



# Updatable Public-Key Encryption, Revisited

Joël Alwen<sup>1(✉)</sup>, Georg Fuchsbauer<sup>2</sup>, and Marta Mularczyk<sup>1</sup>

<sup>1</sup> AWS Wickr, New York, USA  
{alwenjo,mulmarta}@amazon.com  
<sup>2</sup> TU Wien, Vienna, Austria  
Georg.Fuchsbauer@tuwien.ac.at

**Abstract.** We revisit Updatable Public-Key Encryption (UPKE), which was introduced as a practical mechanism for building forward-secure cryptographic protocols. We begin by observing that all UPKE notions to date are neither syntactically flexible nor secure enough for the most important multi-party protocols motivating UPKE. We provide an intuitive taxonomy of UPKE properties – some partially or completely overlooked in the past – along with an overview of known (explicit and implicit) UPKE constructions. We then introduce a formal UPKE definition capturing all intuitive properties needed for multi-party protocols.

Next, we provide a practical pairing-based construction for which we provide concrete bounds under a standard assumption in the random oracle and the algebraic group model. The efficiency profile of the scheme compares very favorably with existing UPKE constructions (despite the added flexibility and stronger security). For example, when used to improve the forward security of the Messaging Layer Security protocol [RFC9420], our new UPKE construction requires less than 1.5% of the bandwidth of the next-most efficient UPKE construction satisfying the strongest UPKE notion considered so far.

## 1 Introduction

Spurred on by the seemingly never-ending procession of data breaches, 0-day exploits and system compromises, it is becoming ever more important in to design protocols with the ability to automatically limit the blast radii of key and state compromises. Among other techniques, this has lead to interest in primitives designed to provide cheap but effective *forward security*, namely the property that security holds despite possible future compromises.

A naïve (though not ineffective) approach to providing forward security for, say, public-key encryption (PKE) is for the owner of a key pair  $(pk, sk)$  to periodically sample a fresh and independent key pair  $(pk', sk')$  that replaces its old keys. While this does provide forward security – old ciphertexts encrypted to  $pk$  remain secure even if the adversary learns  $sk'$  – it comes with a serious drawback from the protocol perspective. After each key rotation the receiver must first inform prospective senders of the new public key before new messages can be

sent privately to the receiver again.<sup>1</sup> Besides increasing communication complexity, the biggest issue with this is that it forces potentially onerous coordination requirements on protocol participants.

Avoiding this cost motivated the study of *Puncturable Public-Key Encryption* [GM15] (PPKE) as a stand-alone primitive. PPKE provides essentially the same security as the naïve approach but without further coordination between parties beyond the initial public key distribution. After that, any number of senders may independently send any number of ciphertexts to the receiver which can be delivered in any order (or not at all). Despite the lack of coordination between parties, PPKE guarantees that at any point, leaking the receiver’s secret key reveals nothing about messages in ciphertexts it had already received and decrypted.

Clearly a powerful tool for building forward-secure protocols, PPKE lies at the heart of recent forward-secure 0-round trip key agreement protocols [GHJL17]. But minimizing round and communication complexity for forward-secure key agreement underpins other classes of cryptographic protocols. Notably, these include 2-party ratcheting [JS18, PR18, JMM19, DV19, CCD+20], the multi-party analogue: continuous group key agreement (CGKA) [ACDT20, AAN+22, ACJM20] and secure group and 2-party messaging [ACDT21]. In this work, we are especially interested in CGKA and secure group messaging (SGM) applications of forward-secure encryption primitives as these demand new, and hitherto seemingly overlooked, properties of the underlying primitive.

**Updatable Public Key Encryption.** Unfortunately, despite its wide-ranging practical applications, to date, PPKE constructions are not practically efficient for many real-world use cases, in particular in the ratcheting and messaging settings. This has given rise to a new class of “off-brand” forward-secure encryption schemes in the messaging literature called *Updatable Public-Key Encryption* (UPKE). They aim for a happy middle ground between forward secrecy with minimal interaction and truly practical efficiency.

Intuitively, UPKE is public-key encryption where senders can also generate *update tokens*. Applying a token  $up$  to a public key  $pk$  produces an updated public key  $pk \rightarrow_{up} pk'$ . Similarly, applying  $up$  to the secret key  $sk$  of  $pk$  yields the secret key  $sk \rightarrow_{up} sk'$  corresponding to  $pk'$ . The essential promise of UPKE is that ciphertexts encrypted to  $pk$  remain secure even when an adversary learns  $pk$ , the token  $up$  and the updated *secret key*  $sk'$ . Thus, a protocol in which parties update receivers’ key pairs whenever encrypting to them can achieve relatively strong forward secrecy properties. Indeed, no secret key is ever used more than once by a party and is immediately deleted (and replaced) upon first use.

However, there is a caveat to this. While using UPKE this way doesn’t require as much coordination between parties as the naïve approach, it does require more than PPKE. To ensure a receiver has the correct secret key available, a sender must encrypt to the *most recent version* of the receiver’s public key. In other words, senders must see each others’  $up$  tokens (or at least the most recently updated public key) before they can send. Otherwise, two senders may

<sup>1</sup> Note that new keys cannot be prepared and distributed too far in advance since this only extends the window of time during which forward secrecy is *not* provided.

concurrently produce update tokens  $up_0$  and  $up_1$  for one public key  $pk$  giving rise to two sibling key pairs  $(pk_0, sk_0)$  and  $(pk_1, sk_1)$ . We refer to this as a “fork”. When a fork occurs, a receiver will typically only derive one of the forked secret keys  $sk_b$  since it must then immediately delete  $sk$  to ensure forward security. Thus, when it later receives  $up_{1-b}$ , it can no longer produce  $sk_{1-b}$  meaning it can’t decrypt anything sent to  $pk_{1-b}$  (or any of its descendent keys). A similar restriction is that the receiver must decrypt ciphertexts in the same order they were sent (even when sent by different senders).

Still, compared to the naïve technique this represents a qualitative reduction in coordination since the receiver can essentially stay silent after initial public key distribution. Crucially, this makes asynchronous communication (as understood in asynchronous (group) messaging) possible, because senders need not wait for a receiver to announce new public keys before they can encrypt new messages to them. Thus, UPKE provides to secure messaging protocol designers the benefits of strong forward secrecy without forcing them to compromise on the ability of parties to privately message each other despite receivers potentially being off-line for extended periods of time.

Unfortunately, no UPKE scheme to date is sufficiently flexible, nor has all of the requisite security properties for natural use in CGKA and SGM applications which UPKE was partly designed for. Indeed, the initial academic work [ACDT20] in this area introduced rTreeKEM, a CGKA protocol which provides strong forward security by using UPKE in place of the PKE. The goal was to provide a more secure CGKA upon which to re-base the IETF’s Messaging Layer Security (MLS) protocol, an open SGM standard specified in RFC9420 [BBR+23]. However, rTreeKEM (and the resulting SGM based on rTreeKEM [ACDT21]) were only analyzed in a restricted model, which lead to relatively lightweight demands being placed on the underlying UPKE (both in terms of functionality and security).

Since then, however, the much more realistic “insider security” paradigm [AJM22] has established itself as a standard in the CGKA and SGM literature [HKP+21, AHKM22, AMT23]. Unlike the security models of [ACDT20, ACDT21], which assume authenticated channels, insider security only uses an insecure network. More challengingly maybe, insider security also provides meaningful security guarantees to parties joining “fake” groups; that is, sessions created arbitrarily by the adversary. These additions mean that insider security better captures the practical security concerns for SGM and CGKA. However, they also mean that to date, all UPKE schemes lack either the flexibility or security necessary for a CGKA (or SGM) application like rTreeKEM to be insider-secure.

**Fake-Group Security.** One such missing security property of existing UPKE notions is the (intuitive) property we call “joiner” security. When UPKE is used in higher-level CGKA/SGM protocols as a forward-secure replacement for PKE (as in rTreeKEM, for example), the joiner security of the UPKE scheme plays a central role in ensuring that the resulting CGKA/SGM protocol provides the “fake group” security aspect of insider security.

In more detail, CGKA and SGM protocols allow for dynamic groups (i.e. groups with evolving membership). Thus, a party  $P$  might receive an invitation

to join an existing group mid-session. To join the group, P also receives the group state including the signature verification keys for each group member (authenticated by some trusted PKI). Fake-group security (for SGM) considers the case when the invitation (and accompanying group state) were produced maliciously by the adversary (who may also corrupt parties). It mandates that if P validates the invitation and state (as specified by the protocol) and subsequently proceeds with the execution to a point where no corrupt signing keys are left in the group's state, then the session should return to a secure state. For example, P's messages to the group should remain hidden from the adversary. Notably, this should be the case even though the group state could still include (U)PKE keys obtained by P from the adversary.

*Fake-Group Security in MLS.* To date, the only protocol we are aware of that achieves fake-group security is MLS. It does so by including signatures in the public group state, which give P a way to identify which PKE keys in the state were (supposedly) generated by which party and to whom the party sent the decryption keys as part of the protocol execution. Whenever a party is removed from the group, so too are any keys they either (supposedly) generated or were sent. In the insider corruption model, leaking a party's signing key also leaks all other secret keys it knows. Thus, if at some point only secure verification keys remain in the group state, we can conclude that all remaining public keys were generated by and sent to uncorrupted parties. As a result, under those conditions, MLS can provide P with meaningful security guarantees for the session.

*UPKE Breaks MLS's Fake-Group Security Mechanism.* When [ACDT20] proposed replacing PKE with UPKE to improve MLS's forward security, the authors left as an open problem how to adapt MLS's mechanism for fake-group security accordingly (at least without growing the group state in the number of updates to UPKE keys). This was one of the primary barriers to adopting UPKE in MLS.

Indeed, in general, the state of a group mid-session would include UPKE keys  $pk$  that are (nominally) the result of updates to some prior original key  $pk_0$ . So, to guarantee that  $pk$  is still secure, a new member must validate that (i)  $pk_0$  was generated by an honest party, and (ii) that  $pk$  is the result of honestly using the update algorithm starting from  $pk_0$ .

One approach to providing (ii) could be to include in the group state all update tokens  $up$  leading from  $pk_0$  to  $pk$  along with proofs that they were generated by the update algorithm. But this results in a state size and computational cost of joining that grow linearly in the number of updates between  $pk_0$  and  $pk$ , which is prohibitive in practice. (MLS sessions can be expected to last for years and have, say,  $n = 50,000$  group members; so some of the  $2n$  public keys in an MLS state could have been updated  $n/2$  times by the time a new member joins.) It is also not an adequate solution to have receivers (i.e., members who can compute the updated  $sk$ , which could be as few as a single party) sign the

updated  $pk$  to attest to its correctness, as it conflicts with the asynchronous nature of MLS.<sup>2</sup>

This motivates the *joiner security* property of UPKE identified in this work. It provides a joiner  $P$  with a concise tag for validating that some UPKE public key  $pk$  is the result of an (unknown) sequence of honest updates to a given “origin” UPKE public key  $pk_0$ . Thus, if an uncorrupted honest party attests to having generated  $pk_0$  via a signature (just as with the PKE keys in MLS) then we can again conclude, in the insider security model, that  $pk$  must be secure.

**Our Proposal: UPKE Allowing for Fake-Group Security.** These issues show that there seems to be no easy way to efficiently adapt MLS’s fake-group security mechanism to UPKE. So instead, we ask the UPKE scheme to directly provide a comparable public key validation mechanism for new members (and a matching security guarantee). A *joiner-secure* UPKE scheme thus includes an algorithm  $\text{Verify}_{jt}$  with 3 inputs: (i) a UPKE public key  $pk$  to be validated, (ii) an original public key  $pk_0$  and (iii) a “joiner tag”  $jt$ . The tag must be constant-size, in particular, independent of how many updates might have lead to from  $pk_0$  to  $pk$ .

The UPKE security game chooses the initial  $pk_0$  honestly at the start of the game (reflecting that in the application we only expect security from  $pk$  if an honest party attested to having generated  $pk_0$ , e.g. via a signature). Then, the UPKE adversary may update  $pk_0$  with honest (i.e., generated by the challenger) or potentially malicious tokens  $up$ . The adversary wins if it can come up with  $pk^*$  and  $jt^*$  which pass  $\text{Verify}_{jt}$  and for which it can break privacy (IND-CCA) of a ciphertext  $c^*$  encrypted to  $pk^*$ . However, the adversary loses if it corrupts a secret key created before requesting  $c^*$ .

This restriction excludes trivial attacks in which  $pk^*$  is an updated version of a corrupted key. On the other hand, the restriction is not tight in the sense that it also excludes corruptions that do not lead to trivial attacks. We believe that our joiner security is a good compromise for the following reasons. First, defining UPKE security that only excludes trivial attacks would require UPKE schemes with additional functionality, which seems to require inefficient constructions.<sup>3</sup> Second, our joiner security is sufficient to prove that MLS with UPKE achieves the same fake-group security as today’s MLS with PKE. In fact, the above can be proven even using UPKE joiner security with *no corruptions at all*. This means that our joiner security notion with corruption could enable an even stronger flavor of fake-group security for MLS with UPKE. Indeed, in the full version we give an example of an MLS execution where MLS with UPKE satisfying our stronger joiner security is secure, but would not be so if its UPKE only satisfied

<sup>2</sup> Indeed, after an update by one group member, new members could only join the group after a different (receiving) group member comes online to validate and sign the updated key. This would mean that at least 2 existing group members are needed to invite a new member to the group.

<sup>3</sup> Essentially, the challenger needs some way to identify which  $pk$ ’s are old versions of  $pk^*$  provided by the adversary. This seems to require storing the whole update history in  $pk^*$  or  $jt^*$ .

a notion disallowing corruptions. Such a stronger notion for MLS has not been defined yet, and we leave this as an interesting open problem.

**UPKE Taxonomy.** Hiding beneath the term “UPKE” and the high-level intuition above, we actually find a series of concrete schemes in the literature (e.g. [JMM19, ACJM20, EJKM22, HLP22, DKW21, AMT23, AW23, HPS23]) that differ in their syntax, security properties and even the purposes they serve in the applications they were conceived for. To better interpret the results in our work, it is instructive to categorize these differences.

*Long vs. Short Syntax:* The most obvious differences between UPKE schemes are their various syntaxes. UPKE was first introduced in [JMM19] using an (*asymmetric*) *long* syntax also used in [AAN+22, EJKM22]. Here, “long syntax” means that key updates are generated and applied using stand-alone algorithms. In contrast, in this work (as in [ACDT20, ACJM20, ACDT21]) we use a *short syntax*, where keys are updated as a side-effect of encryption and decryption, thereby obviating the need for explicit update algorithms. We opted for the simpler syntax as it suffices for the dynamic group protocol applications we focus on and converting to long syntax is trivial.

