

Cryptographic Security of the MLS RFC, Draft 11

Chris Brzuska, Eric Cornelissen, Konrad Kohbrok
Aalto University, Finland

Abstract—Cryptographic communication protocols provide confidentiality, integrity and authentication properties for end-to-end communication under strong corruption attacks, including, notably, post-compromise security (PCS). Most protocols are designed for one-to-one communication. Protocols for group communication are less common, less efficient, and tend to provide weaker security guarantees. This is because group communication poses unique challenges, such as coordinated key updates, changes to group membership and complex post-compromise recovery procedures.

We need to tackle this complex challenge as a community. Thus, the Internet Engineering Task Force (IETF) has created a working group with the goal of developing a sound standard for a continuous asynchronous key-exchange protocol for dynamic groups that is secure and remains efficient for large group sizes. The current version of the *Messaging Layer Security* (MLS) security protocol is in a feature freeze, i.e., no changes are made in order to provide a stable basis for cryptographic analysis. The key schedule and TreeKEM design are of particular concern since they are crucial to distribute and combine several keys to achieve PCS.

In this work, we provide a computational analysis of the MLS key schedule, TreeKEM and their composition, as specified in Draft 11 of the MLS RFC. The analysis is carried out using the State Separating Proofs methodology [9], and showcases the flexibility of the approach, enabling us to provide a full computational analysis shortly after Draft 11 was published.

I. INTRODUCTION

We expect modern-day messaging applications to provide end-to-end security, guaranteeing confidentiality and authenticity of the transmitted messages. This expectation has been made a reality over the course of 25 years through the design of cryptographic protocols, which nowadays guarantee secure communication, potentially even after a participant’s key is compromised.

Our everyday communication protocols inherit their security from the properties of the keys they use, which, in turn, are established via key-exchange protocols.

Key exchange protocols therefore have been designed to achieve strong security guarantees: *Entity Authentication* [4], [23] ensures that the key is shared only with the intended recipient. *Key Indistinguishability* [16], [23], in turn, guarantees that no one but the participants involved has information about the key. A stronger variant of key indistinguishability called *Forward Secrecy* (FS) ensures that short-term communication keys exchanged in past sessions are secure if a party’s current state or long-term key gets compromised [23].

With the introduction of *continuous* key-exchange, for example by protocols such as OTR [7] or ZRTP [28], parties are able to recover from compromise if the adversary remains passive for a brief period of time. This counterpart to FS was first described as *Post-Compromise Security* (PCS) by Cohn-Gordon et al. [11], [18].

Modern messaging protocols require that messages can be sent even if a recipient is offline, requiring an *asynchronous* protocol. As a consequence, protocol sessions don’t naturally end when one participant goes offline, leading to sessions that remain active for months and years. To cover this use-case, non-interactive key exchange protocols were introduced that perform key-exchange and entity authentication continuously even in an asynchronous setting. Most notably, the *Signal protocol* [22] (formerly *TextSecure*) achieves strong FS and PCS guarantees [13], [11] due to the use of double ratcheting (formerly *Axolotl*).

As users of messaging protocols tend to use more than one device and/or communicate with more than one party at a time, messaging protocols are commonly required to be group messaging protocols. This gives rise to the definition of continuous group key exchange protocols by Alwen et al. [2].

The first modern, continuous group key exchange protocol is based on an *Asynchronous Ratcheting Tree* (ART) and was proposed by Cohn-Gordon et al. [12]. Similar to an earlier protocol proposed by Steer et al. [26] (see Figure 1), it is based on a binary tree where each leaf represents a group member. The secrets of the members are recursively combined in lower nodes of the tree such that the secret in the root node is shared by all group members. In contrast to the protocol proposed by Steer et al., ART supports FS and PCS guarantees.

MLS. The *Message Layer Security* (MLS) Working Group of the Internet Engineering Task Force (IETF) aims to create a new continuous, asynchronous group key-exchange protocol [3] that is efficient even for large groups [25]. To achieve this, the current draft makes use of the TreeKEM [6] protocol.

In short, TreeKEM, which is inspired by ART, is a continuous group key distribution based on a tree structure where each leaf node represents a member’s Key Encapsulation Method (KEM) key pair and all other nodes represent a secret value and KEM key pair shared by the members in the node’s subtree. As a result, secrets can be shared efficiently by encrypting it to individual subtrees. The secret value at the root node of the tree is a secret shared by the entire group.

Each member can update their keying material by updating the key pairs and values of their leaf-to-root path, and transfer the new key material to all other members efficiently by encrypting it to the set of subtrees below the path nodes. Each

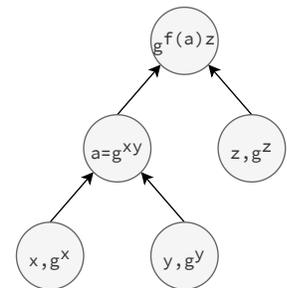


Figure 1: Key derivation in group key exchange protocol [26]. f encodes a as a valid exponent.

member can also remove a member or add a new member, and update the key material so that previous group members cannot access key material any longer. In Section IV, we explain TreeKEM in more detail. The updating of keys allows members to achieve PCS and FS similar to ART.

Note that TreeKEM does not have an internal mechanism for members to reconcile differing views of individual groups and thus relies on members agreeing on a total order of group operations. One straightforward way of achieving this is a central distribution service enforcing that order.

Once a new secret value is distributed via TreeKEM, MLS combines the new secret with existing key material used for communication. From these two keys, a new communication key is derived for the next *epoch*. In addition, key derivation includes information about the group and, optionally, allows to inject one or more external keys. The combination of keys allows MLS to achieve PCS and FS guarantees, and the inclusion of group parameters such as membership ensures agreement on said parameters. The processing and combining of the key material and context information is described in the *MLS key schedule*.

Related Work. Earlier versions of the MLS RFC have been analyzed in several works [1], [2], [15], [27]. Alwen et al. [1] propose an alternative to TreeKEM, named *Tainted TreeKEM*, which replaces blank nodes (node without keys, see Section IV) by so-called *tainted* nodes. In contrast to a blank node, the keys of a tainted node are still used. A tainted node will be blank after the entity that tainted the node updates their keying material. In doing so, depending on the distribution of add and remove operations, updates are more efficient compared to TreeKEM.

Weidner [27] proposes an alternative version of TreeKEM, named *Causal TreeKEM*, that does not require a strict total order of group operations. Where the MLS RFC assumes a trusted server that can enforce ordering of messages because key updates overwrite existing keys, Causal TreeKEM does not have this requirement because it uses keying material that can be combined in an arbitrary order. However, this requires keying material that can be combined in an associative and commutative manner, such as Diffie-Hellman shares. MLS instead opts for higher modularity with a KEM approach, where the underlying encryption scheme can be changed if needed. Specifically, TreeKEM can achieve post-quantum security by choosing an appropriate KEM scheme, which is not possible in causal TreeKEM.

In [15], a comparison between MLS and a straightforward group messaging protocol made of 1:1 Signal sessions, reveals that in the 1:1 setting, authentication can be healed, while, in the group setting, MLS never heals authentication thus weakening PCS guarantees in the case of signature key corruptions.

Relatedly, Alwen et al. [2] show that the Forward Secrecy (FS) guarantees of MLS are relatively weak. An update to the keying material of one client will leave the group vulnerable w.r.t. FS from other corrupted clients. Therefore, in the worst case, FS is only achieved if every group member updates their keying material. To mitigate this shortcoming, Alwen et al. propose to use a PKE scheme in TreeKEM where keys are updated with every encryption and decryption operation. The

proposal was not integrated into MLS due to concerns that supporting homomorphism is not compatible with considering the KEM as a black-box primitive.

Contributions. In this work, we study the security claims made in Draft 11 of the MLS RFC via a cryptographic analysis of TreeKEM, the MLS key schedule and their composition. Key schedule analysis was implicitly suggested in Krawczyk’s study of Sigma protocols [19] and used for the analysis of TLS 1.3 [8]. We develop security models for TreeKEM, the key schedule and their composition. Interestingly, our models themselves are composable, and the composed model is derived from the composition of the TreeKEM model and the key schedule model. This composability feature is enabled via the use of the State Separating Proofs methodology [9] which specifies security models via modular pseudocode and enables code reuse (see Section V for details).

Our model allows the adversary to corrupt all base secrets statically, and to decide to corrupt derived keys upon derivation (CUD). This corruption model simplifies the security analysis (in comparison to fully adaptive corruption) and models the intended security of MLS in a meaningful way. Firstly, adaptive corruption of long-term secrets is most relevant for signing keys, which are outside the scope of this work. Secondly, as session states are intended to be rather short-lived, dynamically corrupting old sessions should be difficult in practice. Thirdly, in light of the impossibility result of non-committing encryption [24], our model seems close to the strongest possible security which can be provably achieved without random oracles.

We establish that the keys produced by TreeKEM and the key schedule of MLS are pseudorandom in the CUD corruption model. We make standard key indistinguishability and collision-resistance assumptions on the key derivation functions (KDF) and assume indistinguishability under chosen-ciphertext attacks (IND-CCA) secure public-key encryption, as well as that the Extract function in Krawczyk’s HKDF design [20] is a dual pseudorandom function and thus, HKDF is a dual KDF, which has also been assumed in the analysis of Noise [17] and TLS 1.3 [8]. Additionally, we assume that all primitives are collision-resistant. Following the approach of [10], [8], we model the security of our symmetric primitives as multi-instance primitives with static corruptions. Jumping ahead, based on our analysis, we suggest to include the group context into the derivation of the joiner secret [14] in order to strengthen its uniqueness properties.

We rely on modular proofs and pseudocode-reuse as specified by the state-separating proofs (SSP) methodology [9]. SSPs allow us to iterate the proof quickly as the MLS draft is developed. Given an initial model of the protocol, we make and verify incremental changes to the model with relative ease. This is what allows us to release this analysis of the protocol shortly after the current feature freeze and it will also allow us to quickly analyse the security of future features.

Overview. Section II introduces notation. Section III provides syntax for continuous key distribution and key schedule. Section IV explains the TreeKEM, the MLS key schedule and their composition. Section V introduces the state-separating proofs methodology and our assumptions. Section VI explains

our security model, Section VII states our three main theorems, and Section VIII and the appendix contain the proofs.

II. PRELIMINARIES AND NOTATION

Binary Trees. A *binary tree* is a finite, connected, and directed graph with a single source — called *root* — such that each node has at most two outgoing edges. For two nodes n_0 and n_1 , if there is an edge $n_0 \rightarrow n_1$, then we call n_0 the *parent* of n_1 and n_1 the *child* of n_0 . All nodes which can be reached from n_0 are its *descendants*, and all nodes from which n_1 can be reached are its *ancestors*. Every node can be reached starting from the root, and we call the set of ancestors of a leaf node its *direct path*. We allow for nodes to be marked as *blank* (the marking will later take on the meaning of no associated data). We call the *co-path* a list of the closest non-blank descendent nodes of the nodes along the direct path.

We leave the encoding of the tree abstract in this paper and only assume that given the size of a tree and the index i of a node, we can find the indices of its direct path.

Notation. We use pseudocode to describe algorithms. We denote sets and tables by capital letters, e.g. S or T . We denote algorithms in lowercase and sans serif. $x \leftarrow a$ assigns value a to variable x , and $x \leftarrow \text{algo}(a)$ runs algo on value a and assigns the result to variable x . When algo is a randomized algorithm, $x \leftarrow_{\$} \text{algo}(a)$ runs algo on value a with fresh randomness and assigns the result to variable x . When S is a set, $x \leftarrow_{\$} S$ samples a uniformly random value from S and assigns that value to variable x . A common set is $\{0, 1\}^\lambda$, the set of bitstrings of length λ . We use \bar{a} for a vector of variables, and $\bar{x} \stackrel{\text{vec}}{\leftarrow} \text{algo}(\bar{a})$ means, we run algo separately on each value in the vector \bar{a} and then assign the results to the corresponding indices in vector \bar{x} . Analogously, $\bar{x} \stackrel{\text{vec}}{\leftarrow}_{\$} S$ means that we sample each value in the vector \bar{x} independently and uniformly at random. For two tables T and U , the operation $T \stackrel{\text{merge}}{\leftarrow} U$ overwrites values $T[i]$ by values $U[i]$ whenever $U[i]$ is defined, and the operation $T \stackrel{\text{rem}}{\leftarrow} U$ removes all values $T[i]$ from T where $U[i]$ is defined. `assert cond` abbreviates that a special error message is sent unless *cond* holds.

III. CONTINUOUS GROUP KEY DISTRIBUTION

A *continuous group key distribution* protocol (CGKD) allows a group of users to maintain a continuous session which

- distributes a secret value among group members in each epoch;
- allows for efficient updates to the key material of a member and the group;
- allows for efficiently adding new clients to the group;
- allows for efficiently removing members from the group.

Assuming pseudorandomness and uniqueness of input keys, a CGKD achieves pseudorandomness and uniqueness of output keys. It does not model signatures and attacks against authentication, but instead allows the adversary to combine keys (almost) arbitrarily. This way, we can model the important security property of post-compromise security (PCS), which demands that if key material of a user gets corrupted (and the adversary learns the session key), the security of the protocol

Input	Description	U	P	J
PK_{mem}	current public view of the group	•	•	•
PK_{rem}	public key(s) to remove	•	•	
PK_{join}	public key(s) to add	•	•	
PK_{upd}	new public key(s) of updater		•	
SK_{own}	own current secret key(s)	•	•	•
i_{own}	own index	•	•	•
i_{upd}	updater's index		•	•
k_{upd}	seed for updating	○		
CP	indices in PK_{mem} to encrypt to	•		
pk_{ext}	external public key	○		
s_{ext}	external secret randomness(es)	•	•	•
s_{int}	internal secret randomness(es)	•	•	
ctx	public context	•	•	•
c	ciphertext from a member		○	•/○
c_{ext}	ciphertext from an outsider		○	
Output	Description	U	P	J
C_{CP}	ciphertext(s) for members	•		
C_{join}	ciphertext(s) for joiners	•		
c_{ext}	ciphertext from outsider	○		
PK'_{upd}	new public key(s) of updater	•		
PK'_{mem}	new view of the group	•	•	•
SK'_{own}	new own secret key(s)	•	•	•
k	new shared secret key(s)	•	•	•
s'_{int}	new s_{int} for the next update	•	•	•
pk'_{ext}	new external public key	•	•	•

Figure 2: • and ○ indicate the value can be an input/output of the algorithms of `cgkd` and `wcgkd` according to our syntax, with grey rows only available in `cgkd`. Values marked by ○ are optional; the algorithm might behave differently if provided. Input c is required for join but optional for `wjoin`.

can recover. This section introduces the syntax of a CGKD, the syntax of a weak variant of continuous group key distribution (wCGKD) and the syntax of a key schedule (KS). The main difference between a CGKD and a wCGKD is that the group keys of a wCGKD are corrupt whenever one group member's current private keys are corrupted. This property is, indeed, rather weak, since in large groups, individual members might use weak randomness, and for insecurity, it suffices if a single member of the group relies on weak randomness for generating their current own keys. In turn, a (strong) CGKD relies on key material which is chained across *epochs*. If the key material from a previous epoch was secure, then security in the current epoch is maintained—unless the key material is leaked by encrypting it to a client with corrupt keys that joins the group. As we will see for MLS, a wCGKD can be combined with a key schedule to obtain a CGKD. We now first introduce the syntax of a CGKD, a wCGKD and a KS and then elaborate how this syntax reflects MLS.

