combining messages.

adversary may also only interact to \mathcal{F} , though it is still permitted to corrupt parties.

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

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

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

By [Can01, Theorem 11], we may assume \mathcal{A} is the *dummy adversary* \mathcal{A}_D that simply delivers messages generated by the environment to the specified recipients, and delivers to the environment all messages generated by the protocol parties.

Corruption. There are two widely used corruption models in UC. In the *static corruption model*, the adversary may not corrupt any parties once execution has started. In the *adaptive corruption model*, the adversary may corrupt parties at any time. In this work, we will consider adversaries who are allowed adaptive corruptions.

SID and SSID. To represent different instances of the same scheme, the UC model uses *session identifiers*. Each party is given a session identifier sid on activation and will only interact with other parties with the same sid . To represent individual protocol executions, we will use *sub-session identifiers*, denoted ssid (these can be established by out-of-band means, or, e.g., by exchanging nonces [BLR04]). A party in session sid may be engaged in multiple simultaneous protocol executions $\text{ssid}_1, \dots, \text{ssid}_n$. As in [GMR06], we may assume that every protocol execution is uniquely identified by its $(\text{sid}, \text{ssid})$ pair. For clarity of presentation, we will assume in our protocol definitions that session and sub-session identifier establishment has already occurred.

2.3 PAKE and iPAKE

PAKE. We describe the function of a PAKE protocol and its intended security properties. A *(balanced) password authenticated key-exchange protocol* is a two-party key-exchange protocol where parties use mutual knowledge of a low-entropy password pw to establish a high-entropy *session key* K . The security goals for a PAKE are (1) to establish a high-entropy shared session key when both parties are honest, (2) to prevent passive adversaries from learning anything about the password, and (3) to limit active adversaries to one (or another small constant) password guess(es) per protocol instance, even if given access to session keys.

We give the corresponding ideal functionality $\mathcal{F}_{\text{PAKE}}$ in Figure 1. We include the small modifications to adversary-controlled keying from [AHH21], as do the authors of CHIP/CRISP. We use the *multi-session* variant of the PAKE functionality [CR03] (i.e., using ssid in addition to sid), which, as noted in [BGHJ24], is equivalent to the base single-session functionality.

PAKE protocols permit an attacker who corrupts a party to learn pw and subsequently impersonate *any* party using the same password. In a setting

Session management	Key generation and authentication
On (NewSession, sid, ssid, \mathcal{P}_j , pw_i) from \mathcal{P}_i	On (NewKey, sid, ssid, \mathcal{P}_i , K') from \mathcal{A}
send (NewSession, sid, ssid, \mathcal{P}_i , \mathcal{P}_j) to \mathcal{A}	retrieve (Session, sid, ssid, \mathcal{P}_i , \mathcal{P}_j , pw_i)
if \nexists (Session, sid, ssid, \mathcal{P}_i , \mathcal{P}_j , pw_i):	not marked completed
record (Session, sid, ssid, \mathcal{P}_i , \mathcal{P}_j , pw_i)	if session is compromised:
Mark it fresh	$K_i := K'$
Active session attack	elif session is fresh and
On (TestPwd, sid, ssid, \mathcal{P}_i , pw') from \mathcal{A}	\exists (Key, sid, ssid, \mathcal{P}_j , pw_j , K_j) s.t. $\text{pw}_i = \text{pw}_j$:
retrieve (Session, sid, ssid, \mathcal{P}_i , \mathcal{P}_j , pw_i)	$K_i := K_j$
marked fresh	else: $K_i \leftarrow \{0, 1\}^\lambda$
if $\text{pw}_i = \text{pw}'$:	if session is fresh:
Mark session compromised	record (Key, sid, ssid, \mathcal{P}_i , pw_i , K_i)
send “correct” to \mathcal{A}	Mark session completed
else	send (sid, ssid, K_i) to \mathcal{P}_i
Mark session interrupted	
send “wrong” to \mathcal{A}	

Figure 1: The (multi-session) $\mathcal{F}_{\text{PAKE}}$ ideal functionality, with the keying modifications suggested in [AHH21].

where multiple devices all share the same password, this results in *catastrophic impersonation* following the compromise of even a single device.

Identity-binding PAKE (iPAKE). iPAKE [CNPR22] prevents catastrophic impersonation attacks. Each party runs a one-time initialization procedure `StorePwdFile` that takes as input a (common) password `pw` and its (public) identity `id`, and outputs a *password file* `pwfile`. That party then saves `pwfile`, and deletes `pw`. At a later point in time, two parties who wish to authenticate run the iPAKE protocol using their respective password files. The output is a tuple (K, id) containing the session key and the identity of the other party.

As with PAKE, iPAKE provides a high-entropy shared key to the two parties running the protocol, and limits active attackers to a single password guess per instance. In contrast to PAKE, however, iPAKE also ensures some measure of robustness following compromise of parties. Specifically, an attacker who compromises Alice only learns Alice’s `pwfile`. With this, the attacker can impersonate Alice to anyone, but does not immediately learn enough information to impersonate any other party to anyone else (including Alice). The best the attacker can do—which is clearly unavoidable in our setting—is to mount a brute-force attack against Alice’s `pwfile` to derive `pw`; this can be made prohibitive if the `pw` has moderate entropy and/or if `StorePwdFile` is relatively slow. The latter can be achieved by incorporating into `StorePwdFile` a time-, memory-, and/or cache-intensive password hashing function, such as Argon2 [BDK15] or `bcrypt` [Tho22a].

We briefly describe the function of the remaining procedures of $\mathcal{F}_{\text{iPAKE}}$. The formal definition is given in Figure 2.

<p>Password Registration On (StorePwFile, sid, id, pw) from \mathcal{P} if \nexists a record (File, \mathcal{P}, \cdot, \cdot) : record (File, \mathcal{P}, id, pw)</p> <p>Password Authentication On (NewSession, sid, ssid, \mathcal{P}_j) from \mathcal{P}_i retrieve (File, \mathcal{P}_i, id$_i$, pw$_i$) send (NewSession, ssid, \mathcal{P}_i, \mathcal{P}_j, id$_i$) to \mathcal{A} if \nexists a record (Session, ssid, \cdot, \cdot, \cdot) : record (Session, ssid, \mathcal{P}_i, \mathcal{P}_j, pw$_i$) and mark it fresh</p> <p>Active Session Attacks On (OnlineTestPw, sid, ssid, \mathcal{P}_i, pw') from \mathcal{A} retrieve (Session, ssid, \mathcal{P}_i, \cdot, pw$_i$) marked fresh or compromised if pw' = pw$_i$: record (Imp, ssid, \mathcal{P}_i, *) Mark the session compromised send "correct" to \mathcal{A} else : Mark the session interrupted send "wrong" to \mathcal{A}</p> <p>On (Impersonate, sid, ssid, \mathcal{P}_i, \mathcal{P}_k) from \mathcal{A} retrieve (Session, ssid, \mathcal{P}_i, \cdot, pw$_i$) marked fresh or compromised retrieve (File, \mathcal{P}_k, id$_k$, pw$_k$) marked compromised if pw$_i$ = pw$_k$: record (Imp, ssid, \mathcal{P}_i, id$_k$) Mark the session compromised send "correct" to \mathcal{A} else : Mark the session interrupted send "wrong" to \mathcal{A}</p>	<p>Key Generation and Authentication On (NewKey, sid, ssid, \mathcal{P}_i, id', K') from \mathcal{A} retrieve (Session, ssid, \mathcal{P}_i, \mathcal{P}_j, pw$_i$) not marked completed retrieve (File, \mathcal{P}_j, id$_j$, pw$_j$) if (session is fresh and id' \neq id$_j$) or (session is compromised and \nexists(Imp, ssid, \mathcal{P}_i, id') and \nexists(Imp, ssid, \mathcal{P}_i, *)) : return // ignore invalid queries if session is compromised : $K_i := K'$ elif session is fresh and \exists a record (Key, ssid, \mathcal{P}_j, pw$_j$, K_j) with pw$_i$ = pw$_j$: $K_i := K_j$ else : $K_i \leftarrow_{\\$} \{0, 1\}^\lambda$ if session is fresh : record (Key, ssid, \mathcal{P}_i, pw$_i$, K_i) Mark session completed send (Key, ssid, id', K_i) to \mathcal{P}_i</p> <p>Stealing Password Data On corruption query (StealPwFile, \mathcal{P}) from \mathcal{A}: retrieve (File, \mathcal{P}, id, pw) Mark the file compromised if \exists a record (Offline, \mathcal{P}, pw) : send (Stolen, id, pw) to \mathcal{A} else : send (Stolen, id, \perp) to \mathcal{A}</p> <p>On (OfflineTestPw, \mathcal{P}, pw') from \mathcal{A} retrieve (File, \mathcal{P}, \cdot, pw) if file marked compromised : if pw' = pw : send "correct" to \mathcal{A} else : send "wrong" to \mathcal{A} else : record (Offline, \mathcal{P}, pw')</p> <p>On (OfflineComparePw, sid, \mathcal{P}_i, \mathcal{P}_j) from \mathcal{A} retrieve (File, \mathcal{P}_i, \cdot, pw$_i$) marked compromised retrieve (File, \mathcal{P}_j, \cdot, pw$_j$) marked compromised if pw$_i$ = pw$_j$: send "match" to \mathcal{A} else : send "no match" to \mathcal{A}</p> <div style="border: 1px dashed black; padding: 5px;"> <p>On (OnlineComparePw, sid, ssid, \mathcal{P}_i, \mathcal{P}_j) from \mathcal{A} retrieve (Session, ssid, \mathcal{P}_i, \cdot, pw$_i$) marked fresh or compromised retrieve (File, \mathcal{P}_j, \cdot, pw$_j$) marked compromised if pw$_i$ = pw$_j$: send "match" to \mathcal{A} else : Mark the session interrupted send "no match" to \mathcal{A}</p> </div>
--	---

Figure 2: The $\mathcal{F}_{\text{iPAKE}}$ and $\mathcal{F}_{\text{ocp-iPAKE}}$ (abbreviated $\mathcal{F}'_{\text{iPAKE}}$) ideal functionalities. $\mathcal{F}'_{\text{iPAKE}}$ includes the code in the dashed box, while $\mathcal{F}_{\text{iPAKE}}$ does not.

Corruption: iPAKE permits two kinds of corruption queries:

- **StealPwdfFile** returns the password file of a specific party \mathcal{P} , and does *not* mark \mathcal{P} corrupted; i.e., the adversary does not get any ephemeral keys nor does it gain control over \mathcal{P} .⁵
- **Corrupt** returns the output of **StealPwdfFile** plus all the ephemeral data held by \mathcal{P} , sends a message to the environment that the action was performed on \mathcal{P} , and *does* mark \mathcal{P} corrupted.

A party who is not corrupted but whose password file has been stolen is called *compromised*.

OfflineTestPwdf: This is called by the simulator \mathcal{S} (acting on behalf of the adversary) to guess the password of a stolen password file.⁶

OfflineComparePwdf: This is called by \mathcal{S} to compare the password files of compromised parties. Through this, it learns which parties' passwords are equal, but does not learn the password itself.

NewSession: This is called by a party \mathcal{P}_i to initiate a new iPAKE session with \mathcal{P}_j . At some point, \mathcal{P}_j will have to call **NewSession** with \mathcal{P}_i and the same sid and ssid if they wish to complete the protocol execution.

NewKey: This is called by \mathcal{S} with a session key K' and ID id' to finalize the key exchange in an active session. If the session is **fresh**, i.e., has not already been interrupted with a different attack, then this outputs a uniformly random key (independent of K') and actual peer identity to the relevant party, and it will output the same key and actual peer identity to the other party when called for them. The values K' and id' are sent to the relevant party when certain compromise conditions are met. These are described below.

OnlineTestPwdf: This is called by \mathcal{S} to guess the targeted party's password in an active session. On success, the session is marked **compromised**, meaning the adversary has full control over the targeted party's session key and the perceived identity of its peer. On failure, the session is marked **interrupted**, i.e., the exchange will fail.

Impersonate: This is called by \mathcal{S} to impersonate a compromised user in an active session. If the compromised user's **pw** matches the session peer's **pw**, the adversary is given full control over the peer's session key (*not* the impersonated user's session key), and the identity is set to the impersonated party's identity.

⁵In prior work [GMR06] this is not explicitly called a corruption query, but—as noted by Hesse [Hes20]—that is the only definitionally sound way to treat this query.

⁶In order to prevent trivial simulators with arbitrary brute-force attack capability, we bound the simulator's usage of **OfflineTestPwdf** by the runtime of the environment \mathcal{Z} . See Hesse [Hes20] for more detail.

In order to permit analysis of LATKE, we slightly weaken the $\mathcal{F}_{\text{iPAKE}}$ functionality by adding a procedure `OnlineComparePwd`. This allows the adversary to check if a compromised user’s password matches the password used in an active session. We call the new functionality $\mathcal{F}_{\text{ocp-iPAKE}}$. For brevity, we refer to this as $\mathcal{F}'_{\text{iPAKE}}$.

We argue that this additional procedure is natural in the sense that similar attacks are already allowed in other PAKE extensions. The Ω -method aPAKE [GMR06] begins with a PAKE over $H(\text{pw})$, and requires the client to store `pw` and the server to store $H(\text{pw})$. Thus, an active attacker who compromises a server can use the stolen hash to intercede in unrelated PAKE sessions and thus determine whether or not the parties of a given session share a password. The reason an `OnlineComparePwd` procedure is not included in aPAKE is because aPAKE doesn’t have a strong notion of identity—to perform the described attack, it suffices to use `Impersonate` using the generic identities of “client” or “server”. This isn’t so in iPAKE: an `Impersonate` query would force the attacker to use the unique identity of the compromised user.

The procedure is natural in another sense. It is possible to emulate `OnlineComparePwd` by calling `Impersonate` as described above, tearing down the session with `OnlineTestPwd(sid, ssid, \mathcal{P} , \perp)`, and forcing the parties to attempt a new session. The only step that is outside of the UC simulator’s capabilities is the last. In reality, though, this is already possible. For example, WPA2-protected WiFi connections can be injected with a deauthentication packet, thus forcing the client to tear down the session and reconnect [SRV23]. Similar attacks in other protocols such as Bluetooth are also documented [Lou21]. Thus, it is not unreasonable to imagine ideal attacker being able to force a retry in just the key exchange portion of a protocol.

2.4 Key Exchange

We review notions of authenticated key exchange (AKE) and identity-based key exchange (IBKE). We include definitions for *post-specified peer* protocols, i.e., protocols where the parties do not know the identity of their peer in advance.

Authenticated key exchange. An AKE protocol allows two parties interact over a public channel and produce a shared secret key. The procedures of a post-specified peer AKE are as follows:

`KeyGen(1^λ)` \rightarrow (`upk`, `usk`) Generates a long-term keypair.

`Execute(skA)` \Leftrightarrow `Execute(skB)` Executes the interactive key exchange protocol. The output is a shared session key K and the public key of the other party, `pkA` or `pkB`.

Identity-based key exchange. Identity-based key exchange (IBKE) is a generalization of AKE which introduces a trusted third party, called the *key generation center* (KGC). The KGC is responsible for generating the *main keypair* (`mpk`, `msk`), and for *extracting* secret keys for users. A user with an identifier string `id` will request a secret key `skid` from the KGC that corresponds to

id . With this secret key, the user can participate in key establishment protocols wherein the other participant knows only the global mpk and their id (and optionally their own sk_{id} if mutual authentication is desired).

Concretely, the procedures of a post-specified peer IBKE are as follows:

$\text{Setup}(1^\lambda) \rightarrow (\text{mpk}, \text{msk})$ Generates the main keypair.

$\text{Extract}_{\text{msk}}(\text{mpk}, \text{id}, \text{aux}) \rightarrow \text{sk}_{\text{id}}$ Extracts a secret key for the given ID under the given mpk , with auxiliary data aux .

$\text{Execute}(\text{sk}_{\text{id}_A}, \text{id}_A, \text{mpk}) \Leftrightarrow \text{Execute}(\text{sk}_{\text{id}_B}, \text{id}_B, \text{mpk})$ Executes the interactive key exchange protocol. The output, on success, is the shared session key K and the ID of the other party, id_A or id_B .

The purpose of aux in extraction is to allow additional data to be included in the extraction request. For example, if extraction requires the KGC to sign a requester-generated public key, aux is where the public key is encoded.