Further, [EJKM22] defines two variants of a long syntax. “Asymmetric” long syntax means an update  $up = (pu, su)$  includes a public component  $pu$  for updating public keys and a private component  $su$  for updating corresponding secret keys. “Symmetric” long syntax uses a single value to update both public and private keys. The notions in [DKW21, HLP22, AW23, HPS23] can be viewed as having a symmetric long syntax where the random coins used by the public key update algorithm are also the update token used for the private key.

*CPA vs. CCA:* The first UPKE applications needed only CPA-style UPKE as they either included additional mechanisms reducing the role of UPKE in their protocol [JMM19, AAN+22] or their application was analyzed in a model that disables all attacks that might leverage honest parties as decryption oracles. (For example, the use of ideal authenticated channels in [ACDT20] trivially prevents the adversary from injecting ciphertexts to honest parties.) However, subsequently, the stronger and more realistic “insider security” model [AJM22] has become the standard in the field [HKP+21, AHKM22, AMT23]. This motivated the need for CCA-style UPKE. Indeed, all subsequent UPKE constructions (including in this work) are now regularly proven secure with CCA-style security games.

*Forking Security:* Almost all UPKE applications in the group setting involve multiple parties using the same UPKE secret key. An adversary that, say, controls the network can easily cause such parties to have diverging views of a protocol session’s transcript. This can result in forked UPKE keys (i.e., the initial key is updated using different sequences of updates). Thus, for such settings UPKE schemes must provide security in the face of forks. To date, we know of no (explicitly defined) UPKE scheme with this property, including those in

**Table 1.** Comparison of security properties of different UPKE schemes. The last two columns indicates whether they are practically efficient and in which model they are proven secure. AGM stands for the algebraic group model [FKL18].

Scheme	Syntax	Privacy	Forking	Agnos-tic	Update validation	Joinersec.	PQ	Practical	Model
[JS18]	long	CCA		✓			✓		ROM
[PR18]	long	CCA		✓			✓		ROM
[JMM19]	long	CPA		✓				✓	ROM
[ACDT20]	short	CPA		✓				✓	ROM
[EJKM22]	long	CPA		✓			✓		standard
[DKW21]	long	CCA			✓		✓		standard
[HLP22]	long	CCA			✓			✓	ROM
[AW23]	long	CCA					✓		ROM
[HPS23]	long	CCA			✓		✓	✓	ROM
[ACJM20]	long	CCA	✓	✓	✓		✓		standard
[AMT23]	long	CCA	✓	✓	✓		✓		standard
This work	short	CCA	✓		✓	✓		✓	ROM+AGM

[JMM19, EJKM22, HLP22, AW23, DKW21, HPS23] making them, a priori, insufficient for such applications.<sup>4</sup>

Notable exceptions are the schemes of [ACJM20, AMT23] that are (implicitly) based on hierarchical identity-based encryption (HIBE). Unfortunately, owing to their use of unbounded-depth HIBE, these are decidedly impractical for real-world applications leaving the state of UPKE for the group setting unsatisfactory.

*Decryption Oracles for Old Keys:* Even assuming there are no forks, in a setting with multiple parties using the same UPKE secret key, one has to account for parties not seeing some of the updates (yet) and hence holding old versions of the secret key. Accordingly, UPKE security notions should account for the attacker trying to inject ciphertexts to such parties. More precisely, assume we want to prove that an SGM scheme using UPKE is secure against adversaries who can inject ciphertexts but can *not* create forks. Even this weaker notion requires a UPKE security notion where, even after receiving the challenge ciphertext, the adversary can use the decryption oracle for any old secret key. However, this is not covered by any CCA-style UPKE definition we know of, in particular, not for [AW23, DKW21, HLP22, HPS23] (Table 1).

*Agnostic Updates:* The applications of UPKE considered in [JS18, PR18, JMM19, AAN+22] require update tokens to be generated without knowing the public key to which they will ultimately be applied which we refer to as “agnostic” updates.

<sup>4</sup> This seems to have happened because initial applications of UPKE are either in the 2-party setting, where forking is inherently not possible [JMM19] or they used very restricted models that artificially avoided forking by definition. Later UPKE constructions relied on UPKE security notions inspired by these early works but were not analyzed in their motivating applications using newer models. We provide a concrete scheme in the full version satisfying the definition [DKW21] but which leads to simple attacks when plugged into rTreeKEM.



Consequently, the UPKE schemes in those works are agnostic (as is the one in [EJKM22] and the implicit ones in [ACJM20, AMT23], although this is not necessary for the applications in those works). Conversely, the constructions of [AW23, DKW21, HLP22, HPS23] create updates for a target key.

*Protocol Usage:* While UPKE is usually used as a tool for achieving forward security in an application, the work of [AAN+22] applies updates to a possibly leaked secret key to refresh it to a new *secure* secret key. In other words, their protocol also relies on UPKE updates to ensure *post-compromise security* (PCS).<sup>5</sup> Thus, unlike any other use for UPKE we are aware of, [AAN+22] needs the additional intuitive property that secret keys of updated public keys have high (computational) entropy given the old secret key and updated public key. Fortunately, to the best of our knowledge, most UPKE schemes already have this property with the exception of the HIBE-based implicit schemes in [JS18, PR18, ACJM20, AMT23]. For the purpose of this work we focus on using UPKE for forward secrecy, so we leave such an entropy requirement for future work.

*Publicly Verifiable Updates:* In multi-party protocols like MLS and rTreeKEM, a common feature is that more than one user might encrypt messages to a particular public key. Suppose we use UPKE in this setting and a party  $P_1$  updates a public key  $pk$  to  $pk'$ . It is important that everyone in the group is convinced that  $pk'$  was generated via an honest update. Otherwise, a corrupt group member  $P_1$  (called an *insider*) might generate a key pair  $(pk^*, sk^*)$  using KeyGen and then convince someone that  $pk^*$  is the updated key. Clearly this would make all future ciphertext sent to the “updated key”  $pk^*$  insecure.

For group members that know  $sk$ , avoiding this is usually not too difficult. For example,  $P_1$  could encrypt to  $pk$  the coins used to produce the update [ACDT20]. However, revealing those coins to members who do *not* know  $sk$  would be problematic since UPKE security notions only ensure forward secrecy for updated keys if the coins used to update  $sk$  to  $sk'$  are kept secret.

So, to prevent an insider from tricking parties that don't know  $sk$  into accepting arbitrary new public keys, the UPKE scheme should provide a method to publicly verify that  $pk'$  was produced from  $pk$  via the update algorithm. To achieve this, the verification procedure can also take as input a *validation tag* provided by  $P_1$  as part of the message it sends to the group to announce the update. Intuitively, UPKE security should guarantee that if  $pk$  is secure and the pair  $(pk, pk')$  passes validation (with some tag), then  $pk'$  is also secure. Accordingly, the UPKE constructions [HLP22, DKW21, HPS23] include a special `VerifyUpdate` algorithm. For the implicit HIBE-based schemes of [ACJM20, AMT23], update verification is quite trivial and the step is left implicit.

To summarize, no UPKE scheme to date is known to satisfy the (CCA and) forking security properties needed to use UPKE in a CGKA protocol

---

<sup>5</sup> PCS is the mirror image of forward security where *future* keys should be secure despite *past* compromises.



like rTreeKEM [ACDT20] and meet the standard insider security for CGKA. (See the full version for a toy scheme that satisfies the UPKE security notion of [DKW21, HLP22, HPS23], yet leads to a trivial insider security attack when used in place of PKE in MLS as proposed in [ACDT20]. The attack leverages the lack of forking security in those UPKE notions.)

## Our Contributions

**New Model.** In this work, we study CCA-secure Updatable Key Encapsulation Mechanisms (UKEM); the KEM analogue of UPKE. Note that building UPKE from a UKEM is straightforward (for both the long and short syntax) e.g. using a standard KEM/DEM construction of CCA-secure PKE from a CCA-secure KEM and a CCA-secure authenticated encryption scheme, as done for example in Hybrid Public Key Encryption (HPKE) [BBLW22].

We present a new UKEM syntax and security definition designed to meet the needs of dynamic group protocols such as MLS and rTreeKEM of [ACDT20]. In particular, it captures CCA-type confidentiality with forks and joiner security. Our notion for UKEM can be easily extended to model UPKE security.

The new syntax does not require agnostic updates as this is not needed for these applications. It is based on the short UPKE syntax augmented with two public key validation algorithms. The first,  $\text{Verify}_{\text{jt}}$ , lets new members joining a group validate the public keys they download as part of the group’s state. It takes as input a public key  $pk_0$ , a public key  $pk_i$  being validated and a *joiner tag*  $jt_i$ . The joiner tag is generated along with  $pk_i$ . In particular, the tag  $jt_0$  is generated alongside  $pk_0$  by KeyGen and for  $i > 0$ , the tag  $jt_i$  is generated by Encaps when encrypting to and updating  $pk_{i-1}$ , given only  $pk_{i-1}$  and  $jt_{i-1}$ .

Joiner tags can be used to provide new-member security in protocols like MLS and rTreeKEM as follows. In addition to each UPKE public key  $pk_i$ , the group state contains the associated tag  $jt_i$ , as well as the original key  $pk_0$  signed by the group member who generated it.<sup>6</sup> Whenever a group member encrypts to  $pk_{i-1}$ , they replace  $pk_{i-1}$  and  $jt_{i-1}$  by  $pk_i$  and  $jt_i$ . Note that this can be done by all members, including new ones who did not see  $pk_1, \dots, pk_{i-2}$ . Further, new members can verify the signature on  $pk_0$  and verify  $jt_i$ , which convinces them, respectively, that  $pk_0$  was honestly generated and then updated to get  $pk_i$ .

The second algorithm,  $\text{Verify}_{\text{mt}}$  plays the same role as  $\text{VerifyUpdate}$  in the syntax of [HLP22, DKW21, HPS23]. It allows existing group members that do not know the secret keys to validate an updated public key. It takes as input the previous public key  $pk_{i-1}$ , the updated public key  $pk_i$  and a member tag  $mt_i$ , also produced as part of the output when encapsulating to  $pk_{i-1}$ .

One may wonder why  $\text{Verify}_{\text{mt}}$  is needed and why members cannot verify  $\text{Verify}_{\text{jt}}$  instead. Indeed, there may exist schemes for which this is the case. However, constructing  $\text{Verify}_{\text{mt}}$  is much easier. Intuitively, this is because the creator of  $mt_i$  can use the actual “witness” (i.e., secret randomness) for updating  $pk_{i-1}$

<sup>6</sup> The number of signatures can be reduced by half using the same “hashing down the path” optimization as in the parent hash mechanism of MLS.

to  $pk_i$ . On the other hand,  $jt_i$  must be generated without knowledge of the witnesses of the updates from  $pk_0$  up to  $pk_{i-1}$ . As a result, our efficient construction achieves better security for  $\text{Verify}_{\text{mt}}$ . On the other hand, joiners cannot profit from this additional security.

**Our Construction.** We provide a practically efficient construction of UKEM satisfying our model based on pairing-friendly elliptic curves. We prove it secure in the combination of the random oracle model (ROM) and the algebraic group model (AGM) [FKL18] (see below) under the co-discrete-log assumption, which in the AGM directly implies the co-CDH assumption [BLS01].<sup>7</sup>

Our starting point is the ElGamal-based KEM of DHIES [ABR98]. Public keys are of the form  $u = g^x \in \mathbb{G}$  in a group  $\mathbb{G}$  of prime order  $p$  with secret key  $x \in \mathbb{Z}_p$ . To encapsulate a symmetric key  $K$ , one chooses  $r \leftarrow \mathbb{Z}_p$ , computes the ciphertext  $v := g^r$  and sets  $K := H(u, u^r)$ , where  $H$  is treated as a random oracle.

To update a public key  $u$  in our scheme, we choose a random  $d \leftarrow \mathbb{Z}_p$ , which defines a new key  $u' := u \cdot g^d$ . The associated member tag  $mt$  is a *proof of knowledge* of  $d$ . Intuitively, this proof guarantees that if  $u$  was “secure” then so is  $u'$ . Indeed, suppose an adversary could update a random key  $u = g^x$  to  $u' = g^y$  for which it knows the secret key  $y$  while also proving knowledge of  $d$  such that  $u' = u \cdot g^d$ . Then by extracting  $d$  from the PoK we can use the adversary to compute the discrete log  $x = y - d$  for a random  $u$ . For our scheme, this intuition about the one-wayness of  $u$  and  $u'$  also extends to CCA-security. To allow receivers to update their secret keys accordingly,  $d$  is encrypted under  $u$ . Decrypters can thus recover  $d$  and update secret key  $x$  to  $x' := x + d$  for  $u'$ .

In fact, in our construction,  $d$  is actually derived via a random oracle (like the encapsulated key  $K$ ). This achieves three goals. First, it allows us to deal with adaptive corruptions, a problem resulting from forks (see below). Second, unlike in [JMM19, ACDT20] we can use the KEM ciphertext directly to transmit  $d$ , which saves on encrypting  $d$  explicitly. Third, using encryption would require key-dependent message security.

Our UKEM *member security* notion requires CCA-security for any public key whose member tag is valid. The notion is strong in that it allows the adversary to adaptively corrupt any secret key  $sk$  as long as  $sk$  does not let the adversary learn the challenge secret key in a trivial way. We achieve this leveraging the random oracle and by devising a careful guessing strategy: the security reduction guesses the first key  $u^*$  on the path of key updates leading from an initial honestly generated public key  $u_0$  to the challenge public key for which (i) the adversary breaks an encryption (which, as in DHIES, corresponds to solving CDH) or (ii) it breaks an encryption of any key the adversary derived from  $pk^*$ . Note that the reduction does not know this path and so it simply guesses a key.

<sup>7</sup> The co-DL assumption in groups  $\mathbb{G}$  and  $\hat{\mathbb{G}}$ , both of prime order  $p$  and generated by  $g$  and  $h$ , respectively, states that given  $g^x \in \mathbb{G}$  and  $h^x \in \hat{\mathbb{G}}$  for  $x \leftarrow \mathbb{Z}_p$ , it is hard to compute  $x$ . The co-CDH assumption states that given  $(g^x, g^r, h^x)$  for  $x, r \leftarrow \mathbb{Z}_p$ , it is hard to compute  $g^{xr}$ .

Despite allowing adaptive corruption, our reduction achieves a security loss of only the number of ciphertexts (and thus new keys) the adversaries asks for. For this to work, we need to assume that the proofs of knowledge of  $d$  (i.e.,  $mt$ ) are *simulation-sound*, that is, even after the adversary has seen simulated proofs (which the reduction creates when embedding its CDH challenge as a key), we can extract from an adversarial proof  $mt$ . This lets us “translate” a CDH solution for a key the adversary derived from the embedded key  $u^*$  to a solution for  $u^*$ .

