ETK: External-Operations TreeKEM and the Security of MLS in RFC 9420

Cas Cremers[1], Esra Günsay[1], Vera Wesselkamp, and Mang Zhao

[1]CISPA Helmholtz Center for Information Security

Version 1.0, February 14, 2025

Abstract

The Messaging Layer Security protocol MLS is standardized in IETF's RFC 9420 and allows a group of parties to securely establish and evolve group keys even if the servers are malicious. Its core mechanism is based on the TreeKEM protocol, but has gained many additional features and modifications during the development of the MLS standard. Over the last years, several partial security analyses have appeared of incomplete drafts of the protocol. One of the major additions to the TreeKEM design in MLS RFC 9420 (the final version of the standard) are the external operations, i.e., external commits and proposals, which interact deeply with the core TreeKEM protocol. These operations have not been considered in any previous security analysis, leaving their impact on the protocol's overall security unclear.

In this work, we formalize ETK: External-Operations TreeKEM that includes external commits and proposals. We propose a corresponding ideal functionality $\mathcal{F}_{\text{ECGKA}}$ and prove that ETK realizes $\mathcal{F}_{\text{ECGKA}}$.

Our work is the first cryptographic analysis that considers both the final changes to the standard's version of TreeKEM as well as external proposals and external commits. Compared to previous works that considered MLS draft versions, our ETK protocol is by far the closest to the final MLS RFC 9420 standard. Our analysis implies that the core of MLS's TreeKEM variant as defined in RFC 9420 is an ETK protocol that realizes $\mathcal{F}_{\text{ECGKA}}$, when used with an $\mathsf{SUF\text{-}CMA}$ secure signature scheme, such as the IETF variant of Ed25519. We show that contrary to previous claims, MLS does not realize $\mathcal{F}_{\text{CGKA}}$ [4] when used with signature schemes that only guarantee $\mathsf{EUF\text{-}CMA}$, such as ECDSA.

Moreover, we show that the security of the protocol could be further strengthened by adding a functionality to insert PSKs, allowing another form of healing, and give a corresponding construction $\mathsf{ETK}^{\text{PSK}}$ and ideal functionality $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$.

1 Introduction

Messaging Layer Security (MLS) is a messaging protocol standardized by the Internet Engineering Task Force (IETF) in [14]. MLS was designed as a highly-secure asynchronous group messaging protocol that scales efficiently for large groups. MLS allows for dynamic membership and provides strong security guarantees such as Forward Secrecy (FS) and Post-Compromise Security (PCS). Informally, FS states that an attacker cannot decrypt messages exchanged before a compromise, while PCS ensures that after a compromise, messages are secure after a *healing period*, in which clients exchange a few secure messages [26]. MLS is already deployed, or in the process of becoming so, in real-world messaging applications such as Cisco WebEx [6], Wickr [33] and Matrix [31]. It also provides the base of an IETF protocol for interoperable messaging services called More Instant Messaging Interoperability (MIMI) [34].

The early development of the MLS standard has been accompanied by rigorous security analyses by the scientific community, e.g., [1, 2, 3, 4, 22, 28, 36]. Most of this work has focused on two sub-protocols of MLS: *TreeKEM*, which encompasses the operations needed for computing a shared group secret via a tree structure, and the subsequent *key schedule* for message encryption. These two components are considered the core of MLS, providing PCS through regular updates to clients' key material, and FS through ratcheting the message encryption keys.

More importantly, the full MLS protocol specifies operations far beyond the basic TreeKEM operations analyzed in previous works. One of the main categories of additional operations are the external operations, such as external commit and external proposal. For example, applications can allow group members

who lost access to the group secrets to *resync* with the group by rejoining it. Further use cases of external proposals are automated services to remove inactive members, or propose adding new staff members. Moreover, some groups may be open, and allow external members join directly by themselves. Although not considered in previous academic work, these, and further operations are used in real-world implementations of MLS [25, 33].

In this work, we set out to analyze and prove formal security guarantees for the core TreeKEM protocol in the final MLS RFC 9420 including external operations. To this end, we build on a previous state-of-the-art analysis of MLS draft 12 [4] that considered malicious insiders.

Contributions. This work presents the first computational security analysis of the TreeKEM variant used in MLS RFC 9420, and includes extended operations that were not addressed in previous studies. Our main contributions are:

1. We introduce ETK: Extended-Operations TreeKEM, which models the core TreeKEM protocol of the final version of the MLS standard in RFC 9420. Whereas previous cryptographic analyses analyzed Draft 12 (October 2021) or earlier, we consider the final version [14] (July 2023), which is based on Draft 20 [30]. Our ETK protocol captures the latest changes to the standard, such as encrypting `groupInfo`, separate keys for leaf nodes, and allowing multiple welcome messages. Notably, ETK is the first TreeKEM variant to include external operations, including proposals, commits, and `resync`.
2. We formalize the security guarantees of the ETK protocol in the ideal functionality $\mathcal{F}_{\mathrm{ECGKA}}$, and prove that ETK realizes $\mathcal{F}_{\mathrm{ECGKA}}$ with respect to active adversaries and potentially malicious insiders, yielding fine-grained security guarantees.

   On the positive side, our results imply that the TreeKEM version in RFC 9420, when instantiated with an SUF-CMA-secure signature scheme, provides strong security guarantees, even with complex operations such as external proposals, commits, and `resync`.

   On the negative side, and contrary to previous claims, we show that MLS can only realize $\mathcal{F}_{\mathrm{CGKA}}$ [4] if the signature scheme is SUF-CMA-secure. RFC 9420 requires implementation of the IETF version of Ed25519, which satisfies SUF-CMA [21], but also allows for the use of ECDSA, which only satisfies EUF-CMA [35]. While we do not know a practical attack on the protocol, we show that when using a signature scheme that is not SUF-CMA secure, the core consistency property of TreeKEM can be violated: the adversary can make two parties that would expect to agree on the state unable to communicate. The underlying reason is that the key schedule includes transcripts, which include the signature data, which could be mangled by the adversary when using, e.g., ECDSA. This also implies that the main theorem in [4] is technically incorrect, and the stated EUF-CMA assumption should be replaced by SUF-CMA. A further corollary is that in FIPS-compliant MLS implementations, the standard's `AuthenticatedContent` objects can be mangled to cause divergence of keys.
3. We formally show that the use of PSKs during a resync improves security guarantees against attacks based on bad or compromised randomness. Concretely, we include the use of PSKs in $\mathsf{ETK}^{\mathrm{PSK}}$, and show that the resulting protocol meets a stronger security notion $\mathcal{F}_{\mathrm{ECGKA}^{\mathrm{PSK}}}$.

Outline. We recall background on MLS and its latest formal security analysis in Section 2. We define our protocol syntax in Section 3 and our security functionality in Section 4. We introduce MLS RFC 9420 and its core operations in Section 5 and prove that it realizes our functionality in Section 5.2, where we also comment on the signature scheme requirements. We formally show the security improvements of PSK injections in Section 6. We recall preliminaries and provide proofs for all theorems in the appendix.

2 Background

We first recall some major changes in the evolution of the Messaging Layer Security (MLS) protocol in Section 2.1. In Section 2.2, we recall the latest existing formal security analysis for MLS. In Section 2.3, we introduce our notations.

2.1 The Evolution of the MLS Protocol

MLS is an asynchronous group key exchange protocol that aims to provide both high efficiency and strong security guarantees. MLS was initiated by the Internet Engineering Task Force (IETF) in February 2018 [16], based on the *Asynchronous Ratchet Tree* (ART) protocol [27]. ART allows a static group of n

parties, each corresponding to a leaf node in a tree, to generate a shared group key, corresponding to the root in a tree, by recursively computing Diffie-Hellman Exchanges (DHE) from leaf nodes to the tree root. Compared to conventional group communication via pairwise channels, ART reduces the computational and communication efforts for every group member from $O(n)$ to $O(\log(n))$. Moreover, every group member can refresh their DHE keys along the path to the root, which allows ART to achieve several strong security guarantees, e.g., *forward secrecy* (FS) and *post-compromise security* (PCS).