Updates of the key material are initiated by one group member and are modeled in the (w)update algorithm in (w)CGKD. In order for other group members to keep their key material in sync with an updating party, they need to process the messages generated by the update algorithm. Since new group members might perform processing in different way

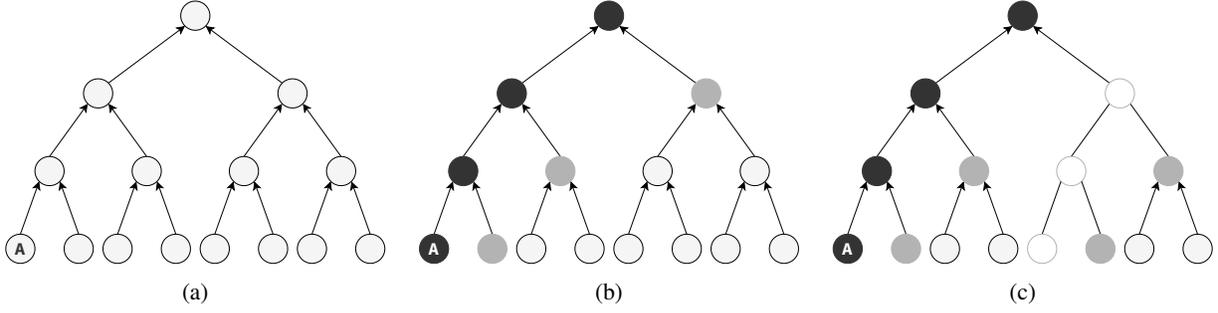


Figure 5: Visualisation of a member A updating their their key material. (a) shows the tree for the group. (b) shows the tree during the update, highlighting the *direct path* in black and the *co-path* in dark grey. (c) shows the tree during the update if there are *blank nodes* in the tree.

the group and indeed derived from the value ps which is stored at the root of the tree. Figure 3 illustrates the data associated with each node as well as the members to which the secret node data is known. Every node in the tree has associated with it a KEM key pair known to the members represented by the leaves in the sub-tree. To share a secret value with some members, a member can encrypt the value either to the public key of a node whose secret key all intended members know or encrypt to several public keys for a more fine-grained information-sharing.

The main sub-procedure of the TreeKEM is the Derive Node procedure dn , see Figure 3 and Figure 4. dn uses an expand function (XPD) and a key encapsulation method (KEM) as building blocks. dn takes as input the secret of a node, referred to as the *path secret* ps_i of node i , and uses XPD to derive two random values: The path secret $ps_{p(i)}$ (where $p(i)$ refers to the parent of node i) and randomness for (DKP) to compute a KEM key pair (pk_i, sk_i) associated with node i .

Roughly, in the tree of Figure 3, each arrow can be interpreted as a possible application of dn with the restriction that if a node i has two non-blank children, then its path secret was only derived from one of them. The keys associated with each node evolve over time via *updates* which start from a leaf node and proceed along the direct path to the root. The path secret ps_i of a node i is derived from the child which was last contained in the direct path of an update.

We next explain key updates. In addition, one can also combine an update with the *adding* and/or the *removing* of one or more members. While adding and removing is implemented simultaneously—see Figure 9 for the pseudocode of update—we describe the conceptual idea for each separately below.

Updating Key Material. Consider the group represented by the tree in Figure 5a. The group member Alice (A) can update the group secret by the following process. First, A samples a random value k_{upd} for the key update and uses it as the new path secret associated with A 's leaf node, i.e., $ps \leftarrow k_{upd}$. Then, Alice runs $(ps', pk, sk) \leftarrow dn(ps)$ and stores pk in $PK'_{mem}[i_{cur}]$ and $PK_{upd}[i_{cur}]$, where i_{cur} is the index of Alice's leaf node, PK'_{mem} is a table which represents Alice's view of the public keys of the group members, and PK_{upd} is a table which contains the public-keys that Alice will tell

other group members to change. Additionally, Alice also stores the secret key in $SK'_{own}[i_{cur}]$. SK is a table of secret keys which Alice knows. Then, Alice computes the parent index of i_{cur} and proceeds in the same way for the parent index: First, she computes $(ps'', pk, sk) \leftarrow dn(ps')$, then stores pk and sk in the adequate tables, then moves to the parent node and applies dn to ps'' etc. until reaching the root node. This loop is captured by $dopath$ in Figure 9—it is only executed if k_{upd} is non-empty. In the end, Alice has replaced all the path secrets, public-keys and secret-keys of the direct path, marked in black in Figure 5b.

As only Alice knows the path secrets used along this path, she now also shares the path secret with the other group members. Thus, for each node i_{cur} on her direct path, she considers the child i_{child} of i_{cur} which is not on her direct path (but, instead, is on the *co-path*) and encrypts the path secret of i_{cur} under $PK_{mem}[i_{child}]$, the public-key associated with this node. In this way, the entire sub-tree below $PK_{mem}[i_{child}]$ can decrypt the ciphertext. In Figure 5b, the co-path is marked in grey, i.e., for each grey node, Alice encrypts to the public-key associated with it.

Figure 5a and Figure 5b now illustrate a case where there are no *blank nodes* in the tree—blank nodes are nodes which do not have data associated with it. Figure 5c depicts the case that the tree contains some blank nodes. In this case, Alice encrypts to the nodes marked in grey in Figure 5c. Note that we nevertheless refer to these nodes as co-path even though our co-path definition does not coincide with the graph-theoretic notion of co-path. In order to simplify the computation of the co-path, we give a description of the co-path CP as input to update. CP is a list of sets, i.e., for each tree index i , $CP[i]$ contains a set of the indices to which Alice needs to encrypt.

Adding a New Member. There are two ways to add a new member. One can either add a member in an *add-only* operation, or one can add a member and perform an update at the same time. (In this case, one can also remove members simultaneously, but let us ignore this operation for now.) When performing an *add* operation, we either find the left-most empty leaf node (there might be empty leaf nodes due to previous removals) or, if the tree is full, we replace a leaf node by a 3-node sub-tree. E.g., consider the original group in Figure 6a. Here, to add Henry (H) to this group, the group

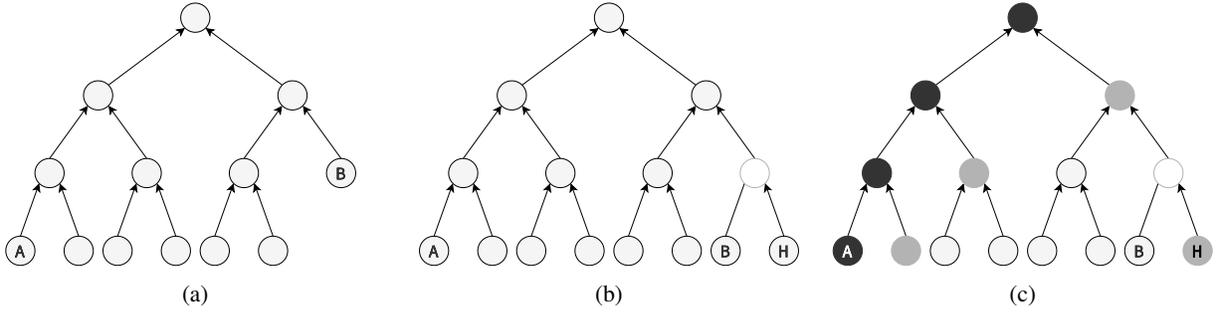


Figure 6: Visualisation of adding a new member to a group. (a) shows the tree for the original group. (b) shows the tree with H added to the group in an *add-only* operation, highlighting blanked nodes in white. (c) shows the tree for the update performed by A , highlighting the *direct path* in black and the *co-path* in dark grey. Note that H 's unmerged leaf is part of the co-path.

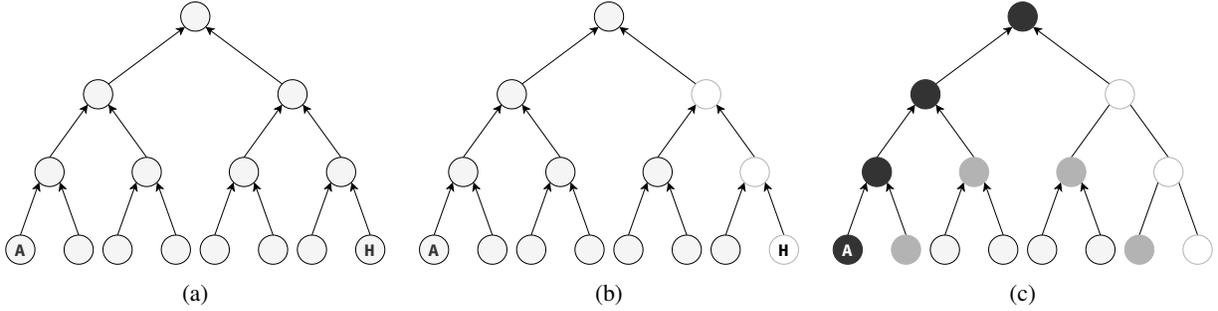


Figure 7: Visualisation of removing a member from a group. (a) shows the tree for the original group. (b) shows the tree with H removed. (c) shows the tree for the update after the removal, highlighting the *direct path* in black and the *co-path* in dark grey. Blanked nodes are highlighted in white in both (b) and (c).

member Alice (A) blanks Bob's (B) node, creates below it a node for H and a new node for B and associates H 's node with a KEM public-key which Alice knows belong to H , and associates B 's new node with the same public-key which B had before. Moreover, Alice runs the key schedule with an empty k_{upd} and encrypts k_{joi} to H .

In the case that A wants to simultaneously perform an update of the key material, she now updates her path as described in Section IV and encrypts the relevant path-secrets to the co-path, marked in grey. Again, the co-path is not the graph-theoretic co-path, since A needs to encrypt a message to H separately. In addition to sharing the relevant path secrets with the co-path, A also encrypts the joiner secret k_{joi} derived via the MLS Key Schedule (see Figure 3 and Figure 4) to H . Recall that the dopath procedure in Figure 9 is only executed if k_{upd} is a non-empty variable. If k_{upd} is empty and an add-only is performed, the computations are performed in the same way except that k_{cs} is now a fixed zero value. It is important to execute the ks function also in an add-only operation to ensure that the new joiners cannot read messages from the time before they joined.

A special case is the situation where the add operation is not performed from *inside* the group, but instead, an outside member adds itself to the group. In this case, the external group public-key pk_{ext} is used (marked in red in Figure 9) to transmit s_{int} . Note that the red code is only run if pk_{ext} is non-empty, which requires that k_{upd} is non-empty, too.

Removing a Member. Consider the group represented by the tree in Figure 7a. To remove a group member Henry (H) from this group, the group member Alice (A) blanks the nodes on the path from node H to the root, shown in Figure 7b. As a result, a co-path that included those nodes before, now, instead includes their non-blank child nodes.

Notice that the removal by itself does not update the group secret and therefore, A also needs to perform an update to ensure that H cannot decrypt messages sent by group members anymore. The requirement to perform an update is captured by

$$\text{if } |PK_{rem}| > 0 : \text{assert } k_{upd} \neq \perp$$

in line 2 of the update procedure in Figure 9. As the tree now contains blanked nodes, the path secrets along the new direct path must be encrypted under the keys of a larger co-path, as shown in Figure 7c.

MLS CGKD. Combining TreeKEM and ks as depicted in Figure 3 yields the MLS CGKD. Note that the MLS CGKD achieves significantly stronger security properties than the wCGKD. Namely, freshness in the wCGKD is violated if *any* of the group keys are compromised (and not yet updated), see left side of Figure 8. In turn, the MLS CGKD also yields security if the s_{ext} or the s_{int} from a previous epoch are fresh. See Figure 9 for the pseudocode of the update procedure and Figure 8 for its freshness.

```

own_keys( $PK_{mem}, PK_{rem}, PK_{joi}, SK_{own}, k_{upd}, pk_{ext}, s_{ext}, s_{int}$ )
if !fresh( $k_{upd}$ ):
  return all corrupt
 $PK'_{mem} \leftarrow PK_{mem}$ 
 $PK'_{mem} \stackrel{rem}{\leftarrow} PK_{rem}$ 
 $PK'_{mem} \stackrel{merge}{\leftarrow} PK_{joi}$ 
 $F \leftarrow []$ 
fresh  $\leftarrow 1$ 
for  $i_{cur}, _ \in SK_{own}$ :
  for  $i_{nod} \in CP[i_{cur}]$ :
    if !fresh( $PK'_{mem}[i_{nod}]$ ):
      fresh  $\leftarrow 0$ 
       $F[i_{cur}] \leftarrow fresh$ 
  return F

```

```

grp_keys( $PK_{mem}, PK_{rem}, PK_{joi}, SK_{own}, k_{upd}, pk_{ext}, s_{ext}, s_{int}$ )
if fresh( $s_{int}$ )  $\wedge$  fresh( $pk_{ext}$ ):
  return 1
if fresh( $s_{ext}$ ):
  return 1
if  $0 \notin \text{own\_keys}(\dots)$ :
  return 1
return 0