Aiming for efficiency, we instantiate these proofs of knowledge of logarithms with Schnorr proofs, which consist of one element from  $\mathbb{G}$  and one from  $\mathbb{Z}_p$ . These proofs were shown simulation-sound in the ROM and the algebraic group model [FPS20,FO22], which provides “straight-line extractability”. That is, extraction of the witness does not require rewinding the adversary (as in the security proof in the ROM), which means we can extract from several proofs without risking an explosion of the running time due to interleaved rewinds for several proofs.

*Joiner Security.* A trivial construction of a joiner tag  $jt$  would be to include all  $mt$  proofs and intermediary public keys on the path from  $u_0$  to  $u'$ , which guarantee knowledge of  $d_1, \dots, d_k$  s.t.  $u' = u_0 \cdot g^d$  for  $d = d_1 + \dots + d_k$ . However, this is inefficient and our goal is constant-size joiner tags. Since the updater does not know the value  $d$ , we need a way to “aggregate” the proofs  $mt_i$  guaranteeing honest hops from  $u_{i-1}$  to  $u_i$  into a single short proof  $jt$  guaranteeing honest hops from  $u_0$  all the way to  $u'$ . An inherent problem with aggregatable proofs is that aggregation introduces malleability, which conflicts with our requirement that  $mt$  should be simulation-sound. Thus, we cannot hope that an instantiation of  $jt$  can also play the role of  $mt$ .

A very simple proof of knowledge of a logarithm is to assume that there exists a second generator  $h$  of  $\mathbb{G}$  of which no one knows the discrete log. To prove knowledge of the logarithm of  $v = g^d$ , one sets  $\pi := h^d$ . The knowledge-of-exponent assumption [Dam92] states that  $\pi$  can only be computed if one knows  $d$ ; formally, for any algorithm outputting  $(g^d, h^d)$ , there exists an extractor that outputs  $d$ . These proofs can be aggregated: given a proof  $\pi = h^d$  for  $u = g^x$  w.r.t.  $u_0 = g^{x_0}$ , that is  $d = x - x_0$ , a proof for  $u' := u \cdot g^d$  is computed as  $\pi' := \pi \cdot h^d$ .

The problem is how to verify whether  $\pi$  was correctly computed. This is why we embed our scheme in a *bilinear group*. That is, we assume a second group  $\hat{\mathbb{G}}$  and a bilinear map  $e: \mathbb{G} \times \hat{\mathbb{G}} \rightarrow \mathbb{G}_T$  for some target group  $\mathbb{G}_T$ .<sup>8</sup> We can now set the basis  $h$  for the proofs as a generator of  $\hat{\mathbb{G}}$  and use the pairing to verify a proof  $\pi \in \hat{\mathbb{G}}$  for  $v \in \mathbb{G}$  by checking whether  $e(v, h) = e(g, \pi)$ .

We prove joiner security directly in the algebraic group model. This model implies states that after having received elements  $h, \pi_1, \dots, \pi_k \in \hat{\mathbb{G}}$ , whenever the adversary returns some  $\pi \in \hat{\mathbb{G}}$ , it must have computed  $\pi$  as a linear

<sup>8</sup> In particular, we use an *asymmetric* pairing. That is, there are no efficiently computable homomorphisms between  $\mathbb{G}$  and  $\hat{\mathbb{G}}$ . In practice, this type of pairing yields the most efficient constructions. Note also that assuming a pairing lets one prove the security of DHIES from co-CDH instead of the interactive assumption *gap-CDH* [OP01,ABR01] which is also the case for our UKEM (see below).

**Table 2.** Comparison of object sizes in {kilo, mega}-bytes of recent UPKE schemes. By  $\phi$  we denote the bit-length of a NIZK that the update was generated correctly. A similar NIZK is needed to make the CPA scheme [DKW21] CCA-secure, while the CRS for the NIZK is included in public keys. In all UPKE applications considered in this work (e.g. rTreeKEM and MLS) ciphertexts are always sent together with a public key, an update  $up$ , joiner tag  $jt$  and member tag  $mt$ .

Scheme	Security	PQ	ROM	$ sk $	$ pk $	$ ctxt $	$ up $	$ jt $	$ mt $
[DKW21]	CPA	✓		166 B	41 KB	41 KB	52.375 MB		
[HPS23]	CCA	✓				1.8 KB	10.8 KB + $\phi$		
[HLP22]	CCA		✓	589 B	1.15 KB	11.375 KB	13.125 KB		
[AW23]	CCA		✓	32 B	80 B	96 B	128 B		
This work	CCA		✓	48 B	48 B	96 B		96 B	96 B

combination (“algebraically”) of all the  $\hat{\mathbb{G}}$  elements it has received. In particular, the AGM assumes that the adversary outputs  $\alpha_0, \dots, \alpha_k \in \mathbb{Z}_p$  such that  $\pi = h^{\alpha_0} \cdot \pi_1^{\alpha_1} \cdots \pi_k^{\alpha_k}$ . In our security proof,  $h$  and the proofs  $\pi_1, \dots, \pi_k$  computed by the reduction will be all  $\hat{\mathbb{G}}$  elements given to the adversary. As the reduction knows the discrete logarithms of the  $\pi_i$ ’s, it can compute the logarithm of  $\pi$  from  $\alpha_0, \dots, \alpha_k$ .

*Weaker Assumption for Member Security.* It turns out that the proofs  $\pi$  for joiner security also allow us to prove member security of our construction under weaker assumptions. In particular, we only require a notion of simulation-soundness for  $mt$  where extraction is done *after* all simulations. Recall that a co-CDH instance consists of  $u = g^x$ ,  $v = g^r \in \mathbb{G}$  and  $\hat{u} = h^x \in \hat{\mathbb{G}}$  and the goal is to compute  $w = g^{xr}$ . In the security proof of DHIES, the reduction embeds  $u$  as the public key and  $v$  as the ciphertext and searches for  $w$  among the random oracle queries made by the adversary. Using co-CDH (rather than CDH) the reduction can efficiently find  $w = g^{xr}$  the pairing  $e$ , by checking if  $e(v, \hat{u}) \stackrel{?}{=} e(w, h)$ .

Our reduction for UKEM embeds  $u$  as some (honestly updated) public key and  $v$  as some ciphertext it hopes the adversary breaks. However,  $v$  may not be created for  $u$  but for some  $u' = u \cdot g^{d'}$  derived from  $u$  by the adversary, who needs to provide proofs  $mt$  and  $jt$  for  $u'$ . The reduction thus searches the random oracle queries for a value  $w' = g^{(x+d')r}$ . It could do so by extracting  $d'$  from the proof of knowledge  $mt$ . However, using  $\pi = h^{d'}$ , it can directly check  $e(v, \hat{u} \cdot \pi) \stackrel{?}{=} e(w', h)$  without extracting anything at all. Extraction of the value  $d'$  is then only needed when a CDH solution is found (and the reduction stops): computing  $w := w'/v^{d'} = g^{(x+d')r}/g^{r d'}$  yields the co-CDH solution  $g^{xr}$ .

**Efficiency of Our Scheme.** We describe the efficiency profile of our scheme when instantiated with the BLS12-381 curve [SKSW22, Bow], which is a concrete 128-bit-secure instance of a BLS curve [BLS04]. It is equipped with an asymmetric pairing from source groups  $\mathbb{G} \times \hat{\mathbb{G}}$  to target group  $\mathbb{G}_T$ . Elements of  $\mathbb{G}$  and  $\hat{\mathbb{G}}$  are of size 48B and 96B respectively. As a NIZK we use a Schnorr proof of knowl-

edge of the discrete log of elements in  $\mathbb{G}$ , which results in proofs of length 96 B. Based on this, in our scheme, public keys are 48 B, ciphertexts are 48 B and both joiner and member tags are 96 B. As seen in Table 2 this represents a *very* significant improvement over all CCA-secure UPKE (and UKEM) schemes to date (despite the new scheme satisfying a considerably stronger security notion).

For example, using UPKE in rTreeKEM to achieve insider security involves sending multiple tuples of the form  $(pk, ctxt, t)$  where  $t$  is either an update token  $up$  or a joiner and member tag pair  $(jt, mt)$ , depending on which UPKE syntax is used and  $ctxt$  is a ciphertext under the previous key. The tuples of the new UPKE construction in this work are  $< 1.5\%$  the size of those of [HLP22]. For other CCA-secure schemes with publicly verifiable updates, the tuples are orders of magnitude larger still (despite none of these schemes providing forking or joiner security like the new construction).

We note that in our scheme, neither key generation, encapsulation nor decapsulation use pairing operations. One pairing is computed during each of the public key validation algorithms (which is run by parties holding the secret key before decapsulation as well).

**Further Results.** In the full version, we discuss extensions of our security model and efficiency improvements of the construction. We also dive into details of the impact of using variants of UPKE, including ours and less secure ones from the literature, on the security of MLS.

## 2 Preliminaries

**Bilinear Groups.** Our scheme will be defined over a bilinear group with an asymmetric pairing, that is, a tuple  $(p, \mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T, g, h, e)$ , where  $\mathbb{G}$  and  $\hat{\mathbb{G}}$  are groups of prime order  $p$  generated by  $g$  and  $h$ , respectively, and  $e: \mathbb{G} \times \hat{\mathbb{G}} \rightarrow \mathbb{G}_T$  is a non-degenerate (i.e.,  $e(g, h)$  generates  $\mathbb{G}_T$ ) bilinear map (i.e., for all  $a, b \in \mathbb{Z}_p$ :  $e(g^a, h^b) = e(g, h)^{ab}$ ).

The security of our scheme relies on the hardness of the co-discrete-logarithm problem in bilinear groups, defined as follows. We also state co-CDH [BLS01].

**Definition 1 (co-DL).** Let  $\mathcal{G} = (p, \mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T, g, h, e)$  be a bilinear group. The advantage of an adversary  $\mathcal{A}$  in solving the co-DL problem over  $\mathcal{G}$  is defined as

$$\text{Adv}_{\mathcal{G}}^{\text{co-DL}}(\mathcal{A}) := \Pr [y = x \mid x \leftarrow_{\$} \mathbb{Z}_p, u \leftarrow g^x, \hat{u} \leftarrow h^x, y \leftarrow \mathcal{A}(u, \hat{u})].$$

**Definition 2 (co-CDH).** Let  $\mathcal{G} = (p, \mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T, g, h, e)$  be a bilinear group. The advantage of an adversary  $\mathcal{A}$  in solving the co-CDH problem over  $\mathcal{G}$  is defined as

$$\text{Adv}_{\mathcal{G}}^{\text{co-CDH}}(\mathcal{A}) := \Pr \left[ w = g^{xr} \mid \begin{array}{l} x, r \leftarrow_{\$} \mathbb{Z}_p, u \leftarrow g^x, \hat{u} \leftarrow h^x, v \leftarrow g^r \\ w \leftarrow \mathcal{A}(u, \hat{u}, v) \end{array} \right].$$

For any  $u = g^x$ ,  $v = g^r$ , we denote a CDH solution  $w = g^{xr}$  by  $w = \text{DH}(u, v)$ .

**The Algebraic Group Model.** We analyze our scheme in the algebraic group model (AGM) [FKL18], which assumes that an adversary is *algebraic*, meaning that it computes any group element it outputs as a linear combination of the group elements it was given. More precisely, if the adversary, given input  $g := u_0, u_1, \dots, u_k \in \mathbb{G}$ , outputs a group element  $v \in \mathbb{G}$ , then it must have computed  $v$  as  $v = u_0^{\alpha_0} \cdots u_k^{\alpha_k}$  for some  $\alpha_0, \dots, \alpha_k$ . Formally, the AGM assumes that such coefficients  $\alpha_i$ , i.e., the “representation” of  $v$  are output by the adversary. The following is implicit in [FKL18]; we include a proof in the full version.

**Lemma 1.** *In the algebraic group model, co-DL tightly implies co-CDH. In particular, for any algebraic adversary  $\mathcal{A}$  against co-CDH in  $\mathcal{G}$ , there exists  $\mathcal{B}$  against co-DL in  $\mathcal{G}$  with approximately the same running time as  $\mathcal{A}$  s.t.  $\text{Adv}_{\mathcal{G}}^{\text{co-DL}}(\mathcal{B}) \geq \text{Adv}_{\mathcal{G}}^{\text{co-CDH}}(\mathcal{A})$ .*

**Simulation-Extractable Zero-Knowledge Proofs.** Our UKEM scheme uses a proof system PoL (“proof of logarithm”) for statements of the form  $\theta := (u, u')$  proving knowledge of a witness  $d$  s.t.  $u'/u = g^d$ . Formally, PoL may use a random oracle  $H$  and comprises the following algorithms:  $\tau \leftarrow \text{PoL.Prove}_H((u, u'), d)$  outputs a proof  $\tau$  and  $0/1 \leftarrow \text{PoL.Verify}_H((u, u'), \tau)$  verifies  $\tau$ .

We require two security notions: *Zero-knowledge* (in the random oracle model) means that the reduction, which can program the random oracle  $H$ , can create proofs  $\tau_i$  for statements  $\theta_i$  without knowing a witness, using an algorithm  $\text{PoL.Simulate}_H$ . The programmed random oracle and simulated proofs are, together, indistinguishable from a fresh random oracle and proofs computed honestly via  $\text{PoL.Prove}_H$  using a witness. We denote by  $\epsilon_{\text{PoL},n}^{\text{sim}}$  the simulation error of PoL when simulating at most  $n$  proofs.

*Strong Simulation Extractability* (sSE) is an adaptation of *strong simulation soundness* [Sah99] to proofs of knowledge [DP92]. It is defined via the following game: an adversary  $\mathcal{A}$  has access to random oracle  $H$  and an oracle that, on input a statement  $\theta_i$  of  $\mathcal{A}$ 's choice, returns a simulated proof  $\tau_i$  (and programs  $H$  as needed). Eventually,  $\mathcal{A}$  returns a statement/proof pair  $(\theta^*, \tau^*) \notin \{(\theta_i, \tau_i)\}_i$ . If  $\tau^*$  is a valid proof for  $\theta^*$  (using the final programmed version of  $H$ ) then a witness for  $\theta^*$  can be extracted from  $\mathcal{A}$ . (The notion is *strong* since after querying a simulated proof for a statement, a different proof for the same statement must be extractable.) We require a *multi-extraction* version of sSE, in which, after having queried simulated proofs, the adversary returns *several* valid pairs  $(\theta_i^*, \tau_i^*)$  with  $\{(\theta_i^*, \tau_i^*)\}_i \cap \{(\theta_i, \tau_i)\}_i = \emptyset$  and one can extract witnesses for all statements  $\theta_i^*$ . We denote by  $\epsilon_{\text{PoL},n}^{\text{ext}}(\mathcal{A})$  the advantage of the adversary  $\mathcal{A}$  in breaking simulation extractability of PoL when returning at most  $n$  proofs.