Since Draft 2 [17], MLS has replaced ART with TreeKEM. TreeKEM has a similar tree structure to ART, but relies on a generic *Key Encapsulation Mechanism* (KEM) instead of DHE. This enables TreeKEM to achieve post-quantum security by selecting a suitable KEM instantiation. Every group member may rekey the secrets on the path from their leaf node to the root. The secret at the tree root is then used for message encryption. However, TreeKEM cannot deal with conflicts caused by concurrent message sending or group operations (e.g., add/remove members, key update), and impractically assumes that no two group members execute group operations at the same time.

To better handle concurrency, MLS added two mechanisms. First, with Draft 7 [7] MLS decreased the frequency of the key update from every message encryption to every "epoch", and passes the root key to a *key schedule* to initialize a *secret tree*[1] for continuous message encryption within every epoch. Second, with Draft 8 [8] MLS decomposed group operations into *Proposals* and *Commits* (also known as "P&C" mechanism). Whenever a group member wants to execute a group operation (e.g., add/remove members or key update), they first send a corresponding proposal. Any group member can then collect and process the received proposals in a commit. The committer is responsible for ordering the proposals and dealing with potential conflicts. Since Draft 8 until the final RFC 9420 standard, the ratchet tree (i.e., TreeKEM), key schedule, and secret tree comprise the core of MLS. The security of these two mechanisms has been formally analyzed in several works, e.g., [1, 4, 19, 22, 27, 28, 29, 36].

After in total twenty drafts, MLS was standardized as RFC 9420 in July 2023 [14]. Compared to basic TreeKEM, RFC 9420 incorporates several new operations and advanced features. These include PSK injection, group re-initialization, sub-group branching, and joining via external commit since Draft 10 [9]; separate keys for leaf nodes in TreeKEM from identity key package since Draft 13 [11]; external self-add proposal since Draft 15 [12]; and allowing multiple welcomes per commit since Draft 17 [13]. To the best of our knowledge, no existing work covers the formal security analysis for any of the above novel operations. We give an overview over previous formal analysis in Table 1. Beyond studies analyzing TreeKEM, some designs aim to enhance its features. Tainted-TreeKEM [32] improves efficiency and security by using a proxy to update encryption keys without blanking direct paths. Quarantined-TreeKEM [24] introduces a 'quarantine' for inactive users, who randomly blank paths and update keys via a distributed secret. QTK strengthens PCS but offers no FS advantage and is compatible with RFC 9420.

	Version	**Framework**	**Scope**	**Adversary**	**Operations**
[27]	Draft 01	symbolic/game-based	static groups	active	-
[2]	Draft 06	game-based	integrity mechanisms	passive	O_1
[20]	Draft 07	symbolic	Messaging	insider	O_1
[22]	Draft 11	state separating proofs	key derivation schedule	passive	-
[1]	Draft 11	game-based	rTreeKEM +group splitting attack	passive+	O_1
[28]	Draft 11	-	cross-group PCS analysis + concurrent group attacks	active-	update-only
[4]	Draft 12	UC	integrity mechanisms + 3 insider attacks	active-	O_1
[36]	Draft 16	symbolic	integrity mechanisms + signature confusion attack	active	O_1
this work	RFC 9420	UC	integrity mechanisms + randomness leakage attack	active	$O_1 \cup O_2$

Table 1: Overview of previous formal analyses of the MLS protocol. O_1 represents the add, update, and remove proposals with regular commits and O_2 represents the external proposals and commits. *passive+* refers to an adversary capable of injecting messages, provided these messages are ultimately rejected by a client. *active-* denotes an active adversary that allows a recovery on healing phase.

[1] MLS introduced the "Tree-based Application Key Schedule" in Draft 7 [7] and renamed it to "secret tree" in Draft 8 [8].

2.2 Formal Security Analysis of MLS Draft 12

The latest formal computational analysis for MLS was provided by Alwen, Jost, and Mularczyk [4] for Draft 12. The paper [4] isolated a core component in MLS Draft 12 called "*Insider Secure TreeKEM*" (ITK), and proved that it realizes a functionality $\mathcal{F}_{\mathrm{CGKA}}$ in the Universal Composability (UC) framework. ITK assumes MLS Draft 12 to be equipped with two abstract services: an authentication service (AS), which manages long-term identity keys, and a key service (KS), which allows parties to upload single-use key packages, used by group members to non-interactively add them to the group. Below, we recall ITK and $\mathcal{F}_{\mathrm{CGKA}}$.

Identity and Credentials. Each leaf node contains a public signature key and a credential that links the member's identity to the signature key. When using an X.509 credential for example, the public signature key of the leaf node is contained in the `subjectPublicKeyInfo` field. The AS validates that the public signature key matches an accepted identity (called a reference ID). An AS might be implemented as a service that signs the members' credentials, allowing other members to verify the credential by verifying the AS's signature. Members must be able to recognize that two credentials belong to the same member according to the application's policy. This might be trivial, e.g., because they use identical signature keys. An update can replace the credentials used by a client. What constitutes a valid successor to a credential is defined by the application.

Insider Secure TreeKEM ITK. ITK augments TreeKEM with message authentication, tree-signing, computation of tags, and parts of the key schedule.

TreeKEM. Each member in TreeKEM is represented as a leaf node in the tree that contains the member's encryption and signature key. Members suggest changes to the group by sending a *proposal* to all group members. The change is applied upon a member *committing* one or more proposals. A commit moves the group to a new epoch by updating the node secrets along the path of the committer, including the root secret. ITK covers TreeKEM with selected operations that can be invoked by group members, including group members' regular proposals and commits for adding new members, updating leaf node secrets, and removing existing members.

- Add Proposal: To add a new member to the group, the proposer fetches a key package containing a fresh leaf node for the client from the KS and adds it to the proposal. With the commit, the new member's leaf node is added to the tree. The committer sends the new member a *welcome* message that contains all information necessary for the new group member to join the group: the current group secret, encrypted with the new member's public encryption key, as well as the secret key of the lowest node in the ratchet tree that the committer and the new member share. The information in the welcome message is encrypted by the *init key* included in the key package. The add proposal always adds the members to the leftmost free node in the tree.
- Update Proposal: An update proposal allows a member to rotate its encryption and optionally its signature keys. The committer of the proposal replaces the updating member's node with the new leaf node contained in the update proposal. It then "blanks" (i.e., deletes the content of) all intermediate nodes on the path from that leaf to the root (the blue path in Figure 1a). Afterwards, the committer updates its own path and encrypts the path secrets to other group members, in particular to the updated party using the updated leaf node's key. This replaces all knowledge of the updated party about the tree.
- Remove Proposal: A remove proposal specifies a member that should be removed from the group. Upon committing, the committer removes the leaf node of the member and blanks all nodes on the path of the removed member to the root. This guarantees that the removed member does not have access to any secrets in the public ratchet tree anymore. The committer update its path secrets (including the new secret for the root) and encrypts the path secrets only to the existing group members. Figure 1b visualizes an example of the process.
- Regular Commit: A regular commit initializes a new epoch, validates and applies the collected proposals, and optionally rekeys the committer's path if the proposals are neither empty nor add-only. The committer rekeys the path from their leaf node to the tree root: refreshes their key pairs of their leaf node by simply generating a new key pair, and of the nodes on their path to root in the following four steps: (1) generate an initial path secret for the first parent node of the leaf, (2) recursively expand it to the next path secret for the next parent node, along the path to the root, (3) derive key pairs for all intermediate nodes on the path to the root from their path secrets, and (4) encapsulate the path secrets at every node on the path under the KEM public keys of the resolution of their sibling nodes. The resolution of a node is an ordered list of non-blank nodes that collectively cover all non-blank

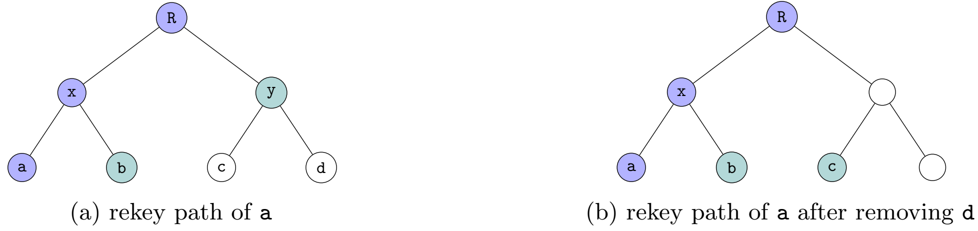

(a) rekey path of `a` (b) rekey path of `a` after removing `d`

Figure 1: MLS group with four members: `a`, `b`, `c`, and `d`. The blue nodes are the secrets known to member `a` that are replaced upon performing a rekey path. The new node secrets are encrypted to the public key of green nodes , i.e., the resolution of their sibling nodes. The right figure shows the removal of node `d`, where the right child node of root is blanked and its resolution is `c`.