Identity concealment. A post-specified peer AKE or IBKE is said to have *passive responder concealment* if the responder’s identity is hidden from a passive adversary, and *active responder concealment* if from an active adversary (i.e., a malicious initiator). Specifically, in a protocol with active responder concealment, the responder will wait until the initiator has authenticated themselves before they send any identifying information. Active initiator concealment is defined similarly. A protocol may have both passive initiator concealment and active responder concealment (as SIGMA-R does; Section 4) or vice-versa, but cannot have active concealment for both parties, since one party must authenticate first.

The Canetti-Krawczyk (CK) model. The CK model [CK01] and its extension, the identity-based CK model (id-CK), are commonly used for proving security of AKEs and IBKEs, respectively. Their purpose is to model everything an adversary can reasonably do in a real-world AKE or IBKE protocol execution. Within these models, it is possible to define notions of ordinary *session-key (SK) security*, key-compromise impersonation resistance (KCIR), maximal exposure resistance (MEX), forward secrecy (FS), post-compromise security (PCS), and more. In order to capture identity concealment, we will use the CK model adapted to the post-specified peer setting, as presented in [CK02a].

For security of LATKE, we require an IBKE that has SK-security with KCIR and *full FS*. In words, KCIR ensures that if Mallory compromises Alice, then Mallory cannot impersonate anyone (aside from other parties Mallory has compromised) to Alice. FS ensures that if Mallory compromises Alice, Mallory remains unable to learn the session keys from Alice’s past completed sessions. For IBKE there is another notion of FS, which we call *KGC-FS*, which requires that if Mallory compromises the KGC, she is still unable to learn the session keys of Alice’s past completed sessions. We say a protocol has *full FS* if it has FS and KGC-FS. More detailed definitions can be found in Appendix A.3.

Finally, we note there are numerous, subtly different, variants of the CK model, including CK_{HMQV} , eCK , and CK^+ [Kra05, LLM07, FSXY12] (their id- variants being defined similarly, by adding KGC extraction and revelation).

In fact, it has been shown that security in the first three variants is pairwise incomparable—security in one model does not imply security in any other model [Cre11]. Nevertheless, these models have enough in common that our main security reduction will be able to use any one of them.

2.5 Symmetric Encryption

LATKE requires a symmetric encryption scheme that is secure under adaptive corruptions (i.e., is *non-committing*). For this, we use the notion of SIM*-AC-CCA security [JT20, Jae23]. Informally, an authenticated encryption scheme AE is SIM*-AC-CCA-secure with respect to some ideal primitive P , e.g., a programmable random oracle or ideal cipher, if there is a simulator S_{cca} such that no PPT adversary can distinguish between S_{cca} and AE, even when given the ability to encrypt messages, decrypt messages, expose keys, and program the random oracle/ideal cipher. We refer the reader to Appendix A.4 for a more complete explanation of this security notion.

For our purposes, it suffices to note that a SIM*-AC-CCA-secure encryption scheme is easily instantiable via the encrypt-then-MAC construction [JT20], given some reasonable modeling choice for the underlying cipher.

3 The LATKE Framework

In this section we present the LATKE framework for constructing iPAKE protocols. The structure of LATKE can be viewed as a synthesis of the Ω -method [GMR06] for constructing aPAKE protocols and the CHIP iPAKE protocol [CNPR22]. We introduce both those constructions and then present ours.

3.1 Ω -method

The Ω -method [GMR06] is a generic construction of aPAKE from PAKE. The transformation works as follows (see Figure 3). A server with password pw computes one-way functions of the password $h_0 := \mathsf{H}_0(\text{sid}, \text{pw})$ and $h_1 := \mathsf{H}_1(\text{sid}, \text{pw})$. Next, it generates public and private keys (pk, sk) for a signature scheme, and encrypts the signing key sk using h_1 to obtain ciphertext c . The password file includes pk , h_0 , and c .

Execution of the protocol proceeds by having the parties run an underlying PAKE protocol using $h_0 := \mathsf{H}_0(\text{sid}, \text{pw})$ to obtain two keys (K, K') . (The client derives h_0 using pw , while the server has h_0 in its password file.) The parties use K to establish a secure channel over which the server sends c . The client uses $h_1 := \mathsf{H}_1(\text{sid}, \text{pw})$ to decrypt c , obtain sk , and then generate a signature σ of the transcript using sk . The client sends σ to the server, which verifies the signature. On success, both parties output K' .

A key feature of the Ω -method is that it uses a PAKE to set up a secure channel, via which the client can then “prove” knowledge of the password to the

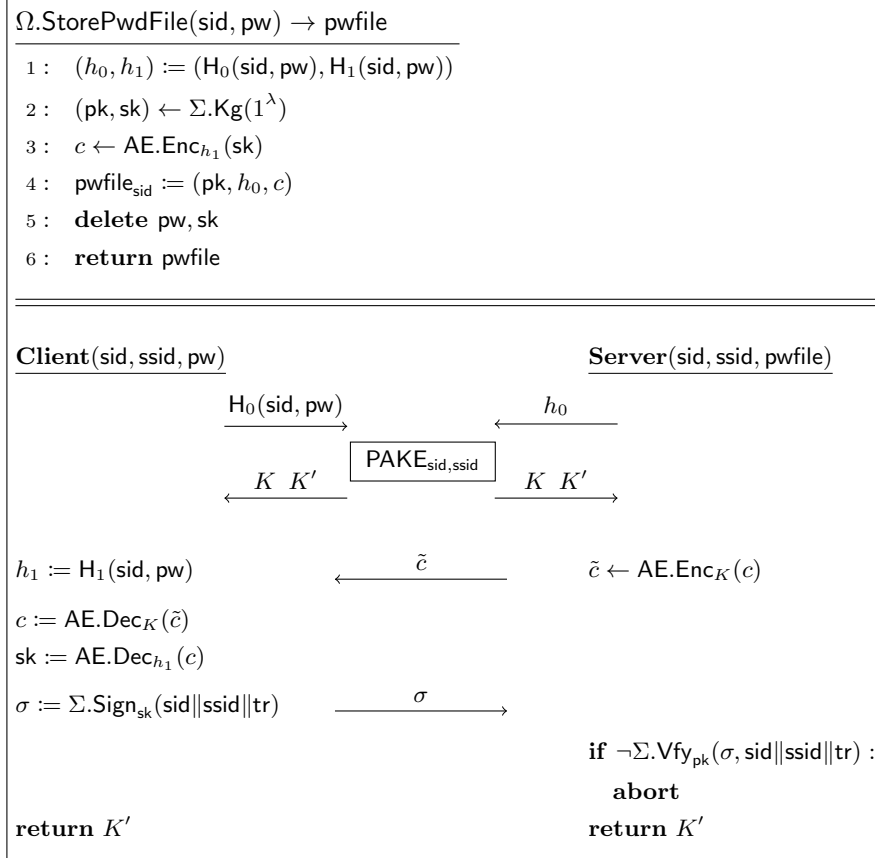


Figure 3: The Ω -method aPAKE [GMR06]. tr represents the protocol transcript up to that point in time. Σ is a signature scheme. AE is an authenticated encryption scheme.

server. This allows the protocol to be secure even though what is transmitted over the secure channel allows for an offline brute-force attack. Intuitively, this holds since any adversary who can learn information about what is sent over the secure channel must already know $\text{H}_0(\text{sid}, \text{pw})$ (since, to see the messages, it must have successfully attacked the PAKE). We use a similar idea in our framework. This will be essential to the flexibility of our iPAKEs.

3.2 CHIP

Recall the purpose of an iPAKE is to permit parties with mutual knowledge of a password to establish a secure connection while identifying themselves to each other based on their unique identities. Further, an attacker who compromises a device should not have the ability to impersonate any identity other than that

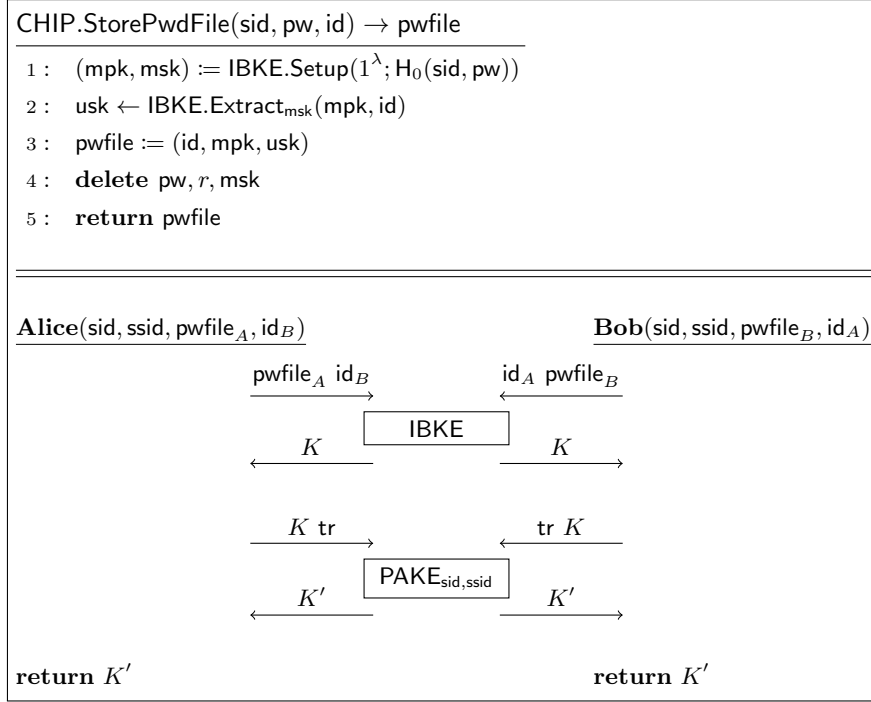


Figure 4: The CHIP iPAKE, using a generic IBKE and PAKE. *tr* represents the protocol transcript up to that point.

of the compromised device (barring a successful brute-force attack).

We give an overview of CHIP in Figure 4. Like the Ω -method, it uses a generic PAKE as a building block. The other building block it requires is an *msk*-independent identity based key-exchange (IBKE) protocol with KCIR. More formally, the messages of the IBKE must be statistically independent of not just *msk*, but also the random coins used in the IBKE’s *Setup* procedure.

The protocol works as follows. A party with ID string *id* and password *pw* generates its password file by hashing *pw* and using the output as random coins to generate a fresh IBKE keypair (*msk*, *mpk*). The party then extracts a secret key *usk* bound to its own identity *id*. The password file is (*id*, *mpk*, *usk*).

An execution of the CHIP protocol consists of two phases. First, the parties execute the IBKE protocol using their extracted secrets and their identities. Next, they use the resulting key (plus the transcript of the protocol) as the “password” for an ordinary PAKE. The output is their final session key.

The crucial element of CHIP is its novel reliance on IBKE, a primitive that typically involves a trusted third party but here is executed by the parties themselves based on their shared password. The IBKE is used to establish a key that simultaneously binds the ID and the password, while ensuring that (1) the *pwfile* is bound to a specific ID, and (2) the password is not directly stored. We

will use this idea in LATKE.

A drawback of CHIP is that it requires an IBKE with strong properties. Specifically, CHIP’s message flow is required to be statistically independent of the password (and hence the random coins of *Setup*). This is because, unlike the Ω -method, CHIP’s key exchange is done “in the clear,” so must ensure that passive adversaries cannot perform an offline brute-force search for the password.

3.3 LATKE

LATKE combines the key features of the Ω -method and CHIP. Password-file generation is identical to that of CHIP, using an IBKE to generate main keys (mpk, msk) that depend on the password, as well as a derived key usk bound to the user’s identity. The protocol itself is similar to the Ω -method, beginning with a PAKE using a one-way function of the password, specifically mpk .⁷ The resulting key is used to set up a secure channel for the rest of the protocol. The parties execute an IBKE protocol inside the secure channel and output the result. The full definition of LATKE can be found in Figure 5. The pre- and post-specified peer variants only differ in what is included in the initial PAKE step.

Since the framework permits a wider range of IBKEs than CHIP, we can now achieve post-quantum iPAKE from post-quantum IBKE and PAKE, and identity-concealing iPAKE from identity-concealing IBKE.

There are a handful of details to work out in order to determine the security and efficiency of LATKE.

Secure channel: The EUE transform. We must specify what we mean by *secure channel*. In the Ω -method, the payload of the channel was a single message. In the case of LATKE, the payload is the entire transcript of an IBKE. In order to achieve security under adaptive corruptions, we must pick our notion of authenticated encryption carefully, and design a secure channel protocol that can be simulated even for sessions between parties whose password files are unknown to the UC simulator. We will design a simple transform, called *encrypt-and-unconditionally execute* (EUE) to meet this specification. We postpone further discussion to Section 3.4.

Necessary security properties. Another detail is precisely what kind of security we need from the PAKE and IBKE. The PAKE requirement does not change from the Ω -method—all we require is a protocol that UC-realizes $\mathcal{F}_{\text{PAKE}}$. The IBKE must have KCIR and full FS, which we note are common properties of key exchange protocols. In the lead-up to our theorem statement, we will explain why each property is necessary.

Building the IBKE from AKE. To further demonstrate the practicality of LATKE, we show in Section 3.5 that it is possible to generically achieve these properties given just an AKE with KCIR (and no FS).

⁷It would suffice to perform the PAKE using a hash of pw as in the Ω -method. Since the parties already hold mpk , however, we simply use that.

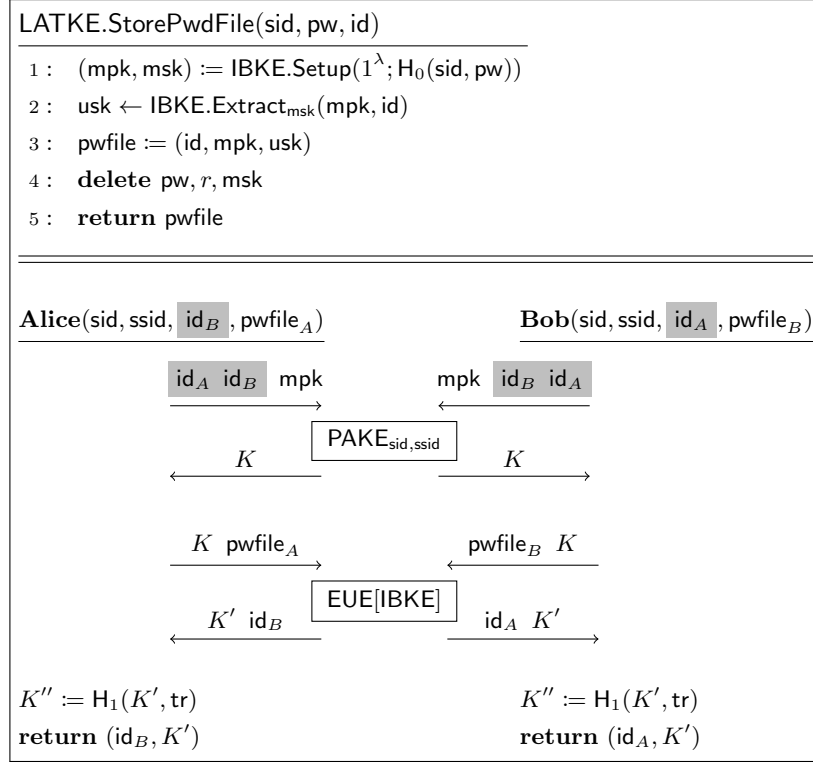


Figure 5: The LATKE iPAKE framework. With the grey text excluded, the above defines LATKE^{post}. With the grey text included, the above defines LATKE^{pre}. EUE[IBKE] is encrypt-and-unconditionally-execute transform of the identity-based key agreement protocol IBKE (Section 3.4), and uses a separate random oracle H_2 for key derivation. tr is the transcript of the entire protocol.

Round complexity. PAKE and IBKE constructions can both be achieved using one round of communication each. So, like CHIP, the minimum round complexity for the LATKE framework is two rounds. While one-round PAKEs and IBKEs exist (e.g., EKE [LLHG23] and Fiore-Gennaro [FG10], resp.), most rely on a Diffie-Hellman-type assumption. In Section 5, we discuss ideas to achieve post-quantum LATKE in two rounds.

3.4 The Encrypt-and-Unconditionally-Execute Transform

Recall we must use the PAKE-derived key to establish a secure channel in which to perform the IBKE protocol. To hide IBKE messages and ensure correct execution, we must carefully apply a symmetric encryption scheme so as to avoid potential replay attacks. In addition, our security proof requires that the protocol transcript look the same to a passive attacker regardless of whether

the initial PAKE succeeded. To achieve this, we must send ciphertexts *even on protocol failure*. Finally, since we must be able to open ciphertexts to arbitrary values, we require a simulation-secure notion of symmetric encryption. We call the combination of these techniques the *encrypt-and-unconditionally-execute transform* EUE[IBKE].