**Schnorr Signatures.** (Key-prefixed) Schnorr signatures are defined over a group  $\mathbb{G}$  of order  $p$  and a hash function  $H: \{0, 1\}^* \rightarrow \mathbb{Z}_p$ , modeled as a random oracle. Using signing key  $x \in \mathbb{Z}_p$ , a signature on a message  $m \in \{0, 1\}^*$  is computed by sampling  $r \leftarrow_s \mathbb{Z}_p$  and returning

$$(v := g^r, s := (r + cx) \bmod p) \quad \text{with} \quad c := H(v, g^x, m).$$

A signature  $(v, s)$  is valid for message  $m$  under public key  $u = g^x$  iff  $g^s = v \cdot u^c$  with  $c = H(v, u, m)$ .

In the combination of the random oracle model and the algebraic group model, [FO22] show that Schnorr signatures are sSE zero-knowledge proofs of knowledge of the logarithm of the public key. That is, they are proofs of knowledge (of the witness) for the NP-relation  $\{((u, m), x) \mid u = g^x, m \in \{0, 1\}^*\}$ .

Proofs for statements  $(u_i, m_i)$  can be simulated by programming the random oracle (as done in the original security proof for Schnorr [PS00]). Suppose an algebraic adversary  $\mathcal{A}$  receives simulated proofs  $(v_i, s_i)$  for statements  $(u_i, m_i)$  of its choosing and then outputs a valid statement/proof pair  $((u^*, m^*), (v^*, s^*)) \notin \{((u_i, m_i), (v_i, s_i))\}$ . Then, [FO22] showed that from the representations for the group elements  $u_1, u_2, \dots, u^*$  and  $v^*$ , which  $\mathcal{A}$  outputted during the game, one can efficiently compute a witness for the statement  $(u^*, m^*)$  with overwhelming probability.<sup>9</sup> In particular, extraction is straight-line and we can extract witnesses for *multiple* proofs produced during a single execution of an adversary. Thus, Schnorr signatures are *multi-extraction* sSE proofs in the ROM and AGM, which we formally prove in the full version.

The proof system PoL for member tags is defined as taking input a statement  $(u, u')$  and a witness  $d = \log(u'/u)$  and returning a Schnorr signature under key  $u'/u$  on the message  $(u, u')$ . Then, sSE guarantees that after receiving simulated proofs for pairs  $(u_i, u'_i)$ , if the adversary returns a new valid statement/proof pair  $((u_*, u'_*), (v_*, s_*))$ , we can extract  $d$  such that  $u'_*/u_* = g^d$ .

### 3 Updatable Key Encapsulation (UKEM)

#### 3.1 Functionality

Intuitively, a UKEM scheme is a key encapsulation mechanism with the following modifications. First, on input a public key  $pk_i$ , the `Encaps` algorithm outputs – in addition to the key  $K$  and the ciphertext  $c$  – the updated public key  $pk_{i+1}$ . Accordingly, on input  $sk_i$ , the `Decaps` algorithm outputs – in addition to  $K$  – the

<sup>9</sup> One might wonder why extraction is not trivial in the AGM anyway: an algebraic adversary that has only seen the generator  $g$  and returns  $u^*$  must know a representation  $\alpha$  s.t.  $u^* = g^\alpha$ . In the context of security proofs, this is not the case: Consider e.g., an *algebraic* reduction  $\mathcal{R}$  to the DL problem. This means that  $\mathcal{R}$  receives a DL instance  $g^*$  and simulates the game to an adversary  $\mathcal{A}$ , providing it with group elements it computes *as linear combinations* of  $g$  and  $g^*$ . When  $\mathcal{A}$  outputs a group element  $z$ , it accompanies it by a representation in basis all group elements received from  $\mathcal{R}$ . From this,  $\mathcal{R}$  can compute a representation  $(\alpha_0, \alpha_1)$  in basis  $(g, g^*)$ , that is,  $z = g^{\alpha_0} \cdot (g^*)^{\alpha_1}$ . To argue that  $\mathcal{R}$  can extract from proofs of knowledge made by  $\mathcal{A}$ , we need to turn  $\mathcal{R}$  together with  $\mathcal{A}$  into an adversary against simulation-extractability. This adversary is algebraic, but only in the sense that it can give representations in basis  $(g, g^*)$  where  $g^*$  is a group element of which the extractor will not know the discrete logarithm. Therefore, [FO22] (and our proof in the full version) actually show that even in the presence of an “auxiliary-input”  $g^*$ , one can extract the witness from a Schnorr proof.



updated secret key  $sk_{i+1}$ . This is analogous to any UKEM/UPKE with short syntax from the literature.

Second, **Encaps** also outputs a “member tag”  $mt_{i+1}$  which can be used by entities holding  $pk_i$  to validate  $pk_{i+1}$ . In particular, running  $\text{Verify}_{\text{mt}}(pk_i, pk_{i+1}, mt_{i+1})$ , such entities can verify that if  $pk_i$  is “honest” then  $pk_{i+1}$  is so, too. In MLS (more precisely, rTreeKEM [ACDT20]),  $\text{Verify}_{\text{mt}}$  is run by members (not joiners) who do not know  $sk_i$  but know and have validated  $pk_i$ .

Third, **Encaps** also generates a “joiner tag”  $jt_{i+1}$  which can be used by entities holding  $pk_0$  to validate  $pk_{i+1}$ : running  $\text{Verify}_{\text{jt}}(pk_0, pk_{i+1}, jt_{i+1})$ , such entities can verify that if  $pk_0$  is “honest” then  $pk_{i+1}$  is so, too. In MLS,  $\text{Verify}_{\text{jt}}$  is run by joiners after checking that  $pk_0$  was signed by the member who generated it using **KeyGen**. Moreover, **Encaps** takes the last joiner tag  $jt_i$  as input.

**Decaps** takes additional input  $pk_{i+1}$  and should output  $\perp$  if it does not “match”  $sk_{i+1}$ . In MLS, members who *do* know  $sk_i$  can thus reject “incorrect” (e.g. adversarially chosen)  $pk_{i+1}$ .

Formally, a UKEM scheme consists of the following algorithms:

**Key Generation.**  $\text{KeyGen}(\kappa) \rightarrow (pk_0, sk_0, jt_0)$ , on input the security parameter, outputs a key pair  $(pk_0, sk_0)$  and the first joiner tag  $jt_0$ .

**Encapsulation.**  $\text{Encaps}(pk_i, jt_i) \rightarrow (K, c, pk_{i+1}, mt_{i+1}, jt_{i+1})$  takes as input the current public key and joiner tag and returns an encapsulated key  $K$ , a ciphertext  $c$ , an updated public key  $pk_{i+1}$ , a new member tag  $mt_{i+1}$  and an updated joiner tag  $jt_{i+1}$ .

**Verification of member tags.**  $\text{Verify}_{\text{mt}}(pk_i, pk_{i+1}, mt_{i+1}) \rightarrow 0/1$  verifies the update from  $pk_i$  to  $pk_{i+1}$  using the tag  $mt_{i+1}$ .

**Verification of joiner tags.**  $\text{Verify}_{\text{jt}}(pk_0, pk_{i+1}, jt_{i+1}) \rightarrow 0/1$  verifies the update from  $pk_0$  to  $pk_{i+1}$  using the tag  $jt_{i+1}$ .

**Decapsulation.**  $\text{Decaps}(sk_i, c, pk_{i+1}) \rightarrow (K, sk_{i+1})/\perp$  outputs the decapsulated key  $K$  and the updated secret key  $sk_{i+1}$ , but only if  $pk_{i+1}$  matches  $sk_{i+1}$ .

*Using UKEM Schemes.* Importantly, **Decaps** does not validate any tags. Therefore, applications using a UKEM scheme should always run  $\text{Verify}_{\text{mt}}$  and  $\text{Verify}_{\text{jt}}$  before **Decaps**. This is reflected in our security notion.

## 3.2 Security

The IND-CCA security of UKEM schemes is formalized by the experiment in Fig. 1.

Intuitively, during the experiment, a tree is created where each node is identified by an integer  $i$  and has a public key  $pk_i$  and a joiner tag  $jt_i$ . The root is identified by  $i = 0$ . Each non-root node has a parent  $par_i$  and a member tag  $mt_i$ . Further, some nodes have a secret key  $sk_i$ . If a node has a secret key, we call it *full*, and otherwise we call it a *half node*.

The root node  $i = 0$  is created by the challenger at the beginning of the experiment. Its public key  $pk_0$ , secret key  $sk_0$  and joiner tag  $jt_0$  are generated using **KeyGen** (the root is thus a full node). All other nodes  $j$  are created by updating existing nodes in one of three ways:

Game UKEM IND-CCA Security	
<p><b>Exp</b> <math>\text{IND-CCA}(\mathcal{A})</math></p> <p><math>(pk_0, sk_0, jt_0) \leftarrow \text{KeyGen}(\kappa)</math>  <math>(mt_0, par_0, rev_0, j) \leftarrow (\epsilon, \epsilon, 0, 0)</math>  <math>b \leftarrow_s \{0, 1\}</math>  <math>b' \leftarrow \mathcal{A}^{\text{Enc, Dec, Rev, MChal, JChal}}(pk_0, jt_0)</math>  <math>S \leftarrow \text{chall-set}(chall)</math>  <b>return</b> <math>b = b' \wedge \forall j \in S : \neg rev_j</math></p> <hr/> <p><b>Oracle Enc</b>(<math>i</math>)</p> <p><math>(K, c) \leftarrow \text{create-honest-node}(i)</math>  <b>return</b> <math>(K, c, pk_j, mt_j, jt_j)</math></p> <p><b>Oracle MChal</b>(<math>i</math>)</p> <p><b>req</b> <math>chall = \perp</math>  <math>K^{(0)} \leftarrow_s \{0, 1\}^\kappa</math>  <math>(K^{(1)}, c^*) \leftarrow \text{create-honest-node}(i)</math>  <math>chall \leftarrow (\text{"member"}, i, c^*, pk_i)</math>  <b>return</b> <math>(K^{(b)}, c^*, pk_j, mt_j, jt_j)</math></p> <p><b>Oracle Rev</b>(<math>i</math>)</p> <p><b>req</b> <math>sk_i \neq \perp</math>  <math>rev_i \leftarrow 1</math>  <b>return</b> <math>sk_i</math></p>	<p><b>Oracle JChal</b>(<math>pk', jt'</math>)</p> <p><b>req</b> <math>chall = \perp</math>  <b>req</b> <math>\text{Verify}_{jt}(pk_0, pk', jt')</math>  <math>K^{(0)} \leftarrow_s \{0, 1\}^\kappa</math>  <math>(K^{(1)}, c^*, pk, mt, jt) \leftarrow \text{Encaps}(pk', jt')</math>  <math>chall \leftarrow (\text{"joiner"}, j, c^*, pk')</math>  <b>return</b> <math>(K^{(b)}, c^*, pk, mt, jt)</math></p> <hr/> <p><b>Oracle Dec</b>(<math>i', c', pk', mt', jt'</math>)</p> <p><b>req</b> <math>pk_{i'} \neq \perp</math> // <math>i</math>-th node exists  <b>if</b> <math>chall = (*, *, c^*, pk^*)</math> <b>then</b>  <b>req</b> <math>c^* \neq c' \vee pk^* \neq pk_{i'}</math>  <b>req</b> <math>\text{Verify}_{mt}(pk_{i'}, pk', mt')</math>  <b>req</b> <math>\text{Verify}_{jt}(pk_0, pk', jt')</math>  <math>j++</math>  <math>(pk_j, mt_j, jt_j, sk_j, par_j, rev_j)</math>  <math>\leftarrow (pk', mt', jt', \perp, i', 0)</math>  <b>if</b> <math>sk_{i'} \neq \perp</math> <b>then</b>  <math>out \leftarrow \text{Decaps}(sk_{i'}, c', pk')</math>  <b>if</b> <math>out \neq \perp</math> <b>then</b>  <math>(K, sk_j) \leftarrow out</math>  <b>return</b> <math>K</math>  <b>return</b> <math>\perp</math></p>
<p><b>Helper create-honest-node</b>(<math>i</math>)</p> <p><b>req</b> <math>pk_i \neq \perp</math> // <math>i</math>-th node exists  <math>j++</math>  <math>(K, c, pk_j, mt_j, jt_j)</math>  <math>\leftarrow \text{Encaps}(pk_i, jt_i)</math>  <b>if</b> <math>sk_i \neq \perp</math> <b>then</b>  // <math>i</math>-th node is full  <math>(*, sk_j) \leftarrow \text{Decaps}(sk_i, c, pk_j)</math>  <b>else</b> <math>sk_j \leftarrow \perp</math>  <math>(par_j, rev_j) \leftarrow (i, 0)</math>  <b>return</b> <math>(K, c)</math></p>	<p><b>Helper chall-set</b>(<math>chall</math>)</p> <p><b>if</b> <math>chall = (\text{"member"}, i^*, *, *)</math> <b>then</b>  <math>base \leftarrow \{i_0, \dots, i_\ell\}</math> where <math>i_0, \dots, i_\ell</math> is  the path from <math>i_0 = 0</math> to <math>i_\ell = i^*</math>  <b>else if</b> <math>chall = (\text{"joiner"}, i^*, *, *)</math> <b>then</b>  <math>base \leftarrow \{0, \dots, i^*\}</math>  <b>else return</b> <math>\emptyset</math>  <math>extd-base \leftarrow \{i' \mid \exists i \in base : (pk_{i'}, mt_{i'}, jt_{i'}) = (pk_i, mt_i, jt_i)\}</math>  // include duplicates  <b>return dec-closure</b>(<math>extd-base</math>)</p> <p><b>Helper dec-closure</b>(<math>S</math>)</p> <p>Return the set of all <math>j</math> reachable from  some <math>i \in S</math> via only edges created by  <b>Dec</b> queries.</p>

**Fig. 1.** The experiment formalizing UKEM IND-CCA security. By default, all variables are initialized to  $\perp$ . We use **req condition** to denote that if *condition* is false, then the current function, and any function calling it, stops and returns  $\perp$ .