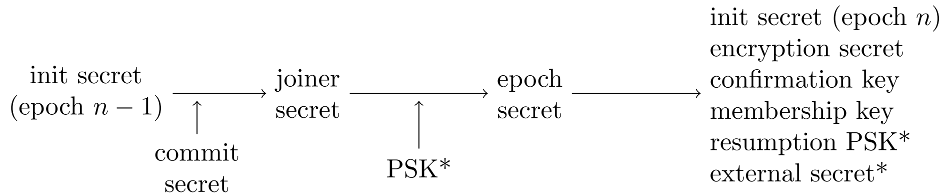

Figure 2: Key derivation schedule for epoch n. *: The PSK injection and the derivation of resumption PSK and external secret are excluded from ITK but included in MLS since Draft 10 [9].

descendants of the node. After the optional rekey path execution, the committer runs "tree-signing", signs the content of the commit, executes the key schedule to derive the secrets for the new epoch, and compute tags. A regular commit message includes the information about the current epoch, identifier of the committer, content of the commit, and above signatures and tags, is output together with a new welcome message for new parties to join the group if available.

Tree-Signing. The tree-signing guarantees newly joining parties that each internal node has been sampled by one of the parties contained in its sub-tree. Upon committing, the committer computes a "parent hash" value, which stores a hash of its parent node and the resolution of its sibling node, for every node on the committer's path from top (i.e., the tree root) to bottom (i.e., the leaf node) in order. By including the parent node in the hash, the freshly sampled key pairs at leaf nodes get recursively bound with all freshly sampled nodes on the path to the root together. By including the siblings node in the hash, the freshly sampled key pairs at every node on the committer's path gets bound with the nodes, to whom every freshly sampled public key should be distributed.

Key Schedule. TreeKEM's ratchet tree provides an efficient way to share a new secret with all group members. The root secret of the ratchet tree is called the commit secret, as it is derived newly in each commit. The commit secret of the new epoch is then combined with a secret from the last epoch, called the *init secret*. In MLS Draft 12, another secret called a "pre-shared key" (PSK) can also be injected into the key derivation process, but these are not modeled in ITK. The derivation of the above secrets is called the *epoch secret*. As the name indicates, ITK derives from the epoch secret a number of group secrets for the computation within this epoch. We depict the key derivation schedule in Figure 2. The most important group secrets are:

- init secret: used for key schedule in the next epoch.
- encryption secret: used to derive the secret tree within this epoch.
- confirmation key: used for computing confirmation tags.
- membership key: used for computing membership tags.
- resumption PSK: used to ensure the membership consistency when the group is branched or reinitialized.
- external secret: used to derive a key pair whose private key is held by the entire group and whose public key can be shared with outsiders to enable them to join the group via external commit.

Computation of Tags. Every commit includes at most two tags: a confirmation tag and a optional membership tag if the commit includes any remove proposal. The confirmation tags guarantee the

succession of the ratchet tree in every new epoch from the one in the past epoch, and further authenticate the group's history. The ratchet tree is equipped with two hashes on the communication transcript: *confirmed transcript hash* and *interim transcript hash.* The confirmed transcript hash is computed by hashing the previous epoch's interim transcript hash, the content of the commit message, and its signature. The interim transcript hash is computed by hashing the confirmed transcript hash with the confirmation tag. Upon committing, the committer uses the confirmation key, which is derived from key schedule, to MAC the confirmed transcript hash, for a confirmation tag, which allows the receiving members to immediately verify whether they agree on the new epoch's key-schedule.

Note that the committer will not encrypt new path secrets to any party that is included in any remove proposal in the commit. Thus, the removed members cannot verify the confirmation tag, because they are unable to execute key schedule to derive the new epoch's confirmation key. In this case, the committer additionally MAC the content of the commit under the current epoch's membership key. Every group member can verify the commit under the same membership key, to determine whether they will be removed in the next epoch.

UC Security and Functionality $\mathcal{F}_{\text{CGKA}}$. We briefly recall the UC framework [23] and the functionality $\mathcal{F}_{\text{CGKA}}$ in [4].

Universal Composability (UC). In the UC framework there exists an environment $\mathcal{Z}$ which can generate the inputs to, read all outputs from, and interact with two worlds: the *real world* and the *ideal world*. In the real world, an instance of a protocol P interacts with an adversary $\mathcal{A}$. In the ideal world, an ideal functionality $\mathcal{F}$, which perfectly simulates the behaviors of P without involving any cryptography, interacts with a simulator $\mathcal{S}$ (also called *ideal adversary*), which mimics every attack of $\mathcal{A}$. We say that the protocol P securely realizes the functionality $\mathcal{F}$ if for any adversary $\mathcal{A}$, there exists a simulator $\mathcal{S}$ such that no environment $\mathcal{Z}$ can distinguish whether it is interacting with the real world (i.e., the instance of P and the adversary $\mathcal{A}$) or with the ideal world (i.e., the ideal functionality $\mathcal{F}$ and the simulator $\mathcal{S}$).

Functionality $\mathcal{F}_{\text{CGKA}}$ in [4]. The functionality $\mathcal{F}_{\text{CGKA}}$ in [4] captures a very strong adversary, which can (1) fully control the network of message and key packages distribution (except for the long-time identity keys on the authentication service AS), (2) adaptively corrupt parties' states, (3) adaptively corrupt and manipulate parties' randomness, and (4) potentially collude with malicious group members. In order to avoid the so-called commitment problem caused by adaptive corruptions in simulation-based frameworks, $\mathcal{F}_{\text{CGKA}}$ restricts the environment not to corrupt parties at certain times to prevent "trivial attacks", e.g., no corruption that yields the change of the safety of a prior "secure" epoch is allowed.

$\mathcal{F}_{\text{CGKA}}$ captures three core security properties of ITK: consistency, confidentiality, and authenticity. Informally, consistency means that all parties in the same epoch agree on the same group state, including the history of the group's evolution. Confidentiality means that the adversary learns nothing about the important secrets in any epoch that is marked "safe", evaluated by a predicate **safe**. The secrets in the "safe" epochs will be derived from a random and independent secret, while the others will be chosen by the simulator $\mathcal{S}$. Finally, authenticity means that the environment cannot forge any message on behalf of a group member as long as the evaluation on a predicate **inj-allowed** is false.

The core theorem of [4] is that ITK, when used with an EUF-CMA signature scheme, realizes $\mathcal{F}_{\text{CGKA}}$ in the random oracle model under the assumption of a modified version of *Generalized Selective Decryption* (GSD) security of the underlying PKE, which can be further reduced to the standard IND-CCA security. We will show in Remark 1 that EUF-CMA is not sufficient for their main theorem.

2.3 Notation

Our notation follows [4]. Let $v \leftarrow x$ denote assigning the value x to v while $v \stackrel{\$}{\leftarrow} S$ denote sampling a random element from a set S. For a set V, $V \mathrel{+}\leftarrow x$ and $V \mathrel{-}\leftarrow x$ denotes adding and removing an element x. Let $A[i] \leftarrow x$ and $y \leftarrow A[i]$ respectively denote the assignment and retrieval of i-th element from an array A, while $A[*] \leftarrow v$ denote setting all array positions to the value v. For brevity, we sometimes omit set or list brackets when there is only one element and the type is clear from the context. Concatenation is denoted by $x +\!\!+ y$. We denote invoking an algorithm A with explicitly supplied randomness r by $A(.;r)$.