EUE definition. We require an IBKE with a fixed number of rounds and fixed message sizes, an authenticated encryption scheme AE with SIM*-AC-CCA security using ideal primitive \mathcal{P} , and a random oracle \mathcal{H} . We use the variable `rmode` to denote whether the party is in *real mode*, i.e., is still encrypting real messages. At the beginning of the protocol, `rmode` := `true`. We define the behavior of a party \mathcal{P} interacting with its peer \mathcal{P}' in session $(\text{sid}, \text{ssid})$. We let K represent the encryption key that \mathcal{P} intends to use for the session.

1. Let the *chain key* $K_{\text{ch}} := \mathcal{H}(\text{sid}, \text{ssid}, 0 \| K)$. We will use the chain key to derive all subsequent keys using a symmetric ratchet similar to the Signal protocol [Mar16, CJSV22]. Once a chain key has been ratcheted forward, the old keys are erased.
2. Protocol step i , where \mathcal{P} is sending: Compute the step key and new chain key $(K_{\text{ch}}, k_i) := \mathcal{H}(\text{sid}, \text{ssid}, 1 \| K_{\text{ch}})$. If `rmode` = `true`, the IBKE hasn't aborted yet, so \mathcal{P} has a specific IBKE message m_i it wishes to send. Compute $c_i := \text{AE.Enc}_{k_i}^{\mathcal{P}}(m)$. Otherwise if `rmode` = `false`, let $c_i := \text{AE.Enc}_{k_i}(0^{\ell_i})$, where 0^{ℓ_i} is the all-zeros string of the known step i message length ℓ_i . Finally send c_i to \mathcal{P}' .
3. Protocol step i , where \mathcal{P} is receiving: Compute the step key as described above. If `rmode` = `false`, the IBKE has aborted, so simply return. Otherwise, proceed as follows. To decrypt the incoming ciphertext c_i , compute $m'_i := \text{AE.Dec}_{k_i}^{\mathcal{P}}(c_i)$. If a decryption failure occurred, set `rmode` := `false` and return. Otherwise pass m'_i to \mathcal{P} to continue the protocol execution. If \mathcal{P} aborts, then abort the protocol.
4. At the end of the protocol, if `rmode` = `false`, output \perp . Otherwise, output \mathcal{P} 's output.

We will not prove any standalone security properties about the EUE protocol, as it only makes sense when used in the broader context of LATKE.

3.5 Building the IBKE

We briefly recall results which make it possible to construct a satisfactory IBKE from the smaller building block of authenticated key exchange (AKE). We start with an AKE with just key compromise impersonation resistance (KCIR). We will use these techniques to build several IBKEs which we benchmark in Section 4.

Forward secrecy for AKE. LATKE requires full forward secrecy from its key exchange. There are a few well known methods to endow a generic AKE with forward secrecy. For generic AKEs, Boyd and Nieto [BN11] describe a

transform wherein the AKE is used to set up an authenticated channel, and through that channel, the parties perform an ordinary (unauthenticated) key exchange to derive the final key. The resulting AKE has forward secrecy as long as the adversary is not given the ability to perform ephemeral key revelations. This security meets the preconditions of Theorem 1.

Another transform achieves forward secrecy from *weak forward secrecy*, i.e., forward secrecy when the adversary does not tamper with the messages in the test session. The transform simply adds *key confirmation* to the end of the protocol—each party sends a MAC whose key is derived from the shared secret [Kra05, Section 8] [GGJJ23].

Both of these transforms add at least one round of communication.

AKE to IBKE. To convert an AKE into an identity-based AKE, we follow a well-known schema, the *certification approach*, used for building identity-based signatures from ordinary signatures [KN09]. In this construction, every user has an AKE public key pk , and some ID string id . The key generation center (KGC) issues identities by signing (id, pk) pairs with its signing key msk . The user’s *certificate* is thus $(\text{id}, \text{pk}, \sigma)$, where σ is the signature they received from the KGC. At the very beginning of key exchange, users exchange and verify each other’s certificates with respect to the KGC’s public key mpk . On success, they proceed with AKE using the user-provided public keys.

We describe the transform more formally. Let Σ be an EUF-CMA-secure signature scheme, and let AKE be an authenticated key exchange protocol in the post-specified peer model. The post-specified peer IBKE is defined as follows:

$\text{Setup}(1^\lambda) \rightarrow (\text{mpk}, \text{msk})$ Generates a signing keypair $(\text{mpk}, \text{msk}) \leftarrow \Sigma.\text{KeyGen}(1^\lambda)$.

$\text{Extract}_{\text{msk}}(\text{id}, \text{upk}) \rightarrow \sigma$ Computes the signature $\sigma \leftarrow \Sigma.\text{Sign}_{\text{msk}}(\text{id} \parallel \text{upk})$.

$\text{Execute}(\text{usk}_A, \text{id}_A, \sigma_A, \text{mpk}) \Leftrightarrow \text{Execute}(\text{usk}_B, \text{id}_B, \sigma_B, \text{mpk})$ Each party first sends each other $\text{cert}_X := (\text{upk}_X, \text{id}_X, \sigma_X)$ where $X \in \{A, B\}$. Then, each party verifies cert with respect to mpk . On success, the parties proceed with the AKE. At the end, the parties check that the outputted upk matches the one they received.

For a pre-specified peer IBKE from a pre-specified peer AKE, it suffices to remove the last check and add a check that the received certificate’s ID matches the expected ID.

We make some remarks on the certification approach. Firstly, the security properties we desire, namely KCIR and FS, are preserved in this transform. Secondly, this transform produces an IBKE with KGC-FS. This is because the KGC keys are strictly used for authentication, and never for key derivation. Thirdly, as stated above the transform is slightly stricter than necessary regarding the ordering of events. In particular, certificates can be sent at any point in the protocol. If AKE is responder-concealing, then the transformed protocol may have the initiator send its certificate at the very beginning, and the responder

send their certificate at the very end of the protocol, encrypted with the session key, thus preserving the responder-concealing property.⁸

3.6 Identity Concealment

Informally, due to LATKE’s initial PAKE, passive attackers cannot observe the identity of either party in a protocol execution. This holds even for active attackers who have not compromised any users or guessed the password. However, the PAKE alone is insufficient to prevent an active attacker in possession of a compromised password file. If Mallory steals Alice’s password file, she can initiate a session with Bob (whose identity she does not yet know), successfully complete the PAKE, and learn Bob’s identity if the underlying IBKE protocol is non-concealing. Thus, even if Alice is not the ID Bob expected or wanted to interact with, he may be forced to reveal his identity.

To address this, we look to IBKE protocols with *active identity concealment*—where one party is forced to identify themselves before the other party does. As mentioned by the authors of the SIGMA-I/SIGMA-R identity-concealing protocols [Kra03], it is not possible to provide active identity concealment to both parties; one must pick either the initiator or the responder.

When $\text{LATKE}^{\text{post}}$ is instantiated with a responder-concealing IBKE, the above scenario is no longer possible: if Bob does not wish to communicate with Alice (currently impersonated by Mallory), he may abort the iPAKE protocol before revealing his identity. We note that an attack is still possible if Mallory successfully guesses pw , since then she may impersonate any party, including those Bob wishes to communicate with. This is an unavoidable property of a password-based identity system.

3.7 Security

We informally discuss the security of LATKE before stating the main security theorem. To build intuition, we explain why each security assumption is necessary for LATKE.

SK-security. The base security of the IBKE prevents the most basic attacks, e.g., determining the session key of a passively observed session, impersonating a party who has not been compromised, or forcing two unrelated sessions to have the same session key.

Key-compromise impersonation resistance. If an adversary, Mallory, compromises Alice, then the only party Mallory should be able to impersonate

⁸It is possible to go even further with this idea. Some protocols which are responder-concealing (resp. initiator-concealing) are also passively initiator-concealing (resp. passively responder concealing). An example is SIGMA-R [Kra03]. The transform described here does not preserve the passive concealing property. But if the protocol is modified so certificates are sent at the exact same time and encrypted under the same key as the user-identifying information in the underlying protocol, the both these properties are preserved. We will do exactly this in our experiments (Section 4).

(barring a brute-force attack) is Alice herself. In particular, she should not be able to impersonate Bob to Alice. This is captured by our KCIR requirement.

Full forward secrecy. Mallory can use a stolen mpk to corrupt the initial PAKE and passively observe the *cleartext* IBKE execution between two uncompromised parties, Alice and Bob. If Mallory records this execution, cracks the password associated with mpk (thus getting msk), and later corrupts Alice and Bob, she should still not be able to later determine the session key. This is captured by our full FS requirement.

OnlineComparePwd. The above scenario also elucidates why `OnlineComparePwd` is necessary in our ideal functionality. If the UC simulator is to simulate the cleartext IBKE session, it must know whether the stolen mpk matches one of the parties' PAKE inputs. This is ordinarily captured by the `Impersonate` procedure, which forces the target party to believe the stolen mpk 's owner is their peer. However, it is not immediately clear when Mallory performs the attack if she intends to impersonate a party or to merely pass messages transparently. Since the latter case would not result in a change of perceived peer, the UC simulator cannot call `Impersonate` at that moment, and must use `OnlineComparePwd` instead.

SIM*-AC-CCA security and programmable ROM. As with `OnlineComparePwd`, we require this symmetric security notion and ideal primitive for UC simulatability reasons. Suppose Alice and Bob are both uncompromised and executing the protocol. In this case, the simulator does not know anything about the parties' mpk : they may be unequal, equal to each other, equal to other parties with known passwords/ mpk , or none of the above. Suppose Mallory waits until the first party, wlog Bob, outputs a session key K'' and sends the final (encrypted) message c to Alice. Then Mallory may immediately corrupt Alice, forcing the simulator to create an internal state which can explain c in the context of the protocol, and also demonstrate that the session key is indeed K'' . The difficulty here is twofold: first, the simulator must find a message m that makes sense and results in the output K'' , and second, the simulator must be able to provide a decryption key that opens c to m . The first issue is resolved by the random oracle at the end of the protocol: the simulator can construct any transcript it wants, and simply program $H_1(K', \text{tr}) := K''$.⁹ The second issue is resolved by our choice of a non-committing encryption primitive, which achieves SIM*-AC-CCA security.

Main theorem. In addition to three random oracles, H_0, H_1, H_2 , we require in our model the ideal primitive P that is used by the underlying SIM*-AC-CCA-

⁹This is resolved differently by Canetti and Krawczyk in their paper tying UC authenticated key exchange to SK-security [CK02b]. They define the *ACK property*—an AKE has the ACK property iff, once the first party outputs the session key, the internal states of both parties are simulatable using only the session key and public information. They show that this property is necessary to achieve UC simulatability, and describe a simple, generic transform to endow an SK-secure AKE with the ACK property, at the cost of one communication round. We are able to circumvent this overhead by using a programmable random oracle on our session key outputs. Since we must use a programmable model for iPAKE regardless [Hes20], this is essentially free.

secure authenticated encryption scheme (for AES-CTR + HMAC, for example, this is an ideal cipher and a random oracle). A proof of the following can be found in Appendix B.

Theorem 1. *Let IBKE be a post-specified peer IBKE with SK-security with KCIR and full FS in the identity-based CK, eCK, CK_{HMQV}, or CK⁺ model. Let AE^P be a SIM*-AC-CCA-secure symmetric encryption scheme with ideal primitive P. Then the LATKE^{post} protocol realizes $\mathcal{F}'_{\text{iPAKE}}$ in the $(\mathcal{F}_{\text{PAKE}}, \mathcal{F}_{\text{RO}}, \mathcal{P})$ -hybrid model with erasure and adaptive corruptions.*

Similarly, if IBKE is a pre-specified peer IBKE with SK-security in these models, then LATKE^{pre} realizes $\mathcal{F}'_{\text{iPAKE}}$ in the $(\mathcal{F}_{\text{PAKE}}, \mathcal{F}_{\text{RO}}, \mathcal{P})$ -hybrid model with erasure and adaptive corruptions.

We note that the theorem accepts a range of IBKE models whose security statements are known to be incomparable [Cre11]. In our proof, we avoid issues by using the narrowest set of adversary powers possible, a set which all models happen to share. In fact, the proof does not even require an IBKE security model with a notion of ephemeral key revelation. This permits a wider class of IBKE constructions, e.g., those constructed from AKEs with the forward secrecy transform from [BN11]. Our reduction is helped by the coarseness of the UC iPAKE model, which has strong session identifiers and only two types of corruption—long-term key corruption and total corruption.

Finally, our proof is written for LATKE^{post}, but it is not dependent on the pre-specified peer model, and transfers with little modification (given in Appendix B). Thus, we achieve a general result for any peer model and any commonly used notion of IBKE security.

4 Experiments

To demonstrate the practicality of LATKE, we instantiate it using a variety of PAKEs and IBKEs with different cryptographic assumptions and round complexities, and perform microbenchmarks on the resulting protocols. In particular, we instantiate the first post-quantum, identity-concealing iPAKE. For fair comparison, we also instantiate CHIP with the same underlying primitives where applicable. All our primitives are chosen to meet 128-bit security.

4.1 Choosing Primitives

In this section we list the PAKEs, IBKEs, and other primitives we chose for experimentation. For hashing, message authentication, and key derivation we use the Blake2b [ANWW13], HMAC [BCK96], and HKDF [Kra10] functions, respectively. For SIM*-AC-CCA-secure encryption, we use the encrypt-then-MAC construction [JT20] with the ChaCha20 stream cipher [Ber08] and the HMAC-Blake2b MAC. For simple Diffie-Hellman key agreement, we use the X25519 key-exchange protocol [Ber06]. For any protocol requiring prime-order group operations, we use the Ristretto255 group [VGH⁺23].

PAKEs. We now describe the PAKEs we used to instantiate LATKE and CHIP.

For our pre-quantum PAKE, we use CPace [HL19, AHH21], a one-round Diffie-Hellman-based construction. We also implemented and benchmarked KC-SPAKE2 [Sho20], but found it had higher communication cost, higher round complexity, and worse runtime efficiency. Thus, we omit KC-SPAKE2 entirely from our analysis.

For our post-quantum PAKE, we use CAKE [BCP⁺23], a three-round generic construction which can be instantiated from any fuzzy, key-anonymous IND-CPA-secure KEM and any keyed permutation (standing in for an ideal cipher). For the keyed permutation, we use Kravatte, an instantiation of the Farfalle wide block cipher over the Keccak permutation [BDH⁺17]. For the KEM, we use Saber [DKRV18],¹⁰ whose ciphertexts and public keys pack perfectly into bytes, and so is compatible with CAKE with a wide-block cipher-based keyed permutation.

IBKE. We now describe the IBKEs we used to instantiate LATKE and CHIP. We remark that we benefit greatly from the generality of Theorem 1. The protocols used below are proven in the id-CK, CK, eCK, and CK_{HMQV} models, respectively.

For more direct comparison with CHIP, we use the same Fiore-Gennaro IBKE that CHIP uses [FG10]. However, the IBKE only provides weak forward secrecy. To make it usable for LATKE, we add key confirmation to both sides using a MAC, as described in Section 3.5. We call the key confirmation variant FgC.

For identity-concealing and post-quantum security, we instantiate LATKE with the SIGMA-R responder-concealing protocol, due to Krawczyk [Kra03]. This is the only post-specified peer AKE we test. We apply the modification suggested by Peikert [Pei14] to base the protocol on a generic IND-CPA-secure KEM, rather than Diffie-Hellman. Finally, we apply our AKE-to-IBKE transform to the protocol, while preserving responder concealment. We illustrate the protocol in full in Figure 6. SIGMA-R and the IBKE transform both require signatures, so we benchmark with respect to the Ed25519 [BDL⁺12] and Dilithium [BDK⁺] signature schemes.¹¹

We also instantiate LATKE with the one-round AKE protocol by Bergsma et al. [BJS15] that, at its core, is signed Diffie-Hellman. We apply the same AKE-to-IBKE transform, using Ed25519 signatures, to this protocol.

Finally, we instantiate LATKE with the HMQV-C protocol (i.e., HMQV with key confirmation), described by Krawczyk [Kra05]. We apply the same

¹⁰Saber is very similar to Kyber [BDK⁺18] in construction, with the only essential difference being its reliance on the Module Learning with Rounding (MLWR) problem rather than Module Learning with Errors (MLWE). Their similarity allows CAKE’s required fuzziness and key-anonymity properties proven about Kyber in [BCP⁺23, Lemma 2] to transfer to Saber.