```

```

update( $PK_{mem}, PK_{rem}, PK_{joi}, SK_{own}, i_{own}, k_{upd}, CP, pk_{ext}, s_{ext}, s_{int}, ctx$ )
if  $|PK_{rem}| > 0 \vee pk_{ext} \neq \perp$ :
  assert  $k_{upd} \neq \perp$ 
  ( $C_{cp}, PK'_{upd}, PK'_{mem}, SK'_{own}, k_{cs}$ )
   $\leftarrow$  wupdate( $PK_{mem}, PK_{rem}, PK_{joi}, SK_{own}, i_{own}, k_{upd}, CP$ )
   $c_{ext} \leftarrow$  pke.enc( $pk_{ext}, s_{int}$ )
  ( $s'_{int}, \bar{k}_{\overline{tbl}}, k_{joi}, pk'_{ext}, sk'_{ext}$ )
   $\leftarrow$  ks( $k_{cs}, s_{int}, s_{ext}, \bar{tbl}, ctx$ )
   $SK'_{own}[-1] \leftarrow sk'_{ext}$ 
   $C_{joi} \leftarrow []$ 
  for  $i_{nod}, pk \in PK_{joi}$ :
     $C_{joi}[i_{nod}] \leftarrow$  s pke.enc( $pk, k_{joi}$ )
  return ( $C_{cp}, C_{joi}, c_{ext}, PK'_{upd}, PK'_{mem}, SK'_{own}, \bar{k}_{\overline{tbl}}, s'_{int}, pk'_{ext}$ )

dopath( $PK_{mem}, i_{cur}, k_{upd}, CP$ )
 $C_{cp} \leftarrow []$ 
 $PK'_{own}, SK'_{own} \leftarrow [], []$ 
 $ps \leftarrow k_{upd}$ 
while  $i_{cur} \neq \perp$ :
  ( $ps', pk, sk$ )  $\leftarrow$  dn( $ps$ )
   $PK'_{own}[i_{cur}] \leftarrow pk$ 
   $SK'_{own}[i_{cur}] \leftarrow sk$ 
  if  $i_{cur} \neq i_{own}$ :
    for  $i_{nod} \in CP[i_{cur}]$ :
       $pk_{nod} \leftarrow PK_{mem}[i_{nod}]$ 
       $c \leftarrow$  s pke.enc( $pk_{nod}, ps'$ )
       $C_{cp}[i_{nod}] \leftarrow c$ 
   $i_{cur} \leftarrow$  parent( $i_{cur}, |PK_{mem}|$ )
   $ps \leftarrow ps'$ 
return ( $C_{cp}, PK'_{own}, SK'_{own}, ps$ )

process( $PK_{mem}, PK_{rem}, PK_{joi}, PK_{upd}, SK_{own}, i_{nod}, i_{upd}, s_{ext}, s_{int}, ctx, c, c_{ext}$ )
( $PK'_{mem}, SK'_{own}, k_{cs}$ )  $\leftarrow$ 
wprocess( $PK_{mem}, PK_{rem}, PK_{joi}, PK_{upd}, SK_{own}, i_{nod}, i_{upd}, c$ )
 $s_{int} \leftarrow$  pke.dec( $SK_{own}[-1], c_{ext}$ )
assert  $s_{int} \neq \perp$ 
( $s'_{int}, \bar{k}_{\overline{tbl}}, \dots, pk'_{ext}, sk'_{ext}$ )
 $\leftarrow$  ks( $k_{cs}, s_{int}, s_{ext}, \bar{tbl}, ctx$ )
 $SK'_{own}[-1] \leftarrow sk'_{ext}$ 
return ( $PK'_{mem}, SK'_{own}, \bar{k}_{\overline{tbl}}, s'_{int}, pk'_{ext}$ )

join( $PK_{mem}, SK_{own}, i_{own}, i_{upd}, s_{ext}, ctx, c$ )
 $SK'_{own} \leftarrow SK_{own}$ 
 $sk_{own} \leftarrow SK'_{own}[i_{own}]$ 
 $k_{joi} \leftarrow$  pke.dec( $sk_{own}, c$ )
assert  $k_{joi} \neq \perp$ 
( $s'_{int}, \bar{k}_{\overline{tbl}}, \dots, pk'_{ext}, sk'_{ext}$ )
 $\leftarrow$  subks( $k_{joi}, s_{ext}, \bar{tbl}, ctx$ )
 $SK'_{own}[-1] \leftarrow sk'_{ext}$ 
return ( $PK_{mem}, SK'_{own}, \bar{k}_{\overline{tbl}}, s'_{int}, pk'_{ext}$ )

wupdate( $PK_{mem}, PK_{rem}, PK_{joi}, SK_{own}, i_{own}, k_{upd}, CP$ )
 $k_{cs} \leftarrow 0$ 
 $PK'_{mem} \leftarrow PK_{mem}$ 
 $PK'_{mem} \stackrel{rem}{\leftarrow} PK_{rem}$ 
 $PK'_{mem} \stackrel{merge}{\leftarrow} PK_{joi}$ 
 $C_{cp}, PK'_{upd} \leftarrow [], []$ 
 $SK'_{own} \leftarrow SK_{own}$ 
for  $i \in \text{directpath}(i_{own}, |PK_{mem}|)$ :
   $PK_{upd}[i] \leftarrow PK'_{mem}[i]$ 
  ( $C_{cp}, PK'_{upd}, SK'_{own}, k_{cs}$ )
   $\leftarrow$  dopath( $PK'_{mem}, i_{own}, k_{upd}, CP$ )
   $PK'_{mem} \stackrel{merge}{\leftarrow} PK'_{upd}$ 
  return ( $C_{cp}, PK'_{upd}, PK'_{mem}, SK'_{own}, k_{cs}$ )

```

Figure 8: Freshness conditions of the MLS Key Distribution.

The algorithms process and join update their key material consistently with update. I.e., the ciphertexts C_{cp} and the list of changes to the public-key information PK_{rem} , PK_{joi} and PK_{upd} suffice for process to derive the same values as update. Similarly, the ciphertexts C_{add} and the view of the entire new tree PK_{mem} suffice to allow the join procedure to derive the same values. See Figure 9 for their code.

V. ASSUMPTIONS

Game-Based Security Notions. We formulate security properties via indistinguishability between games, a *real* game which models the real behaviour of a cryptographic primitive, and an *ideal* game which models an ideally secure variant of the primitive. For example, security of a pseudorandom function is captured by demanding that the input-output behaviour of the (keyed) pseudorandom function is indistinguishable from the input-output behavior of a truly random function which draws a uniformly random output for each input.

Especially since Bellare and Rogaway promoted code-based game-hopping [5], it has been popular to formulate games in pseudocode, and we follow this approach here. I.e., a game G provides a set of *oracles* to the adversary \mathcal{A} , these oracles are specified by pseudocode and operate on a common state which is secret to the adversary. This hidden state might contain, e.g., a secret key. In this work, we denote by $\mathcal{A} \circ G$ that the adversary \mathcal{A} interacts with the oracles of game G . I.e., the adversary is the *main procedure* that makes queries to the oracles of G , which returns an answer to \mathcal{A} , and then, in the end, the adversary \mathcal{A} returns a bit indicating their guess whether the game is indeed real or ideal. The *advantage* $\text{Adv}_{G^0, G^1}(\mathcal{A})$ measures an adversary \mathcal{A} 's ability to distinguish between G^0 (real) and G^1 (ideal).

Definition V.1 (Advantage). *For an adversary \mathcal{A} and two games G^0, G^1 , the advantage $\text{Adv}_{G^0, G^1}(\mathcal{A})$ denotes the term*

$$\left| \Pr[1 = \mathcal{A} \circ G^0] - \Pr[1 = \mathcal{A} \circ G^1] \right|.$$

We formulate security statements by relating advantages. E.g., we upper bound an adversary \mathcal{A} 's advantage against a game pair for MLS key distribution security by an adversary

Figure 9: Procedures update, process, join, and wupdate.

\mathcal{B} 's advantage against a game pair for the base primitive security. Importantly, the runtime of \mathcal{A} and \mathcal{B} will be similar. Namely, we often write \mathcal{B} as $\mathcal{A} \circ \mathcal{R}$ where \mathcal{R} is referred to as a *reduction* which translates \mathcal{A} 's oracle queries to its own game (in this case, the game pair for MLS key distribution security) into oracle queries to the game for the base primitive.

State-Separating Proofs. The previous description lacks a natural way to compose games. However, if games are not described as a single block of pseudocode but rather sliced into individual packages of code, then the same package of code can be reused. This notion of code-reuse, together with a separate state for each package results in a natural notion of package composition to form games.

A frequent example of a natural, re-usable code package is the KEY package described in Figure 10. Cryptographic primitives often share key material (e.g., in the MLS key schedule, KDF produces outputs which become keys of KDF which produces outputs which become keys of XPD etc., see Figure 3). Now, if KDF computes keys and stores them in a KEY package using the SET oracle, then XPD can retrieve

<u>KEY^{b,λ}</u> <u>Package Parameters</u> <i>b</i> : idealization bit <i>λ</i> : key length <u>Package State</u> <i>K</i> [<i>h</i> ↦ <i>k</i>]: key table <i>H</i> [<i>h</i> ↦ <i>b</i>]: honesty table <u>SET(<i>h, hon, s</i>)</u> <u>REG(<i>hon, s</i>)</u> assert <i>s</i> = <i>λ</i> assert <i>s</i> = <i>λ</i> <i>i</i> ← index(KEY) <i>i</i> ← index(KEY) <i>h</i> ← reg (<i>K</i> , <i>i</i>) if <i>K</i> [<i>h</i>] ≠ ⊥: if <i>K</i> [<i>h</i>] ≠ ⊥: return <i>h</i> return <i>h</i> if <i>b</i> ∧ <i>hon</i> : if <i>hon</i> : <i>s</i> ← _s {0, 1} ^λ <i>s</i> ← _s {0, 1} ^λ UNQ(<i>h, hon, s</i>) UNQ(<i>h, hon, s</i>) <i>K</i> [<i>h</i>] ← <i>s</i> <i>K</i> [<i>h</i>] ← <i>s</i> <i>H</i> [<i>h</i>] ← <i>hon</i> <i>H</i> [<i>h</i>] ← <i>hon</i> return <i>h</i> return <i>h</i> <u>GET(<i>h</i>)</u> <u>CGET(<i>h</i>)</u> assert <i>K</i> [<i>h</i>] ≠ ⊥ assert <i>H</i> [<i>h</i>] = 0 return <i>K</i> [<i>h</i>] return <i>K</i> [<i>h</i>] <u>HON(<i>h</i>)</u> assert <i>H</i> [<i>h</i>] ≠ ⊥ return <i>H</i> [<i>h</i>]	<u>LOG^b</u> <u>Package Parameters</u> <i>b</i> : idealization bit <u>Package State</u> <i>L</i> : log table <u>UNQ(<i>h, hon, s</i>)</u> if <i>L</i> [<i>h</i>] ≠ ⊥ : (<i>h'</i> , <i>hon'</i> , <i>s'</i>) ← <i>L</i> [<i>h</i>] assert <i>s'</i> = <i>s</i> if <i>hon'</i> = <i>hon</i> : return <i>h'</i> for (<i>h'</i> , <i>hon'</i> , <i>s'</i>) ∈ <i>L</i> with <i>s</i> = <i>s'</i> : <i>r</i> ← (<i>level</i> (<i>h</i>) = <i>set</i>) <i>r'</i> ← (<i>level</i> (<i>h'</i>) = <i>set</i>) if <i>r</i> ≠ <i>r'</i> ∧ <i>M</i> [<i>s</i>] = ⊥ ∧ <i>hon</i> = <i>hon'</i> = 0: <i>M</i> [<i>s</i>] ← 1 <i>L</i> [<i>h</i>] ← (<i>h'</i> , <i>hon</i> , <i>s</i>) return <i>h'</i> if <i>hon</i> = <i>hon'</i> = 0 ∧ <i>r</i> = <i>r'</i> : abort if <i>b</i> ∧ <i>r</i> = <i>r'</i> = 1: abort <i>Log</i> [<i>h</i>] ← (<i>h, hon, s</i>) return <i>h</i>
--	--

Figure 10: Definition of the KEY and LOG package, where **reg**(\cdot) is an injective handle constructor, creating a tuple from the inputs.

them via the GET oracle.

We now explain how to use KEY packages to model security. Figure 12 and Figure 13 describe the security games for KDF, XPD and HPKE, decomposed into several packages. For KDF and XPD, the upper KEY models the input key, the lower KEY package models the output keys. KDF and XPD are stateless packages, (see Figure 12d) which retrieve an input key from the upper KEY package, perform a computation, and store the result in the lower KEY package. The lower KEY package models security, namely pseudorandomness of the output keys.

Recall that our security notion aims to capture that output keys are indistinguishable from uniformly random values of the same length. Thus, if $b = 0$ in KEY^{b,λ}, the concrete key values are stored, and if $b = 1$ in KEY^{b,λ}, then uniformly random keys of the same length are drawn. Finally, the adversary can retrieve the output keys via the GET oracle.

Note, that we want to model multi-instance security, where some of the keys might be known to the adversary. We track the adversary's knowledge of a key by marking the key as *dishonest* (sometimes also called corrupt). In that case we will say that the honesty value *hon* is 0. In turn, for honest keys, *hon* is 1. In the case that *hon* = 0, the concrete keys (i.e. the real result of the key computations) are stored regardless of whether $b = 1$ or $b = 0$. Concretely, we model multi-instance security with static corruptions on the input keys and corrupt-

<u>PKEY^{b,pke}</u> <u>Package Parameters</u> <i>b</i> : idealization bit <i>pke</i> : pub. key enc. sch. <u>Package State</u> <i>K</i> [<i>h</i> ↦ <i>k</i>]: key table <i>H</i> [<i>h</i> ↦ <i>b</i>]: honesty table <i>Q</i> [<i>h</i> ↦ <i>h</i>]: handle table <u>GET(<i>h</i>)</u> assert <i>K</i> [<i>h</i>] ≠ ⊥ return <i>K</i> [<i>h</i>]	<u>SET(<i>h, hon, pk, sk</i>)</u> assert <i>pke.valid</i> (<i>pk, sk</i>) if <i>Q</i> [<i>h</i>] ≠ ⊥: return <i>Q</i> [<i>h</i>] if <i>b</i> ∧ <i>hon</i> : (<i>pk, sk</i>) ← _s <i>pke.kgen</i> () <i>h'</i> ← (<i>h, pk</i>) UNQ(<i>h', hon, sk</i>) <i>K</i> [<i>h'</i>] ← <i>sk</i> <i>H</i> [<i>h'</i>] ← <i>hon</i> return <i>h'</i> <u>CGET(<i>h</i>)</u> assert <i>H</i> [<i>h</i>] = 0 return <i>K</i> [<i>h</i>]	<u>REG(<i>hon, pk, sk</i>)</u> assert <i>pke.valid</i> (<i>pk, sk</i>) if <i>hon</i> : (<i>pk, sk</i>) ← _s <i>pke.kgen</i> () <i>h'</i> ← reg ((<i>K</i> , <i>pk</i>)) UNQ(<i>h', hon, sk</i>) <i>K</i> [<i>h'</i>] ← <i>sk</i> <i>H</i> [<i>h'</i>] ← <i>hon</i> return <i>h'</i> <u>HON(<i>h</i>)</u> assert <i>H</i> [<i>h</i>] ≠ ⊥ return <i>H</i> [<i>h</i>]
--	--	---

Figure 11: Definition of the PKEY package. **reg**(\cdot) is an injective handle constructor, creating a tuple from the inputs.