We use three keywords to encode the expectation of expressions:

- **req** ⟨**condition**⟩: if the condition following the **req** keyword fails, the function returns $\perp$ and reverts all changes that were applied so far. The keyword **req** is used to control the allowed function inputs.
- **try** $y \leftarrow f(x)$ is equivalent to $y \leftarrow f(x)$ followed by **req** $y \neq \perp$.

Color	Meaning
●	Identical to [4] for syntax, protocols, and functionalities
●	Updates from Draft 12 (modeled in [4]) to the final RFC 9420
●	External commit / resync
●	External self-add proposal
●	(Resumption) PSK proposals

Table 2: We use color coding to highlight and differentiate our new elements in the syntax, functionalities, and protocols.

- **assert** ⟨**condition**⟩: if the condition following the **assert** keyword fails, the functionality permanently halts, making the ideal and real-world trivially indistinguishable. The keyword **assert** encodes security properties of the protocol.

We use pattern matching to simply the definitions of algorithms in our protocol and security model. For example, in $p \leftarrow (\texttt{Propose}, \texttt{rem-id}_t)$, `rem` is a keyword, while $\texttt{id}_t$ is the actual identity signifier of the removed party.

Throughout this work, we indicate the differences and commonalities between our syntax, functionalities, and protocols and the corresponding ones in [4] using the color coding in Table 2.

3 Extended Continuous Group Key Agreement: Syntax

In this section, we define our syntax for *Extended Continuous Group Key Agreement* (ECGKA) protocols. Our definition extends the definition in [4, Section 3.3] to include the additional operations: *external commits/resync*, *external self-add proposals*, and *PSK proposals*, each of which we will introduce in detail in Sections 4 and 5. We assume each party in ECGKA has a unique identifier `id`. For simplicity, we assume that every algorithm includes the `id` of the invoker (i.e., the party who runs this algorithm) as an implicit input.

Definition 1. *An ECGKA protocol includes the following algorithms:*

Group Creation: $(\texttt{Create}, \texttt{spk})$ *creates a new group with signature key* `spk` *and* `id` *as the only member.*

Add Proposals: $p \leftarrow (\texttt{Propose}, \texttt{add-id}_t)$ *proposes to add the party* $\texttt{id}_t$*. The output is a proposal message* p *if* $\texttt{id}_t$ *is not yet a group member or* $\perp$ *if it is.*

Remove Proposals: $p \leftarrow (\texttt{Propose}, \texttt{rem-id}_t)$ *proposes to remove the party* $\texttt{id}_t$*. The output is a proposal message* p *if* $\texttt{id}_t$ *is a group member or* $\perp$ *if it is not.*

Update Proposals: $p \leftarrow (\texttt{Propose}, \texttt{up-spk})$ *proposes to update the member's leaf key and (if different from the current) signature key* `spk`*. The output is a proposal* p *or* $\perp$ *if* `id` *is not a member.*

Commit: $(c, \texttt{vec_w}, g) \leftarrow (\texttt{Commit}, \vec{p}, \texttt{spk}, \texttt{force-rekey})$ *commits the vector of proposals* $\vec{p}$ *(that were previously generated) and outputs the commit message* c*. If* $\vec{p}$ *contains an add proposal, it outputs a vector of welcome messages* `vec_w`*. It additionally outputs the group information (a.k.a. groupInfo)* g *for external joiners.* `force-rekey` *determines if the committer's keying material is enforced to be replaced. Optionally, the committer's signature key is updated to* `spk`*.*

Process: $(\texttt{id}_c, \texttt{propSem}) \leftarrow (\texttt{Process}, c, \vec{p})$ *processes the commit* c *that includes proposals* $\vec{p}$ *and advances the epoch to the next one. The output is the committer* $\texttt{id}_c$ *and* `propSem`*, a vector of the proposers and actions of* $\vec{p}$*. If* $\vec{p} = \perp$*, it assumes that* c *is an external commit and the proposals are directly included in the commit message.*

Join: $(\texttt{roster}, \texttt{id}_c) \leftarrow (\texttt{Join}, \texttt{vec_w})$ *allows a new member* `id` *to use a vector of welcome messages* `vec_w` *to join the group. The output is* `roster`*, the set of all group members' identities and signature key, and* $\texttt{id}_c$*, the identity of the committer of the add proposal.*

Key: $\texttt{K} \leftarrow \texttt{Key}$ *outputs the current application secret, i.e., the secrets that* `id` *derives for the current epoch.*

External Self-Add Proposal: $p \leftarrow (\texttt{Propose}, \texttt{extAdd-spk}, \texttt{epoch})$ *proposes to add the proposing member* `id` *with the signature key* `spk` *to the group. It outputs an external proposal* p*, or* $\perp$ *if the member is already in the group.*

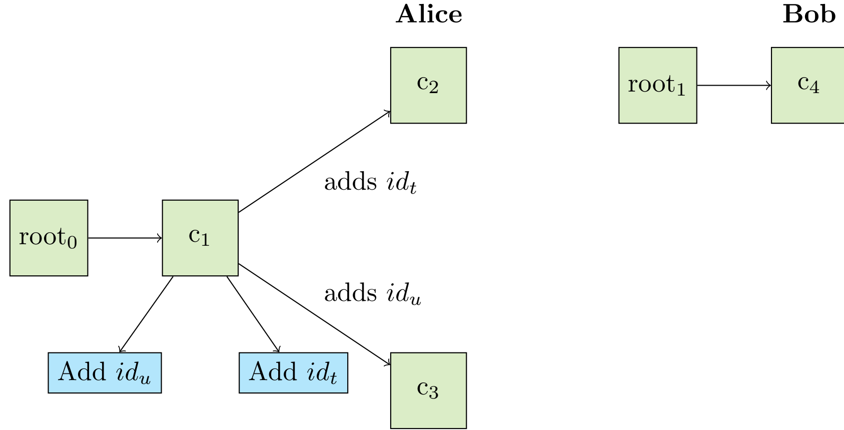

Figure 3: An example history graph similar to [4]. The green and blue boxes respectively denote commit and proposal nodes. In the epoch created by the commit $c1$, two proposals that respectively add parties $\mathtt{id}_u$ and $\mathtt{id}_t$ are generated and respectively committed by $c2$ and $c3$. Now the group splits and Alice chooses to process the commit $c2$. Additionally, Bob is invited to join the group via a welcome injected by the adversary, which does not correspond to any existing commit. The history graph hence creates a new detached node root_1, and appends a new node $c4$ to this root, to which Bob then joins.

PSK Proposals: $p \leftarrow (\mathtt{Propose}, \mathtt{psk\text{-}epoch\text{-}use})$ *proposes to inject the (internal) resumption PSK from an epoch into the key derivation schedule. It outputs a PSK proposal message p, or $\perp$ if the PSK is invalid.*

External Commit: $(c, g') \leftarrow (\mathtt{Commit}, g, \mathtt{spk}, e_{\vec{psk}}, \mathtt{resync})$ *creates an external commit using groupInfo g and injecting PSKs from a vector of epochs $e_{\vec{psk}}$ that adds the committer to the group and (optionally) removes an old representation of the member (if $\mathtt{resync} = \mathtt{true}$). The operations return the commit c and the group info g' of the next epoch for external joiners.*

4 Security Model and Functionality

We first define our extended notion of *history graph* in Section 4.1, and use this to present our extended generic functionality $\mathcal{F}_{\mathrm{ECGKA}}$ in Section 4.2. Afterwards, we highlight core differences with the functionality $\mathcal{F}_{\mathrm{CGKA}}$ [4] in Section 4.3.

4.1 History Graph

A history graph is a labeled directed graph that can symbolically represent a group's evolution. The concept was introduced in [1] and adapted in [4]. A history graph consists of two types of nodes: *commit* and *proposal nodes*, which respectively represent all executed commit and proposal operations. We show an example history graph in Figure 3. For our work, we define an extended notion of history graphs, in which every node stores the following variables:

- orig: the creator $\mathtt{id}$ (i.e., the proposal or commit sender) of this node.
- $\mathtt{stat} \in \{\mathtt{good}, \mathtt{bad}, \mathtt{adv}\}$: a status flag indicating whether the secrets in this node is secure (i.e., $\mathtt{stat} = \mathtt{good}$), leaked to the adversary (i.e., $\mathtt{stat} = \mathtt{bad}$), or injected by the adversary (i.e., $\mathtt{stat} = \mathtt{adv}$).
- pars: the list of parent commit nodes, representing the node's evolution history.