¹¹There is good reason to consider using a pre-quantum signature scheme in an otherwise post-quantum protocol. The *store now, decrypt later* threat model only applies to encryption that may be broken in the future, not authentication. Thus, it suffices to use pre-quantum signatures until a cryptographically relevant quantum computer is imminent.

iPAKE	PQ?	Con.?	Rounds	Comm.	Setup	Online	
Chip[Cspace,Fg] [CNPR22]	x	x	2	208B	284µs	5.33ms	(1×)
Latke ^{pre} [Cspace,FgC]	x	x	4	404B	314µs	5.56ms	(1.04×)
Latke ^{pre} [Cspace,IdHmqvC]	x	x	4	532B	467µs	5.62ms	(1.05×)
Latke ^{pre} [Cspace,IdSigDh]	x	x	2	616B	615µs	8.32ms	(1.56×)
Latke ^{post} [Cake,IdEd25519]	enc.	✓	6	3.53kB	813µs	5.46ms	(1.03×)
Latke ^{post} [Cake,IdEd1lith2]	enc.+auth.	✓	6	15.5kB	2.55ms	10.1ms	(1.89×)

Table 1: Performance characteristics of CHIP and LATKE, instantiated with varying IBKES and PAKES. The *Con.* column indicates whether the iPAKE has identity concealment. Reported online latencies are for the full protocol, excluding communication times.

transformation, also using Ed25519 signatures.

4.2 Experimental Setup

We now describe the hardware, software, and methodology of our benchmarks.

Hardware. Since the intended use case for iPAKE is in mesh networking protocols, we chose to conduct benchmarks on a commodity WiFi router. The router, a Linksys E8450 AX3200, has a 64-bit ARM Cortex-A53 CPU (late 2012), and runs OpenWrt [ope] snapshot r17758-b118efa0d2 (late 2021).

Software. All benchmarks were written in Rust, in a total of 2.9kloc.¹² We used the Criterion benchmarking framework to measure performance, and disabled SHA2 hardware acceleration. SIMD is not supported on the target chipset, and we did not use any other form of parallel execution.

Methodology. For each combination of CHIP / LATKE with PAKE and IBKE, we measured *online runtime*, i.e., how long it takes to run the online portion of the protocol from beginning to end, ignoring any communication costs. We separately measured communication costs. To have a fair comparison of the included pre- and post-specified peer protocols, we exclude the communication costs that come from sending identifier strings (this is at most 64B per protocol execution). We also measure *setup runtime*, i.e., the time it takes to run `StorePwdFile`, which is a one-time offline cost per device. To facilitate comparison between schemes, we do not use any hard password hashing functions in `StorePwdFile` (e.g., Argon2 or PBKDF2), though any realistic implementation must use one.

We report the medians of the recorded latencies. Across all benchmarks, then maximum observed relative standard error of the median was 0.5%.

4.3 Discussion

We show the results of our benchmarks in Table 1.

As expected, LATKE^{pre} with FgC performs closely to CHIP with Fg, both in setup and online time. The additional communication overhead of 196B can be

¹²Code is available at <https://github.com/rozbb/latke-ipake>

attributed to the additional key confirmation information that CHIP does not require, plus the ciphertext expansion due to EUE.

The `HmqvC` instantiation of LATKE performs similarly well, albeit with a slightly higher setup time and communication cost. The round-optimal instantiation with `IdSigDh` performs noticeably worse due to its reliance on signatures for authentication, rather than the implicit authentication of `FgC` and `HmqvC`. This is one area in which CHIP has a strong advantage—since CHIP does not require any form of forward secrecy from its IBKE, it can use efficient, one-round, implicitly authenticated IBKEs, while LATKE must use digital signatures or else increase its round complexity.

The `IdΣEd25519` post-quantum encryption is surprisingly close to CHIP in online runtime, albeit with nearly $17\times$ the communication cost. And, as expected, `IdΣDilith2` post-quantum encryption and authentication instantiation performed the worst across the board, with over $74\times$ the communication cost compared to CHIP, but surprisingly only $1.89\times$ the online runtime.

We also measured an optimization in which users save the certificate of the other party so that in future exchanges the user does not have to perform a signature verification. This saves $822\mu\text{s}$ on average, making `Latkepre[Cpace, IdHmqvC]` and `Latkepost[Cake, IdΣEd25519]` faster than `Chip[Cpace, Fg]` on every execution after the first (again, ignoring communication costs).

Finally, we note there are still performance gains to be had for the post-quantum SIGMA-R instantiation. The reason we use Saber for IBKE is simply for convenience in testing, since it is already used in CAKE. But there are post-quantum KEMs faster than Saber, such as NTTRU [LS19] and Kyber [BDK⁺18].

5 Future Work

Post-quantum siPAKE. Recall the CRISP protocol is a *strong* iPAKE (siPAKE), meaning that it is robust to precomputation attacks. In CHIP and LATKE, an adversary who precomputes a table of `mpk` values from various password guesses can use it to speed up a brute-force attack, should they ever retrieve an `mpk` from a user device. The CRISP construction uses algebraic techniques which rely on bilinear pairings in order to locally rerandomize its `pwfile` without breaking authenticity. It is not immediately clear if these properties can be obtained generically with an AKE and PAKE, nor if there exists a post-quantum analogue to the local rerandomization and pairing steps.

In addition, we reiterate the call for future work by the authors of CHIP/CRISP, that it would be useful to find a local rerandomization step that produces a problem for an attacker with tweakable hardness. Currently, CRISP forces the attacker to compute 1 pairing per guess—on the order of 1ms—whereas an OPRF-based strong construction like OPAQUE force an attacker to compute a hard hash per guess—on the order of 100ms.

Hybrid iPAKE. We have described LATKE using entirely post-quantum and entirely pre-quantum primitives, but the common path taken recently by standardization bodies is to define *hybrid* schemes, which combine two schemes

and enjoy the security of the harder of the two, even if we don't yet know which one that is. Since there already exist hybrid IBKEs with the properties necessary for LATKE, all that remains to build a hybrid iPAKE is to find a hybrid PAKE.

To the authors' knowledge, there is no known hybrid PAKE, let alone a generic method for creating one. Consider the concatenation transform, used in KEMs [BCD⁺24], whereby two KEMs are hybridized by simply running them in parallel, concatenating the final shared secrets, and hashing the concatenation. Suppose we do the same with two PAKEs which rely on a computational indistinguishability assumption for passive security, as CAKE does (specifically, LWE for fuzziness and key anonymity). Then the resulting hybrid PAKE's passive security is only as secure as the *least* secure of the two underlying PAKEs, since a passive attacker may simply pick which subprotocol transcript they wish to attack, and derive the password from that.

Round-optimal post-quantum iPAKE. As demonstrated in Section 4, pre-quantum LATKE can be achieved using only 2 rounds of communication. The same is true for post-quantum LATKE, but there are limited options for the underlying primitives.

For post-quantum one-round PAKE, the only efficient scheme that satisfies our security requirements appears to be EKE-NIKE [BGHJ24]. For post-quantum security, this can be instantiated with the CPA-secure variant of the Swoosh [GdKQ⁺23] NIKE, which is based on LWE (note, though, that communication costs are high because it requires high lattice dimension), or the CSIDH [CLM⁺18] NIKE, which is based on an isogeny walk assumption. Beyond this, there are PAKEs from isogeny assumptions [AEK⁺22, IY23], though, they are only proven in (an extension of) the BPR model [BPR00], not UC. Finally, there is a PAKE from (approximate) smooth projective hash functions and simulation sound NIZKs, which can both be realized from LWE assumptions [BBDQ18].

One-round post-quantum IBKEs are also difficult to instantiate. We may consider just AKE in our round complexity analysis, since our AKE-to-IBKE transform does not add rounds. Note a one-round AKE is equivalently a non-interactive key exchange (NIKE), since we can simply make the first round an exchange of the parties' public keys.

The only post-quantum AKE considered in Section 4 is SIGMA-R [Kra03], which has four rounds of communication. Schemes with two rounds are well known [BCNS15, ADPS16], but there are few schemes with just one round of communication. CSIDH appears to be the only plausibly efficient protocol in this space. The only other option in, the CCA-secure variant of Swoosh, requires each party to compute a non-interactive zero-knowledge proof over such vectors for every key exchange, and, as in the CPA-secure version, has relatively high communication costs.

Finally, as noted also by the CHIP/CRISP authors, it would be useful to find a one-round iPAKE, or else prove that one cannot be constructed.

Acknowledgements

This work was supported in part by NSF award CNS-2154705. Many thanks to Mark Schultz for inspiring conversations. Thanks also to Levi Schuck, Greg Rubin, Sanketh Menda, Doruk Gür, Cas Cremers, Shahar Paz, Eyal Ronen, Julia Hesse, and Steve Thomas for providing feedback and answering questions.

References

- [ADPS16] Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange - A new hope. In Thorsten Holz and Stefan Savage, editors, *USENIX Security 2016*, pages 327–343. USENIX Association, August 2016.
- [AEK⁺22] Michel Abdalla, Thorsten Eisenhofer, Eike Kiltz, Sabrina Kunzweiler, and Doreen Riepel. Password-authenticated key exchange from group actions. In Dodis and Shrimpton [DS22], pages 699–728. doi:10.1007/978-3-031-15979-4_24.
- [AHH21] Michel Abdalla, Björn Haase, and Julia Hesse. Security analysis of CPace. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part IV*, volume 13093 of *LNCS*, pages 711–741. Springer, Heidelberg, December 2021. doi:10.1007/978-3-030-92068-5_24.
- [All22a] Connectivity Standards Alliance. Matter Specification v1.0, September 2022. URL: https://csa-iot.org/wp-content/uploads/2022/11/22-27349-001_Matter-1.0-Core-Specification.pdf.
- [All22b] The Wifi Alliance. WPA3 Specification v3.1. 2022. URL: <https://www.wi-fi.org/system/files/WPA3%20Specification%20v3.1.pdf>.
- [ANWW13] Jean-Philippe Aumasson, Samuel Neves, Zooko Wilcox-O’Hearn, and Christian Winnerlein. BLAKE2: Simpler, smaller, fast as MD5. In Michael J. Jacobson Jr., Michael E. Locasto, Payman Mohassel, and Reihaneh Safavi-Naini, editors, *ACNS 13*, volume 7954 of *LNCS*, pages 119–135. Springer, Heidelberg, June 2013. doi:10.1007/978-3-642-38980-1_8.
- [App] Apple. Apple Platform Security, May 2022. URL: https://help.apple.com/pdf/security/en_US/apple-platform-security-guide.pdf.
- [Bas22] Andrea Basso. Poster: A post-quantum oblivious prf from isogenies. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, CCS ’22*, page 3327–3329, New

- York, NY, USA, 2022. Association for Computing Machinery. doi: 10.1145/3548606.3563542.
- [Bas23] Andrea Basso. A post-quantum round-optimal oblivious PRF from isogenies. Cryptology ePrint Archive, Report 2023/225, 2023. <https://eprint.iacr.org/2023/225>.
- [BBDQ18] Fabrice Benhamouda, Olivier Blazy, Léo Ducas, and Willy Quach. Hash proof systems over lattices revisited. In Michel Abdalla and Ricardo Dahab, editors, *Public-Key Cryptography – PKC 2018*, pages 644–674, Cham, 2018. Springer International Publishing.
- [BCD⁺24] Manuel Barbosa, Deirdre Connolly, João Diogo Duarte, Aaron Kaiser, Peter Schwabe, Karolin Varner, and Bas Westerbaan. X-wing. *IACR Communications in Cryptology*, 1(1), 2024. doi: 10.62056/a3qj89n4e.
- [BCK96] Mihir Bellare, Ran Canetti, and Hugo Krawczyk. Keying hash functions for message authentication. In Neal Koblitz, editor, *CRYPTO'96*, volume 1109 of *LNCS*, pages 1–15. Springer, Heidelberg, August 1996. doi:10.1007/3-540-68697-5_1.
- [BCNS15] Joppe W. Bos, Craig Costello, Michael Naehrig, and Douglas Stebila. Post-quantum key exchange for the TLS protocol from the ring learning with errors problem. In *2015 IEEE Symposium on Security and Privacy*, pages 553–570. IEEE Computer Society Press, May 2015. doi:10.1109/SP.2015.40.
- [BCP⁺23] Hugo Beguinet, Céline Chevalier, David Pointcheval, Thomas Ricosset, and Mélissa Rossi. Get a cake: Generic transformations from key encapsulation mechanisms to password authenticated key exchanges. In Mehdi Tibouchi and XiaoFeng Wang, editors, *Applied Cryptography and Network Security*, pages 516–538, Cham, 2023. Springer Nature Switzerland.
- [BDFH24] Ward Beullens, Lucas Dodgson, Sebastian Faller, and Julia Hesse. The 2Hash OPRF framework and efficient post-quantum instantiations. Cryptology ePrint Archive, Paper 2024/450, 2024. <https://eprint.iacr.org/2024/450>. URL: <https://eprint.iacr.org/2024/450>.
- [BDH⁺17] Guido Bertoni, Joan Daemen, Seth Hoffert, Michaël Peeters, Gilles Van Assche, and Ronny Van Keer. Farfalle: parallel permutation-based cryptography. *IACR Trans. Symm. Cryptol.*, 2017(4):1–38, 2017. doi:10.13154/tosc.v2017.i4.1-38.
- [BDK⁺] Shi Bai, Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS-Dilithium. URL: <https://pq-crystals.org/>

- dilithium/data/dilithium-specification-round3-20210208.pdf.
- [BDK15] Alex Biryukov, Daniel Dinu, and Dmitry Khovratovich. Argon2: the memory-hard function for password hashing and other applications. 2015.
- [BDK⁺18] Joppe Bos, Leo Ducas, Eike Kiltz, T Lepoint, Vadim Lyubashevsky, John M. Schanck, Peter Schwabe, Gregor Seiler, and Damien Stehle. Crystals - kyber: A cca-secure module-lattice-based kem. In *2018 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 353–367, 2018. doi:10.1109/EuroSP.2018.00032.
- [BDL⁺12] Daniel J. Bernstein, Niels Duif, Tanja Lange, Peter Schwabe, and Bo-Yin Yang. High-speed high-security signatures. *Journal of Cryptographic Engineering*, 2(2):77–89, September 2012. doi:10.1007/s13389-012-0027-1.
- [Ber06] Daniel J. Bernstein. Curve25519: New Diffie-Hellman speed records. In Moti Yung, Yevgeniy Dodis, Aggelos Kiayias, and Tal Malkin, editors, *PKC 2006*, volume 3958 of *LNCS*, pages 207–228. Springer, Heidelberg, April 2006. doi:10.1007/11745853_14.
- [Ber08] Daniel J. Bernstein. The ChaCha family of stream ciphers, January 2008. URL: <https://cr.yp.to/chacha.html>.
- [BGHJ24] Manuel Barbosa, Kai Gellert, Julia Hesse, and Stanislaw Jarecki. Bare pake: Universally composable key exchange from just passwords. Cryptology ePrint Archive, Paper 2024/234, 2024. <https://eprint.iacr.org/2024/234>. URL: <https://eprint.iacr.org/2024/234>.
- [BJS15] Florian Bergsma, Tibor Jager, and Jörg Schwenk. One-round key exchange with strong security: An efficient and generic construction in the standard model. In Jonathan Katz, editor, *PKC 2015*, volume 9020 of *LNCS*, pages 477–494. Springer, Heidelberg, March / April 2015. doi:10.1007/978-3-662-46447-2_21.
- [BKM⁺21] Andrea Basso, Péter Kutas, Simon-Philipp Merz, Christophe Petit, and Antonio Sanso. Cryptanalysis of an oblivious PRF from supersingular isogenies. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part I*, volume 13090 of *LNCS*, pages 160–184. Springer, Heidelberg, December 2021. doi:10.1007/978-3-030-92062-3_6.
- [BKW20] Dan Boneh, Dmitry Kogan, and Katharine Woo. Oblivious pseudorandom functions from isogenies. In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part II*, volume 12492 of *LNCS*, pages 520–550. Springer, Heidelberg, December 2020. doi:10.1007/978-3-030-64834-3_18.