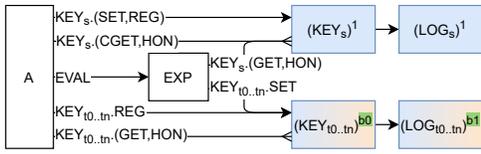
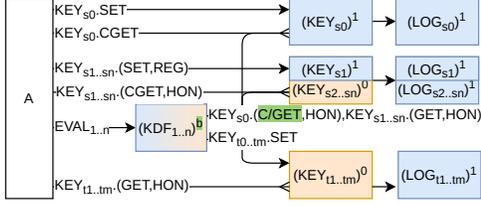
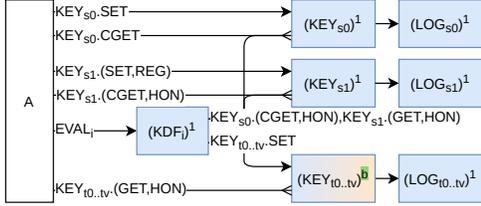
upon-derivation (CUD) on the output keys. If all input keys are corrupt, then the output value is marked as corrupt as well. Moreover, if the adversary uses *hon* = 0 as a parameter in its oracle query to EVAL, then the key will be marked as corrupt, too. Else, keys are marked as honest. Note, that in the upper KEY package, honest keys are sampled uniformly at random and secret from the adversary, as the adversary cannot access them via GET, because the GET query cannot be asked by the adversary—the arrows specify which package may call which oracles of which package. Importantly, this means that the graphs are part of the formal game definitions.

Handles. If an adversary wants to interact with a specific key, e.g., to retrieve it via a GET or have it be the subject of another oracle query such as EVAL, they use the *handle* of that key. We construct handles in such a way that each handle injectively maps to a key. When a key is initially randomly generated via the REG oracle, its handle **reg**(\cdot) is composed of the index of the KEY (denoted by index(KEY)) and a counter, implemented as |*K*| where *K* is the table of keys already stored in the KEY package. We denote by **reg** the corresponding an injective handle constructor, see Figure 10.

If a key is derived from one or more other keys, the handle of the output key is constructed from the handles of the input key(s), as well as any other inputs to the key derivation. See, e.g., the description of the EVAL oracles in Figure 12d. This guarantees that the handle-to-key mapping is *computationally injective*, i.e., will not have collisions unless we found a collision in the key derivation function used to derive the key.

Finally, let us elaborate on the role of the KEY package in HPKE. Here, KEY models a message that shall remain secret. If the HPKE is idealized, the message is instead replaced with an all-zero string before encryption (Figure 13). Note that pairs of secret keys and public keys are stored in the PKEY package, see Figure 11 and Figure 13.

Logging. We now turn to how we model collision resistance and key uniqueness. We call two handles *h* and *h'* such that they correspond to the same key in the *K* table in KEY a *key collision*. For honest, uniformly random keys, such collisions

(a) Game $\text{GEXP}_f^{b_0 b_1}$.(b) Game GKDFL_g^b .(c) Game GKDFR_g^b .EXP

Package Parameters

 f : a PRF \overline{lbl} : labels s : input name \overline{t} : output namesKDF_i^b

Package Parameters

 g : a KDF s_0 : left name \overline{lbl} : labels s_1 : right name b : idealization bit i : epoch indexPackage State*no state*Package State*no state*EVAL($h_s, \overline{ctx}, \overline{hon}$) $s \leftarrow \text{KEY}_s.\text{GET}(h_s)$ $\overline{hon}' \stackrel{\text{vec}}{\leftarrow} \text{KEY}_s.\text{HON}(h_s) \wedge \overline{hon}$ $\overline{t} \stackrel{\text{vec}}{\leftarrow} f(s, (\overline{lbl}, \overline{ctx}))$ $\overline{h}_t \stackrel{\text{vec}}{\leftarrow} (h_s, (\overline{lbl}, \overline{ctx}), \overline{t})$ $\overline{h}_t' \stackrel{\text{vec}}{\leftarrow} \text{KEY}_{\overline{t}}.\text{SET}(\overline{h}_t, \overline{hon}', \overline{t})$ **return** \overline{h}_t EVAL($h_{s_0}, h_{s_1}, \overline{ctx}, \overline{hon}$) $\text{hon}_0 \leftarrow \text{KEY}_{s_0}.\text{HON}(h_{s_0})$ $\text{hon}_1 \leftarrow \text{KEY}_{s_1}.\text{HON}(h_{s_1})$ **if** $b \wedge \text{hon}_0$: $\overline{t} \stackrel{\text{vec}}{\leftarrow} \{0, 1\}^{\lambda'}$ **else** $s_0 \leftarrow \text{KEY}_{s_0}.\text{GET}(h_{s_0})$: $s_1 \leftarrow \text{KEY}_{s_1}.\text{GET}(h_{s_1})$ $\overline{ctx}' \stackrel{\text{vec}}{\leftarrow} (\overline{ctx}, i)$ $\overline{t} \stackrel{\text{vec}}{\leftarrow} g(s_0, s_1, (\overline{lbl}, \overline{ctx}'))$ $\overline{hon}' \stackrel{\text{vec}}{\leftarrow} (\text{hon}_0 \vee \text{hon}_1) \wedge \overline{hon}$ $\overline{h}_t \stackrel{\text{vec}}{\leftarrow} (h_{s_0}, h_{s_1}, (\overline{lbl}, \overline{ctx}'), \overline{t})$ $\overline{h}_t' \stackrel{\text{vec}}{\leftarrow} \text{KEY}_{\overline{t}}.\text{SET}(\overline{h}_t, \overline{hon}', \overline{t})$ **return** \overline{h}_t (d) Code of packages EXP and KDF_i^b .Figure 12: Games $\text{GEXP}_f^{b_0 b_1}$, GKDFL_g^b and GKDFR_g^b .

are unlikely, but for keys registered by the adversary or known to the adversary, such collisions can trivially occur. We thus 1) disallow the adversary from registering the same dishonest key value k twice under two different handles (since identical input keys would lead to identical output keys), 2) remove collisions between registered and set dishonest keys (once) and 3) prove that then derived keys do not have collisions. We model both properties using a LOG package which keeps track of the keys and handles used so far and sends a special abort message when a collision occurs.

This modular style of code-writing has recently been developed by Brzuska, Delignat-Lavaud, Fournet, Kohbrok and Kohlweiss (BDFKK [9]) and been coined *state-separating proofs* (SSP) since the state separation of the games enables proofs which rely on code-separation and state-separation. We now state all our assumptions formally using the composed games style and return to proofs in Section VIII.

Expand. We consider Krawczyk's XPD function as a function $\text{xpd} : (\{0, 1\}^\lambda \times \{0, 1\}^*) \mapsto \{0, 1\}^{\lambda'}$, where $\lambda' = \lambda'(\lambda)$ is hardcoded (rather than provided as an explicit input). The second input is a pair $(\overline{lbl}, \overline{ctx})$, where in some cases, \overline{ctx} might be empty. We parametrize our security game for XPD with a list of labels \overline{lbl} and measure the security of the XPD function by the advantage which adversaries \mathcal{A} have in distinguishing the real game $\text{GEXP}_f^{0,0}$ and the ideal game $\text{GEXP}_f^{1,1}$, defined in Figure 12a, for $f = \text{xpd}$. For an adversary \mathcal{A} , we define the advantage function $\text{Adv}_{\text{EXP}}^f(\mathcal{A}) :=$

$$\left| \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GEXP}_f^{0,0} \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GEXP}_f^{1,1} \right] \right|.$$

In asymptotic terminology, if the advantage $\text{Adv}_{\text{EXP}}^{\text{xpd}}(\mathcal{A})$ is negligible for all PPT adversaries, then xpd is a secure pseudorandom function. $\text{Adv}_{\text{EXP}}^f(\mathcal{A})$ incorporates both, pseudorandomness *and* collision-resistance into a single assumption, since the bit in the LOG also changes from 0 to 1. While this is convenient in proofs, in some cases, we need pseudorandomness only, since collision resistance has already been taken care of. For an adversary \mathcal{A} , we write $\text{Adv}_{\text{PREXP}}^f(\mathcal{A}) :=$

$$\left| \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GEXP}_f^{1,1} \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GEXP}_f^{0,1} \right] \right|.$$

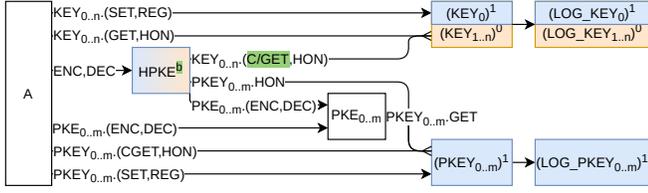
KDF. We model Krawczyk's HKDF as a function $\text{kdf} : (\{0, 1\}^\lambda \times \{0, 1\}^\lambda \times \{0, 1\}^*) \mapsto \{0, 1\}^{\lambda'}$ for some hard-coded $\lambda' = \lambda'(\lambda)$ such that a triple $(k, \overline{lbl}, \overline{ctx})$ is mapped to a string of length λ' . We assume that kdf is a *dual* KDF. I.e., if either of the inputs is random and secret from the adversary, then kdf produces a pseudorandom output. Idealization based on the *left* input is modeled by a bit in the code of KDF. For an adversary \mathcal{A} , we define the advantage $\text{Adv}_{\text{KDFL}}^g(\mathcal{A}) :=$

$$\left| \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GKDFL}_g^0 \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GKDFL}_g^1 \right] \right|,$$

where GKDFL_g^b is defined in Figure 12b. Note that here, right inputs might come from one of many packages. In turn, idealization based on the *right* input is modeled by a bit in the lower KEY and for an adversary \mathcal{A} , is captured by the advantage $\text{Adv}_{\text{KDFR}}^g(\mathcal{A}) :=$

$$\left| \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GKDFR}_g^0 \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GKDFR}_g^1 \right] \right|,$$

where GKDFR_g^b is defined in Figure 12c.

(a) Game $\text{GHPKE}_{\text{pke}}^b$.

<u>PKE</u>	<u>HPKE</u>	
<u>Package Parameters</u>	<u>Package Parameters</u>	
$\text{pke} : \text{pub. key enc. sch.}$	$b : \text{idealisation bit}$	
	$\lambda : \text{encrypted key length}$	
<u>Package State</u>	<u>Package State</u>	
$H[c \mapsto b] : \text{hon. ciph.}$	$Q[c \mapsto h] : \text{handle table}$	
<u>ENC(h_{kp}, h_k)</u>	<u>ENC(h_{kp}, h_k)</u>	<u>DEC(h_{kp}, c)</u>
$pk \leftarrow \text{parse}(h_{kp})$	$j \leftarrow \text{parse}(h_{kp})$	$j \leftarrow \text{parse}(h_{kp})$
$c \leftarrow s \cdot \text{pke.enc}(pk, m)$	$i \leftarrow \text{parse}(h_k)$	if $Q[c] \neq \perp$:
$H[c] \leftarrow 1$	assert $\text{PKEY}_j.\text{HON}(h_{kp})$	return $Q[c]$
return c	$\vee \neg \text{KEY}_i.\text{HON}(h_k)$	$m \leftarrow \text{DEC}(h_{kp}, c)$
	if $b \wedge \text{PKEY}_j.\text{HON}(h_{pk})$:	return m
<u>DEC(h_{kp}, c)</u>	$m \leftarrow 0^\lambda$	
$j \leftarrow \text{parse}(h_{kp})$	else if b :	
if $H[c] = 1$:	$m \leftarrow \text{KEY}_j.\text{CGET}(h_k)$	
return \perp	else :	
$sk \leftarrow \text{PKEY}_j.\text{GET}(h_{kp})$	$m \leftarrow \text{KEY}_i.\text{GET}(h_k)$	
$m \leftarrow \text{pke.dec}(sk, c)$	$c \leftarrow s \cdot \text{PKEY}_j.\text{ENC}(h_{kp}, m)$	
return m	$Q[c] \leftarrow h_k$	
	return c	

(b) Packages PKE and HPKE

Figure 13: Game $\text{GHPKE}_{\text{pke}}^b$.

Public-Key Encryption. Lastly, we assume the existence of a public-key encryption (PKE) scheme pke , which consists of three algorithms pke.kgen , pke.enc and pke.dec with standard correctness and confidentiality properties, namely indistinguishability under chosen ciphertext attacks (IND-CCA2). IND-CCA2 captures that encryptions of (adversarially chosen) messages m are computationally indistinguishable from encryptions of $0^{|m|}$, even when the adversary can use a decryption oracle (except on the challenge ciphertexts). We use the PKE scheme exclusively to model Hybrid Public Key Encryption (HPKE), i.e. we only encrypt symmetric keys with the PKE scheme. The length of the symmetric keys is λ . For an adversary \mathcal{A} , we define the advantage as $\text{Adv}_{\text{HPKE}}^{\text{pke}}(\mathcal{A}) :=$

$$\left| \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GHPKE}_{\text{pke}}^0 \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GHPKE}_{\text{pke}}^1 \right] \right|,$$

where the game $\text{GHPKE}_{\text{pke}}^b$ is defined in Figure 13 and n and m refers to an upper bound on the number of secret-key public-key pairs, and symmetric-keys to be encrypted, respectively. Since we will need to run pke.kgen with explicit randomness, we introduce a *derive key pair function* (DKP), which maps the random coins to the output of pke.kgen .

Collision-Resistance. In addition to assuming that xpd and kdf produce pseudorandom outputs, we also assume their collision-resistance. In particular, we encode collision-

resistance by assuming that xpd and kdf pass on the uniqueness of their input keys to their output keys, i.e., if the bits in the logs of their input keys have bit 1, then the bits in the log packages of their output keys can also be flipped from 0 to 1. For an adversary \mathcal{A} , we denote these advantages by $\text{Adv}_{\text{CRXPD}}^{\text{xpd}}(\mathcal{A})$ and $\text{Adv}_{\text{CRKDF}}^{\text{xpd}}(\mathcal{A})$, respectively, and we refer by $\text{Adv}_{\text{CRDKP}}^{\text{xpd}}(\mathcal{A})$ to the collision-resistance advantage of the derive key pair function dkp .