Every *proposal* node further stores the following variables:

- $\mathsf{act} \in \{\mathtt{up\text{-}spk}, \mathtt{add\text{-}id}_t\text{-}\mathtt{spk}_t, \mathtt{rem\text{-}id}_t, \mathtt{extAdd\text{-}spk}, \mathtt{psk\text{-}epoch\text{-}use}\}$: the proposed action with related information. We further explain these actions in Section 4.2.
- c_{psk}: specifies the commit, whose PSK secret this PSK proposal refers to.[2] This variable is only defined for PSK proposal $\mathtt{psk}$.

[2]If the adversary injects a PSK proposal on behalf of any party, then this PSK proposal might point to a node in a detached tree, where the commit itself and the resumption PSK value have not been fixed. In this case, we use the variable c_{psk} to denote the commit message that the PSK proposal refers to.

- psk: a status flag indicating whether the resumption PSK secrets in this node is secure (i.e., stat = good), leaked to the adversary (i.e., stat = bad), or injected by the adversary (i.e., stat = adv). This variable is only defined for PSK proposal psk.

Every *commit* node further stores the following values:

- pro: the ordered list of committed proposals. This variable is only defined for regular commits.
- mem: the list of group members and their signature public keys.
- sender_type $\in$ {'member', 'new_member'}: the role of sender that distinguishes regular (i.e., 'member') and external (i.e., 'new_member') commit nodes.
- ExtCommitProps: the ordered list of committed proposals. This variable is only defined for external commits.
- epoch: the sender's current epoch.
- has_psk: the list of group members who have derived the resumption PSK from the key schedule, see Figure 2.
- key: the group key.
- chall $\in$ {true, false}: a flag indicating whether the group key has been challenged (i.e., random, if chall = true) or not (i.e., chosen by the adversary or not generated yet, if chall = false).
- exp: a set that includes the corrupted parties in this node.

Comparison with the History Graph in [4]: Our history graph definition differs in three main ways from [4]. First, our graph incorporates external proposals that can be applied to all commit nodes with the same epoch. To this end, the pars variable stores a list of parent commit nodes rather than unique one, and every commit node additionally stores a new epoch variable.

Second, our history graph incorporates the external commit operation. To distinguish these node from the regular commits, our history graph includes a sender_type variable. While regular proposals can be generated and sent independently before they are committed, external commits and the included proposals are closely connected and sent together. To capture this, our external commit nodes have an additional ExtProps variable for all included proposals.

Third, our history graph can also be used to incorporate the resumption PSK proposal operation, which we will introduce in Section 5. To this end, our history graph includes a new PSK proposal node that has two additional variables: c_{psk} and psk, which respectively specify the commit whose PSK this proposal refers to, and the security of this PSK. Moreover, every commit node has an additional has_psk variable that stores the list of group members who have processed this commit and derived its resumption PSK.

4.2 Functionality $\mathcal{F}_{\text{ECGKA}}$

We give our new functionality $\mathcal{F}_{\text{ECGKA}}$ in Figures 4 and 5. Note that we define two functionalities in these figures: $\mathcal{F}_{\text{ECGKA}}$ is defined in these figures by omitting the orange (●) code. Similar to [4], our $\mathcal{F}_{\text{ECGKA}}$ is also equipped with two additional (ideal) functionalities $\mathcal{F}_{\mathcal{AS}}^{\mathcal{IW}}$ and $\mathcal{F}_{\mathcal{KS}}^{\mathcal{IW}}$ that capture the behaviors of authentication and key services (See Section 2.2). We formally describe $\mathcal{F}_{\mathcal{AS}}^{\mathcal{IW}}$, $\mathcal{F}_{\mathcal{KS}}^{\mathcal{IW}}$, and all helper functions in Appendix C.

Threat Model. Following [4], our model considers an adversary with the following strong capabilities:

(1) Fully control the network of message and key packages distribution, except for the long-time identity keys on the authentication service AS. This is captured by allowing the adversary (which is controlled by the environment $\mathcal{Z}$) to invoke each algorithm and to provide their inputs in $\mathcal{F}_{\text{ECGKA}}$, $\mathcal{F}_{\mathcal{AS}}^{\mathcal{IW}}$, $\mathcal{F}_{\mathcal{KS}}^{\mathcal{IW}}$, except for having the algorithm verify-cert directly return to the invoking party id.

(2) Adaptively corrupt the authentication and key service, i.e., obtain the secrets of any party id's identity or single-use key packages. This is captured by the expose algorithm with any id as input in $\mathcal{F}_{\mathcal{AS}}^{\mathcal{IW}}$ and $\mathcal{F}_{\mathcal{KS}}^{\mathcal{IW}}$.

(3) Adaptively corrupt any party's state through the Expose algorithm in $\mathcal{F}_{\text{ECGKA}}$.

(4) Adaptively control any party's randomness. This is captured by CorRand in $\mathcal{F}_{\text{ECGKA}}$ and corRand in $\mathcal{F}_{\mathcal{AS}}^{\mathcal{IW}}$ and $\mathcal{F}_{\mathcal{KS}}^{\mathcal{IW}}$ with $b = \texttt{bad}$ as input.

(5) Potentially collude with malicious group members and malicious servers. This is captured in all the above points.

$\mathcal{F}_{\text{ECGKA}}$ restricts the corruption in the following scenarios: once an epoch is "created" securely, and hence the group secret is generated randomly by $\mathcal{F}_{\text{ECGKA}}$, a corruption of this epoch is no longer possible, similar to [4].

Security Guarantees. $\mathcal{F}_{\text{ECGKA}}$ captures three core security guarantees: *consistency (and correctness)*, *confidentiality*, and *authenticity*.

Consistency. Consistency ensures that all parties in the same epoch agree on the group state, including the group membership and the history of group evolution (i.e., proposals and commits). We capture this by checking **assert** **cons-invariant** at the end of every algorithm that causes an incremented epoch, i.e., `Commit`, `Process`, `Join`, and `ExtCommit`. We depict our **cons-invariant** in Figure 6. Consistency holds (i.e., **cons-invariant** = `true`) if and only if all the following conditions hold:

a) every (non-detached) commit c and all included proposals must be generated at the same parent commit node. This ensures that all committed proposals must be applied to the same group state.

b) every party `id` that can enter a new epoch created by any commit c only if `id` belongs to the membership specified by c. This ensures that all group members at the same epoch must agree on the same group membership.

c) the history graph contains no cycles. This ensures that the epoch always moves forward and never reverts.

d) every external self-add proposal must have the same epoch as all its parent commit nodes. This ensures that external self-add proposals will be only applied to a group with the epoch that this proposal aims at.

Confidentiality. Confidentiality means that the environment must have no information about the important secrets (which we call *application secrets*) in any "safe" epoch, evaluated by a predicate **safe**. We capture this by deriving the application secrets in the "safe" epochs randomly, while having the simulator $\mathcal{S}$ choose them in other "unsafe" epochs. We depict the **safe** predicate in Figure 7. We consider that an epoch of a commit c is safe if and only if the epoch secret can be neither directly be exposed, indicated by $(*, \texttt{true}) \in \texttt{Node}[c].\mathsf{exp}$, nor indirectly leaked through other secrets, indicated by ***can-traverse**(c). Our ***can-traverse**(c) considers five cases.

Case a) the regular commit node c is an orphan root with a corrupted signature key. Identical to [4], an orphan root indicates a commit node that is injected by the adversary either directly via a commit message, or indirectly via a welcome message. In this case, a regular commit c is secure unless the adversary can inject any commit or welcome message with any group member's corrupted signature key.

Case b) a regular commit c adds a corrupted member. The only difference is that our definition also considers a corrupted member that is added by an external self-add proposal. In this case, a regular commit c is secure unless no added signer in this epoch is corrupted.

Case c) the secrets encrypted in the welcome message are exposed by corrupting an added member in the future. The only difference is that our definition additionally considers the welcome message for adding new group members via external self-add proposal. In this case, a regular commit c is secure unless a member `id` is added in this epoch and corrupted in the future, without doing any update in-between.