1. When the adversary  $\mathcal{A}$  calls the oracle  $\text{Enc}(i)$ , the challenger creates a child  $j$  of  $i$  by running  $\text{Encaps}$ . If  $i$  is a full node,  $j$  is also a full node with secret key generated by running  $\text{Decaps}$ .
2. A child of  $i$  with a possibly “adversarial” public key may be created when  $\mathcal{A}$  calls the oracle  $\text{Dec}(i, c, pk', mt', jt')$ . In such case, the challenger verifies  $mt'$  and  $jt'$  and, if the check passes, creates the node  $j$  using these values. If  $i$  is a full node and  $\text{Decaps}(sk_i, c, pk')$  outputs  $(K, sk_j)$  (and not  $\perp$ ), then  $j$  is also a full node with secret key  $sk_j$ ; in that case,  $\mathcal{A}$  also receives  $K$ , which reflects CCA-security. Otherwise,  $j$  is a half node. Observe that  $j$  is a half node if  $\mathcal{A}$  provides correct (publicly verifiable) tags but  $c$  inconsistent with  $pk'$  (which is not publicly verifiable).
3. A node can be created during a challenge call. We address such calls next. There are two challenge oracles: member challenge  $\text{MChal}$  and joiner challenge  $\text{JChal}$ . Without loss of generality,  $\mathcal{A}$  can only call one of them, and only once.

**Member Security.** Consider the case that  $\mathcal{A}$  calls  $\text{MChal}$ , which means that the notion implies security for group members when used in a secure messaging application. On query  $\text{MChal}(i^*)$ , the challenger creates a child  $j^*$  of  $i^*$  just like during an  $\text{Enc}$  query creating a “real” key  $K^{(1)}$ .  $\mathcal{A}$  gets either  $K^{(1)}$  or a random and independent key  $K^{(0)}$  and has to decide which is the case. It also receives the resulting tags, public key and the ciphertext  $c^*$ . To disable trivial wins, on inputs  $i$  and  $c$  the  $\text{Dec}$  oracle returns  $\perp$  if  $pk_i = pk_{i^*}$  and  $c = c^*$ .

Furthermore, our notion implies forward secrecy by giving  $\mathcal{A}$  access to an oracle  $\text{Rev}$ , which reveals secret keys (of full nodes). In particular,  $\mathcal{A}$  can ask for the secret key of any node outside the *challenge set of  $i^*$* , which consists of three parts. First, the *base* of the challenge set, which is the path from the root 0 to  $i^*$ . Clearly, revealing the secret key for any such node would allow  $\mathcal{A}$  to trivially win by computing the secret key of  $i^*$  by running  $\text{Decaps}$  sequentially on the ciphertexts between the corrupted and the challenged node, and then decapsulating  $c^*$ . This *base* is extended to *extd-base*, which also includes *duplicates*, i.e., any nodes that have the same public key and tags as a node in *base*.<sup>10</sup>

Finally, the challenge set contains *branches*, which are nodes reachable from *extd-base* via nodes created by  $\text{Dec}$  queries. This is where our notion does not formalize optimal security: there exist UKEM schemes, notably the ones based on HIBE that achieve security even when  $\mathcal{A}$  can corrupt keys on branches. However, we are not aware of any *efficient* schemes that achieve this. Observe that the secret keys of nodes on branches are generated by updating a secret key on the challenge path (or a duplicate node) with updates generated by  $\mathcal{A}$ . Therefore, for optimal security we would need a mechanism that does not allow  $\mathcal{A}$  to undo its updates, which resembles PPKE.

<sup>10</sup> This restriction prevents trivial attacks, as in the following example:  $\mathcal{A}$  queries  $\text{Enc}(0)$ , which creates node 1 with  $(pk_1, mt_1, jt_1)$  and ciphertext  $c_1$ . It next queries  $\text{Rev}(1)$ , to obtain the corresponding  $sk_1$ . It then queries  $\text{Dec}(0, c_1, pk_1, mt_1, jt_1)$ , which creates node 2 with  $sk_2 = sk_1$ , and finally  $\text{MChal}(2)$ , to receive  $(c^*, K^*, pk_3, mt_3, jt_3)$  and checks whether for  $(K', sk_3) \leftarrow \text{Decaps}(sk_1, c^*, pk_3)$  it holds that  $K' = K^*$ .

We note that  $\mathcal{A}$  is allowed to ask for the secret key for  $j^*$  created by MChal, which corresponds to the fact that in typical UPKE security notions [DKW21, HPS23, HLP22, AW23] the challenge oracle returns the updated secret key. However,  $\mathcal{A}$  can also obtain many other keys, e.g., any node created by Enc and not on the challenge path (and all their children).

**Joiner Security.** Next, consider the case that  $\mathcal{A}$  calls JChal, formalizing a notion that implies security for joiners when used in a secure messaging application. On query JChal( $pk', jt'$ ), the challenger verifies  $jt'$  for  $pk'$  w.r.t. the (honest)  $pk_0$  and, if the check passes, runs Encaps on  $pk'$  to generate the “real” key  $K^{(1)}$ . As for member security,  $\mathcal{A}$  is also given the resulting ciphertext, public key and tags.  $\mathcal{A}$ 's goal is to distinguish  $K^{(1)}$  from a random and independent  $K^{(0)}$ . To disable trivial wins, on inputs  $i$  and  $c$  the Dec oracle returns  $\perp$  if  $pk_i = pk'$  and  $c = c^*$ .

Reveal queries are more restricted for joiner security than for member security. In particular, the challenge set *base* now contains all nodes generated before the call to JChal was made (which is thus a superset of the set *base* in the MChal setting). Analogously to member security,  $\mathcal{A}$  is not allowed to corrupt keys for nodes in the set *base*, any duplicates of such nodes and *branches* (i.e., nodes derived from these via Dec queries).

The above restriction cannot be relaxed without enabling “trivial” attacks against any correct scheme (with our syntax). To illustrate this, consider the following adversary  $\mathcal{A}$ . By calling Enc(0) twice,  $\mathcal{A}$  generates two children of node 0 with keys  $pk_1, pk_2$  and tags  $jt_1, jt_2$ . Then by running Encaps( $pk_1, jt_1$ ) (possibly repeating this to create a longer path),  $\mathcal{A}$  computes a new pair ( $pk', jt'$ ) on its own and submits it to its JChal oracle. If  $\mathcal{A}$  was allowed to query Rev(1), it could then, by running Decaps (possibly consecutively), compute the secret key for  $pk'$ .

In general,  $pk'$  may have been derived via Encaps from any  $pk_i$  that  $\mathcal{A}$  saw before generating  $pk'$ . Our restriction thus disallows Rev( $i$ ) for all such  $pk_i$ , including  $pk_0, pk_1$  and  $pk_2$  in the above example, even though corrupting  $pk_2$  would not lead to an attack. However, the challenger cannot identify keys that can be revealed, as the UKEM syntax does not allow to decide, given the challenger’s information, whether  $pk'$  could *not* have been derived from them.

*Remark 1.* One could consider relaxing the above restriction on reveal queries for a UPKE with modified syntax, e.g. with an additional algorithm that decides, given  $pk', jt', pk_i$  and  $sk_i$  (and any other information the challenger has), whether  $pk_i$  is an ancestor of  $pk'$ . However, implementing such an algorithm seems to require inefficient techniques, such as storing all ancestor public keys in  $jt'$ .

*Remark 2.* One could imagine achieving stronger joiner security by having JChal( $pk', jt'$ ) create an (incomplete) node  $i'$  with  $pk_{i'} = pk'$  and  $jt_{i'} = jt'$  and allowing the adversary  $\mathcal{A}$  to create a (detached) tree rooted at  $i'$ . (Note that, by the arguments in Remark 1, we cannot define a parent of  $i'$ .) However, the resulting notion would be equivalent to our notion. Since  $i'$  has no parent, its sub-tree contains only half-nodes without secret keys. So no oracle call related

to such nodes uses any secrets unknown to  $\mathcal{A}$  (which are the secret keys of full nodes and the bit  $b$ .) Thus,  $\mathcal{A}$  could emulate such oracle calls itself.

**Definition 3 (UKEM Security).** Let  $\text{Exp}^{\text{IND-CCA}}(\mathcal{A})$  be as defined in Fig. 1. The advantage of an adversary  $\mathcal{A}$  against the IND-CCA security of a UKEM scheme is defined as

$$\text{Adv}^{\text{IND-CCA}}(\mathcal{A}) := 2 \Pr [\text{Exp}^{\text{IND-CCA}}(\mathcal{A}) = 1] - 1.$$

## 4 Construction

The basis of our construction is the KEM part of DHIES [ABR98], which is basically “hashed ElGamal” for a hash function (modeled as a random oracle)  $H: \{0, 1\}^* \rightarrow \mathcal{K}$ , the symmetric key space. We use groups  $\mathbb{G}$  and  $\hat{\mathbb{G}}$  of order  $p$  with a pairing  $e$  from  $\mathbb{G} \times \hat{\mathbb{G}}$  and define the KEM in  $\mathbb{G}$ : Public keys are of the form  $u = g^x \in \mathbb{G}$  and symmetric keys  $K$  are encapsulated by choosing  $r \leftarrow_s \mathbb{Z}_p$ , defining the ciphertext as  $v := g^r$  and deriving  $K := H(u, u^r)$ . Using the secret key  $x$ , keys are decapsulated from  $v$  as  $K := H(g^x, v^x)$ .

We extend this to derive updated public keys as follows: using a second random oracle  $H_1$ , we define  $d := H_1(u, u^r)$  and set the new public key as  $u' := u \cdot g^d$ . Decapsulation now takes as additional argument the updated key  $u'$ , derives  $d := H_1(g^x, v^x)$ , updates the secret key to  $x' := x + d$  and checks if  $u' = g^{x'}$ . To guarantee that  $u'$  was derived correctly (and not chosen freshly with a known secret key), we add a proof of knowledge (PoK)  $\tau$  of  $d$ , that is, a PoK of the discrete log of  $u'/u$ . (For our security notion allowing adaptive corruption,  $\tau$  needs to be simulation-sound.) This  $\tau$  corresponds to  $mt$  in the UKEM model.

The tag  $jt$  given to joiners will be a PoK of  $D' := x' - x_0$ , with  $x_0$  the secret key of the root key  $u_0$  and  $x'$  the secret key of the updated key  $u'$ . This guarantees that  $u'$  is linked to the root key  $u_0$ . A straightforward solution would be to define  $jt_j := (u_1, mt_1, \dots, u_{j-1}, mt_{j-1}, mt_j)$ . To avoid a growth in size depending on the number of updates, we would require a direct proof of knowledge of  $D' = x' - x_0$ , but the updater will not know  $D'$ . Our solution is to use “aggregatable” proofs, that is, given a PoK of  $D = x - x_0$  corresponding to key  $u$ , and deriving  $u'$  from  $u$  using  $d$ , one should be able to derive a PoK of  $D' := D + d$ .

We use the second pairing source group  $\hat{\mathbb{G}}$ , generated by  $h$ , to instantiate these aggregatable proofs. A proof  $\pi$  proving knowledge of the logarithm of an element  $u = g^x \in \mathbb{G}$  is defined as  $\pi := h^x \in \hat{\mathbb{G}}$ . Using the pairing, a proof can be verified by checking  $e(u, h) = e(g, \pi)$ . Making “knowledge-of-exponent”-type assumptions (in our security proof we will directly rely on the algebraic group model), we get that from any algorithm that returns  $u$  and  $\pi$  satisfying the above equation, one can extract  $x = \log_g u = \log_h \pi$ , meaning  $\pi$  is a proof of knowledge.

Using these proofs for  $jt$  allows the updater to transform a proof  $\pi$  for  $u$  into a proof  $\pi' := \pi \cdot h^d$  for  $u' = u \cdot g^d$ . A proof  $\pi'$  for  $u'$  w.r.t.  $u_0$  is verified by checking  $e(u'/u_0, h) = e(g, \pi')$ . Our UKEM scheme is formally defined in Fig. 2.

## 5 Security of the Construction

Security of our construction is expressed by the following theorem.

**Theorem 1.** *If PoL is a strongly simulation-extractable proof system and co-CDH holds for  $\mathcal{G}$ , and assuming adversary  $\mathcal{A}$  is algebraic, then the UKEM construction from Fig. 2 is IND-CCA secure in the ROM. More precisely, for any adversary  $\mathcal{A}$ , there exist reductions  $\mathcal{B}$  and  $\mathcal{B}'$  such that*

$$\text{Adv}^{\text{IND-CCA}}(\mathcal{A}) \leq (n_e + 2)(\epsilon_{\text{PoL}, n_e+1}^{\text{sim}} + \epsilon_{\text{PoL}, n_d}^{\text{ext}}(\mathcal{B}') + \text{Adv}_{\mathcal{G}}^{\text{co-CDH}}(\mathcal{B})),$$

where  $n_e$  ( $n_d$ , resp.) are upper bounds on the number of Enc (Dec, resp.) queries made by  $\mathcal{A}$ , and  $\epsilon_{\text{PoL}, n}^{\text{sim}}$  ( $\epsilon_{\text{PoL}, n_d}^{\text{ext}}(\cdot)$ , resp.) are the probabilities that simulation of  $n$  proofs (extraction from  $n_d$  proofs, resp.) fails for PoL.

**Construction UKEM** $[H_1, H_2, H_3, \mathcal{G}, \text{PoL}]$

<p><u>KeyGen()</u></p> <p><math>x \leftarrow \mathbb{Z}_p</math>  <b>return</b> <math>(u \leftarrow g^x, x, \pi_0 \leftarrow h^0)</math></p> <p><u>Encaps</u><math>(u_i, \pi_i)</math></p> <p><math>r \leftarrow \mathbb{Z}_p</math>  <math>v \leftarrow g^r</math> // ciphertext  <math>d_{i+1} \leftarrow H_1(u_i, u_i^r)</math> // key update  <math>u_{i+1} \leftarrow u_i \cdot g^{d_{i+1}}</math> // update public key  <math>\tau_{i+1} \leftarrow \text{PoL.Prove}_{H_3}((u_i, u_{i+1}), d_{i+1})</math>  <math>\pi_{i+1} \leftarrow \pi_i \cdot h^{d_{i+1}}</math>  <math>K \leftarrow H_2(u_i, u_i^r)</math> // output key  <b>return</b> <math>(K, v, u_{i+1}, \tau_{i+1}, \pi_{i+1})</math></p>	<p><u>Verify</u><sub>mt</sub><math>(u_i, u_{i+1}, \tau_{i+1})</math></p> <p><b>return</b> <math>\text{PoL.Verify}_{H_3}((u_i, u_{i+1}), \tau_{i+1})</math></p> <p><u>Verify</u><sub>jt</sub><math>(u_0, u_{i+1}, \pi_{i+1})</math></p> <p><b>return</b> <math>e(u_{i+1}/u_0, h) = e(g, \pi_{i+1})</math></p> <p><u>Decaps</u><math>(x_i, v, u_{i+1})</math></p> <p><math>d_{i+1} \leftarrow H_1(g^{x_i}, v^{x_i})</math>  <math>x_{i+1} \leftarrow x_i + d_{i+1}</math> // update secret key  <b>if</b> <math>u_{i+1} \neq g^{x_{i+1}}</math> <b>then return</b> <math>\perp</math>  <math>K \leftarrow H_2(g^{x_i}, v^{x_i})</math> // output symm. key  <b>return</b> <math>(K, x_{i+1})</math></p>
---	---

**Fig. 2.** The UKEM construction. Here  $H_1, H_2$  and  $H_3$  are hash functions modeled as random oracles,  $\mathcal{G} = (p, \mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T, g, h, e)$  is a bilinear group, and PoL is a proof of knowledge system for discrete logarithm statements in  $\mathbb{G}$ , which might use  $H_3$ .