- [BLR04] Boaz Barak, Yehuda Lindell, and Tal Rabin. Protocol initialization for the framework of universal composability. Cryptology ePrint Archive, Report 2004/006, 2004. <https://eprint.iacr.org/2004/006>.
- [BM92] Steven M. Bellovin and Michael Merritt. Encrypted key exchange: Password-based protocols secure against dictionary attacks. In *1992 IEEE Symposium on Security and Privacy*, pages 72–84. IEEE Computer Society Press, May 1992. doi:10.1109/RISP.1992.213269.
- [BN11] Colin Boyd and Juan González Nieto. On forward secrecy in one-round key exchange. In Liqun Chen, editor, *Cryptography and Coding*, pages 451–468, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.
- [Bon03] Dan Boneh, editor. *CRYPTO 2003*, volume 2729 of *LNCS*. Springer, Heidelberg, August 2003.
- [BPR00] Mihir Bellare, David Pointcheval, and Phillip Rogaway. Authenticated key exchange secure against dictionary attacks. In Bart Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 139–155. Springer, Heidelberg, May 2000. doi:10.1007/3-540-45539-6_11.
- [Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145. IEEE Computer Society Press, October 2001. doi:10.1109/SFCS.2001.959888.
- [CC07a] Zhaohui Cheng and Liqun Chen. On security proof of mccullaghbarreto’s key agreement protocol and its variants. *International Journal of Security and Networks*, 2(3-4):251–259, 2007. doi:10.1504/IJSN.2007.013178.
- [CC07b] Sherman S. M. Chow and Kim-Kwang Raymond Choo. Strongly-secure identity-based key agreement and anonymous extension. In Juan A. Garay, Arjen K. Lenstra, Masahiro Mambo, and René Peralta, editors, *Information Security*, pages 203–220, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg.
- [CJSV22] Ran Canetti, Palak Jain, Marika Swanberg, and Mayank Varia. Universally composable end-to-end secure messaging. In Dodis and Shrimpton [DS22], pages 3–33. doi:10.1007/978-3-031-15979-4_1.
- [CK01] Ran Canetti and Hugo Krawczyk. Analysis of key-exchange protocols and their use for building secure channels. In Birgit Pfizmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 453–474. Springer, Heidelberg, May 2001. doi:10.1007/3-540-44987-6_28.

- [CK02a] Ran Canetti and Hugo Krawczyk. Security analysis of IKE’s signature-based key-exchange protocol. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 143–161. Springer, Heidelberg, August 2002. <https://eprint.iacr.org/2002/120/>. doi:10.1007/3-540-45708-9_10.
- [CK02b] Ran Canetti and Hugo Krawczyk. Universally composable notions of key exchange and secure channels. In Lars R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 337–351. Springer, Heidelberg, April / May 2002. doi:10.1007/3-540-46035-7_22.
- [CLM⁺18] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. CSIDH: An efficient post-quantum commutative group action. In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part III*, volume 11274 of *LNCS*, pages 395–427. Springer, Heidelberg, December 2018. doi:10.1007/978-3-030-03332-3_15.
- [CNPR22] Cas Cremers, Moni Naor, Shahar Paz, and Eyal Ronen. CHIP and CRISP: Protecting all parties against compromise through identity-binding PAKEs. In Dodis and Shrimpton [DS22], pages 668–698. doi:10.1007/978-3-031-15979-4_23.
- [CR03] Ran Canetti and Tal Rabin. Universal composition with joint state. In Boneh [Bon03], pages 265–281. doi:10.1007/978-3-540-45146-4_16.
- [Cre11] Cas Cremers. Examining indistinguishability-based security models for key exchange protocols: the case of CK, CK-HMQV, and eCK. In Bruce S. N. Cheung, Lucas Chi Kwong Hui, Ravi S. Sandhu, and Duncan S. Wong, editors, *ASIACCS 11*, pages 80–91. ACM Press, March 2011.
- [DFG⁺23] Gareth T. Davies, Sebastian H. Faller, Kai Gellert, Tobias Handirk, Julia Hesse, Máté Horváth, and Tibor Jager. Security analysis of the WhatsApp end-to-end encrypted backup protocol. In Handschuh and Lysyanskaya [HL23], pages 330–361. doi:10.1007/978-3-031-38551-3_11.
- [DKRV18] Jan-Pieter D’Anvers, Angshuman Karmakar, Sujoy Sinha Roy, and Frederik Vercauteren. Saber: Module-LWR based key exchange, CPA-secure encryption and CCA-secure KEM. In Antoine Joux, Abderrahmane Nitaj, and Tajjeeddine Rachidi, editors, *AFRICACRYPT 18*, volume 10831 of *LNCS*, pages 282–305. Springer, Heidelberg, May 2018. doi:10.1007/978-3-319-89339-6_16.
- [Dod23] Lucas Dodgson. Post-Quantum Building Blocks for Secure Computation – the Legendre OPRF, September 2023. URL: <https://ethz>.

- ch/content/dam/ethz/special-interest/infk/inst-infsec/appliedcrypto/education/theses/Master_Thesis_Post_Quantum_Building_blocks_for_secure_computation.pdf.
- [DS22] Yevgeniy Dodis and Thomas Shrimpton, editors. *CRYPTO 2022, Part II*, volume 13508 of *LNCS*. Springer, Heidelberg, August 2022.
- [FG10] Dario Fiore and Rosario Gennaro. *Identity-Based Key Exchange Protocols without Pairings*, pages 42–77. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010. doi:10.1007/978-3-642-17499-5_3.
- [Fil18] Rick Fillion. Secure Remote Password (SRP): How 1Password uses it, 2018. Section: 1Password. URL: <https://blog.1password.com/developers-how-we-use-srp-and-you-can-too/>.
- [FOO23] Sebastian Faller, Astrid Ottenhues, and Johannes Ottenhues. Composable oblivious pseudo-random functions via garbled circuits. In Abdelrahman Aly and Mehdi Tibouchi, editors, *Progress in Cryptology – LATINCRYPT 2023*, pages 249–270, Cham, 2023. Springer Nature Switzerland.
- [FSXY12] Atsushi Fujioka, Koutarou Suzuki, Keita Xagawa, and Kazuki Yoneyama. Strongly secure authenticated key exchange from factoring, codes, and lattices. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012*, volume 7293 of *LNCS*, pages 467–484. Springer, Heidelberg, May 2012. doi:10.1007/978-3-642-30057-8_28.
- [GdKQ+23] Phillip Gajland, Bor de Kock, Miguel Quaresma, Giulio Malavolta, and Peter Schwabe. Swoosh: Practical lattice-based non-interactive key exchange. *Cryptology ePrint Archive*, Report 2023/271, 2023. <https://eprint.iacr.org/2023/271>.
- [GGJJ23] Kai Gellert, Kristian Gjøsteen, Håkon Jacobsen, and Tibor Jäger. On optimal tightness for key exchange with full forward secrecy via key confirmation. In Handschuh and Lysyanskaya [HL23], pages 297–329. doi:10.1007/978-3-031-38551-3_10.
- [GMR06] Craig Gentry, Philip MacKenzie, and Zulfikar Ramzan. A method for making password-based key exchange resilient to server compromise. In Cynthia Dwork, editor, *CRYPTO 2006*, volume 4117 of *LNCS*, pages 142–159. Springer, Heidelberg, August 2006. doi:10.1007/11818175_9.
- [Gon93] Li Gong. Lower bounds on messages and rounds for network authentication protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *ACM CCS 93*, pages 26–37. ACM Press, November 1993. doi:10.1145/168588.168592.

- [Gün90] Christoph G. Günther. An identity-based key-exchange protocol. In Jean-Jacques Quisquater and Joos Vandewalle, editors, *EUROCRYPT'89*, volume 434 of *LNCS*, pages 29–37. Springer, Heidelberg, April 1990. doi:10.1007/3-540-46885-4_5.
- [Ham15] Mike Hamburg. [curves] SPAKE2 and SPAKE2 Elligator Edition, 2015. URL: <https://moderncrypto.org/mail-archive/curves/2015/000424.html>.
- [Har08] Dan Harkins. Simultaneous authentication of equals: A secure, password-based key exchange for mesh networks. In *2008 Second International Conference on Sensor Technologies and Applications (sensorcomm 2008)*, pages 839–844, 2008. doi:10.1109/SENSORCOMM.2008.131.
- [HC09] Hai Huang and Zhenfu Cao. An ID-based authenticated key exchange protocol based on bilinear Diffie-Hellman problem. In Wanqing Li, Willy Susilo, Udaya Kiran Tupakula, Reihaneh Safavi-Naini, and Vijay Varadharajan, editors, *ASIACCS 09*, pages 333–342. ACM Press, March 2009.
- [Hes20] Julia Hesse. Separating symmetric and asymmetric password-authenticated key exchange. In Clemente Galdi and Vladimir Kolesnikov, editors, *SCN 20*, volume 12238 of *LNCS*, pages 579–599. Springer, Heidelberg, September 2020. doi:10.1007/978-3-030-57990-6_29.
- [HL19] Björn Haase and Benoît Labrique. AuCPace: Efficient verifier-based PAKE protocol tailored for the IIoT. *IACR TCHES*, 2019(2):1–48, 2019. <https://tches.iacr.org/index.php/TCHES/article/view/7384>. doi:10.13154/tches.v2019.i2.1-48.
- [HL23] Helena Handschuh and Anna Lysyanskaya, editors. *CRYPTO 2023, Part IV*, volume 14084 of *LNCS*. Springer, Heidelberg, August 2023.
- [IY23] Ren Ishibashi and Kazuki Yoneyama. Compact password authenticated key exchange from group actions. In Leonie Simpson and Mir Ali Rezazadeh Bae, editors, *ACISP 23*, volume 13915 of *LNCS*, pages 220–247. Springer, Heidelberg, July 2023. doi:10.1007/978-3-031-35486-1_11.
- [Jab97] David P. Jablon. Extended password key exchange protocols immune to dictionary attacks. In *6th IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE 1997)*, pages 248–255, Cambridge, MA, USA, June 18–20, 1997. IEEE Computer Society.
- [Jae23] Joseph Jaeger. Let attackers program ideal models: Modularity and composability for adaptive compromise. In Carmit Hazay

- and Martijn Stam, editors, *EUROCRYPT 2023, Part III*, volume 14006 of *LNCS*, pages 101–131. Springer, Heidelberg, April 2023. doi:10.1007/978-3-031-30620-4_4.
- [JKX18] Stanislaw Jarecki, Hugo Krawczyk, and Jiayu Xu. OPAQUE: An asymmetric PAKE protocol secure against pre-computation attacks. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 456–486. Springer, Heidelberg, April / May 2018. doi:10.1007/978-3-319-78372-7_15.
- [JT20] Joseph Jaeger and Nirvan Tyagi. Handling adaptive compromise for practical encryption schemes. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part I*, volume 12170 of *LNCS*, pages 3–32. Springer, Heidelberg, August 2020. doi:10.1007/978-3-030-56784-2_1.
- [KN09] Eike Kiltz and Gregory Neven. Identity-Based Signatures. 2009.
- [Kra03] Hugo Krawczyk. SIGMA: The “SIGn-and-MAC” approach to authenticated Diffie-Hellman and its use in the IKE protocols. In Boneh [Bon03], pages 400–425. doi:10.1007/978-3-540-45146-4_24.
- [Kra05] Hugo Krawczyk. HMQV: A high-performance secure Diffie-Hellman protocol. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 546–566. Springer, Heidelberg, August 2005. doi:10.1007/11535218_33.
- [Kra10] Hugo Krawczyk. Cryptographic extraction and key derivation: The HKDF scheme. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 631–648. Springer, Heidelberg, August 2010. doi:10.1007/978-3-642-14623-7_34.
- [LLHG23] Xiangyu Liu, Shengli Liu, Shuai Han, and Dawu Gu. EKE meets tight security in the Universally Composable framework. In Alexandra Boldyreva and Vladimir Kolesnikov, editors, *PKC 2023, Part I*, volume 13940 of *LNCS*, pages 685–713. Springer, Heidelberg, May 2023. doi:10.1007/978-3-031-31368-4_24.
- [LLM07] Brian A. LaMacchia, Kristin Lauter, and Anton Mityagin. Stronger security of authenticated key exchange. In Willy Susilo, Joseph K. Liu, and Yi Mu, editors, *ProvSec 2007*, volume 4784 of *LNCS*, pages 1–16. Springer, Heidelberg, November 2007.
- [Lou21] Karim Lounis. Cut it: Deauthentication attack on bluetooth. In *2021 14th International Conference on Security of Information and Networks (SIN)*, volume 1, pages 1–8, 2021. doi:10.1109/SIN54109.2021.9699265.

- [LS19] Vadim Lyubashevsky and Gregor Seiler. NTTRU: Truly fast NTRU using NTT. *IACR TCHES*, 2019(3):180–201, 2019. <https://tches.iacr.org/index.php/TCHES/article/view/8293>. doi:10.13154/tches.v2019.i3.180-201.
- [Mar16] Moxie Marlinspike. The Double Ratchet Algorithm, 2016. URL: <https://signal.org/docs/specifications/doubleratchet/>.
- [Oka88] Eiji Okamoto. Key distribution systems based on identification information. In Carl Pomerance, editor, *CRYPTO'87*, volume 293 of *LNCS*, pages 194–202. Springer, Heidelberg, August 1988. doi:10.1007/3-540-48184-2_15.
- [ope] Welcome to the OpenWrt Project. URL: <https://openwrt.org>.
- [Pei14] Chris Peikert. Lattice cryptography for the internet. In Michele Mosca, editor, *Post-Quantum Cryptography - 6th International Workshop, PQCrypto 2014*, pages 197–219. Springer, Heidelberg, October 2014. doi:10.1007/978-3-319-11659-4_12.
- [SE15] Jae Hong Seo and Keita Emura. Revocable hierarchical identity-based encryption: History-free update, security against insiders, and short ciphertexts. In Kaisa Nyberg, editor, *CT-RSA 2015*, volume 9048 of *LNCS*, pages 106–123. Springer, Heidelberg, April 2015. doi:10.1007/978-3-319-16715-2_6.
- [Shi03] Kyungah Shim. Efficient id-based authenticated key agreement protocol based on weil pairing. *Electronics Letters*, 39:653–654(1), April 2003. URL: https://digital-library.theiet.org/content/journals/10.1049/el_20030448.
- [Sho20] Victor Shoup. Security analysis of SPAKE2+. In Rafael Pass and Krzysztof Pietrzak, editors, *Theory of Cryptography*, pages 31–60, Cham, 2020. Springer International Publishing.
- [SPT13] Craig A. Shue, Nathanael Paul, and Curtis R. Taylor. From an IP address to a street address: Using wireless signals to locate a target. In *7th USENIX Workshop on Offensive Technologies (WOOT 13)*, Washington, D.C., August 2013. USENIX Association. URL: <https://www.usenix.org/conference/woot13/workshop-program/presentation/shue>.
- [SRV23] Domien Schepers, Aanjhan Ranganathan, and Mathy Vanhoef. Framing frames: Bypassing Wi-Fi encryption by manipulating transmit queues. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 53–68, Anaheim, CA, August 2023. USENIX Association. URL: <https://www.usenix.org/conference/usenixsecurity23/presentation/schepers>.

- [Tho22a] Steve Thomas. bscript: A Cache Hard Password Hash, 2022. URL: <https://tobtu.com/files/bsideslv2022.pdf>.
- [Tho22b] Steve Thomas. Demystifying Key Stretching and PAKEs. Black Hat 2022, Aug 2022. URL: <https://www.blackhat.com/us-22/briefings/schedule/#demystifying-key-stretching-and-pakes-27615>.
- [Thr15] Thread Group. Thread commissioning, July 2015. URL: https://www.threadgroup.org/Portals/0/documents/support/CommissioningWhitePaper_658_2.pdf.
- [VGH⁺23] Henry de Valence, Jack Grigg, Mike Hamburg, Isis Lovecruft, George Tankersley, and Filippo Valsorda. The ristretto255 and decaf448 Groups. Request for Comments RFC 9496, Internet Engineering Task Force, December 2023. URL: <https://datatracker.ietf.org/doc/rfc9496>.
- [Wan13] Yongge Wang. *Efficient Identity-Based and Authenticated Key Agreement Protocol*, pages 172–197. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013. doi:10.1007/978-3-642-35840-1_9.

A Deferred Preliminaries

Here we include the preliminaries necessary for our main proof in Appendix B.