Case d) the adversary has corrupted sufficient state secrets from members' states to recover the epoch secret.

Case e) the external commit is injected by the adversary or involves bad randomness. In this case, an external commit c is secure, unless the committer's randomness or the commit itself is chosen by the adversary.

Authenticity. Authenticity means that the environment cannot forge any message on behalf of any group member if the predicate **inj-allowed** is false. We consider the following four cases, corresponding to different message types:

a) Every regular commit c can be injected by the adversary on behalf of any group member `id` with its valid leaf key `spk` only if both `spk` and the epoch secret at the parent node $c_p = \texttt{Node}[c].\mathsf{pars}$ has been exposed.

Functionality $\mathcal{F}_{\text{ECGKA}}$ and $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ - Part 1

The functionality expects as part of the instance's session identifier **sid** the group creator's identity $\text{id}_{\text{creator}}$. It is parameterized in the predicates **safe**(c), specifying which keys are confidential.

```
Initialization
  Ptr[*], Node[*], Prop[*], Wel[*] ← ⊥
  RndCor[*] ← good; HasKey[*] ← false
  rootCtr ← 0
  GroupInfo[*] ← ⊥ // groupInfo object for external commits, indexed by
  groupInfo g
  ExtProps[*] ← ⊥ // external self-add, indexed by p

Inputs from id_creator
Input (Create, spk)
  req Node[root_0] = ⊥ ∧ *valid-spk(id_creator, spk)
  mem ← {(id_creator, spk)}
  Node[root_0] ← *create-root(id_creator, mem,
        RndCor[id_creator], epoch = 0)
  HasKey[id_creator] ← true; Ptr[id_creator] ← root_0

Inputs from a party id

Input (Propose, act), act ∈ {up-spk, add-id_t, rem-id_t, psk-epoch-use}
  req Ptr[id] ≠ ⊥
  Send (Propose, id, act) to the adversary and
       receive (p, spk_t, ack).
  if ¬*req-correctness('prop', id, act) then req ack
  if act = up-spk then assert *valid-spk(id, spk)
  if act = add-id_t then act ← add-id_t-spk_t
  if Prop[p] = ⊥ then
    Prop[p] ← *create-prop(Ptr[id], id, act, RndCor[id])
  else
    *consistent-prop(p, id, act)
  if RndCor[id] = bad then
    Send (exposed, id, spk) to F_AS.
  return p

Input (Propose, extAdd-id-spk, epoch)
  req Ptr[id] = ⊥
  Send (Propose, id, extAdd-spk, epoch) to the adversary and receive
  (p, ack).
  if ¬*req-correctness('ext-prop', id, extAdd-spk) then req ack
  assert *valid-spk(id, spk)
  act ← extAdd-id-spk
  if Prop[p] = ⊥ then
    Prop[p] ← *create-ext-prop(id, act, epoch,
          RndCor[id])
    ExtProps[p] ← new node with epoch ← epoch
  else
    *consistent-ext-prop(p, id, act, epoch)
  if RndCor[id] = bad then
    Send (exposed, id, spk) to F_AS.
  return p

Input (Commit, p⃗, spk, force-rekey, wel_type)
  req Ptr[id] ≠ ⊥
  Send (Commit, id, p⃗, spk, force-rekey, wel_type) to the adversary
      and receive (ack, c, vec_w, g, rt, c⃗_psk).
  if ¬*req-correctness('comm', id, p⃗, spk, force-rekey) then req ack
  *fill-props(id, p⃗)
  res = *fixate-psk-refs((Prop[p] : p ∈ p⃗), c⃗_psk)
  assert res ≠ ⊥
  if *keep-spk(p⃗, force-rekey) then
    spk ← Node[Ptr[id]].mem[id]
  assert *valid-spk(id, spk)
  mem ← *members(Ptr[id], id, [Prop[p] : p ∈ p⃗], spk)
  assert mem ≠ ⊥ ∧ (id, spk) ∈ mem
  for p ∈ p⃗ s.t. Prop[p].act = psk-* do
    req id ∈ Node[Prop[p].c_psk].has_psk
  if Node[c] = ⊥ ∧ rt = ⊥ then
    if *keep-spk(p⃗, force-rekey) then
      stat ← bad
    else stat ← RndCor[id]
    Node[c] ← *create-child(Ptr[id], id, p⃗, mem, stat)
  else
    if rt ≠ ⊥ then
      c' ← root_rt
      assert RndCor[id] ≠ good
    else
      c' ← c
      if ¬*keep-spk(p⃗, force-rekey) then
        assert RndCor[id] ≠ good
    assert Node[c'].*origRChild = id
    assert *valid-successor(c', id, p⃗, mem)
    if c ≠ c' then *attach(c, c', id, p⃗)
  assert vec_w ≠ ⊥ iff ∃p ∈ p⃗ : Node[p].act ∈ {add-*, extAdd-*}
  if vec_w ≠ ⊥ then
    c⃗_psk = Prop[p].c_psk for all p ∈ p⃗ s.t. Prop[p].act = psk-*
    assert Wel[vec_w] ∈ {⊥, (c, c⃗_psk)}
    Wel[vec_w] ← (c, c⃗_psk)
  assert g ≠ ⊥
  assert groupInfo[g] ∈ {⊥, c}
  groupInfo[g] ← c
  assert cons-invariant ∧ auth-invariant
  if RndCor[id] = bad then
    Send (exposed, id, Node[Ptr[id]].mem[id]) to F_AS.
  return (c, vec_w, g)

Input Key
  req Ptr[id] ≠ ⊥ ∧ HasKey[id]
  if Node[Ptr[id]].key = ⊥ then *set-key(Ptr[id])
  HasKey[id] ← false
  return Node[Ptr[id]].key
```

Figure 4: Part 1 of our $\mathcal{F}_{\text{ECGKA}}$ and $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ definition, using the color coding from Table 2. The orange code (●) is included in $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$, but not in $\mathcal{F}_{\text{ECGKA}}$.