Together with Lemma 1, Theorem 1 implies that the security of our construction can be reduced to co-DL. Moreover, using the fact that Schnorr proofs, against algebraic adversaries, are strongly simulation-(multi-)extractable (as we show in the full version) with simulation error  $\epsilon_n^{\text{sim}} := n/(p - n_h - n)$  and (multi-)extraction error  $\epsilon_n^{\text{ext}} = n/p$ , yields the following:

**Corollary 1.** *Let  $\mathcal{G}$  be an asymmetric bilinear group. If PoL is instantiated using Schnorr (cf. Sect. 2) and co-DL holds for  $\mathcal{G}$ , then the UKEM construction from Fig. 2 is IND-CCA secure in the ROM and the AGM. More precisely, for any algebraic adversary  $\mathcal{A}$ , there exist a reduction  $\mathcal{B}$  such that*

$$\text{Adv}^{\text{IND-CCA}}(\mathcal{A}) \leq (n_e + 2) \left( \frac{n_e + 1}{p - n_h - n_e - 1} + \frac{n_d}{p} + \text{Adv}_{\mathcal{G}}^{\text{co-DL}}(\mathcal{B}) \right),$$

where  $n_e$ ,  $n_d$  and  $n_h$  are upper bounds on the number of, respectively, *Enc*, *Dec* and *RO* queries made by  $\mathcal{A}$ .

**Proof of Theorem 1.** We split the security notion IND-CCA into two: CCA-M, in which the JChal oracle is disabled, and CCA-J, in which the MChal oracle is disabled. The advantages  $\text{Adv}^{\text{CCA-M}}$  and  $\text{Adv}^{\text{CCA-J}}$  are defined accordingly. In Lemmas 2 and 3 we then bound these advantages. Theorem 1 then follows by summing them and letting  $\mathcal{B}$  and  $\mathcal{B}'$  be those adversaries from Lemma 2 or Lemma 3 that have the greater advantage.

## 5.1 Member Security

We start with the following lemma, which formalizing member security, CCA-M, of our UKEM scheme. For space reasons, we defer the full proof to the full version.

**Lemma 2.** *If PoL is a strongly simulation-extractable proof system and co-CDH holds for  $\mathcal{G}$ , then the UKEM construction from Fig. 2 is CCA-M-secure in the ROM. More precisely, for any adversary  $\mathcal{A}$ , there exist reductions  $\mathcal{B}$  and  $\mathcal{B}'$  such that*

$$\text{Adv}^{\text{CCA-M}}(\mathcal{A}) \leq (n_e + 1)(\epsilon_{\text{PoL}, n_e+1}^{\text{sim}} + \epsilon_{\text{PoL}, n_d}^{\text{ext}}(\mathcal{B}') + \text{Adv}_{\mathcal{G}}^{\text{co-CDH}}(\mathcal{B})),$$

where  $n_e$  and  $n_d$  are upper bounds on the number of  $\mathcal{A}$ 's *Enc* and *Dec* queries, resp.

**Proof Intuition.** Let  $\mathcal{A}$  be any adversary against the CCA-M security of our UKEM scheme. We will construct a reduction  $\mathcal{B}$  against the co-CDH problem, i.e., given  $u^*$ ,  $\hat{u}^*$  and  $v^*$ ,  $\mathcal{B}$  must compute  $w^* = \text{DH}(u^*, v^*)$ .

We start by adapting the proof idea for the security of the KEM of DHIES in the ROM.  $\mathcal{B}$  embeds  $u^*$  as some  $u_j$  generated by the challenger, that is, either as  $u_0$  or some  $u_j$  returned by an  $\text{Enc}(i)$  query, hoping that  $\mathcal{A}$  calls  $\text{MChal}(j)$ . If this happens,  $\mathcal{B}$  embeds  $v^*$  as the ciphertext returned by the oracle. Now as long as  $\mathcal{A}$  never queries  $(u^*, w^*)$  to the RO  $H_2$  with  $w^* = \text{DH}(u^*, v^*)$ , the challenge key  $K^{(b)}$  is independently random in both the real and the ideal game, and so no information on  $b$  is revealed. On the other hand, querying  $(u^*, w^*)$  means  $\mathcal{A}$  solved CDH; moreover,  $\mathcal{B}$  can test this by checking if  $e(w^*, h) = e(v^*, \hat{u}^*)$ .

*Embedding  $u^*$ .* Consider embedding  $u^* = g^x$  as  $u_{i^*}$  during a query  $\text{Enc}(p^*)$  (with  $p^*$  the parent of  $i^*$ ), which returns ciphertext  $v_{i^*}$ . Recall that  $\text{Encaps}$  would compute  $d_{i^*} = H_1(u_{p^*}, w_{i^*})$  with  $w_{i^*} := \text{DH}(u_{p^*}, v_{i^*})$  and define  $u_{i^*} := u_{p^*} \cdot g^{d_{i^*}}$  and  $\pi_{i^*} := \pi_{p^*} \cdot h^{d_{i^*}} = h^{x_{i^*} - x_0}$ . So when setting  $u_{i^*} := u^*$ , the reduction  $\mathcal{B}$  does not know  $d_{i^*} = \log(u^*/u_{p^*})$ . It thus generates the proof  $\tau_{i^*}$  using the simulator guaranteed by zero knowledge of PoL. To compute  $\pi_{i^*}$ , it uses  $\hat{u}^* = h^x$  from its co-CDH challenge and sets  $\pi_{i^*} := \hat{u}^*/h^{x_0}$  (and  $\pi_{i^*} := h^0$  if  $j = 0$ ).



While  $\mathcal{B}$  can simulate the proofs, not knowing  $d_{i^*}$ , it cannot consistently answer if  $\mathcal{A}$  queries  $H_1$  on  $(u_{p^*}, w_{i^*})$ . On the other hand, as long as this query has not been made, the simulation is consistent. Now, to make this query,  $\mathcal{A}$  would have to solve CDH w.r.t.  $u_{p^*}$  and  $v_{i^*}$ . But if  $\mathcal{A}$  ever does so, then  $\mathcal{B}$  should have guessed differently and embedded  $u^*$  as  $u_{p^*}$  and  $v^*$  as  $v_{i^*}$  (assuming for the moment there are no Dec queries).  $\mathcal{B}$ 's guessing strategy will therefore be to guess the index  $i^*$  of the *first* key  $u_{i^*}$  generated during a query  $\text{Enc}(p^*)$  on the path to the challenge for which  $\mathcal{A}$  will solve CDH via an RO query. (Note that  $\mathcal{B}$  does not know the path; it simply guesses the index of an Enc query.)

For now we only considered the case that  $\mathcal{A}$  makes the query  $\text{MChal}(j^*)$  or  $\text{Enc}(j^*)$  assuming  $u_{j^*}$  was itself created during an Enc query; but  $u_{j^*}$  might have been created during a Dec query. That is, the attacked key (i.e., the one for which  $\mathcal{A}$  solves CDH) has been generated by the adversary. Security now relies on the fact that ultimately the attacked key was derived (possibly via many Dec queries) from an honest key, say  $u_{i^*}$  (which might be  $u_0$ ).

Since  $\mathcal{A}$  must provide proofs  $\tau_i$  for the hops from  $u_{i^*}$  to  $u_{j^*}$  (where  $\tau_i$  proves knowledge of  $d_i = x_i - x_{\text{par}_i}$ ),  $\mathcal{B}$  can extract the values  $d_i$  and sum them to  $d_{i^* \rightarrow j^*} := x_{j^*} - x_{i^*}$ , which it can use to “translate” CDH solutions for  $u_{j^*}$  to  $u_{i^*}$ . Thus, it can embed  $u^*$  as  $u_{i^*}$  and embed  $v^*$  as the ciphertext the adversary breaks. A solution  $w = \text{DH}(u_{j^*}, v^*)$  then yields a solution  $w/(v^*)^{d_{i^* \rightarrow j^*}} = g^{x_{j^*}r}/g^{r(x_{j^*} - x_{i^*})} = \text{DH}(u^*, v^*)$ .

Our strategy is thus to guess the following index  $i^*$ : if the first attacked key is  $u_{j^*}$ , then  $i^*$  is the closest ancestor of  $j^*$  with a public key generated by the challenger. That is, at the latest  $i^* = j^*$  (if  $j^*$  is generated during an Enc query), and at the earliest  $i^* = 0$ .

*Answering Rev Queries.* Say  $\mathcal{B}$  embeds  $u^*$  as  $u_{i^*}$  and consider a query  $\text{Enc}(i^*)$ , which creates a new key  $u_j$ . If node  $j$  turns out not to lie on the challenge path, then  $\mathcal{A}$  is allowed to query  $\text{Rev}(j)$ . However, if  $\mathcal{B}$  ran  $\text{Encaps}$  to answer the query, setting  $u_j := u^* \cdot g^{d_j}$  with  $d_j := H_1(u_{i^*}, \text{DH}(u_{i^*}, v_j))$ , then it would not know  $x_j = \log u_j$  to answer the Rev query.

But recall that  $\mathcal{B}$  hopes that  $\mathcal{A}$  attacks key  $u_{i^*}$ ! Every time Enc or MChal is queried on  $i^*$ , the reduction thus embeds  $v^*$  from its co-CDH challenge into the ciphertext. In particular, using random self-reducibility,  $\mathcal{B}$  chooses a uniform  $s_j$  and defines the new ciphertext as  $v_j := v^* \cdot g^{s_j}$ . If  $\mathcal{A}$  ever queries  $H_1(u_{i^*}, w_j)$  for  $w_j := \text{DH}(u_{i^*}, v_j)$ , the game stops and  $\mathcal{B}$  returns  $w^* := w_j/(u^*)^{s_j} = g^{x^*(r+s_j)}/g^{x^*s_j} = \text{DH}(u^*, v^*)$ . On the other hand, as long as no such query is made,  $d_j$  is not defined, and thus  $\mathcal{B}$  can simply sample  $x_j$ , set  $u_j := g^{x_j}$  (which implicitly defines  $d_j$ ) and simulate the proofs  $\tau_j$  and  $\pi_j$ . This way,  $\mathcal{B}$  can then answer the query  $\text{Rev}(j)$ .

The case  $\text{Enc}(i)$  for an index  $i$  whose path from  $i^*$  consists of only Dec queries is dealt with similarly:  $\mathcal{B}$  embeds  $v^* \cdot g^{s_j}$  as  $v_j$  and samples  $x_j$  freshly. As long as  $\mathcal{A}$  does not query  $H_1(u_i, w_j)$  with  $w_j = \text{DH}(u_i, v_j)$ , the simulation is perfect. If the adversary makes that query, it can be translated back to a solution  $\text{DH}(u^*, v_j)$ , and thus to  $\text{DH}(u^*, v^*)$ , by extracting  $d_{i^* \rightarrow i} = x_i - x^*$  from the  $\tau$ -proofs provided

by the adversary when making the Dec queries linking  $u_{i^*}$  to  $u_i$ : we have  $w^* := w_j \cdot (u^*)^{-s_j} \cdot v_j^{-d_{i^* \rightarrow i}} = g^{(x^* + d_{i^* \rightarrow i})(r + s_j)} g^{-x^* s_j} g^{-(r + s_j) d_{i^* \rightarrow i}} = \text{DH}(u^*, v^*)$ .

*Extracting from Adversarial Proofs.* Simulation-extractability of  $\tau$ -proofs only lets us extract from proofs computed by the adversary (and not ones created by the simulator). So what happens if the adversary “copies” proofs simulated by the challenger?

In particular, consider the situation where we embedded our challenge key  $u^*$  as  $u_{i^*}$  and the adversary attacked one of its Dec-descendants  $u_{j^*}$ . If none of the key/proof pairs  $(u_i, \tau_i)$  on the path from  $i^*$  to  $j^*$  appear elsewhere in the tree, then the statement/proof pairs are different from those of the simulated proofs, and we can extract their witnesses. On the other hand, assume that on this path, there is a pair  $(u_{k^*}, \tau_{k^*})$  which appears elsewhere as  $(u_{k'}, \tau_{k'})$  in the tree. If (and only if)  $k'$  was created in a query  $\text{Enc}(i')$  and  $i'$  is a Dec-descendant of  $i^*$ , then  $\tau_{k'}$  was simulated, and thus we cannot extract from  $\tau_{k'} = \tau_{k^*}$ . (Note that since for every  $u_k$  there is a unique valid  $\pi_k$ , we have  $(u_{k^*}, \tau_{k^*}, \pi_{k^*}) = (u_{k'}, \tau_{k'}, \pi_{k'})$ .)

However, this just means that we should have guessed differently: assume  $k^*$  is the last “copied” node on the path from  $i^*$  to  $j^*$ . If we had embedded our challenge key  $u^*$  as  $u_{k^*}$  (when we created it as  $u_{k'}$  when answering an Enc query) then we could now solve CDH: since, by assumption, no nodes between  $u_{k^*}$  and  $u_{j^*}$  are copied, we can extract from their  $\tau$ -proofs and thus compute  $d_{k^* \rightarrow j^*} = x_{j^*} - x_{k^*}$ , which lets us shift a CDH solution for  $u_{j^*}$  to one for  $u_{k^*}$ . Note that we would not be able to answer Rev for  $k'$  and its Dec-descendants, but such queries are disallowed (as they are part of *chall-set*, cf. Fig. 1).

Our actual guess strategy is therefore: let  $u_{j^*}$  be the first key the adversary attacks during the game; then what is the index of the Enc query that creates the node  $(u_{k^*}, \tau_{k^*})$  so that when starting from  $u_{j^*}$  and moving up Dec-edges,  $(u_{k^*}, \tau_{k^*})$  is the first key/proof pair created by the challenger during an Enc query (at latest, this is  $u_0$ ).