A.1 Notation

We say a function $f : \mathbb{N} \rightarrow \mathbb{R}$ is *negligible in n* , denoted $f(n) = \text{negl}(n)$, iff $|f(n)| = n^{-\omega(1)}$ or, equivalently, for any $c > 0$, there is an n_0 such that $|f(n)| < n^{-c}$ for all $n > n_0$. We say a function $f : \mathbb{N} \rightarrow \mathbb{R}$ is *polynomial in n* , denoted $f(n) = \text{poly}(n)$, iff for some $M, d > 0$, there is an n_0 such that $|f(n)| \leq Mn^c$ for all $n > n_0$.

A.2 Probability

A *probability ensemble* \mathcal{S} is an infinite sequence of probability distributions $\mathcal{S}_1, \mathcal{S}_2, \dots$, where each \mathcal{S}_i is over a set $A_i \subset \{0, 1\}^{\ell(i)}$, where $\ell(n) = \text{poly}(n)$ is some length function. The *statistical distance* between two probability distributions $\mathcal{S}, \mathcal{S}'$ over sets A, A' is

$$\Delta(\mathcal{S}, \mathcal{S}') = \frac{1}{2} \sum_{x \in A \cup A'} |\Pr[y = x \mid y \leftarrow \mathcal{S}] - \Pr[y = x \mid y \leftarrow \mathcal{S}']|.$$

We say that two probability ensembles $\mathcal{S}, \mathcal{S}'$ are *statistically indistinguishable*, denoted $\mathcal{S} \stackrel{s}{\approx} \mathcal{S}'$ iff $\Delta(\mathcal{S}_n, \mathcal{S}'_n) = \text{negl}(n)$. We say that two probability ensembles

$\mathcal{S}, \mathcal{S}'$ are *computationally indistinguishable*, denoted $\mathcal{S} \stackrel{c}{\approx} \mathcal{S}'$, if for any probabilistic polynomial-time (PPT) distinguisher $|\Pr[\mathcal{D}(\mathcal{S}) = 0] - \Pr[\mathcal{D}(\mathcal{S}') = 0]|$ is negligible. We say an event X in some probability ensemble \mathcal{S} occurs with *overwhelming probability* if the event fails to occur with negligible probability.

A.3 Canetti-Krawczyk Model

The authors of [CK02a] present a game-based approach to modeling adaptive adversaries in an authenticated key agreement protocol in the post-specified peer setting. The adversary is permitted to activate new sessions between honest parties, intercept their messages, inject messages, and corrupt them in various ways. We describe in more detail the post-specified peer CK model. Similar definitions exist for pre-specified peers, and can be found in [CK01].

A session activated at party \mathcal{P}_i (the *owner*) is denoted by the tuple $(\mathcal{P}_i, \text{role}, s, d)$, where *role* is the role (*initiator* or *responder*) of \mathcal{P}_i ; s is the unique session identifier; and d is the *destination address* of the intended peer, i.e., an abstract identifier that may define a network location or shared address where the intended peer may be found. For brevity, we omit *role* and d when referring to a session. At the end of a successful session, the party outputs a public triple $(\mathcal{P}_i, s, \mathcal{P}_j)$ where \mathcal{P}_j is the (discovered) peer to the session, and a private session key K . It also erases all its *session state*, i.e., all the ephemeral data used to conduct the protocol. A session that has output its peer-key tuple is called *completed*. A session that has failed is called *aborted* and returns \perp .

The model permits the adversary \mathcal{A} to query the following oracles regarding the state of protocol sessions and participating parties.

NewSession $(\mathcal{P}_i, s, d, \text{role})$ Creates a new session for party \mathcal{P}_i with destination address d . If *role* = *initiator*, then \mathcal{P}_i will send an initiating message to the given address. If s is given, the new session is given ID s . If $s = \perp$, then a fresh ID is generated and returned to the caller.

Send (\mathcal{P}_i, s, m) Sends a message m to \mathcal{P}_i in session s . Returns the response of \mathcal{P}_i after processing the message according to the protocol.

Corrupt (\mathcal{P}_i) Reveals to \mathcal{A} the long-term keys of \mathcal{P}_i , the session keys of all its (unexpired) completed sessions, and the state of all its incomplete sessions.

RevealKey (\mathcal{P}_i, s) Reveals to \mathcal{A} the session key derived by \mathcal{P}_i in the *completed* session s .

RevealLtk (\mathcal{P}_i, s) Reveals to \mathcal{A} the long-term keys of \mathcal{P}_i .¹³

RevealState (\mathcal{P}_i, s) Reveals to \mathcal{A} the session state of \mathcal{P}_i in the given *incomplete* session s .

Expire (\mathcal{P}_i, s) *Expires* the completed session at the given party, i.e., erases the session key of s from \mathcal{P}_i 's memory.

¹³This oracle is not actually in the CK model. We show our extended model's equivalence to the base model in Appendix C

$\text{Test}(\mathcal{P}_i, s)$ Can only be called on a completed, unexpired session. Flips a coin b .
 If $b = 0$, reveals to \mathcal{A} the session key of s held by \mathcal{P}_i . If $b = 1$, sends to \mathcal{A} a uniform value from the session key space.

In order to rule out trivial attacks, we must prevent the adversary from corrupting a session's peer and using it to win the test session at the owner. We give the definition of *matching session* from [CK02a].

Definition (Matching session). *Let (\mathcal{P}_i, s) be a completed session with public output $(\mathcal{P}_i, s, \mathcal{P}_j)$. The session (\mathcal{P}_j, s) is called the matching session of (\mathcal{P}_i, s) if either*

1. (\mathcal{P}_j, s) is not completed, or
2. (\mathcal{P}_j, s) is completed and its public output is $(\mathcal{P}_j, s, \mathcal{P}_i)$.

Definition (Session exposure). *A completed session (\mathcal{P}_i, s) with public output $(\mathcal{P}_i, s, \mathcal{P}_j)$ is exposed if any of the following holds:*

1. RevealKey was called on the session or its matching session (if it exists),
2. RevealState was called on the session or its matching session (if it exists),
3. Corrupt or RevealLtk was called on \mathcal{P}_i or the owner of its matching session (if it exists) at any point.

We now give the basic notion of security in this model. We will build on this by adding notions of KCIR and full FS.

Definition (Session key (SK) security). *An authenticated key exchange protocol Π has session key security if the following hold*

1. Π is correct, i.e., if two uncorrupted parties complete matching sessions, then their session keys are equal with overwhelming probability.
2. The advantage of an adversary in distinguishing $b = 0$ versus $b = 1$ in an unexposed Test session is negligible.

Definition (Key compromise impersonation resistance (KCIR)). *An adversary \mathcal{A} breaks SK-security with KCIR against an AKE protocol iff it wins the SK-security game with the following extra capability: \mathcal{A} may learn the long-term key of the owner of the test session at any time (via, e.g., RevealLtk or Corrupt outside the lifetime of s).*

Definition (Forward secrecy (FS)). *An adversary \mathcal{A} breaks SK-security with FS against an AKE protocol iff it wins the SK-security game with the following extra capabilities: (1) \mathcal{A} may learn the long-term key of the owner of the test session after it is expired; (2) similarly, if the test session has a matching session, \mathcal{A} may learn the long term key of the owner of that session after it is expired.*

SK-security with KCIR and full FS is the combination of the above three definitions. Note that the combination of capabilities afforded by KCIR and FS gives the attacker the novel ability to get the long-term key of the owner of the test session, expire the session, and call Corrupt on the test session peer.

A.3.1 The id-CK Model

The CK model has been extended in prior works to apply to identity-based authenticated key exchange protocols [SE15, HC09, FSXY12]. The changes to CK are minimal:

1. Every party is given a unique identifier id , chosen by the adversary. In protocol outputs and session tuples, id_i replaces \mathcal{P}_i .
2. At the beginning of the game, the main keypair (mpk, msk) is generated. mpk is given to the adversary.
3. There is an additional oracle, RevealMsk , which reveals to \mathcal{A} the value of msk . The notion of exposure is extended to include: RevealMsk was called at any point.

Finally, since there is a new type of secret, msk , we extend FS to the id-CK setting.

Definition (Key generation center forward secrecy (KGC-FS)). *An adversary \mathcal{A} breaks SK-security with KGC-FS against an IBKE protocol iff it wins the SK-security game with the following extra capability: \mathcal{A} may call RevealMsk after both the test session is expired, and its matching session (if it exists) are expired.*

Note that the above definition is equivalent to the ordinary FS, treating RevealMsk as if it were a call to RevealLtk on both the owner of the test session and the owner of its matching session (if it exists).

Finally, we combine our notions of FS:

Definition (Full FS). *An IBKE achieves full FS iff it has FS and KGC-FS.*

A.4 Symmetric Encryption

LATKE requires a symmetric encryption scheme that is secure under adaptive corruptions (i.e., *non-committing*). For this, we use SIM*-AC-CCA encryption [JT20, Jae23]. Informally, an authenticated encryption scheme AE is SIM*-AC-CCA-secure with respect to some ideal primitive P (e.g., programmable random oracle or ideal cipher), if there is a simulator S_{cca} such that no efficient adversary can distinguish between S_{cca} and AE, even when given the ability to encrypt messages, decrypt messages, expose keys, and program the ideal primitive.

The procedures exposed by the symmetric encryption scheme AE are:

$\text{KeyGen}(1^\lambda) \rightarrow k$ Generates a fresh symmetric key.

$\text{Enc}_k^{\mathsf{P}}(m) \rightarrow c$ Produces an authenticated ciphertext for message m under key k . This procedure has access to ideal primitive P .

$\text{Dec}_k^{\mathsf{P}}(c) \rightarrow m$ Decrypts the given ciphertext using key k , returning the plaintext m or an error \perp . This also has access to ideal primitive P .

<p>Game $G_{F, S_{cca}, \mathcal{A}}^{\text{sim}^* \text{-ac-cca}}(\lambda)$</p> <p>$k_{(\cdot)} \leftarrow \text{AE.KeyGen}(1^\lambda)$</p> <p>$\sigma_P \leftarrow \text{P.Init}(1^\lambda)$</p> <p>$\sigma \leftarrow S_{cca}.\text{Init}(1^\lambda)$</p> <p>$b \leftarrow_{\\$} \{0, 1\}$</p> <p>$b' \leftarrow \mathcal{A}^{\text{Enc, Dec, Exp, PPrim}}(1^\lambda)$</p> <p>return $(b = b')$</p> <p>Exp(u)</p> <p>if $b = 1 : k' := k_u$</p> <p>else $: k' \leftarrow S_{cca}.\text{Exp}^{\text{PPrim}}(1^\lambda, u, M_u : \sigma)$</p> <p>$X_u := \text{true}$</p> <p>return k'</p>	<p>PPrim(Op, k, x, y)</p> <p>assert $\text{Op} \in \{\text{Ls}, \text{Prog}\}$</p> <p>$y \leftarrow \text{P.Op}(1^\lambda, k, x, y : \sigma_P)$</p> <p>return y</p> <p>Enc(u, m)</p> <p>if $\neg X_u : \ell := m , \text{else} : \ell := m$</p> <p>if $b = 1 : c \leftarrow \text{AE.Enc}_{k_u}^{\text{P}}(1^\lambda, m)$</p> <p>else $: c \leftarrow S_{cca}.\text{Enc}^{\text{PPrim}}(1^\lambda, u, \ell : \sigma)$</p> <p>$M_u.\text{add}(c, m)$</p> <p>return c</p> <p>Dec(u, c)</p> <p>if $M_u[c] \neq \perp : \text{return } M_u[c]$</p> <p>if $b = 1 : m \leftarrow \text{AE.Dec}_{k_u}^{\text{P}}(1^\lambda, c)$</p> <p>else $: m \leftarrow S_{cca}.\text{Dec}^{\text{PPrim}}(1^\lambda, u, c : \sigma)$</p> <p>return m</p>
--	---

Figure 7: The SIM*-AC-CCA security game, for authenticated encryption scheme AE, ideal primitive P, simulator S_{cca} , and adversary \mathcal{A} .

The simulator S_{cca} associated with AE exposes the following procedures. We use σ to denote state that S_{cca} maintains. We mark the separation between this state and other inputs using ‘:’, as in [JT20, Jae23] (this is done for technical reasons beyond the scope of this overview).

$\text{Init}(1^\lambda) \rightarrow \sigma$ Creates the simulator’s state

$\text{Dec}(u, c : \sigma) \rightarrow m$ Simulates a party u decrypting a ciphertext c .

$\text{Exp}(u, M_u : \sigma) \rightarrow k$ Simulates the corruption, i.e., exposure of a key, of a party u . M_u is the set of ciphertext-plaintext pairs that the simulator must behave consistently with.

In LATKE, specifically EUE, each round of each subsession has a unique key. Thus, we will let the game’s user identifier u be a tuple containing the party identifier, the subsession identifier, and the current step in the protocol execution: $u = (\mathcal{P}, \text{ssid}, \text{step})$. When \mathcal{P} is corrupted in our LATKE simulator, it will trigger the exposure of all its incomplete subsessions.

We reproduce the SIM*-AC-CCA security game from [Jae23] in Figure 7. The game permits the adversary to encrypt messages, decrypt ciphertexts, expose keys, and query and program the ideal primitive P. When $b = 0$, the adversary \mathcal{A} sees the real world, i.e., the behavior of the authenticated encryption scheme

AE. When $b = 1$, \mathcal{A} sees the ideal world, i.e., the behavior of the simulator S_{cca} . AE is *SIM*-AC-CCA-secure* if there exists a simulator S_{cca} such that for any PPT adversary \mathcal{A} , \mathcal{A} 's advantage in distinguishing these two worlds is negligible.

Concrete instantiation. We can instantiate a SIM*-AC-CCA-secure encryption scheme via the encrypt-then-MAC (EtM) transform [JT20]. If SE is SIM*-AC-CPA-secure with primitive P and MAC is UF-CMA*-secure (a weaker notion than SIM*-AC-PRF-secure) with primitive P, then EtM[SE, MAC] is SIM*-AC-CCA-secure with primitive P [Jae23, Lemma 2].

Ideal ciphers and random oracles are both SIM*-AC-PRF-secure [Jae23, Lemma 3]. Thus, if we model AES as an ideal cipher or the ChaCha20 block function [Ber08] as a random oracle, and use them in CTR mode, by Lemma 2, we have a SIM*-AC-CPA-secure encryption scheme. If the hash function H is modeled as a random oracle, then HMAC[H] is UF-CMA*-secure. Putting these together, Lemma 2 says that AES-CTR + HMAC and ChaCha20 + HMAC are SIM*-AC-CCA-secure authenticated encryption algorithms.

B Proof of Main Theorem

We prove the main theorem of the paper, Theorem 1, which claims LATKE UC-realizes \mathcal{F}_{iPAKE} given an IBKE and a symmetric encryption scheme. The proof consists of multiple game hops, beginning with the real world, i.e., the protocol itself, and ending in the ideal world, i.e., where the simulator may only make calls to \mathcal{F}_{iPAKE} . For brevity we do not write out each game's simulator in full, but we include the final simulator in Figure 8. We will use $tr_{sid,ssid}^{pake}$ to denote the transcript of the PAKE in the appropriate subsession.¹⁴

We state the proof below for LATKE^{post}. The argument for LATKE^{pre} for pre-specified peer IBKEs is nearly identical. The only difference is that the simulator must use the IDs given in the initial PAKE to activate IBKE sessions in its SK-security reduction. This is a procedural change, and does not affect the soundness of any reduction.

Beyond CK. The below proof references the *id-CK** model as a stand-in for any of the identity-based variants of CK, eCK, CK_{HMQV}, or CK⁺ (adding the RevealLtk oracle as necessary, as explained in Appendix C). It is known that these models all differ [Cre11], but the differences are subtle and are finer than the somewhat coarse UC definitions we target. More specifically, the models differ in their notions of matching sessions (and hence, who is allowed to be compromised and which session keys can be revealed) and permissiveness of attack scenarios.

The differences in matching sessions lie in whether sessions match when their transcripts match, or their *ssid*'s and roles match, or some combination thereof. Our reduction does not corrupt or reveal the session key to sessions that satisfy

¹⁴This transcript exists in practice, but in our proof this is a fictive variable. This is because the real world is in the \mathcal{F}_{iPAKE} -hybrid model, and thus has no PAKE protocol transcript. To address this, it suffices to imagine an extension of \mathcal{F}_{iPAKE} that also outputs a transcript variable to the environment.

any of these definitions. We also avoid a correctness issue—whether sessions with matching `ssid` but non-matching transcripts will output the same key—by simply hashing the transcript into the final key.