VI. SECURITY MODELS

We now introduce our security model for the MLS CGKD. We start with a high-level overview which considers the CGKD, defined as a composition of TreeKEM and the key schedule and then turn to modeling TreeKEM and the key schedule each on their own.

A. CGKD

A crucial component in the modeling are different keys, inputs and outputs as described in Figure 2. Based on the rationale outlined in Section V, storing keys in separate packages is useful for composition, and so, our model stores each type of key in a separate KEY package, i.e., there are two main packages wCGKD (modeling TreeKEM) and KS , each reading and writing keys from/into the different KEY and PKEY packages, see Figure 14a for the composed game GWCGK and Figure 14c for GKS . Figure 14b and 14c respectively show the security definition of TreeKEM and the Key Schedule on their own. Note that each KEY package has a different subscript and so have their oracles to ensure that queries are unique across the entire graph. We now map the inputs and outputs described in Figure 2 to the KEY packages.

The values (pk, sk) of the KEMs associated with different nodes in the tree are stored in PKEY_NKP packages, the k_{upd} variable, conceptually, corresponds to a path secret and thus is stored in the KEY_PS packages with other path secrets. Finally, values for the s_{ext} variable are stored in KEY_PSK packages, values for the s_{int} variable are stored in KEY_IS packages, and values of each of the k variables in $\overline{k}_{\text{tbl}}$ are stored in KEY_GR .

Our base keys are modeled as uniformly random, unique keys (for honest keys) or adversarially chosen (but still unique) keys, e.g., the KEY_PSK starts with an idealization bit $b = 1$, reflecting that keys are sampled uniformly at random and that the adversary cannot register the same dishonest key twice. We can then rely on the assumptions stated in Section V to prove that a game where the idealization bits of the other KEY packages are all 0 is computationally indistinguishable from a game where the idealization bits of the other KEY packages are all 1. This establishes that the MLS CGKD provides pseudorandom and unique keys.

Definition VI.1 (CGKD advantage). *For an expand function xpd , a KDF kdf , a PKE pke , for all polynomials d and t and all adversaries \mathcal{A} which generate trees of depth at most d and run groups for at most t epochs, we denote by $\text{Adv}_{\text{CGKD}}^{\text{xpd}, \text{kdf}, \text{pke}}(\mathcal{A})$ the advantage*

$$\left| \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GCGKD}_{\text{xpd}, \text{kdf}, \text{pke}}^{0, d, t} \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GCGKD}_{\text{xpd}, \text{kdf}, \text{pke}}^{1, d, t} \right] \right|.$$

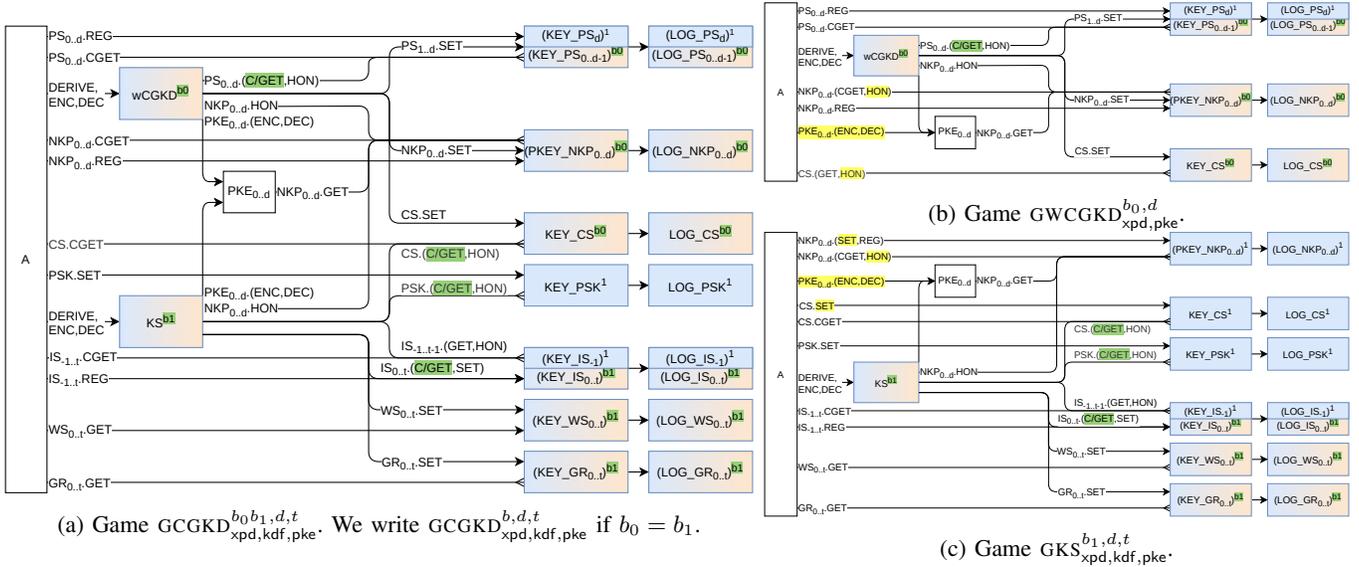


Figure 14: Composed (left) and individual (right) TreeKEM and Key Schedule games.

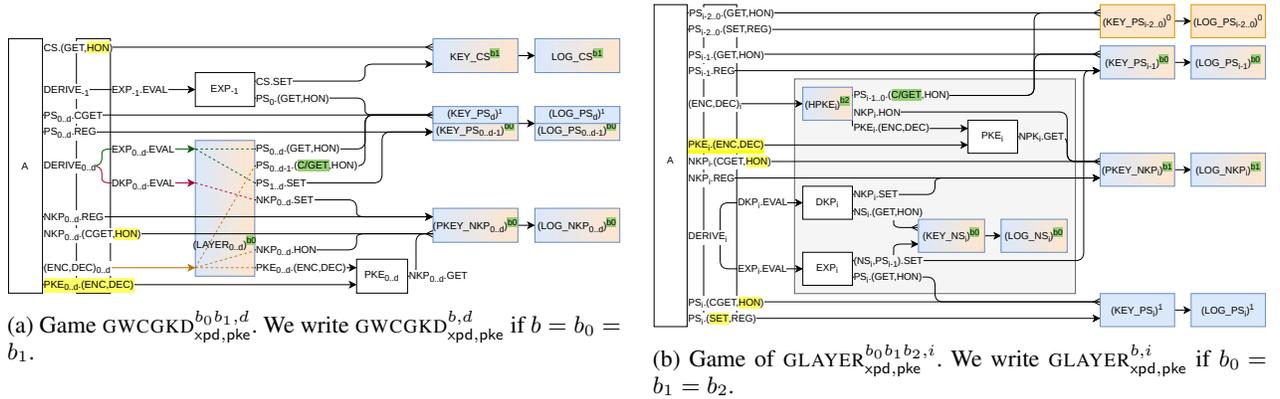


Figure 15: TreeKEM subgames.

B. TreeKEM model

In Figure 14b, the $wCGKD$ package represents the composition of packages shown in the center of Figure 15a ($LAYER_{0..d}$ and EXP). Note that the KEY , $PKEY$, and PKE packages correspond to those in Figure 14b. The $LAYER$ packages represent an at most d -layer tree structure where each layer is encoded as a separate composition of packages. Note that for each layer, the KEY_PS package, storing path secrets, has a REG query that allows nodes on that layer to function as a leaf node by generating a random (secret) value from scratch.

We define each $LAYER$ package as composition of packages, see Figure 15b. Again, the KEY , $PKEY_NKP$, and PKE packages correspond to those in Figure 14b. Figure 15b encodes both the dn function (Figure 4) through the EXP and DKP packages, and the encryption of path secrets to other members through the $HPKE$, PKE , $PKEY$, and KEY_PS packages. The GET access of $HPKE$ to all layers above it encodes that path secrets at various positions in the tree might be encrypted by $HPKE$.

This model gives significantly more power to the adversary than the MLS TreeKEM interface would (making our security

statements only stronger). The model allows the adversary to set up derivations and encryptions that encode arbitrary tree structures. The model merely enforces the concept of layers.

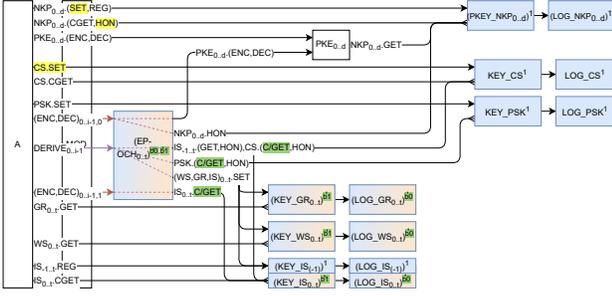
Definition VI.2 ($wCGKD$ advantage). *For an expand function xpd , a PKE pke , for all polynomials d and all adversaries \mathcal{A} which generate trees of depth at most d , we denote by $\text{Adv}_{wCGKD}^{xpd,pke}(\mathcal{A})$ the advantage*

$$\left| \Pr \left[1 \leftarrow \mathcal{A} \circ GWCGKD_{xpd,pke}^{0,d} \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ GWCGKD_{xpd,pke}^{1,d} \right] \right|.$$

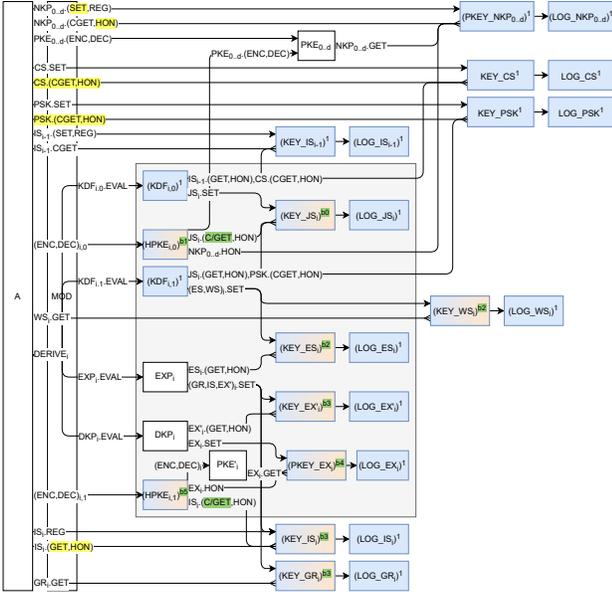
C. Key Schedule model

We define the KS package in Figure 14c by the composition of packages shown in the center of Figure 16a ($EPOCH_{0..t}$). The KEY , $PKEY$, and PKE packages correspond to those in Figure 14c. The $EPOCH$ packages represent a t -epoch group where each epoch is encoded as a separate composition of packages. Each $EPOCH$ package is the composition of packages shown in the center of Figure 16b. Recall that the KEY , $PKEY$, and PKE packages correspond to those in Figure 14c.

Figure 16b encodes the ks and $subks$ functions (Figure 4). The KDF and EXP packages correspond to kdf and xpd calls



(a) Game $\text{GGKS}_{xpd,kdf,pke}^{b'_0 b'_1, d, t}$. We write $\text{GGKS}_{xpd,kdf,pke}^{b, d, t}$ if $b'_0 = b'_1$.



(b) Game $\text{GEPOCH}_{xpd,kdf,pke}^{b_0 b_1 b_2 b_3 b_4 b_5, d, t, i}$. We write $\text{GEPOCH}_{xpd,kdf,pke}^{b, d, t, i}$ if $b_0 = b_1 = b_2 = b_3 = b_4 = b_5$. Note that KEY_{IS}_{-1} should technically have a REG query, but for an ideal KEY package REG and SET are functionally equivalent.

Figure 16: Key Schedule subgames.

in ks and subs, and the DKP package corresponds to the dkp call. The HPKE packages encode the encryption of the joiner secret to new members (top), and the encryption of init secrets by external committers (bottom). The REG query allows the adversary to create external init secrets. The values of external key pairs, (pk_{ext}, sk_{ext}) , are stored in the PKEY_EX package.

Definition VI.3 (KS advantage). *For an expand function xpd , a key derivation function kdf , a PKE pke , for all polynomials d, t and all adversaries \mathcal{A} which generate trees of depth at most d and run groups for at most t epochs, $\text{Adv}_{KS}^{xpd,kdf,pke}(\mathcal{A})$ denotes*

$$\left| \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GKS}_{xpd,kdf,pke}^{0,d,t} \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GKS}_{xpd,kdf,pke}^{1,d,t} \right] \right|.$$

VII. THEOREMS

This section relates the security of the MLS CGKD to the security of the underlying primitives XPD, KDF, and HPKE. We first make separate security claims for the wCGKD and the key schedule and then bound the security of their composition. Our security statements are concrete and constructive, i.e., they

transform an adversary against the CGKD into an adversary against the underlying primitives.

Theorem 1. (*wCGKD Security*) *For all polynomials d and all adversaries \mathcal{A} which generate trees of depth at most d , it holds that*

$$\text{Adv}_{wCGKD}^{xpd,pke}(\mathcal{A}) \leq \text{Adv}_{EXP}^{xpd}(\mathcal{A} \circ \mathcal{R}_{exp}) + \sum_{i=0}^{d-1} (\text{Adv}_{EXP}^{xpd}(\mathcal{A} \circ \mathcal{R}_i \circ \mathcal{R}_{exp,i}) + \text{Adv}_{HPKE}^{pke}(\mathcal{A} \circ \mathcal{R}_i \circ \mathcal{R}_{hpke,i}))$$

where $\text{GWCGKD}_{xpd,pke}^{b,d}$ is provided in Figure 14b.

Theorem 2. (*Key Schedule Security*) *For all polynomials d and t and all adversaries \mathcal{A} which generate trees of depth at most d and run groups for at most t epochs, it holds that*

$$\begin{aligned} \text{Adv}_{KS}^{xpd,kdf,pke}(\mathcal{A}) &\leq \text{Adv}_{CR}^{xpd,kdf,dkp}(\mathcal{A} \circ \mathcal{R}_{cr}) \\ &+ \text{Adv}_{KDFL}^{kdf}(\mathcal{A} \circ \mathcal{R}_{int}) + \text{Adv}_{KDFL}^{kdf}(\mathcal{A} \circ \mathcal{R}_{joi}) \\ &+ \sum_{i=0}^t \text{Adv}_{KDFR}^{kdf}(\mathcal{A} \circ \mathcal{R}_{hyb,i} \circ \mathcal{R}_{a,i}) + \text{Adv}_{HPKE}^{pke}(\mathcal{A} \circ \mathcal{R}_{hyb,i} \circ \mathcal{R}_{b,i}) \\ &+ \text{Adv}_{KDFR}^{kdf}(\mathcal{A} \circ \mathcal{R}_{hyb,i} \circ \mathcal{R}_{c,i}) + \text{Adv}_{EXP}^{xpd}(\mathcal{A} \circ \mathcal{R}_{hyb,i} \circ \mathcal{R}_{d,i}) \\ &+ \text{Adv}_{HPKE}^{pke}(\mathcal{A} \circ \mathcal{R}_{hyb,i} \circ \mathcal{R}_{e,i}), \end{aligned}$$

where $\text{GKS}_{xpd,kdf,pke}^{b,d,t}$ is defined in Figure 14c.