*Answering Dec Queries.* We address answering decryption queries  $\text{Dec}(i)$  for nodes whose secret key is not known to the reduction. These are all nodes whose public key  $u^*$  is the embedded co-CDH instance, or any Dec-descendant of such nodes. Here, we again follow the ideas for proving CCA-security of DHIES, namely to inspect the random-oracle table. We moreover use the fact that CDH solutions can be checked via the pairing using the associated proof  $\pi_i$ : given a ciphertext  $v_j$  for key  $u_i = g^{x_i}$ , we have  $K_j = H_2(u_i, w_j)$  with  $w_j := \text{DH}(u_i, v_j)$  and the latter can be efficiently checked: setting  $\hat{u}_i := \pi_i \cdot h^{x_0} = h^{x_i}$  (where  $h^{x_0} := \hat{u}^*$  if  $i^* = 0$ ), check if  $e(w_j, h) \stackrel{?}{=} e(v_j, \hat{u}_i)$ .

So to decrypt ciphertext  $v_j$  for key  $u_i$  we do the following: if there has been a query  $(u_i, \text{DH}(u_i, v_j))$  to  $H_2$ , then we return the same key again; if there has not been such a query, we sample a fresh key  $K_j$  and (implicitly) program the random oracle: store an entry  $(u_i, v_j, \perp, K_j)$ , meaning that  $(u_i, \text{DH}(u_i, v_j))$  gets mapped to  $K_j$ . To detail how the Dec queries are answered, we first address programming of the random oracles.

*Programming the Random Oracles.* Answering Enc, Dec and MChal queries results in defining the entries of the random oracle tables for  $H_1$  and  $H_2$ . The inputs are of the form  $(u, w)$ , on which  $H_1$  outputs  $d$  and  $H_2$  which outputs  $K$ . For certain queries, these entries are partial, since the reduction does not know all inputs/outputs, i.e., the RO is programmed implicitly. The reduction thus stores RO entries of the form  $(u, \hat{u}, v, w, u', d, K)$ , some of whose components can be  $\perp$ . For  $u = g^x$ , the (non- $\perp$ ) values are:  $\hat{u} = h^x$ ,  $w = v^x = \text{DH}(u, v)$ ,  $u' = u \cdot g^d$  and  $d$  and  $K$  are the outputs of, respectively,  $H_1$  and  $H_2$ , on input  $(u, w)$ . Note that  $\hat{u}$ ,  $w$  and  $u'$  are determined by the other values. During Enc and MChal queries, implicit programming happens at the following positions:

1. When embedding the key  $u^*$  as  $u_{i^*}$  for  $i^* \neq 0$ , letting  $p^* := \text{par}_{i^*}$ , the reduction implicitly defines the oracles at  $(u_{p^*}, v_{i^*}^{x_{p^*}})$  (where  $x_{p^*}$  was chosen by the reduction);  $H_1$  is set to  $d_{i^*} := \log(u_{i^*}/u_{p^*})$  (unknown to the reduction) and  $H_2$  is set to  $K_{i^*}$  (chosen by the reduction). When answering this query, the reduction thus stores the following entry (where  $\hat{u}_{p^*} := h^{x_{p^*}}$ ):

$$(u_{p^*}, \hat{u}_{p^*}, v_{i^*}, \perp, u_{i^*}, \perp, K_{i^*})$$

2. For any call of Enc or MChal at position  $i$  with  $u_i = u^*$  or  $i$  being a Dec-descendant of a node with public key  $u^*$ , the reduction creates  $v_j$  (embedding  $v^*$  from its co-CDH instance) and  $u_j$  ( $:= g^{x_j}$  for fresh  $j$ ) and defines  $H_1$  and  $H_2$  at position  $(u_i, \text{DH}(u_i, v_j))$ , which is unknown to the reduction. While the reduction chooses the value  $K_j$  at this position for  $H_2$  (for the MChal query,  $K_j$  corresponds to “ $K^{(1)}$ ”), it will not know the value  $d_j = \log(u_j/u_i)$  for  $H_1$ . The reduction thus stores  $(u_i, \hat{u}_i, v_j, \perp, u_j, \perp, K_j)$ , where, as above,  $\hat{u}_i = \hat{u}^*$  if  $i^* = 0$  and  $\hat{u}_i := \pi_i \cdot h^{x_0} = h^{x_i}$  otherwise.

For every random-oracle query  $(u, w)$  the adversary makes, the reduction checks if  $(u, w) = (u_i, \text{DH}(u_i, v_j))$  holds when  $i = i^*$ , or  $i = \text{par}_{i^*}$  or  $i$  is a Dec-descendant of  $i^*$ . It does this by checking  $u \stackrel{?}{=} u_i$  and  $e(w, h) \stackrel{?}{=} e(v_j, \hat{u}_i)$ . (Note that such queries to  $H_1$  cannot be answered, since the reduction does not know  $d_j = \log(u_j/u_i)$ .)

If this is the case for  $i = \text{par}_{i^*}$ , the reduction stops, since the guess  $i^*$  was wrong, as  $\text{par}_{i^*}$  would have been the right guess. If it happens for  $i^*$  or any of its Dec-descendants, the reduction stops and returns the co-CDH solution (computed as described above). Otherwise, fresh values  $d$  and  $K$  are sampled and a new entry  $(u, \perp, \perp, w, u \cdot g^d, d, K)$  is created. We say that in this case the RO was *explicitly* programmed.

*Details of Answering Dec Queries.* Let us consider a query  $\text{Dec}(i', v', u', \tau', \pi')$ . If  $\tau'$  and  $\pi'$  are valid, a new (for now: half-)node is created. If  $i'$  is a full node, the oracle would do the following: run **Decaps** on  $sk_{i'}$ , that is, compute  $d' := H_1(u_{i'}, \text{DH}(u_{i'}, v'))$ ; check if  $u' = u_{i'} \cdot g^{d'}$ ; if so, return  $K := H_2(u_{i'}, \text{DH}(u_{i'}, v'))$  and declare the new node a full node; else return  $\perp$ .

If  $i'$  is a half-node, then the reduction can simulate the Dec oracle perfectly, as the latter uses only public values. Moreover, if  $i' \notin \text{chall-set}$ , then the reduction

knows  $sk_{i'}$  and can thus simulate the oracle perfectly as well. For  $i' \in \text{chall-set}$  and  $i'$  being a full node, the reduction uses its extended RO table as follows:

- (i) If there is an entry  $(u_{i'}, *, v', \perp, u'', d, K)$  for some  $u''$ ,  $d$  (where possibly  $d = \perp$ ) and  $K$ , then the RO was already implicitly programmed at  $(u_{i'}, \text{DH}(u_{i'}, v'))$  (either during an Enc or MChal query as described above, or during a Dec query as described below). The reduction checks if  $u' = u''$  (as Decaps does) and if so, it declares the new node a full node and returns  $K$ ; else, it declares the new node a half node and returns  $\perp$ .
- (ii) Else if there is an entry  $(u_{i'}, \perp, \perp, w, u'', d, K)$  for some  $u''$  and  $w = \text{DH}(u_{i'}, v')$ , which can be checked using  $\hat{u}_{i'} = \pi_{i'} \cdot h^{x_0}$ , then the RO was already explicitly programmed at  $(u_{i'}, \text{DH}(u_{i'}, v'))$ . As above, the reduction checks if  $u' = u''$  (as Decaps does); if so, it declares the new node a full node and returns  $K$ ; else, it declares the new node a half node and returns  $\perp$ .
- (iii) If none of the above apply, then sample  $d$  and  $K$ , create new entry  $(u_{i'}, \hat{u}_{i'}, v', \perp, u_{i'} \cdot g^d, d, K)$  and proceed as in Decaps. (Note that, with overwhelming probability, this will return  $\perp$ , since  $d$  will be inconsistent with  $u_{i'}$  and  $u'$ .)

Finally, note that the only RO query that would reveal the challenge bit  $b$  is querying  $H_2$  on  $(u_{i_c}, \text{DH}(u_{i_c}, v_{j_c}))$ , where  $i_c$  is the value queried to MChal and  $j_c$  the current value of  $j$  at that point. “Explicit” queries are dealt with by our guessing strategy: if the guess  $i^*$  was correct then such a query is used to solve co-CDH. On the other hand, “implicit” queries via the Dec oracle cannot occur, since this would correspond to  $\text{Dec}(i', v', u', \tau', \pi')$  with  $u' = u_{i_c}$  and  $v' = v_{i_c}$ , which is forbidden (as trivial wins).

## 5.2 Joiner Security

We next state the following lemma, which formalizes the joiner security, CCA-J, of our UKEM scheme. The full proof is deferred to the full version.

**Lemma 3.** *Let  $\mathcal{G}$  be an asymmetric bilinear group. If PoL is a simulation-extractable proof system, co-CDH holds for  $\mathcal{G}$  and adversary  $\mathcal{A}$  is algebraic in  $\hat{\mathcal{G}}$ , the UKEM construction from Fig. 2 is CCA-J secure in ROM. More precisely, for any algebraic adversary  $\mathcal{A}$ , there exist reductions  $\mathcal{B}$  and  $\mathcal{B}'$  such that*

$$\text{Adv}^{\text{CCA-J}}(\mathcal{A}) \leq \epsilon_{\text{PoL}, n_e}^{\text{sim}} + \epsilon_{\text{PoL}, n_d}^{\text{ext}}(\mathcal{B}') + \text{Adv}_{\mathcal{G}}^{\text{co-CDH}}(\mathcal{B}),$$

where  $n_e$  and  $n_d$  are upper bounds on the number of  $\mathcal{A}$ 's Enc and Dec queries, resp.

**Proof Intuition.** We build upon the proof of Lemma 2 (member security). The only difference is that instead of MChal, the adversary  $\mathcal{A}$  now calls JChal( $u', \pi'$ ). Accordingly, instead of embedding  $v^*$  in the ciphertext returned by the MChal oracle, the reduction  $\mathcal{B}$  against co-CDH now embeds  $v^*$  in the ciphertext  $v'$  returned by JChal; specifically, using random self-reducibility, it sets  $v' := v^* \cdot g^{s'}$  for a random  $s'$ . The security of  $v'$  encrypted to  $u'$  hinges on the link between  $u'$

and the honest  $u_0$  via the associated proof  $\pi'$ . More precisely, unless  $\mathcal{A}$  queries  $H_2$  on  $w' := \text{DH}(u', v')$ , both the “random” key  $K^{(0)}$  and the “real” key  $K^{(1)} = H_2(u', w')$  are random and independent, so  $\mathcal{A}$ 's advantage is 0. On the other hand, if  $\mathcal{A}$  makes such an RO query,  $\mathcal{B}$  can compute the co-CDH solution by extracting from the proof  $\pi'$  as follows.

*Extracting from the  $\pi'$  Proof.* Since  $\mathcal{A}$  is algebraic, when it calls  $\text{JChal}(u', \pi')$ ,  $\mathcal{B}$  can extract the representation of  $\pi'$  as a linear combination of all  $\hat{\mathbb{G}}$  elements given to  $\mathcal{A}$  so far, which are (precisely) the  $\pi_j$  proofs returned by the Enc oracle.  $\mathcal{B}$  knows the logarithm of each such  $\pi_j$  because it emulates the Enc oracle by running Encaps honestly. Thus,  $\mathcal{B}$  can use the representation of  $\pi'$  and the known logarithms to compute the logarithm  $d'$  of  $\pi'$ . Since  $\pi'$  is valid,  $d'$  is equal to the logarithm of  $u'/u_0 = u'/u^*$ . Thus  $\mathcal{B}$  can use  $d'$  and  $s'$  chosen when embedding  $v'$  as  $v' = v^* \cdot g^{s'}$  to translate  $w' = \text{DH}(u', v') = \text{DH}(u^* \cdot g^{d'}, v^* \cdot g^{s'})$  from  $\mathcal{A}$ 's RO query to the solution  $w^* = \text{DH}(u^*, v^*)$  analogously to the reduction for member security:  $w^* = w' \cdot (u^*)^{-s'} \cdot (v')^{-d'}$ .

*Answering Rev Queries.* If, after the JChal call,  $\mathcal{A}$  makes a query  $\text{Enc}(i)$  creating a node  $j$ , then it is allowed to query  $\text{Rev}(j)$ .  $\mathcal{B}$  deals with such queries the same way as the reduction for member security: It samples the secret key  $x_j$  itself. This implies that  $\mathcal{B}$  cannot answer a query  $(u_i, \text{DH}(u_i, v_j))$  to  $H_1$ , which should return  $d_j = \log(g^{x_j}/u_i)$ . But again, such a query would allow  $\mathcal{B}$  to solve its co-CDH instance, had it embedded  $v^*$  in  $v_j$ .

In more detail, recall that Rev queries are allowed for nodes outside of *chall-set* which is the *dec-closure* of all nodes (and their duplicates) created before the JChal call. If after the JChal query  $\mathcal{A}$  queries  $\text{Enc}(i)$  for some  $i \in \text{chall-set}$ ,  $\mathcal{B}$  samples  $x_j$  itself and returns  $v_j = v^* \cdot g^{s_j}$  for a random  $s_j$  together with simulated proofs  $\pi_j$  and  $\tau_j$ . If  $\mathcal{A}$  later “breaks”  $v_j$  by making an RO query  $(u_i, w_j)$  with  $w_j = \text{DH}(u_i, v_j)$  then  $\mathcal{B}$  translates  $w_j$  to the co-CDH solution  $\text{DH}(u^*, v^*)$ : it does so by collecting and summing all  $d$  values on the path from node 0 to node  $i$ : for any Dec edge,  $\mathcal{B}$  extracts  $d$  from the  $\tau$  proof provided by  $\mathcal{A}$ ; for any Enc edge,  $\mathcal{B}$  generated  $d$  itself, as all Enc edges in *chall-set* were created before JChal and hence by running Encaps. Knowing  $d_{0 \rightarrow i}$  s.t.  $u_i = u_0 \cdot g^{d_{0 \rightarrow i}} = u^* \cdot g^{d_{0 \rightarrow i}}$  and  $s_j$  s.t.  $v_j = v^* \cdot g^{s_j}$ , the reduction can compute  $\text{DH}(u^*, v^*) = \text{DH}(u_i, v_j) \cdot (u^*)^{-s_j} \cdot v_j^{-d_{0 \rightarrow i}}$ .

Observe that for the above “simulated” edges,  $\mathcal{B}$  does not know the logarithm of the simulated  $\pi_j$ . This does not affect extraction from the proof  $\pi'$  above, since extraction is done when JChal is called, thus before any proofs are simulated.

*Extracting from Adversarial  $\tau$  Proofs.* Say  $\mathcal{A}$  breaks some  $v_j$  embedded in the response to  $\text{Enc}(i)$  as described above. Simulation-extractability allows  $\mathcal{B}$  to extract from  $\tau$ -proofs for Dec edges as long as these were not simulated. However,  $\mathcal{A}$  could copy a node with a simulated proof via a Dec query. Therefore, we need to modify  $\mathcal{B}$ 's strategy, similarly to the proof of member security.