Functionality $\mathcal{F}_{\text{ECGKA}}$ and $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ - Part 2

```
Input (ExtCommit, g, spk, e_psk→, resync)
  if resync = true then req Ptr[id] ≠ ⊥
  else req Ptr[id] = ⊥
  Send (ExtCommit, id, g, spk, e_psk→, resync) to the adversary
      and receive (ack, c, g', p_exi, p_rem, p_psk→,
            c_gm, orig_g, mem_g, epoch, c_psk→).
  if ¬*req-correctness('ext-comm', id, g, spk
            resync = true) then req ack
  assert *valid-spk(id, spk)
  c_g ← groupInfo[g]
  if c_g = ⊥ then
    if Node[c_gm] ≠ ⊥ then
      c_g ← c_gm
    else
      rootCtr++; c ← root_rootCtr
      Node[c_g] ← *create-root(orig_g, mem_g, 'adv', epoch)
    groupInfo[g] ← c_g
  ExtCommitProps[p_exi] += *create-ext-commit-props(
            groupInfo[g], 'exi-id', id)
  if p_rem ≠ ⊥ then
    ExtCommitProps[p_rem] +=
         *create-ext-commit-props(groupInfo[g],
            'rem-id', id)
  for e, p in zip(e_psk→, p_psk→) do
    ExtCommitProps[p_psk] +=
         *create-ext-commit-props(groupInfo[g],
            'psk-e-application', id)
  try ExtCommitProps =
            *fixate-psk-refs(ExtCommitProps, c_psk→)
  for prop ∈ ExtCommitProps s.t. prop.act = psk-* do
    req id ∈ Node[prop.c_psk].has_psk
  p→ = [p_exi, p_rem, p_psk]
  mem ← *members(groupInfo[g], id, ExtCommitProps, spk)
  assert mem ≠ ⊥ ∧ (id, spk) ∈ mem
  if Node[c] = ⊥ then
    stat ← RndCor[id]
    Node[c] ← *create-child(groupInfo[g], id, p→, mem, stat, 'new_member',
    ExtCommitProps)
  else
    *consistent-ext-comm(c, g, id, p, mem, ExtCommitProps)
  assert g' ≠ ⊥
  assert groupInfo[g'] ∈ {⊥, c}
  groupInfo[g'] ← c
  assert cons-invariant ∧ auth-invariant
  if RndCor[id] = bad then
    Send (exposed, id, Node[Ptr[id]].mem[id]) to F_AS.
  return (c, g)

Corruptions
Input (Expose, id)
  if Ptr[id] ≠ ⊥ then
    Node[Ptr[id]].exp +← (id, HasKey[id])
    *update-stat-after-exp(id)
    Send (exposed, id, Node[Ptr[id]].mem[id]) to F_AS.
  Send (get-sk) to F_KS and receive SK and SPK.
  for each kp s.t. SK[id, kp] ≠ ⊥ ∧ SPK[id, kp] = spk do
    for each c s.t. ∃p ∈ Node[c].pro :
                    Prop[p].act = add-id-spk do
      Node[c].exp +← (id, true)
  This input is disallowed if ∃c : Node[c].chall ∧ ¬safe(c)

Input (CorrRand, id, b), b ∈ {good, bad}
  RndCor[id] ← b

Input (Process, c, p→)
  if p→ = ⊥ then return *process-ec(c)
  Send (Process, id, c, p→) to the adversary and
      receive (ack, rt, orig', spk', c_psk→).
  if ¬*req-correctness('proc', id, c, p→) then
    req ack
  *fill-props(id, p→)
  try res = *fixate-psk-refs(Prop[p] : p ∈ p→, c_psk→)
  if Node[c] = ⊥ ∧ rt = ⊥ then
    mem ← *members(Ptr[id], orig', p→, spk')
    assert mem ≠ ⊥
    for p ∈ p→ s.t. Prop[p].act = psk-* do
      req id ∈ Node[Prop[p].c_psk].has_psk ∨ ∄(id, *) ∈ mem
    Node[c] ← *create-child(Ptr[id], orig', p→, mem, adv)
  else
    if Node[c] = ⊥ then c' ← root_rt
    else c' ← c
    id_c ← Node[c'].orig; spk_c ← Node[c'].mem[id_c]
    mem ← *members(Ptr[id], id_c, p→, spk_c)
    assert mem ≠ ⊥
    for p ∈ p→ s.t. Prop[p].act = psk-* do
      req id ∈ Node[Prop[p].c_psk].has_psk ∨ ∄(id, *) ∈ mem
    *valid-successor(c', id, p→, mem)
    if c ≠ c' then *attach(c, c', id, p→)
  if ∃p ∈ p→ : Prop[p].act = rem-id then
    Ptr[id] ← ⊥
  else
    assert id ∈ Node[c].mem
    Ptr[id] ← c; HasKey[id] ← true
    Node[c].has_psk +← id
  assert cons-invariant ∧ auth-invariant
  return *output-proc(c)

Input (Join, vec_w)
  Send (Join, id, vec_w) to the adversary and
      receive (ack, c', g, orig', mem', epoch, c_psk→').
  req ack
  (c, c_psk→) ← Wel[vec_w]
  if c = ⊥ then
    if Node[c'] ≠ ⊥ then
      c ← c'; c_psk→ ← c_psk→'
    else
      rootCtr++; c ← root_rootCtr; c_psk→ ← c_psk→'
      Node[c] ← *create-root(orig', mem', adv, epoch)
    Wel[vec_w] ← (c, c_psk→)
  for p ∈ Node[c].pro s.t. Prop[p].act = psk-* do
    req id ∈ Node[Prop[p].c_psk].has_psk
  Ptr[id] ← c
  HasKey[id] ← true
  Node[c].has_psk +← id
  assert g ≠ ⊥
  assert groupInfo[g] ∈ {⊥, c}
  groupInfo[g] ← c
  assert id ∈ Node[c].mem ∧ cons-invariant
        ∧ auth-invariant
  return *output-join(c)
```

Figure 5: Part 2 of our $\mathcal{F}_{\text{ECGKA}}$ and $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ definition, using the color coding from Table 2. The orange code (●) is included in $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$, but not in $\mathcal{F}_{\text{ECGKA}}$.

Functionality $\mathcal{F}_{\text{ECGKA}}$ and $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$: Invariants

// The history graph is consistent.

helper cons-invariant

return true iff

a) $\forall c$ s.t. $\texttt{Node}[c].\texttt{pars} \neq \bot$: $\texttt{Node}[c].\texttt{pro} \neq \bot$ and $\forall p \in \texttt{Node}[c].\texttt{pro}$: $\texttt{Prop}[p].\texttt{pars} = \texttt{Node}[c].\texttt{pars}$, and

b) $\forall \texttt{id}$ s.t. $\texttt{Ptr}[\texttt{id}] \neq \bot$: $\texttt{id} \in \texttt{Node}[\texttt{Ptr}[\texttt{id}]].\texttt{mem}$, and

c) the graph contains no cycles, and

d) $\forall ep$ s.t. $\texttt{ExtProps}[ep] \neq \bot$: $(\forall c \in \texttt{Prop}[ep].\texttt{par}$: $\texttt{Node}[c].\texttt{epoch} = \texttt{ExtProps}[ep].\texttt{epoch})$

// No injections when authenticity guaranteed.

helper auth-invariant

return true iff

a) $\forall c$ with $\texttt{Node}[c].\texttt{sender_type} = \text{'member'}$ and $c_p = \texttt{Node}[c].\texttt{pars} \neq \bot$ and $\texttt{id} = \texttt{Node}[c].\texttt{orig}$ and $(\texttt{id}, \texttt{spk}) \in \texttt{Node}[c_p].\texttt{mem}$, if $\texttt{Node}[c].\texttt{stat} = \texttt{adv}$ then **inj-allowed**$(c_p, \texttt{spk})$, and

b) $\forall p$ with $p \in \texttt{Prop} \setminus \texttt{ExtProps}$ and $c_p = \texttt{Prop}[p].\texttt{pars} \neq \bot$ and $\texttt{id} = \texttt{Prop}[p].\texttt{orig}$ and $(\texttt{id}, \texttt{spk}) \in \texttt{Node}[c_p].\texttt{mem}$, if $\texttt{Prop}[p].\texttt{stat} = \texttt{adv}$ then **inj-allowed**$(c_p, \texttt{spk})$, and

c) $\forall c$ with $\texttt{Node}[c].\texttt{sender_type} = \text{'new_member'}$ and $\texttt{id} = \texttt{Node}[c].\texttt{orig}$ and $(\texttt{id}, \texttt{spk}) \in \texttt{Node}[c].\texttt{mem}$, if $\texttt{Node}[c].\texttt{stat} = \texttt{adv}$ then **inj-allowed**$(\bot, \texttt{spk})$, and

d) $\forall p$ with $p \in \texttt{ExtProps}$ and $\texttt{id} = \texttt{Prop}[p].\texttt{orig}$ and $\texttt{ExtProps}[p].\texttt{act} = \texttt{extAdd-id-spk}$, if $\texttt{Prop}[p].\texttt{stat} = \texttt{adv}$ then **inj-allowed**$(\bot, \texttt{spk})$.

Figure 6: The cons- and auth-invariant predicates for the $\mathcal{F}_{\text{ECGKA}}$ and $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ functionalities, adapted from [4] using the color coding from Table 2. The codes highlighted with ● color are only included in $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$.

b) Every regular proposal p can be injected by the adversary on behalf of any group member `id` with its valid leaf key `spk` only if both `spk` and the epoch secret at the parent node $c_p = \texttt{Prop}[p].\texttt{pars}$ has been exposed.

c) Every external commit c can be injected by the adversary on behalf of any group member `id` with its valid leaf key `spk` only if `spk` has been exposed.

d) Every external proposal p can be injected by the adversary on behalf of any group member `id` with its valid leaf key `spk` only if `spk` has been exposed.