Theorem 3. (*CGKD*) *For all polynomials d and t and all adversaries \mathcal{A} which generate trees of depth at most d and run groups for at most t epochs, it holds that*

$$\text{Adv}_{CGKD}^{xpd,kdf,pke}(\mathcal{A}) \leq \text{Adv}_{wCGKD}^{xpd,pke}(\mathcal{A} \circ \mathcal{R}_p) + \text{Adv}_{KS}^{xpd,kdf,pke}(\mathcal{A} \circ \mathcal{R}_q),$$

where $\text{GCGKD}_{xpd,kdf,pke}^{b,d,t}$ is given in Figure 14a.

We prove Theorem 1 in Section VIII, provide the proof of Theorem 3 in Appendix A and the proof of Theorem 2 in Appendix B, C and D.

VIII. PROOF OF THEOREM 1 (wCGKD SECURITY)

The proof of Theorem 1 consists of Lemma 1 and a hybrid argument. We start with the former which shows that a TreeKEM layer with an ideal path secret can be completely idealized such that the path secret of its parent is ideal as well as the encryptions it performs.

Lemma 1. (*Layer Security*) *Let \mathcal{B} be an adversary and $\text{GLAYER}_{xpd,pke}^{b,i}$ the indistinguishability game defined by Figure 15b. Then it holds that*

$$\text{Adv}_{LAYER}^{xpd,pke,i}(\mathcal{B}) \leq \text{Adv}_{EXP}^{xpd}(\mathcal{B} \circ \mathcal{R}_{exp,i}) + \text{Adv}_{HPKE}^{pke}(\mathcal{B} \circ \mathcal{R}_{hpke,i}),$$

where $\mathcal{R}_{exp,i}$ is defined as the grey area in Figure 17a and $\mathcal{R}_{hpke,i}$ is defined as the grey area in Figure 17c.

The proof of Lemma 1 consists of three steps, each of which modifies one or two bits in the overall game. We denote

$$\text{GLAYER}_{xpd,pke}^{b,i} := \text{GLAYER}_{xpd,pke}^{bbb,i}$$

i.e., $\text{GLAYER}_{xpd,pke}^{b_0 b_1 b_2, i}$ with $b = b_0 = b_1 = b_2$.

In the first proof step, we apply the expand assumption for $\overline{bl} = \{\text{path}, \text{node}\}$ to idealize the node secret of layer i and the path secret of layer $p(i)$. We obtain the following bound:

$\frac{\text{wprocess}(PK_{mem}, PK_{rem}, PK_{joi}, PK_{upd}, SK_{own}, i_{nod}, i_{upd}, c)}{k_{cs} \leftarrow 0$ $PK'_{mem} \leftarrow PK_{mem}$ $PK'_{mem} \stackrel{rem}{\leftarrow} PK_{rem}$ $PK'_{mem} \stackrel{merge}{\leftarrow} PK_{joi}$ $PK'_{mem} \stackrel{yc}{\leftarrow} PK_{upd}$ $SK'_{own} \leftarrow SK_{own}$ $sk_{nod} \leftarrow SK_{own}[i_{nod}]$ $ps \leftarrow \text{pke.dec}(sk_{nod}, c)$ $\text{assert } ps \neq \perp$ $i_{cur} \leftarrow \text{lcp}(i_{own}, i_{upd}, PK'_{mem})$ $(_, _, SK''_{own}, k_{cs}) \leftarrow \text{dopath}(PK'_{mem}, i_{cur}, ps, \perp)$ $SK'_{own} \stackrel{yc}{\leftarrow} SK''_{own}$ $\text{return } (PK'_{mem}, SK'_{own}, k_{cs})$	$\frac{\text{wjoin}(PK_{mem}, SK_{own}, i_{own}, i_{upd}, c)}{k_{cs} \leftarrow 0}$ $sk_{own} \leftarrow SK_{own}[i_{own}]$ $ps \leftarrow \text{pke.dec}(sk_{own}, c)$ $\text{assert } ps \neq \perp$ $i_{cur} \leftarrow \text{lcp}(i_{own}, i_{upd}, PK_{mem})$ $(_, _, _, k_{cs}) \leftarrow \text{dopath}(PK_{mem}, i_{cur}, ps, \perp)$ $\text{return } (PK_{mem}, SK_{own}, k_{cs})$
--	--

Figure 19: Procedures wprocess and wjoin.

where \mathcal{R}_i is defined as the grey area in Figure 18a. We apply this hybrid argument until KEY_PS_0 is ideal. At this point the entire TreeKEM tree is ideal, except for the commit secret derived at the root. Therefore, as visualized in Figure 18b, we apply the expand assumption and obtain

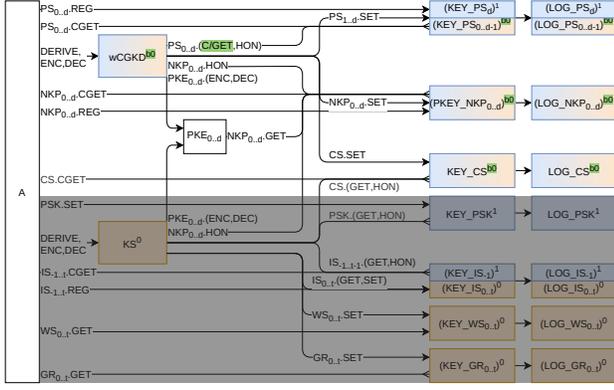
$$\left| \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GWCGKD}_{\text{upd,pke}}^{1\dots 10,d} \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GWCGKD}_{\text{upd,pke}}^{1\dots 11,d} \right] \right| \leq \text{Adv}_{\text{EXP}}^{\text{upd,pke}}(\mathcal{A} \circ \mathcal{R}_{\text{exp}}),$$

where \mathcal{R}_{exp} is defined as the grey area in Figure 18b. Using $\mathcal{B} = \mathcal{A} \circ \mathcal{R}_i$, we derive the desired bound for Theorem 1:

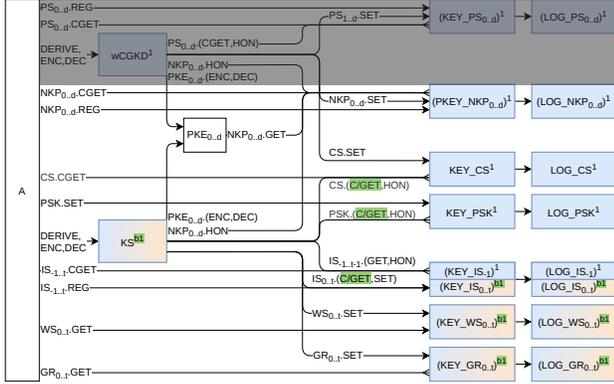
$$\begin{aligned} \text{Adv}_{\text{WCGKD}}^{\text{upd,pke}}(\mathcal{A}) &\leq \text{Adv}_{\text{EXP}}^{\text{upd,pke}}(\mathcal{A} \circ \mathcal{R}_{\text{exp}}) + \sum_{i=0}^{d-1} \text{Adv}_{\text{LAYER}}^{\text{upd,pke},i}(\mathcal{A} \circ \mathcal{R}_i) \\ &= \text{Adv}_{\text{EXP}}^{\text{upd,pke}}(\mathcal{A} \circ \mathcal{R}_{\text{exp}}) + \sum_{i=0}^{d-1} (\text{Adv}_{\text{EXP}}^{\text{upd,pke}}(\mathcal{A} \circ \mathcal{R}_i \circ \mathcal{R}_{a,i}) \\ &\quad + \text{Adv}_{\text{HPKE}}^{\text{pke}}(\mathcal{A} \circ \mathcal{R}_i \circ \mathcal{R}_{b,i})) \end{aligned}$$

REFERENCES

- [1] Joël Alwen, Margarita Capretto, Miguel Cueto, Chethan Kamath, Karen Klein, Iliia Markov, Guillermo Pascual-Perez, Krzysztof Pietrzak, Michael Walter, and Michelle Yeo. Keep the dirt: Tainted TreeKEM, adaptively and actively secure continuous group key agreement. Cryptology ePrint Archive, Report 2019/1489, 2019. <https://eprint.iacr.org/2019/1489>.
- [2] Joël Alwen, Sandro Coretti, Yevgeniy Dodis, and Yiannis Tselekounis. Security analysis and improvements for the IETF MLS standard for group messaging. Cryptology ePrint Archive, Report 2019/1189, 2019. <https://eprint.iacr.org/2019/1189>.
- [3] Richard Barnes, Benjamin Beurdouche, Jon Millican, Katriel Cohn-Gordon Emad Omara, and Raphael Robert. Message layer security rfc draft 11. 2019. <https://datatracker.ietf.org/doc/draft-ietf-mls-protocol/>.
- [4] Mihir Bellare and Phillip Rogaway. Entity authentication and key distribution. In Annual international cryptography conference, pages 232–249. Springer, 1993.
- [5] Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In Serge Vaudenay, editor, EUROCRYPT 2006, volume 4004 of LNCSS, pages 409–426. St. Petersburg, Russia, May 28 – June 1, 2006. Springer, Heidelberg, Germany.
- [6] Karthikeyan Bhargavan, Richard Barnes, and Eric Rescorla. TreeKEM: Asynchronous Decentralized Key Management for Large Dynamic Groups A protocol proposal for Messaging Layer Security (MLS). PhD thesis, Inria Paris, 2018. <https://prosecco.gforge.inria.fr/personal/karthik/pubs/treekem.pdf>.
- [7] Nikita Borisov, Ian Goldberg, and Eric Brewer. Off-the-record communication, or, why not to use pgp. In Proceedings of the 2004 ACM Workshop on Privacy in the Electronic Society, WPES '04, page 77–84. New York, NY, USA, 2004. Association for Computing Machinery.
- [8] Chris Brzuska, Antoine Delignat-Lavaud, Cédric Fournet, Konrad Kohbrok, and Markulf Kohlweiss. Handshake security for the TLS 1.3 standard. Preprint.
- [9] Chris Brzuska, Antoine Delignat-Lavaud, Cédric Fournet, Konrad Kohbrok, and Markulf Kohlweiss. State separation for code-based game-playing proofs. In Thomas Peyrin and Steven Galbraith, editors, ASIACRYPT 2018, Part III, volume 11274 of LNCSS, pages 222–249. Brisbane, Queensland, Australia, December 2–6, 2018. Springer, Heidelberg, Germany.
- [10] Chris Brzuska and Jan Winkelmann. Nprfs and their application to message layer security. Preprint. <http://chrisbrzuska.de/2020-NPRF.html>.
- [11] Katriel Cohn-Gordon, Cas Cremers, and Luke Garratt. On post-compromise security. Cryptology ePrint Archive, Report 2016/221, 2016. <http://eprint.iacr.org/2016/221>.
- [12] Katriel Cohn-Gordon, Cas Cremers, Luke Garratt, Jon Millican, and Kevin Milner. On ends-to-ends encryption: Asynchronous group messaging with strong security guarantees. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, ACM CCS 2018, pages 1802–1819. Toronto, ON, Canada, October 15–19, 2018. ACM Press.
- [13] Katriel Cohn-Gordon, Cas J. F. Cremers, Benjamin Dowling, Luke Garratt, and Douglas Stebila. A formal security analysis of the signal messaging protocol. In 2017 IEEE European Symposium on Security and Privacy, EuroS&P 2017, Paris, France, April 26–28, 2017, pages 451–466. IEEE, 2017.
- [14] Eric Cornelissen. Pull request 453: Use the GroupContext to derive the joiner_secret. <https://github.com/mlswg/mls-protocol/pull/453>.
- [15] Cas Cremers, Britta Hale, and Konrad Kohbrok. Efficient post-compromise security beyond one group. Cryptology ePrint Archive, Report 2019/477, 2019. <https://eprint.iacr.org/2019/477>.
- [16] Cyprien Delpèch de Saint Guilhem, Marc Fischlin, and Bogdan Warinschi. Authentication in key-exchange: Definitions, relations and composition. 2019.
- [17] Benjamin Dowling, Paul Rösler, and Jörg Schwenk. Flexible authenticated and confidential channel establishment (fACCE): Analyzing the noise protocol framework. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, PKC 2020, Part I, volume 12110 of LNCSS, pages 341–373. Edinburgh, UK, May 4–7, 2020. Springer, Heidelberg, Germany.
- [18] Yongdae Kim, Adrian Perrig, and Gene Tsudik. Communication-efficient group key agreement. In Trusted Information: The New Decade Challenge, IFIP TC11 Sixteenth Annual Working Conference on Information Security (IFIP/Sec'01), June 11–13, 2001, Paris, France, volume 193 of IFIP Conference Proceedings, pages 229–244. Kluwer, 2001.
- [19] Hugo Krawczyk. SIGMA: The “SIGN-and-MAC” approach to authenticated Diffie-Hellman and its use in the IKE protocols. In CRYPTO 2003, volume 2729 of LNCSS, pages 400–425. Santa Barbara, CA, USA, August 17–21, 2003. Springer, Heidelberg, Germany.
- [20] Hugo Krawczyk. Cryptographic extraction and key derivation: The HKDF scheme. In Tal Rabin, editor, CRYPTO 2010, volume 6223 of LNCSS, pages 631–648. Santa Barbara, CA, USA, August 15–19, 2010. Springer, Heidelberg, Germany.
- [21] Hugo Krawczyk and Hoeteck Wee. The OPTLS protocol and TLS 1.3. Cryptology ePrint Archive, Report 2015/978, 2015. <http://eprint.iacr.org/2015/978>.
- [22] M. Marlinspike and T. Perrin. Signal specifications. Technical report, 2016. <https://signal.org/docs/>.
- [23] Alfred J. Menezes, Paul C. van Oorschot, and Scott A. Vanstone. Handbook of Applied Cryptography. CRC Press, Boca Raton, Florida, 1996.
- [24] Jesper Buus Nielsen. Separating random oracle proofs from complexity theoretic proofs: The non-committing encryption case. In Moti Yung, editor, CRYPTO 2002, volume 2442 of LNCSS, pages 111–126. Santa Barbara, CA, USA, August 18–22, 2002. Springer, Heidelberg, Germany.
- [25] Emad Omara, Benjamin Beurdouche, Eric Rescorla, Srinivas Inguva, Kwon Albert, and Alan Duric. The messaging layer security (mls) architecture rfc draft 04. 2019. <https://tools.ietf.org/html/draft-ietf-mls-architecture-04>.
- [26] D. G. Steer, L. Strawczynski, Whitfield Diffie, and Michael J. Wiener. A secure audio teleconference system. In Shafi Goldwasser, editor,



(a) First idealization step of composition theorem.