Consider the following example :  $\mathcal{A}$  starts by calling JChal and then queries  $\text{Enc}(0)$  creating node 1, at which point  $\mathcal{B}$  picks  $x_1$  and sets  $u_1 = g^{x_1}$  as described above. Now  $\mathcal{A}$  forwards the outputs of the above Enc query to  $\text{Dec}(0)$ , creating node 2 with  $(u_2, \tau_2) = (u_1, \tau_1)$ . Finally,  $\mathcal{A}$  queries  $\text{Enc}(2)$ , which creates node

3. Since node 2 is in *chall-set*,  $\mathcal{B}$ , following the above strategy, would choose a fresh key  $x_3$  for  $u_3 := g^{x_3}$ . However, if  $\mathcal{A}$  queries  $(u_2, \text{DH}(u_2, v_3))$  to  $H_1$ , then  $\mathcal{B}$  would not be able to answer, since it does not know  $d_3 := x_3 - x_2$ ; moreover, the value  $\text{DH}(u_2, v_3)$  is of no use, as  $\mathcal{B}$  can compute it itself as  $v_3^{x_2} = v_3^{x_1}$ .  $\mathcal{B}$  should thus just have computed  $u_3$  honestly as  $u_3 = u_2 \cdot g^{d_3}$ . It could then still answer  $\text{Rev}(3)$ , as required, since it knows  $x_3 = x_1 + d_3$ .

$\mathcal{B}$ 's strategy is thus the following:  $u^*$  is embedded as  $u_0$ , and before the JChal query, every Enc query is answered by running **Encaps** (and thus  $\mathcal{B}$  does not know the resulting secret key). After the JChal query, every query  $\text{Enc}(i)$  creating node  $j$  must be answered in a way so  $\mathcal{B}$  knows the resulting secret key  $x_j$ . We distinguish two cases: (1)  $\mathcal{B}$  knows  $x_i$ , or  $x_k$  for any Dec-“ancestor”  $k$  of  $i$ : then  $\mathcal{B}$ , knowing  $x_i$ , runs **Encaps**, and will thus know  $x_j$ . (2) Else  $\mathcal{B}$  sets the resulting key as  $u_j := g^{x_j}$  and  $v_j := v^* \cdot g^{s_j}$  for fresh  $x_j, s_j$ , and simulates the proofs.

Note that for every  $i \notin \text{chall-set}$ ,  $\mathcal{B}$  either knows  $x_i$ , or it can compute it by running **Decaps** between the node  $k$  for which it knows  $x_k$  and  $i$  in order to derive  $x_i$ . Therefore,  $\mathcal{B}$  can answer all Rev queries for such  $i$ .

Moreover, when  $\mathcal{A}$  makes an unanswerable RO query,  $\mathcal{B}$  can use it to break co-CDH. Any such query is of the form  $(u_i, w_j = \text{DH}(u_i, v_j))$  where  $j$  is a node created in mode (2) above (for which  $\mathcal{B}$  does not know  $d_j$ ).  $\mathcal{B}$  extracts all  $d$  values from the proofs  $\tau$  on the path from the root to node  $i$ . This must succeed as long as none of the proofs was simulated, which we show next.

Towards a contradiction, assume that for some  $k$  on that path,  $\tau_k$  for the statement  $(u_{\text{par}_k}, u_k)$  was simulated.  $\mathcal{B}$  must have simulated the proof when, after the JChal call,  $\mathcal{A}$  called  $\text{Enc}(\text{par}_{k'})$ , creating node  $k'$  with  $u_{k'} = u_k$ . However, this means that  $\mathcal{B}$  chose  $x_{k'}$  itself, and thus  $i$  has a Dec-ancestor with a known secret key, meaning that node  $j$  was not created in mode (2), which is a contradiction.

Since  $\mathcal{B}$  can extract all values  $d$  and thus compute  $d_{0 \rightarrow i}$  with  $u_i = u^* \cdot g^{d_{0 \rightarrow i}}$ , and since  $v_j = v^* \cdot g^{s_j}$ , it can translate  $\text{DH}(u_i, v_j)$  to  $\text{DH}(u^*, v^*)$ , as done above.

**Acknowledgments.** This work was funded by the Vienna Science and Technology Fund (WWTF) [10.47379/VRG18002] and by the Austrian Science Fund (FWF) [10.55776/F8515-N]. The authors would like to thank the anonymous reviewers for their valuable comments and suggestions.

## References

- [AAN+22] Alwen, J., et al.: CoCoA: concurrent continuous group key agreement. In: Dunkelmann, O., Dziembowski, S. (eds.) EUROCRYPT 2022. LNCS, vol. 13276, pp. 815–844. Springer, Cham (2022). [https://doi.org/10.1007/978-3-031-07085-3\\_28](https://doi.org/10.1007/978-3-031-07085-3_28)
- [ABR98] Abdalla, M., Bellare, M., Rogaway, P.: DHIES: an encryption scheme based on the Diffie-Hellman problem. In: Contributions to IEEE P1363a, September 1998

- [ABR01] Abdalla, M., Bellare, M., Rogaway, P.: The oracle Diffie-Hellman assumptions and an analysis of DHIES. In: Naccache, D. (ed.) CT-RSA 2001. LNCS, vol. 2020, pp. 143–158. Springer, Heidelberg (2001). [https://doi.org/10.1007/3-540-45353-9\\_12](https://doi.org/10.1007/3-540-45353-9_12)
- [ACDT20] Alwen, J., Coretti, S., Dodis, Y., Tselekounis, Y.: Security analysis and improvements for the IETF MLS standard for group messaging. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020. LNCS, vol. 12170, pp. 248–277. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-56784-2\\_9](https://doi.org/10.1007/978-3-030-56784-2_9)
- [ACDT21] Alwen, J., Coretti, S., Dodis, Y., Tselekounis, Y.: Modular design of secure group messaging protocols and the security of MLS. In: Vigna, G., Shi, E. (eds.) ACM CCS 2021, pp. 1463–1483. ACM Press, November 2021
- [ACJM20] Alwen, J., Coretti, S., Jost, D., Mularczyk, M.: Continuous group key agreement with active security. In: Pass, R., Pietrzak, K. (eds.) TCC 2020. LNCS, vol. 12551, pp. 261–290. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-64378-2\\_10](https://doi.org/10.1007/978-3-030-64378-2_10)
- [AHKM22] Alwen, J., Hartmann, D., Kiltz, E., Mularczyk, M.: Server-aided continuous group key agreement. In: Yin, H., Stavrou, A., Cremers, C., Shi, E. (eds.) ACM CCS 2022, pp. 69–82. ACM Press (2022)
- [AJM22] Alwen, J., Jost, D., Mularczyk, M.: On the insider security of MLS. In: Dodis, Y., Shrimpton, T. (eds.) CRYPTO 2022. LNCS, vol. 13508, pp. 34–68. Springer, Cham (2022). [https://doi.org/10.1007/978-3-031-15979-4\\_2](https://doi.org/10.1007/978-3-031-15979-4_2)
- [AMT23] Alwen, J., Mularczyk, M., Tselekounis, Y.: Fork-resilient continuous group key agreement. In: Handschuh, H., Lysyanskaya, A. (eds.) CRYPTO 2023. LNCS, vol. 14084, pp. 396–429. Springer, Cham (2023). [https://doi.org/10.1007/978-3-031-38551-3\\_13](https://doi.org/10.1007/978-3-031-38551-3_13)
- [AW23] Asano, K., Watanabe, Y.: Updatable public key encryption with strong CCA security: security analysis and efficient generic construction. IACR Cryptology ePrint Archive, p. 976 (2023)
- [BBLW22] Barnes, R., Bhargavan, K., Lipp, B., Wood, C.A.: Hybrid public key encryption. RFC 9180, February 2022
- [BBR+23] Barnes, R., Beurdouche, B., Robert, R., Millican, J., Omara, E., Cohn-Gordon, K.: The messaging layer security (MLS) protocol. RFC 9420, July 2023
- [BLS01] Boneh, D., Lynn, B., Shacham, H.: Short signatures from the weil pairing. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 514–532. Springer, Heidelberg (2001). [https://doi.org/10.1007/3-540-45682-1\\_30](https://doi.org/10.1007/3-540-45682-1_30)
- [BLS04] Barreto, P.S.L.M., Lynn, B., Scott, M.: On the selection of pairing-friendly groups. In: Matsui, M., Zuccherato, R.J. (eds.) SAC 2003. LNCS, vol. 3006, pp. 17–25. Springer, Heidelberg (2004). [https://doi.org/10.1007/978-3-540-24654-1\\_2](https://doi.org/10.1007/978-3-540-24654-1_2)
- [Bow] Bowe, S.: Bls12-381: New zk-snark elliptic curve construction
- [CCD+20] Cohn-Gordon, K., Cremers, C., Dowling, B., Garratt, L., Stebila, D.: A formal security analysis of the signal messaging protocol. *J. Cryptol.* **33**(4), 1914–1983 (2020)
- [Dam92] Damgård, I.: Towards practical public key systems secure against chosen ciphertext attacks. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 445–456. Springer, Heidelberg (1992). [https://doi.org/10.1007/3-540-46766-1\\_36](https://doi.org/10.1007/3-540-46766-1_36)



- [DKW21] Dodis, Y., Karthikeyan, H., Wichs, D.: Updatable public key encryption in the standard model. In: Nissim, K., Waters, B. (eds.) TCC 2021, Part III. LNCS, vol. 13044, pp. 254–285. Springer, Cham (2021). [https://doi.org/10.1007/978-3-030-90456-2\\_9](https://doi.org/10.1007/978-3-030-90456-2_9)
- [DP92] De Santis, A., Persiano, G.: Zero-knowledge proofs of knowledge without interaction (extended abstract). In: 33rd FOCS, pp. 427–436. IEEE Computer Society Press, October 1992
- [DV19] Durak, F.B., Vaudenay, S.: Bidirectional asynchronous ratcheted key agreement with linear complexity. In: Attrapadung, N., Yagi, T. (eds.) IWSEC 2019. LNCS, vol. 11689, pp. 343–362. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-26834-3\\_20](https://doi.org/10.1007/978-3-030-26834-3_20)
- [EJKM22] Eaton, E., Jao, D., Komlo, C., Mokrani, Y.: Towards post-quantum key-updatable public-key encryption via supersingular isogenies. In: AlTawy, R., Hülsing, A. (eds.) SAC 2021. LNCS, vol. 13203, pp. 461–482. Springer, Cham (2022). [https://doi.org/10.1007/978-3-030-99277-4\\_22](https://doi.org/10.1007/978-3-030-99277-4_22)
- [FKL18] Fuchsbauer, G., Kiltz, E., Loss, J.: The algebraic group model and its applications. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018. LNCS, vol. 10992, pp. 33–62. Springer, Cham (2018). [https://doi.org/10.1007/978-3-319-96881-0\\_2](https://doi.org/10.1007/978-3-319-96881-0_2)
- [FO22] Fuchsbauer, G., Orrù, M.: Non-interactive mumblewumble transactions, revisited. In: Agrawal, S., Lin, D. (eds.) ASIACRYPT 2022, Part I. LNCS, vol. 13791, pp. 713–744. Springer, Cham (2022). [https://doi.org/10.1007/978-3-031-22963-3\\_24](https://doi.org/10.1007/978-3-031-22963-3_24)
- [FPS20] Fuchsbauer, G., Plouviez, A., Seurin, Y.: Blind Schnorr signatures and signed ElGamal encryption in the algebraic group model. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020, Part II. LNCS, vol. 12106, pp. 63–95. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-45724-2\\_3](https://doi.org/10.1007/978-3-030-45724-2_3)
- [GHJL17] Günther, F., Hale, B., Jager, T., Lauer, S.: 0-RTT key exchange with full forward secrecy. In: Coron, J.-S., Nielsen, J.B. (eds.) EUROCRYPT 2017, Part III. LNCS, vol. 10212, pp. 519–548. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-56617-7\\_18](https://doi.org/10.1007/978-3-319-56617-7_18)
- [GM15] Green, M.D., Miers, I.: Forward secure asynchronous messaging from puncturable encryption. In: 2015 IEEE Symposium on Security and Privacy, pp. 305–320. IEEE Computer Society Press, May 2015
- [HKP+21] Hashimoto, K., Katsumata, S., Postlethwaite, E., Prest, T., Westerbaan, B.: A concrete treatment of efficient continuous group key agreement via multi-recipient PKEs. In: Vigna, G., Shi, E. (eds.) ACM CCS 2021, pp. 1441–1462. ACM Press, November 2021
- [HLP22] Haidar, C.A., Libert, B., Passelègue, A.: Updatable public key encryption from DCR: efficient constructions with stronger security. : Yin, H., Stavrou, A., Cremers, C., Shi, E. (eds.) ACM CCS 2022, pp. 11–22. ACM Press, November 2022
- [HPS23] Haidar, C.A., Passelègue, A., Stehlé, D.: Efficient updatable public-key encryption from lattices. Cryptology ePrint Archive, Paper 2023/1400 (2023). <https://eprint.iacr.org/2023/1400>
- [JMM19] Jost, D., Maurer, U., Mularczyk, M.: Efficient ratcheting: almost-optimal guarantees for secure messaging. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019. LNCS, vol. 11476, pp. 159–188. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-17653-2\\_6](https://doi.org/10.1007/978-3-030-17653-2_6)

- [JS18] Jaeger, J., Stepanovs, I.: Optimal channel security against fine-grained state compromise: the safety of messaging. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018. LNCS, vol. 10991, pp. 33–62. Springer, Cham (2018). [https://doi.org/10.1007/978-3-319-96884-1\\_2](https://doi.org/10.1007/978-3-319-96884-1_2)
- [OP01] Okamoto, T., Pointcheval, D.: The gap-problems: a new class of problems for the security of cryptographic schemes. In: Kim, K. (ed.) PKC 2001. LNCS, vol. 1992, pp. 104–118. Springer, Heidelberg (2001). [https://doi.org/10.1007/3-540-44586-2\\_8](https://doi.org/10.1007/3-540-44586-2_8)
- [PR18] Poettering, B., Rösler, P.: Towards bidirectional ratcheted key exchange. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part I. LNCS, vol. 10991, pp. 3–32. Springer, Cham (2018). [https://doi.org/10.1007/978-3-319-96884-1\\_1](https://doi.org/10.1007/978-3-319-96884-1_1)
- [PS00] Pointcheval, D., Stern, J.: Security arguments for digital signatures and blind signatures. *J. Cryptol.* **13**(3), 361–396 (2000)
- [Sah99] Sahai, A.: Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In: 40th FOCS, pp. 543–553. IEEE Computer Society Press, October 1999
- [SKSW22] Sakemi, Y., Kobayashi, T., Saito, T., Wahby, R.S.: Pairing-friendly curves. internet-draft draft-IRTF-CFRG-pairing-friendly-curves-11, Internet Engineering Task Force, November 2022. Work in Progress