The differences in attack permissiveness in these models comes down to the adversary’s ability to use short-term state revelation (called “state revelation” in CK-type models, and “ephemeral key revelation” in eCK-type models). Since our reduction never uses short-term state revelation, we are able to use the weakest notion of security, i.e., the least permissive to the adversary, in any of these models. In particular, this means that our notion of key compromise impersonation resistance and full forward secrecy are covered by those in any of these models.

Proof. **Game 0.** Real world

Game 1. \mathcal{S} simulates $\mathcal{F}_{\text{PAKE}}$ and \mathcal{P} . \mathcal{S} also simulates \mathcal{F}_{RO} using hash tables H_0, H_1, H_2 . This is perfectly indistinguishable from Game 0.

Game 2. We introduce $\mathcal{F}_{\text{iPAKE}}$ and have parties call $\mathcal{F}_{\text{iPAKE}}.\text{StorePwdFile}$ on initialization instead of directly generating its `mpk`. Instead, a party will generate its `mpk` lazily via `IBKE.Setup` on its first session initiation. After the setup, the simulator will simulate the rest of the session. In this game, the simulator must do two things: (1) consistently respond to corruption queries even for parties that haven’t participated in an active session, and (2) simulate honest parties in a session. For an honest party \mathcal{P}_i with password `pwi`, we denote its `IBKE.Setup` randomness by `coinsi`. \mathcal{S} will read from and program H_0 in order to make consistent choices for `coinsi`.

For `StealPwdFile` (resp. `Corrupt`) queries on parties who have not yet begun a session, \mathcal{S} calls $\mathcal{F}_{\text{iPAKE}}.\text{StealPwdFile}$ (resp. `Corrupt`) and receives `(Stolen, idi, pwi)`, where `pwi` = \perp if the password has not yet been guessed. \mathcal{S} must now determine that party’s randomness, `coinsi`. If `pwi` $\neq \perp$, \mathcal{S} sets `coinsi` := $H_0[\text{sid}, \text{pw}_i]$. Otherwise \mathcal{S} performs $\mathcal{F}_{\text{iPAKE}}.\text{OfflineComparePwd}$ queries to see if the stolen password file matches any other known stolen password files. If there are any matches, then `coinsi` is set to the other party’s `coins`. Otherwise, this a fresh compromise, so `coinsi` $\leftarrow_{\mathcal{S}} \{0, 1\}^m$. \mathcal{S} then saves a record `(KnownCoins, $\mathcal{P}_i, \text{id}_i, \text{coins}_i$)`. Now that `coinsi` is set, \mathcal{S} may generate the main keypair `mski, mpki` and extract the user keys `uski, upki` with respect to `idi`. \mathcal{S} returns all these values to the environment.

To maintain consistency of `coins`, we must also modify how \mathcal{S} simulates H_0 queries. On query `(sid, pw)`, \mathcal{S} calls $\mathcal{F}_{\text{iPAKE}}.\text{OfflineTestPwd}$ for each compromised party to test if `pw` produces their stolen password. If there is a match on party \mathcal{P}_j , then \mathcal{S} sets $H_0[\text{sid}, \text{pw}] := \text{coins}_j$, where `coinsj` comes from the record `(KnownCoins, $\mathcal{P}_j, \cdot, \text{coins}_j$)`.

Finally, \mathcal{S} simulates honest parties in the protocol as follows. For a session with an honest party \mathcal{P}_i , \mathcal{S} receives `mpki` via its simulation of $\mathcal{F}_{\text{PAKE}}.\text{NewSession}$. From `mpki`, \mathcal{S} can find the corresponding `pwi` by testing which `pw'` yields `(mpki, \cdot) = IBKE.KeyGen($1^\lambda; H_0[\text{sid}, \text{pw}']$). This is guaranteed to exist because an honest party in an active session must have called IBKE.KeyGen on a password. If the user keypair uski, upki has not yet been extracted for the party, \mathcal{S} does so`

now via `IBKE.Extract`, and save it for future simulations. \mathcal{S} then compares the `mpk` provided by both parties in the PAKE session. If the PAKE fails, \mathcal{S} does not need to know the parties' secrets in order to simulate. If the PAKE succeeds, then the `mpk` values are the same, and if one of them is honest, then \mathcal{S} knows the `msk` for both, and can thus generate user keypairs for both sides and use them for a complete IBKE simulation $A(\text{mpk}, \text{upk}_A, \text{usk}_A) \Leftrightarrow B(\text{mpk}, \text{upk}_B, \text{usk}_B)$. Since the functionality is symmetric, we assume for simplicity that the first party to `NewSession` is the initiator in the IBKE.

This simulation is perfect.

Game 3. Rather than returning the session key directly to the parties, \mathcal{S} now uses `$\mathcal{F}_{\text{IPAKE}}$.NewKey`. Since this procedure will only use the provided key if the session is marked `compromised`, the simulator must call the appropriate compromising functions, `Impersonate` and `OnlineTestPwd`, when possible.

We define the simulator to mark a compromise at multiple points. Firstly, any time an `mpk` is provided to the protocol by a corrupted party (in the next step uncorrupted parties will not provide `mpk` at all), \mathcal{S} checks if there is a `pw` such that $(\text{mpk}, \cdot) = \text{IBKE.KeyGen}(1^\lambda; \text{coins})$, where $\text{coins} = \text{H}_0[\text{sid}, \text{pw}]$. If there is a match, \mathcal{S} calls `$\mathcal{F}_{\text{IPAKE}}$.OnlineTestPwd` with that password. This is done in the simulation of `$\mathcal{F}_{\text{PAKE}}$.TestPwd` and `$\mathcal{F}_{\text{PAKE}}$.NewKey`. Secondly, if an party uses a compromised `mpk`, i.e., one that is the result of `IBKE.KeyGen` with `coinsk` for some recorded `coinsk` belonging to party \mathcal{P}_k , and the session is tampered with by the adversary, then \mathcal{S} calls `$\mathcal{F}_{\text{IPAKE}}$.Impersonate` on the corresponding party as \mathcal{P}_k .¹⁵ Similarly, if both parties are compromised, then `Impersonate` is called on both. `Impersonate` is also done in the simulation of `$\mathcal{F}_{\text{PAKE}}$.TestPwd` and `$\mathcal{F}_{\text{PAKE}}$.NewKey`. Finally, if an honest party \mathcal{P}_i is corrupted during a session with another honest party, \mathcal{S} calls `$\mathcal{F}_{\text{PAKE}}$.Impersonate` as \mathcal{P}_i , “impersonating” itself.

Game 2 \approx Game 3. We argue that this is indistinguishable from the last game, assuming IBKE is SK-secure with KCIR and full FS.

Let \mathcal{Z} be an environment that can distinguish between the two games. We will use \mathcal{Z} to construct an adversary \mathcal{A} against the SK-security game with KCIR and full FS. We will show the reduction for a sequence of hybrids. In hybrid i , all sessions $j \geq i$ return the key directly to the user, and all sessions $j < i$ call `$\mathcal{F}_{\text{IPAKE}}$.NewKey`. Hybrid 0 is game 2, and hybrid N is game 3, where N is the (polynomial) number of sessions in the experiment.

Note that, between these two games, the only time `NewKey` behaves differently (i.e., sending a random key) is when, at the end of the protocol execution, both parties are uncorrupted and at most one is compromised.¹⁶ In addition, `NewKey` is only different when `mpk` doesn't come from a known password (else \mathcal{S} runs

¹⁵The tampering condition is to handle the following scenario. If Alice is compromised and Mallory breaks the PAKE session between Alice and Bob, but otherwise does not tamper with any messages, then the session should succeed. That is, `NewKey` must produce the same key for both parties. If Mallory modifies or injects messages, though, then Bob's output must be consistent with Mallory's view, so \mathcal{S} must `Impersonate` to Bob. If Bob is not compromised, then this breaks the symmetry of the keys. The final transcript hash helps us break this symmetry.

¹⁶Note that this still permits an adversary to participate in the encrypted portion of the protocol, assuming they passed the PAKE using a stolen `mpk`.

OnlineTestPwd and wins). Finally, NewKey is only different when the initial PAKE succeeds, i.e., both parties submitted the same mpk (otherwise both parties output randomness). Thus, for hybrid i , we may assume wlog that the session in question ends with both parties uncorrupted, at most one party compromised, $\text{H}_0[\text{sid}, \text{pw}]$ belonging to the mpk not queried, and both parties submitted the same mpk to the PAKE. We can safely assume this because, outside of this scenario, the simulator’s success probability is 1.

We reduce hybrid i from id-CK* now, defining an adversary \mathcal{A} . Let mpk^* be the main public key provided by the id-CK* challenger. Let session i be a session between parties \mathcal{P} and \mathcal{P}' . This will be the Test session. We program the mpk of both \mathcal{P} and \mathcal{P}' to be mpk^* . We use the id-CK* Send oracle when simulating the sending of messages in every session using mpk^* as \mathcal{P} and \mathcal{P}' , CK.RevealTk for StealPwdfFile queries, CK.Corrupt for Corrupt queries (by hypothesis, on parties other than \mathcal{P} and \mathcal{P}'), and RevealKey to determine the final key of all other sessions. If the adversary did not modify the IBKE messages during a session (either by knowing the PAKE key and modifying the messages, or not knowing the PAKE key and modifying the ciphertext), then the iPAKE output key for *both* parties is the one given by Test. This corresponds to the condition that matching completed sessions between honest parties must produce the same key, and non-matching sessions overwhelmingly do not produce the same key (otherwise an adversary wins by simply calling RevealKey on the non-matching session). We return these as described above, based on the session number. For the test session (which is clean by hypothesis), we simply return the challenge to \mathcal{P} and \mathcal{P}' . Clearly, $b = 0$ is hybrid $i - 1$ and $b = 1$ is hybrid i . \mathcal{A} returns whatever \mathcal{Z} guesses at the end of the UC game. Thus the advantage is precisely the SK-security with KCIR and full FS advantage.

Game 4. We replace AE with the simulator \mathcal{S}_{cca} whenever a PAKE completed and at most one of the parties is compromised. Recall, rather than using just the party \mathcal{P}_i as a user identifier for \mathcal{S}_{cca} , we use $(\mathcal{P}_i, \text{ssid}, \text{step})$, where step refers to the current step in the protocol. This corresponds to the fact that the real-world protocol uses a unique key for every message.

Recall the simulator knows the msk and usk of any party in a session where at least one compromise has occurred. Thus, it knows the output of the PAKE in these scenarios and does not need to use any indistinguishability property. Similarly, the simulator learns the value of the PAKE if $\mathcal{F}_{\text{PAKE}}.\text{TestPwdf}$ is called on the session. The only case we must handle is where, during the EUE phase of a session with two honest parties and an uncompromised PAKE session, one party is corrupted.

When Corrupt is called on \mathcal{P}_i , \mathcal{S} uses the long term keys to check if the PAKE failed or if any messages were modified by the adversary. Either of these events imply a decryption error in the real world. Thus, the simulator creates a valid IBKE transcript between \mathcal{P}_i and its peer \mathcal{P}_j , up until the point of failure (if at all), and then, per EUE, pads the rest of the transcript with zeros. The simulator selects a random chain key K_{ch} and programs the decryption key $k = \text{AE.Exp}(m, c)$, where c is the last ciphertext to be received by \mathcal{P}_i in the

protocol and m last message to be received by \mathcal{P}_i in the fabricated transcript.

After corruption, \mathcal{S} switches to using AE and the revealed key. If the corruption occurred after the last message was sent, then \mathcal{S} programs $H_1(K', \text{tr})$ to be the random session key previously chosen.

In order to make the keys consistent, we must also program H_2 so that $H_2(\text{sid}, \text{ssid}, 1 \| K_{\text{ch}}^{(i)}) = (K_{\text{ch}}^{(i+1)}, k_{i+1})$ for every chain key $K_{\text{ch}}^{(i)}$ and exposed encryption key k_i , and $H_2(\text{sid}, \text{ssid}, 0 \| K) = K_{\text{ch}}^0$ for the initial key.

Game 3 \approx Game 4. We reduce this hop to the SIM*-AC-CCA security of AE. The only bad situation to handle is if H_2 has already been called on one of the chain keys. Since the PAKE keys are unknown to the adversary by hypothesis, and all keys are defined in a random oracle hash chain, this probability is negligible.

We can show this game reduces to SIM*-AC-CCA using a sequence of hybrids. We define hybrid i as the world where every iPAKE subsession $\leq i$ is defined without the simulator, and each $> i$ is. Each game hop goes from defining a ciphertext $c := \text{AE.Enc}_k(m)$ for a known key k and message m , to defining $c := \text{S}_{\text{cca}}.\text{Enc}$ and later generating k as the opening of c to m . Further, all tampered ciphertexts are assumed by the simulator to trigger a decryption error. This is precisely the SIM*-AC-CCA security game. Thus, the advantage of any adversary in the larger game hop is bounded above by $N \cdot \text{Adv}_A^{\text{sim}^*-\text{ac}-\text{cca}}$, where N is the number of sessions in the game.

Game 5. We now remove the ability of \mathcal{S} to see the PAKE keys for sessions with participants who are honest up to and including $\mathcal{F}_{\text{PAKE}}.\text{NewKey}$. On activation, each party calls $\mathcal{F}_{\text{iPAKE}}.\text{NewSession}$ instead of directly calling $\mathcal{F}_{\text{PAKE}}.\text{NewSession}$. Using NewSession means that OnlineTestPwd can return meaningful answers.

Since the specific value of the PAKE keys no longer matters (since S_{cca} allows \mathcal{S} to generate the appropriate hash chain key), the thing \mathcal{S} loses in this hop is knowing whether a PAKE session succeeded, i.e., both parties got the same key. Specifically, \mathcal{S} must know this for sessions with at least one compromised party in order for the above steps to work. Thus, it will suffice to define a way for \mathcal{S} to determine the success of any such PAKE session.

We define \mathcal{S} to check every possible value to determine whether the PAKE between \mathcal{P}_i and \mathcal{P}_j succeeded. If the parties are both corrupted, then they provided their own mpk values, and those can be compared directly. If only one of them is corrupted, then that supplied mpk value is checked against known compromised users (use OnlineComparePwd) and compromised passwords (use OnlineTestPwd). Finally, if one party is compromised and not corrupted then its mpk is checked against the other's (use OnlineTestPwd). This new simulator covers every edge case using just $\mathcal{F}'_{\text{iPAKE}}$ functionality, and so this hop is perfectly indistinguishable from the last. This completes our proof. \square