(b) Second idealization step of composition theorem.

Figure 20: Theorem 3.

CRYPTO '88, volume 403 of LNCS, pages 520–528, Santa Barbara, CA, USA, August 21–25, 1990. Springer, Heidelberg, Germany.

- [27] Matthew Weidner. Group messaging for secure asynchronous collaboration. PhD thesis, MPhil Dissertation, 2019. Advisors: A. Beresford and M. Kleppmann, 2019.
- [28] P. Zimmermann, A. Johnston, and J. Callas. Zrtp: Media path key agreement for unicast secure rtp. RFC 6189, RFC Editor, April 2011. <http://www.rfc-editor.org/rfc/rfc6189.txt>.

APPENDIX

A. Proof of Theorem 3 (CGKD security)

We prove Theorem 3 using Theorem 1 and Theorem 2. Recall (see Figure 14a) that, if $b = b_0 = b_1$, then

$$\text{GCGKD}_{\text{xpd,kdf,pke}}^{b_0 b_1, d, t} = \text{GCGKD}_{\text{xpd,kdf,pke}}^{bb, d, t} = \text{GCGKD}_{\text{xpd,kdf,pke}}^{b, d, t}$$

In the first step of this proof we modify b_0 from 0 to 1 and reduce to GWCGKD. For all PPT adversaries \mathcal{A} , we obtain

$$\left| \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GCGKD}_{\text{xpd,kdf,pke}}^{00, d, t} \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GCGKD}_{\text{xpd,kdf,pke}}^{10, d, t} \right] \right| \leq \text{Adv}_{\text{wCGKD}}^{\text{xpd,pke}}(\mathcal{A} \circ \mathcal{R}_p),$$

where \mathcal{R}_p is defined as the grey area in Figure 20a. Next, we modify b_1 from 0 to 1 and reduce to GKS. We obtain

$$\left| \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GCGKD}_{\text{xpd,kdf,pke}}^{10, d, t} \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GCGKD}_{\text{xpd,kdf,pke}}^{11, d, t} \right] \right| \leq \text{Adv}_{\text{KS}}^{\text{xpd,kdf,pke}}(\mathcal{A} \circ \mathcal{R}_q),$$

where \mathcal{R}_q is defined as the grey area in Figure 20b. Hence, we obtain the desired bound for Theorem 3:

$$\text{Adv}_{\text{CGKD}}^{\text{xpd,kdf,pke}}(\mathcal{A}) \leq \text{Adv}_{\text{wCGKD}}^{\text{xpd,pke}}(\mathcal{A} \circ \mathcal{R}_p) + \text{Adv}_{\text{KS}}^{\text{xpd,kdf,pke}}(\mathcal{A} \circ \mathcal{R}_q).$$

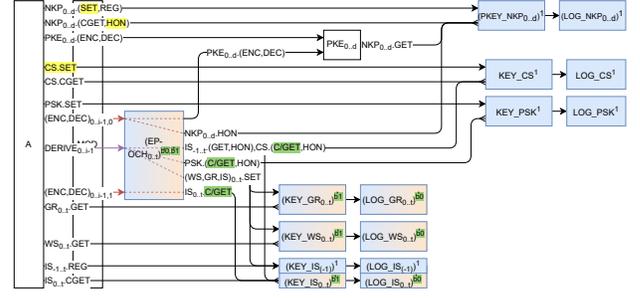
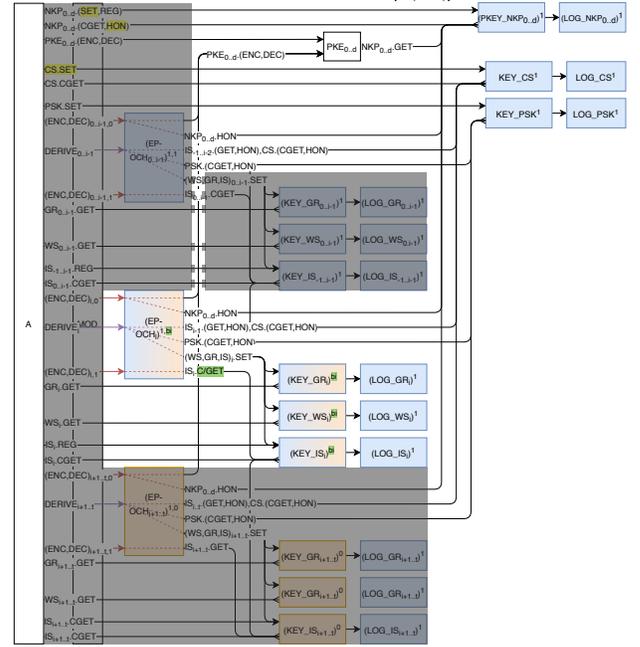
(a) Definition of $\text{GGKS}_{\text{xpd,kdf,pke}}^{b'_0 b'_1, d, t}$.(b) Reducing a hybrid step in the key schedule proof to the security of a single epoch. Reduction $\mathcal{R}_{\text{hyb},i}$ is marked in grey.

Figure 21: Proof of Theorem 2.

B. Proof of Theorem 2 (KS security)

This appendix proves Theorem 2. We first idealize collision-resistance for all primitives (unique inputs translate into unique outputs), then idealize the two KDFs of the key schedule based on the commit secret and the external secret globally for all epochs (since the commit secret and the external secret do not have an epoch) and then finally apply a hybrid argument over the epochs which is captured by Lemma 3.

KEY_IS₀ represents the initial init secret and is ideal from the start, i.e., it samples honest values independently at random and allows the adversary to register corrupt values. Similarly, the pre-shared keys (KEY_PSK) are ideal from the start as they are exchanged out-of-bound. Lastly, note that the public-private key-pairs (PKEY_NKP) and commit secrets (KEY_CS) are ideal from the start, which follows from Theorem 3.

Figure 21a depicts $\text{GGKS}_{\text{xpd,kdf,pke}}^{b'_0 b'_1, d, t}$ for trees up to depth d and groups that run for up to t epochs. For $b = b'_0 = b'_1$, we define

$$\text{GKS}_{\text{xpd,kdf,pke}}^{b, d, t} := \text{GGKS}_{\text{xpd,kdf,pke}}^{b'_0 b'_1, d, t}$$

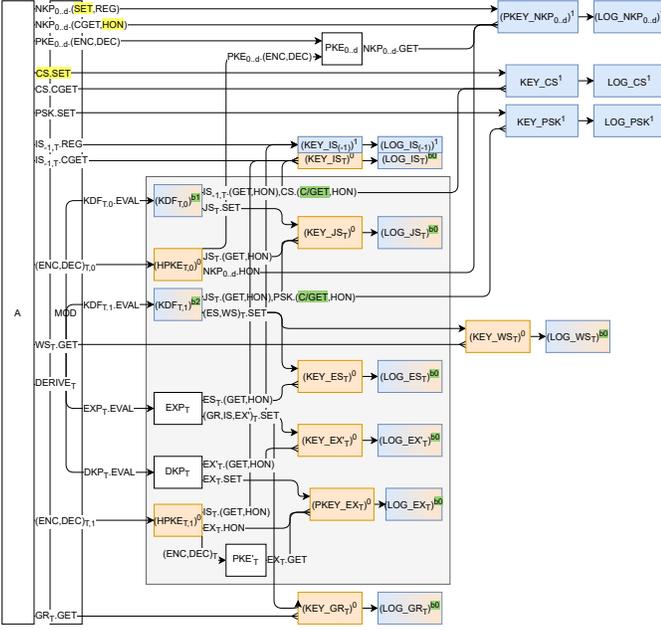


Figure 22: $\text{GPGKS}_{\text{xpd,kdf,pke}}^{b_0 b_1 b_2 0, d, t}$.

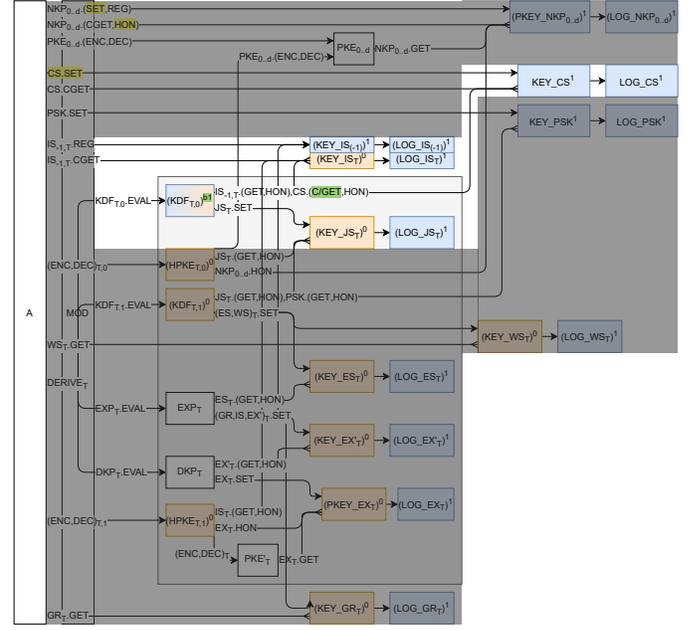


Figure 23: Reducing to KDFL security based on ideal CS. Reduction \mathcal{R}_{CS} is marked in grey.

Here, b'_0 corresponds to *global* proof steps (collision-resistance and KDFs), and b'_1 corresponds to *epoch-wise* proof steps. Note that the second G in GGKS stands for *global*.

We now state Lemma 2 and Lemma 3, then show that they imply Theorem 2 and then prove Lemma 2 and Lemma 3.

Lemma 2. *For all polynomials d and t and all adversaries \mathcal{A} which generate trees of depth at most d and run groups for at most t epochs, it holds that $\text{Adv}_{\text{GGKS}}^{\text{xpd,kdf,pke}, b_0}(\mathcal{A}) :=$*

$$\left| \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GGKS}_{\text{xpd,kdf,pke}}^{00, d, t} \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GGKS}_{\text{xpd,kdf,pke}}^{10, d, t} \right] \right| \leq \text{Adv}_{\text{CR}}^{\text{xpd,kdf,dkp}}(\mathcal{A} \circ \mathcal{R}_{\text{cr}}) + \text{Adv}_{\text{KDFL}}^{\text{kdf}}(\mathcal{A} \circ \mathcal{R}_{\text{cs}}) + \text{Adv}_{\text{KDFL}}^{\text{kdf}}(\mathcal{A} \circ \mathcal{R}_{\text{psk}})$$

Recall that $\text{GGKS}_{\text{xpd,kdf,pke}}^{b'_0 b'_1, d, t}$ is defined in Figure 21a. Analogously to $\text{Adv}_{\text{GGKS}}^{\text{xpd,kdf,pke}, b_0}(\mathcal{A})$, we use the notation

$$\text{Adv}_{\text{GGKS}}^{\text{xpd,kdf,pke}, 1b}(\mathcal{A}) := \left| \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GGKS}_{\text{xpd,kdf,pke}}^{10, d, t} \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GGKS}_{\text{xpd,kdf,pke}}^{11, d, t} \right] \right|.$$

In order to state Lemma 3, we need to split b'_1 into t different bits b_0 to b_t . We write $\text{GGKS}_{\text{xpd,kdf,pke}}^{0b_0 \dots b_t, d, t}$ and note that

$$\text{GGKS}_{\text{xpd,kdf,pke}}^{0b_0 \dots b_t, d, t} = \text{GGKS}_{\text{xpd,kdf,pke}}^{0b, d, t}$$

for $b'_1 = b_0 = \dots = b_t$.

Lemma 3. (Epoch Security) *Let \mathcal{B} be an adversary and $\text{GEOCH}_{\text{xpd,kdf,pke}}^{b, d, t, i}$ the indistinguishability game defined by Figure 25. Then $\forall d, t, i \in \mathbb{N} : 0 \leq i \leq t$*

$$\text{Adv}_{\text{EPOCH}}^{\text{xpd,kdf,pke}, i}(\mathcal{B}) := \left| \Pr \left[1 \leftarrow \mathcal{B} \circ \text{GEOCH}_{\text{xpd,kdf,pke}}^{0, d, t, i} \right] - \Pr \left[1 \leftarrow \mathcal{B} \circ \text{GEOCH}_{\text{xpd,kdf,pke}}^{1, d, t, i} \right] \right| \leq \text{Adv}_{\text{KDFR}}^{\text{kdf}}(\mathcal{B} \circ \mathcal{R}_{a, i}) + \text{Adv}_{\text{HPKE}}^{\text{pke}}(\mathcal{B} \circ \mathcal{R}_{b, i}) + \text{Adv}_{\text{KDFR}}^{\text{kdf}}(\mathcal{B} \circ \mathcal{R}_{c, i}) + \text{Adv}_{\text{EXP}}^{\text{xpd}}(\mathcal{B} \circ \mathcal{R}_{d, i}) + \text{Adv}_{\text{HPKE}}^{\text{pke}}(\mathcal{B} \circ \mathcal{R}_{e, i}),$$

where $\mathcal{R}_{a, i}$ is defined as the grey area in Figure 26, $\mathcal{R}_{b, i}$ as the grey area in Figure 27, $\mathcal{R}_{c, i}$ as the grey area in Figure 28, $\mathcal{R}_{d, i}$ as the grey area in Figure 29, and $\mathcal{R}_{e, i}$ as the grey area in Figure 30.

Proof of Theorem 2. We first apply Lemma 2 and then apply a hybrid argument over the t epochs using that for all \mathcal{B} ,

$$\text{Adv}_{\text{EPOCH}}^{\text{xpd,kdf,pke}, i}(\mathcal{B} \circ \mathcal{R}_{\text{hyb}, i}) = \left| \Pr \left[1 \leftarrow \mathcal{B} \circ \text{GGKS}_{\text{xpd,kdf,pke}}^{1, 1_0, \dots, 1_{i-1}, 0_i, 0_{i+1}, \dots, 0_t, d, t} \right] - \Pr \left[1 \leftarrow \mathcal{B} \circ \text{GGKS}_{\text{xpd,kdf,pke}}^{1, 1_0, \dots, 1_{i-1}, 1_i, 0_{i+1}, \dots, 0_t, d, t} \right] \right|,$$

where $\mathcal{R}_{\text{hyb}, i}$ is defined in Figure 21b. We obtain Theorem 2 as follows:

$$\begin{aligned} \text{Adv}_{\text{KS}}^{\text{xpd,kdf,pke}}(\mathcal{A}) &\leq \text{Adv}_{\text{GKS}}^{\text{xpd,kdf,pke}, b_0}(\mathcal{A}) + \text{Adv}_{\text{GKS}}^{\text{xpd,kdf,pke}, 1b}(\mathcal{A}) \\ &\leq \text{Adv}_{\text{GKS}}^{\text{xpd,kdf,pke}, b_0}(\mathcal{A}) + \sum_{i=0}^{t-1} \text{Adv}_{\text{EPOCH}}^{\text{xpd,kdf,pke}, i}(\mathcal{A} \circ \mathcal{R}_{\text{hyb}, i}). \end{aligned}$$

Plugging in the bounds from Lemma 2 and Lemma 3 with $\mathcal{B} = \mathcal{A} \circ \mathcal{R}_{\text{hyb}, i}$ yields Theorem 2. We now turn to the proof of Lemma 2.

C. Proof of Lemma 2

In order to prove Lemma 2, we split b'_0 into 3 different bits b_0 , b_1 , and b_2 and write $\text{GPGKS}_{\text{xpd,kdf,pke}}^{b_0 b_1 b_2 0, d, t}$ for the game depicted in Figure 22, where the P in GPGKS stands for *proof*. Note that when $b_0 = b_1 = b_2 = b'_0$, then

$$\text{GPGKS}_{\text{xpd,kdf,pke}}^{b_0 b_1 b_2 0, d, t} = \text{GGKS}_{\text{xpd,kdf,pke}}^{b'_0 0, d, t}$$

For collision-resistance, we can construct a reduction \mathcal{R}_{cr} such that the distinguishing advantage between $\text{GPGKS}_{\text{xpd,kdf,pke}}^{0000, d, t}$ and

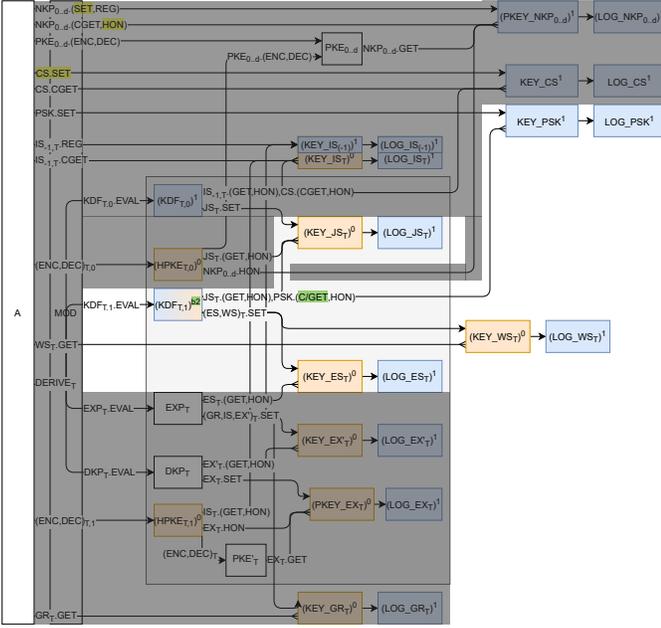


Figure 24: Reducing to KDFL security based on ideal PSK. Reduction \mathcal{R}_{psk} is marked in grey.

$\text{GPGKS}_{\text{xpd,kdf,pke}}^{1000,d,t}$ can be turned into a collision against one of the underlying primitives. This is possible because (a) the initial commit secret, PSK, and level 0 init secrets are unique and (b) we have domain separation between the different epochs due to the inclusion of the group context in the KDFs. Thus, we can argue inductively that

$$\left| \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GPGKS}_{\text{xpd,kdf,pke}}^{0000,d,t} \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GPGKS}_{\text{xpd,kdf,pke}}^{1000,d,t} \right] \right| \leq \text{Adv}_{\text{CR}}^{\text{xpd,kdf,dkp}}(\mathcal{A} \circ \mathcal{R}_{cr}).$$

After idealizing collision-resistance, we can now reduce to the LKDF security of the commit secret, since the respective second inputs are unique. We thus obtain that

$$\left| \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GPGKS}_{\text{xpd,kdf,pke}}^{1000,d,t} \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GPGKS}_{\text{xpd,kdf,pke}}^{1100,d,t} \right] \right| \leq \text{Adv}_{\text{KDFL}}^{\text{kdf}}(\mathcal{A} \circ \mathcal{R}_{cs}),$$

where \mathcal{R}_{cs} is defined in Figure 23. Analogously,

$$\left| \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GPGKS}_{\text{xpd,kdf,pke}}^{1100,d,t} \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GPGKS}_{\text{xpd,kdf,pke}}^{1110,d,t} \right] \right| \leq \text{Adv}_{\text{KDFL}}^{\text{kdf}}(\mathcal{A} \circ \mathcal{R}_{psk}),$$

where \mathcal{R}_{psk} is defined in Figure 24. We obtain

$$\begin{aligned} \text{Adv}_{\text{GKS}}^{\text{xpd,kdf,pke},b_0}(\mathcal{A}) &:= \\ &\left| \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GPGKS}_{\text{xpd,kdf,pke}}^{00,d,t} \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GPGKS}_{\text{xpd,kdf,pke}}^{10,d,t} \right] \right| \\ &\leq \text{Adv}_{\text{CR}}^{\text{xpd,kdf,dkp}}(\mathcal{A} \circ \mathcal{R}_{cr}) + \text{Adv}_{\text{KDFL}}^{\text{kdf}}(\mathcal{A} \circ \mathcal{R}_{cs}) \\ &\quad + \text{Adv}_{\text{KDFL}}^{\text{kdf}}(\mathcal{A} \circ \mathcal{R}_{psk}), \end{aligned}$$

which concludes the proof of Lemma 2.

D. Key Schedule Epoch

We now prove Lemma 3 which captures the security of a single epoch of the key schedule, a core argument in the proof

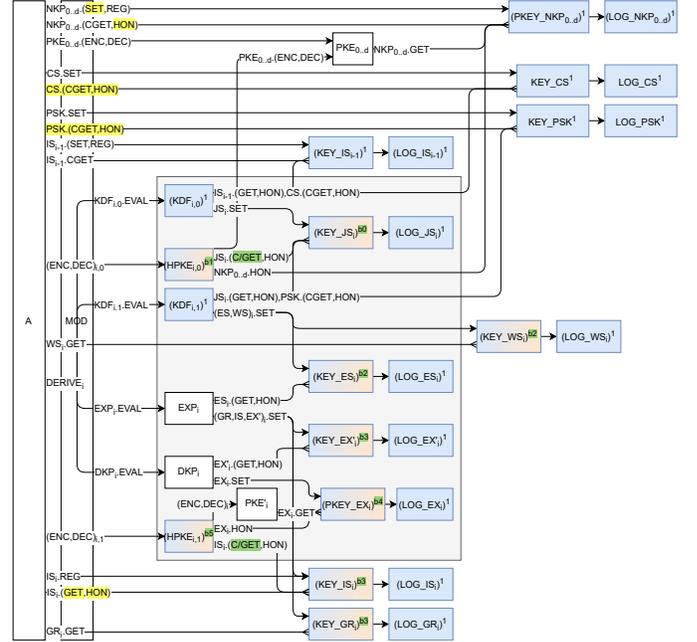


Figure 25: Definition of $\text{GEPOCH}_{\text{xpd,kdf,pke}}^{b_0 b_1 b_2 b_3 b_4 b_5, d, t, i}$. We write $\text{GEPOCH}_{\text{xpd,kdf,pke}}^{b, d, t, i}$ if $b_0 = b_1 = b_2 = b_3 = b_4 = b_5$. Note that $\text{KEY}_{\text{IS}_{-1}}$ should technically have a REG query, but for an ideal KEY package REG and SET are functionally equivalent.

of Theorem 2. We show that when an MLS epoch relies on an ideal internal secret s_{int} derived in a previous epoch or an ideal PSK s_{ext} or an ideal commit secret k_{cs} , then the current epoch provides security guarantees, too, for the new s'_{int} derived in the current epoch, all group secrets for the current epoch and for the current external public-key encryption keys.

The proof of Lemma 3 consists of six steps, each of which flips one or more bits in the overall game. See Figure 25 for the definition of $\text{GEPOCH}_{\text{xpd,kdf,pke}}^{b_0 b_1 b_2 b_3 b_4 b_5, d, t, i}$ and note that

$$\text{GEPOCH}_{\text{xpd,kdf,pke}}^{b, d, t, i} = \text{GEPOCH}_{\text{xpd,kdf,pke}}^{b_0 b_1 b_2 b_3 b_4 b_5, d, t, i}$$

if $b = b_0 = b_1 = b_2 = b_3 = b_4 = b_5$.

In the first step of this proof we use the KDFR assumption to idealize the KEY_{JS} package of epoch i . We obtain

$$\left| \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GEPOCH}_{\text{xpd,kdf,pke}}^{000000, d, t, i} \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GEPOCH}_{\text{xpd,kdf,pke}}^{100000, d, t, i} \right] \right| \leq \text{Adv}_{\text{KDFR}}^{\text{kdf}}(\mathcal{A} \circ \mathcal{R}_{a,i}).$$

Now, in the second game hop, we use the HPKE assumption for $n = 1$ and $m = d$ to idealize the first HPKE package of epoch i and to remove access to the joiner secret. We obtain

$$\left| \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GEPOCH}_{\text{xpd,kdf,pke}}^{100000, d, t, i} \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GEPOCH}_{\text{xpd,kdf,pke}}^{110000, d, t, i} \right] \right| \leq \text{Adv}_{\text{HPKE}}^{\text{pke}}(\mathcal{A} \circ \mathcal{R}_{b,i}).$$

In the third step of this proof we use the KDFR assumption to idealize the KEY_{ES} package of epoch i . This is possible since there is no GET query on the joiner secret anymore. We obtain the bound

$$\left| \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GEPOCH}_{\text{xpd,kdf,pke}}^{110000, d, t, i} \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \text{GEPOCH}_{\text{xpd,kdf,pke}}^{111000, d, t, i} \right] \right| \leq \text{Adv}_{\text{KDFR}}^{\text{kdf}}(\mathcal{A} \circ \mathcal{R}_{c,i}).$$

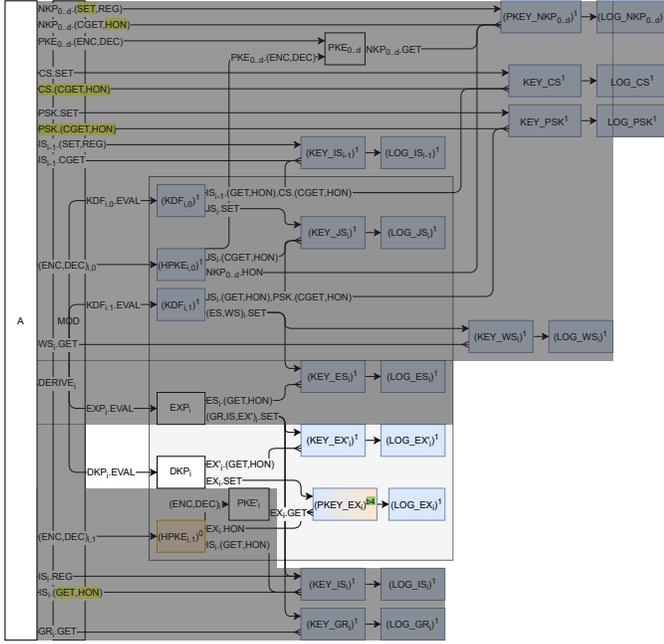


Figure 30: Step 5 of idealizing epoch i . We mark \mathcal{R}_{dkp} in grey.

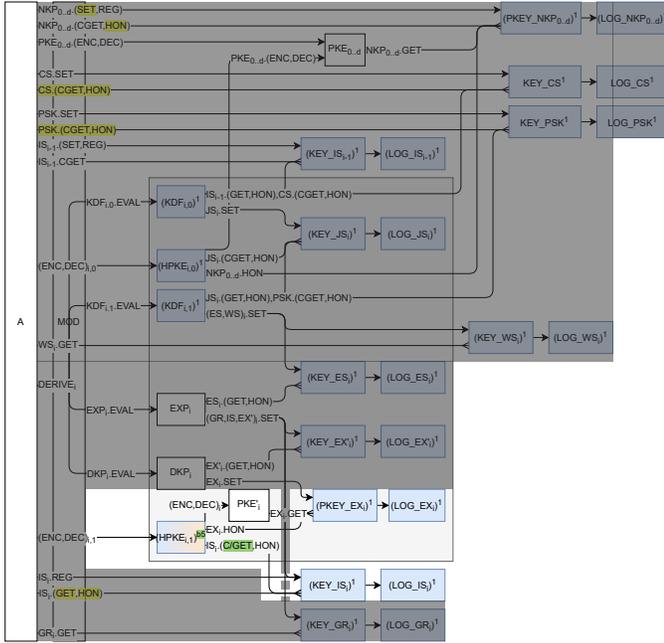


Figure 31: Step 6 of idealizing epoch i . We mark $\mathcal{R}_{e,i}$ in grey.

$$\begin{aligned}
 \text{Adv}_{\text{EPOCH}}^{\text{xpd, kdf, pke, } i}(\mathcal{A}) &\leq \text{Adv}_{\text{KDFR}}^{\text{kdf}}(\mathcal{A} \circ \mathcal{R}_{a,i}) + \text{Adv}_{\text{HPKE}}^{\text{pke}}(\mathcal{A} \circ \mathcal{R}_{b,i}) \\
 &\quad + \text{Adv}_{\text{KDFR}}^{\text{kdf}}(\mathcal{A} \circ \mathcal{R}_{c,i}) + \text{Adv}_{\text{EXP}}^{\text{xpd}}(\mathcal{A} \circ \mathcal{R}_{d,i}) \\
 &\quad + 0 + \text{Adv}_{\text{HPKE}}^{\text{pke}}(\mathcal{A} \circ \mathcal{R}_{e,i}),
 \end{aligned}$$

and Lemma 3 holds.