<p>On corruption query (StealPwdfFile, sid, \mathcal{P}_i) from \mathcal{Z}</p> <pre> send (StealPwdfFile, sid, \mathcal{P}_i) to $\mathcal{F}_{\text{IPAKE}}$ if received “no record” : send “no record” to \mathcal{Z} and exit else received (Stolen, id_i, pw_i) if $\text{pw}_i \neq \perp$, $\text{coins}_i := H_0[\text{sid}, \text{pw}_i]$ else : $\text{coins}_i \leftarrow_{\\$} \{0, 1\}^m$ for each record (KnownCoins, $\mathcal{P}_j, \cdot, \text{coins}_j$) : send (OfflineComparePwdf, $\mathcal{P}_i, \mathcal{P}_j$) to $\mathcal{F}_{\text{IPAKE}}$ if received “match”, $\text{coins}_i := \text{coins}_j$ record (KnownCoins, $\mathcal{P}, \text{id}, \text{coins}_i$) ($\text{msk}_i, \text{mpk}_i$) := IBKE.Setup($1^\lambda$; coins_i) ($\text{usk}_i, \text{upk}_i$) \leftarrow IBKE.Extract$_{\text{msk}_i}$(id_i) ibkeStates[$\cdot, \mathcal{P}_i, 0$] := (begin, $\text{msk}_i, \text{usk}_i$) Mark \mathcal{P}_i compromised for each (Session, ssid, $\mathcal{P}_i, \mathcal{P}_j, \cdot, \cdot$) marked pakecompleted and not complete : EueCatchup(sid, ssid) record (Pwfile, $\mathcal{P}_i, \text{mpk}_i, \text{upk}_i, \text{usk}_i$) send ($\text{mpk}_i, \text{upk}_i, \text{usk}_i$) to \mathcal{Z} </pre> <p>On corruption query (Corrupt, sid, \mathcal{P}_i) from \mathcal{Z}</p> <pre> ($\text{mpk}_i, \text{upk}_i, \text{usk}_i$) := S.StealPwdfFile(\mathcal{P}_i) curKeys := {} for each (Session, sid, ssid, $\mathcal{P}_i, \cdot, \cdot, \cdot$) not complete : curKeys.append((ssid, K_{ch}, k)) Mark \mathcal{P} corrupted and compromised send ($\text{mpk}_i, \text{upk}_i, \text{usk}_i, \text{curKeys}$) to \mathcal{Z} </pre> <p>On (TestPwdf, sid, ssid, $\mathcal{P}_i, \text{mpk}$) from \mathcal{Z} to $\mathcal{F}_{\text{IPAKE}}$</p> <pre> retrieve (Session, ssid, $\mathcal{P}_i, \mathcal{P}_j, \cdot, \text{mpk}_i$) retrieve (Session, ssid, $\mathcal{P}_j, \mathcal{P}_i, \cdot, \text{mpk}_j, \cdot$) succ := false if $\text{mpk} = \text{mpk}_i \neq \perp$: succ := true else : if received “correct” : succ := true if succ : Mark session compromised return “correct” else : Mark session interrupted return “wrong” </pre>	<p>On (NewSession, sid, ssid, $\mathcal{P}_j, \text{mpk}$) from corrupt \mathcal{P}_i to $\mathcal{F}_{\text{IPAKE}}$</p> <pre> if \nexists a record (Session, ssid, $\mathcal{P}_i, \mathcal{P}_j, \cdot, \cdot$) role := if \nexists(Session, ssid, $\mathcal{P}_j, \mathcal{P}_i, \cdot, \cdot$) : init else resp record (Session, ssid, $\mathcal{P}_i, \mathcal{P}_j, \text{id} = \perp, \text{mpk}, \text{role}$) marked fresh </pre> <p>On (NewSession, sid, ssid, $\mathcal{P}_i, \mathcal{P}_j, \text{id}_i$) from $\mathcal{F}_{\text{IPAKE}}$</p> <pre> role := if \nexists(Session, ssid, $\mathcal{P}_j, \mathcal{P}_i, \cdot, \cdot$) : init else resp record (Session, ssid, $\mathcal{P}_i, \mathcal{P}_j, \text{id}_i, \text{mpk} = \perp, \text{role}$) marked fresh </pre> <p>On (NewKey, sid, ssid, \mathcal{P}_i, K') from \mathcal{Z} to $\mathcal{F}_{\text{IPAKE}}$</p> <pre> retrieve (Session, ssid, $\mathcal{P}_i, \mathcal{P}_j, \cdot, \text{mpk}_i, \text{role}$) retrieve (Session, ssid, $\mathcal{P}_j, \mathcal{P}_i, \cdot, \text{mpk}_j, \cdot$) not marked pakecompleted if session is compromised : $K_i := K'$ elif session is fresh and \exists(FreshKey, ssid, \mathcal{P}_j, K_j) : $\text{iseq} := \text{SessMpkEq}(\text{ssid})$ if $\text{iseq} = \text{true}$: $K_i := K_j$ elif $\text{iseq} = \text{false}$: $K_i \leftarrow_{\\$} \mathcal{K}$ else : $K_i := \text{indeterminate}$ else : $K_i \leftarrow_{\\$} \mathcal{K}$ if fresh : record (FreshKey, ssid, \mathcal{P}_i, K_i) Mark session pakecompleted if $K_i \neq \text{indeterminate}$: record (ChoseKey, ssid, \mathcal{P}_i, K_i) if \mathcal{P}_i corrupted : send (ssid, K_i) to \mathcal{P}_i </pre> <p>On (Hash, sid, (i, x)) from \mathcal{Z} to \mathcal{F}_{RO}</p> <pre> if $i = 0$, for each party \mathcal{P} : // program keygen coins $\text{pw} := x$ send (OfflineTestPwdf, ssid, \mathcal{P}, pw) to $\mathcal{F}_{\text{IPAKE}}$ if received “correct” : // “correct” \implies compromised retrieve (Compromised, $\mathcal{P}, \cdot, \text{coins}$) $H_0[\text{sid}, \text{pw}] := \text{coins}$ return $H_i[\text{sid}, x]$ </pre>
--	---

Figure 8: The final UC simulator \mathcal{S} for the LATKE^{post} protocol. \mathcal{S} maintains hash tables, H_0, H_1, H_2 . Hash queries $H_i[x]$ are defined to be uniform and consistent when not otherwise specified.

<pre> On msg ($\mathcal{P}_j \rightarrow \mathcal{P}_i, c, \text{sid}, \text{ssid}$) Check all (Session, ssid, ...) are pokecompleted or abort retrieve (Session, ssid, $\mathcal{P}_i, \mathcal{P}_j, \cdot, \text{id}_i, \text{role}$) step := sentMsgs[sid, ssid, \mathcal{P}_i, \cdot] recvdMsgs[sid, ssid, $\mathcal{P}_i, \text{step}$] := c if $\exists(\text{KnownCoins}, \mathcal{P}_i, \text{id}_i, \text{coins}_i)$: st := $\text{ibkeStates}[\text{ssid}, \mathcal{P}_i]$ inCt := $\text{recvdMsgs}[\text{sid}, \text{ssid}, \mathcal{P}_i, \text{step} - 1]$ or begin if st $\neq \perp$: inMsg := $S_{\text{cca}}.\text{Dec}_k^P(\text{inCt} : \sigma)$ if inMsg $\neq \perp$: (outMsg, st') := $\text{IBKE.next}(\text{st}, \text{inMsg})$ if st' = (Done, \cdot, \cdot) : FinishIbke(sid, ssid, st') return \perp $\text{ibkeStates}[\text{ssid}, \mathcal{P}_i] := \text{st}'$ else : outMsg := $0^{\text{msgSize}_{\text{step}}}$ else : outMsg := $0^{\text{msgSize}_{\text{step}}}$ outCt := $S_{\text{cca}}.\text{Enc}^P((\mathcal{P}_i, \text{ssid}, \text{step}), \text{outMsg} : \sigma)$ else : outCt $\leftarrow \{0, 1\}^{\text{msgSize}_{\text{step}} + \tau}$ send ($\mathcal{P}_i \rightarrow \mathcal{P}_j, \text{ssid}, \text{outCt}$) sentMsgs[sid, ssid, $\mathcal{P}_i, \text{step}$] := outCt </pre>	<pre> Procedure EueCatchup(sid, ssid) // assume wlog i is the compromised party retrieve (Session, ssid, $\mathcal{P}_i, \mathcal{P}_j, \text{id}_i, \text{mpk}_i, \text{role}$) retrieve (Session, ssid, $\mathcal{P}_j, \mathcal{P}_i, \cdot, \text{mpk}_j, \cdot$) steps = sentMsgs[sid, ssid, \mathcal{P}_i, \cdot] if SessMpkEq(ssid) : if $\exists(\text{ChoseKey}, \text{ssid}, \mathcal{P}_i, K_i) : K := K_i$ else : $K \leftarrow \{0, 1\}^\lambda$ $K_{\text{ch}} := H_2(\text{sid}, \text{ssid}, 0 \ K)$ $\text{ibkeStates}[\text{ssid}, \mathcal{P}_j, \text{step}] := (\text{begin}, \text{msk}, \text{usk}_j)$ outMsg := \emptyset for step in $0 \dots \text{steps}$: st := $\text{ibkeStates}[\text{ssid}, \mathcal{P}_k, \text{step}]$ or \perp if step is even : $\ell := \text{if role} = \text{init} : i \text{ else } j$ else : $\ell := \text{if role} = \text{init} : j \text{ else } i$ if st = \perp or $\text{recvdMsgs}[\text{sid}, \text{ssid}, \mathcal{P}_k, \text{step} - 1]$ $\neq \text{sentMsgs}[\text{sid}, \text{ssid}, \mathcal{P}_k, \text{step} - 1]$: outMsg := $0^{\text{msgSize}_{\text{step}}}$ st' := \perp else : outCt := $\text{sentMsgs}[\text{sid}, \text{ssid}, \mathcal{P}_k, \text{step}]$ (outMsg, st') := $\text{IBKE.next}(\text{st})$ $k := S_{\text{cca}}.\text{Exp}((\mathcal{P}_k, \text{ssid}, \text{step}), (\text{outMsg}, \text{outCt}) : \sigma)$ $K'_{\text{ch}} \leftarrow \{0, 1\}^\lambda$ $H_2[\text{sid}, \text{ssid}, 1 \ K_{\text{ch}}] := (K'_{\text{ch}}, k)$ $\text{ibkeStates}[\text{ssid}, \mathcal{P}_k, \text{step}] := \text{st}'$ </pre>	<pre> Procedure SessMpkEq(sid, ssid) retrieve (Session, ssid, $\mathcal{P}_i, \mathcal{P}_j, \text{id}_i, \text{mpk}_i, \text{role}$) retrieve (Session, ssid, $\mathcal{P}_j, \mathcal{P}_i, \cdot, \text{mpk}_j, \cdot$) if $\text{mpk}_i \neq \perp \wedge \text{mpk}_j \neq \perp$: return $\text{mpk}_i = \text{mpk}_j$ out := indeterminate coins := \perp if $\exists \ell \text{ s.t. } \text{mpk}_\ell \neq \perp$: $\hat{\ell} := i \text{ if } \ell = j \text{ else } j$ if mpk_ℓ has owner \mathcal{P}_ℓ : send (OnlineComparePwd, sid, ssid, $\mathcal{P}_\ell, \mathcal{P}_\ell$) out := resp = "correct" : if \mathcal{P}_ℓ honest : coins := coins of mpk_ℓ elif mpk_ℓ is derived from pw_ℓ : if \mathcal{P}_ℓ honest : coins := coins of pw_ℓ send (OnlineTestPwd, sid, ssid, $\mathcal{P}_\ell, \text{pw}_\ell$) out := resp = "correct" : elif $\exists \ell \text{ s.t. } \mathcal{P}_\ell$ compromised : $\hat{\ell} := i \text{ if } \ell = j \text{ else } j$ if \mathcal{P}_ℓ honest : coins := coins of mpk_ℓ send (OnlineComparePwd, sid, ssid, $\mathcal{P}_\ell, \mathcal{P}_\ell$) out := resp = "correct" if out = true : record (KnownCoins, $\mathcal{P}, \text{id}_i, \text{coins}$) return out </pre>
<pre> Procedure FinishIbke(sid, ssid, \mathcal{P}_i, st) retrieve (Session, ssid, $\mathcal{P}_i, \cdot, \cdot, \text{id}_i, \cdot$) (Done, K, id') = st $K' := H_1(K, \text{tr}_{\text{sid}, \text{ssid}}^{\text{poke}} \text{recvdMsgs}[\text{sid}, \text{ssid}, \cdot, \cdot])$ sessTampered := sentMsgs[sid, ssid, \cdot, \cdot] \neq recvdMsgs[sid, ssid, \cdot, \cdot] if $\text{id}' \neq \text{id}_i$ or sessTampered : if id is compromised send (Impersonate, sid, ssid, $\mathcal{P}_i, \mathcal{P}_{\text{id}}$) elif $\exists \text{id}', \text{pw}$ s.t. $\text{mpk}_i = \text{IBKE.Kg}(1^\lambda; H_0(\text{id}', \text{pw}))$: send (OnlineTestPwd, sid, ssid, \mathcal{P}_i, pw) else : abort send (NewKey, sid, ssid, $\mathcal{P}_i, \text{id}', K'$) Mark session complete </pre>		

Figure 9: (cont.) the LATKE^{post} UC simulator

C RevealLtk in the CK Model

In our presentation of the CK model in Appendix A.3, we include the oracle `RevealLtk`, which, unlike `Corrupt`, reveals the long-term key of the given party without revealing any internal state. While `RevealLtk` is included in the extended Canetti-Krawczyk model (eCK) [LLM07], it does not appear in the original CK or CK_{HMV} models [CK01, CK02a, Kra05], which Theorem 1 claims compatibility with.

We claim that this new model is equivalent to the base model. An adversary against the SK-without-`RevealLtk` game is trivially an adversary against the SK-with-`RevealLtk` game. So it only remains to show the converse. The theorem applies to protocols in the CK and CK_{HMV} model.

We also note that the SK-with-`RevealLtk` game was implicitly used in the security proof of CHIP in version 3.0 of the CHIP paper [CNPR22]. Thus, this section is not only necessary for LATKE, but also fills a gap in the security proof of CHIP.

Theorem 2. *Let Π be an authenticated key exchange protocol . If \mathcal{B} is an adversary with advantage ϵ against the SK-security game with access to the `RevealLtk` oracle, then there exists an adversary \mathcal{A} against the ordinary SK-security game with advantage at least $\epsilon/4$. This also applies to SK-security with `KCIR`, `FS`, and `KCIR+FS`.*

Proof. Let $N = \text{poly}(\lambda)$ denote the number of parties that \mathcal{B} initializes in its game. We define \mathcal{A} as follows. At the very beginning of the SK-security game, \mathcal{A} selects $N/2$ parties at random and calls `Corrupt` on them, thus receiving all their long term keys. Let C denote the set of the initially corrupted parties, and U the set of initially uncorrupted parties. Let sk_i denote the long-term key of party \mathcal{P}_i .

Then, \mathcal{A} runs \mathcal{B} . For every oracle query \mathcal{B} makes, \mathcal{A} makes the same query (if it exists), including when choosing the test session. When \mathcal{B} makes a query of the form `RevealLtk`(\mathcal{P}_i), then \mathcal{A} (1) checks if it has stored sk_i and, if so, returns it; or (2) calls `Corrupt` on \mathcal{P}_i , saves the long-term key as sk_i , and returns sk_i . For its test session, \mathcal{A} outputs whichever bit \mathcal{B} outputs. If any of the previous instructions lead \mathcal{A} to expose the test session, a *failsafe* will trigger and \mathcal{A} will instead respond to the test session with a random bit.

It will suffice to consider a specific condition E under which \mathcal{A} is guaranteed to win its game when \mathcal{B} wins its game, i.e.,

$$\Pr[\mathcal{A} \text{ wins} \mid E] = \Pr[\mathcal{B} \text{ wins}].$$

In other words, when E holds, then \mathcal{A} 's test session is unexposed.

In all of the SK-security variants, `RevealLtk` and `Corrupt` are treated equally in terms of exposure except *during* the test session (i.e., after it is initiated and before it is expired). In the `KCIR` variant, the adversary is permitted to call `RevealLtk` on the owner \mathcal{P}_i of the test session during the test session . If \mathcal{B} does this, then \mathcal{A} may have to call `Corrupt` on the owner, exposing the session. Thus, \mathcal{A} 's test session remains unexposed if both

1. $\mathcal{P}_i \in C$ (so \mathcal{A} can respond with sk_i), and
2. the peer of the session \mathcal{P}_j (if it exists) is in U (so at least one party is uncorrupted).

Let E be the event in which both statements are true. It is clear that the above equality holds with respect to E in every SK-security variant. Since both (1) and (2) occur independently and with even odds, $\Pr[E] = 1/4$.

When E does not occur, then \mathcal{A} 's failsafe may trigger. We may lower bound $\Pr[\mathcal{A} \text{ wins} \mid \neg E] \geq 1/2$.

Thus, the advantage of \mathcal{A} is

$$\begin{aligned}
\left| \Pr[\mathcal{A} \text{ wins}] - \frac{1}{2} \right| &= \left| \Pr[\mathcal{A} \text{ wins} \mid E] \Pr[E] + \Pr[\mathcal{A} \text{ wins} \mid \neg E] \Pr[\neg E] - \frac{1}{2} \right| \\
&\geq \left| \frac{\Pr[\mathcal{B} \text{ wins}]}{4} - \frac{1}{8} \right| \\
&= \frac{1}{4} \left| \Pr[\mathcal{B} \text{ wins}] - \frac{1}{2} \right| \\
&= \epsilon/4
\end{aligned}$$

□

D Changelog

v1.0

- Replaced $\mathcal{F}_{\text{PAKE}}$ with its multisession variant to clarify iPAKE presentation
- Renamed the new functionality $\mathcal{F}_{\text{ocp-iPAKE}}$, or $\mathcal{F}'_{\text{iPAKE}}$ for short
- Removed claim that the reductions are plausibly history-free
- Added informal discussion of identity concealment
- Placed expanded overview of SIM*-AC-CCA security in appendix
- Fixed symmetric key schedule in the main proof's simulator
- Updated future work to mention existing one-round PQ PAKEs
- Removed unnecessary *role* field from PAKE ideal functionality
- Miscellaneous cleanup and clarification

v0.5

Initial ePrint submission