Interfaces in $\mathcal{F}_{\text{ECGKA}}$. The functionality $\mathcal{F}_{\text{ECGKA}}$ is initialized with an empty history graph, i.e., all proposal and commit nodes $\texttt{Prop}[p]$ and $\texttt{Node}[c]$ are set to $\bot$ and the root is identified by the label $\texttt{root}_0$. For each party `id`, each vector of welcome messages `vec_w`, and each groupInfo g, $\mathcal{F}_{\text{ECGKA}}$ respectively stores a pointer $\texttt{Ptr}[\texttt{id}]$, $\texttt{Wel}[\texttt{vec_w}]$, $\texttt{GroupInfo}[g]$, to its corresponding commit node in the history graph, initialized with $\bot$. Moreover, for each external self-add proposal p, $\mathcal{F}_{\text{ECGKA}}$ also uses a helper node $\texttt{ExtProps}[p]$ to store the `epoch` that p will apply to. Furthermore, $\mathcal{F}_{\text{ECGKA}}$ uses $\text{RndCor}[\texttt{id}]$ to record whether a party `id`'s randomness source is indeed random (i.e., `good`) or chosen by the adversary (i.e., `bad`), and $\text{HasKey}[\texttt{id}]$ to specify whether the current application secret of `id` can be challenged (i.e., `true`) or not (i.e., `false`).

Below, we explain the interface of external self-add proposal and external commit in $\mathcal{F}_{\text{ECGKA}}$ in detail, and briefly introduce other interfaces that have been studied in [4]. We detail the interfaces of the remaining interfaces in Appendix C.

Interface `Create`*:* This interface $(\texttt{Create}, \texttt{spk})$ models group creation and is invoked only if the node of initial root is not initialized and the public key `spk` of the invoker $\texttt{id}_{\texttt{creator}}$ is valid. Then, the initial root of the history graph is initialized with the only member $\texttt{id}_{\texttt{creator}}$ with its public key `spk`. Finally, $\texttt{Ptr}[\texttt{id}_{\texttt{creator}}]$ is set to the initial root and $\texttt{id}_{\texttt{creator}}$'s application key can be challenged.

Interface `Propose` *(External):* This interface of the form $(\texttt{Propose}, \texttt{extAdd-spk}, \texttt{epoch})$ considers the new operation external self-add proposal. The input `epoch` specifies the epoch of a group that this proposal is intended for. The invoker `id` must not be yet a member of the group. Because the proposal comes from outside the group, it is only signed and does not employ any group internal secret, and there is no indication of the commit node of the history graph p is intended for.

$\mathcal{F}_{\text{ECGKA}}$ forwards this query to the adversary $\mathcal{A}$ and receives a proposal p and an acknowledgment *ack*. $\mathcal{F}_{\text{ECGKA}}$ checks whether the proposal is correct. Otherwise, it must be maliciously injected by the adversary (i.e., $ack = \texttt{true}$). Then, $\mathcal{F}_{\text{ECGKA}}$ checks whether the input `spk` is valid. If the node $\texttt{Prop}[p]$ of the proposal p does not exist, then $\mathcal{F}_{\text{ECGKA}}$ creates it using input of this query. More concretely, the list of parent node $\texttt{Prop}[p].\texttt{pars}$ is set to all existing commit nodes at same epoch as the input `epoch`. Recall that an external self-add proposal is sent to all commit nodes at the epoch `epoch`, including both existing and future ones. If a commit node at epoch `epoch` is created in the future, it will also be added into $\texttt{Prop}[p].\texttt{pars}$. The origin $\texttt{Prop}[p].\texttt{orig}$ is set to the invoker `id`, the action $\texttt{Prop}[p].\texttt{act}$ is set to the input

Functionality $\mathcal{F}_{\text{ECGKA}}$ and $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$: **Predicate safe**

Knowledge of state secrets.

know$(c, \texttt{id}) \iff$

a) // id's state leaks directly e.g. via corruption (see below):
$\ast$**state-directly-leaks**$(c, \texttt{id}) \lor$

b) // know state in the parent:
$(\texttt{Node}[c].\text{pars} \neq \bot \land \neg\ast\textbf{secrets-replaced}(c, \texttt{id}) \land \textbf{know}(\texttt{Node}[c].\text{pars}, \texttt{id})) \lor$

c) // know state in a child:
$(\exists c' : \texttt{Node}[c'].\text{pars} = c \land \neg\ast\textbf{secrets-replaced}(c', \texttt{id}) \land \textbf{know}(c', \texttt{id}))$

$\ast$**state-directly-leaks**$(c, \texttt{id}) \iff$

a) // id has been exposed in c:
$(\texttt{id}, \ast) \in \texttt{Node}[c].\text{exp} \lor$

b) // c is in a detached tree and id's spk is exposed
$\Big((\exists rt : \ast\text{ancestor}(\text{root}_{rt}, c)) \land (\exists \text{spk} : (\texttt{id}, \text{spk}) \in \texttt{Node}[c].\text{mem} \land \text{spk} \in \text{Exposed})\Big) \lor$

c) // id's secrets in c are injected by the adversary:
$((\texttt{id}, \text{spk}) \in \texttt{Node}[c].\text{mem} \land \ast\textbf{secrets-injected}(c, \texttt{id}))$

$\ast$**secrets-injected**$(c, \texttt{id}) \iff$

a) // id is the sender of c and c was injected or generated with bad randomness
$(\texttt{Node}[c].\text{orig} = \texttt{id} \land \texttt{Node}[c].\text{stat} \neq \text{good}) \lor$

b) // c commits an update of id that is injected or generated with bad randomness
$(\exists p \in \texttt{Node}[c].\text{pro} : \texttt{Prop}[p].\text{act} = \text{up-}\ast \land \texttt{Prop}[p].\text{orig} = \texttt{id} \land \texttt{Prop}[p].\text{stat} \neq \text{good}) \lor$

c) // c adds id with corrupted spk
$(\exists p \in \texttt{Node}[c].\text{pro} : \texttt{Prop}[p].\text{act} \in \{\text{add-id-spk}, \texttt{extAdd-id-spk}\} \land \text{spk} \in \text{Exposed})$

$\ast$**secrets-replaced**$(c, \texttt{id}) \iff$

a) // id is the sender of c and c includes an empty or add-(and-psk)-only proposal list
$\Big(\texttt{Node}[c].\text{orig} = \texttt{id} \land ((\exists p \in \texttt{Node}[c].\text{pro} : \texttt{Prop}[p].\text{act} \notin \{\text{add-}\ast, \texttt{extAdd-}\ast, \texttt{psk-}\ast\}) \lor \texttt{Node}[c].\text{pro} = \varnothing)\Big) \lor$

b) // c adds id, removes id, or commits and update of id
$(\exists p \in \texttt{Node}[c].\text{pro} : \texttt{Prop}[p].\text{act} \in \{\text{add-id-}\ast, \text{rem-id}\} \lor (\texttt{Prop}[p].\text{act} = \text{up-}\ast \land \texttt{Prop}[p].\text{orig} = \texttt{id}))$

Knowledge of epoch secrets.

know$(c, \text{'epoch'}) \iff \texttt{Node}[c].\text{exp} \neq \varnothing \lor \ast\textbf{can-traverse}(c)$

// Can the adversary process c using exposed individual secrets and parent's init secret?

$\ast$**can-traverse**$(c) \iff$

a) // orphan root with a corrupted signature public key:
$(\texttt{Node}[c].\text{pars} = \bot \land (\ast, \text{spk}) \in \texttt{Node}[c].\text{mem} \land \text{spk} \in \text{Exposed}) \lor$

b) // commit to an add proposal that uses an exposed signature public key:
$(\exists p \in \texttt{Node}[c].\text{pro} : \texttt{Prop}[p].\text{act} \in \{\text{add-id-spk}, \texttt{extAdd-id-spk}\} \land \text{spk} \in \text{Exposed}) \lor$

c) // secrets encrypted in the welcome message under an exposed secret, i.e., exposed init secret and psk
$\Big(\exists \texttt{id}, p \in \texttt{Node}[c].\text{pro} : \texttt{Prop}[p].\text{act} \in \{\text{add-id-}\ast, \texttt{extAdd-id-}\ast\} \land (\exists c_d : \ast\text{ancestor}(c, c_d) \land (\texttt{id}, \ast) \in$
$\texttt{Node}[c_d].\text{exp} \land$ no node c_h with $\ast\textbf{secrets-replaced}(c_h, \texttt{id})$ on c-c_d path$)\Big) \lor$

d) // know necessary info to traverse the edge:
$\Big(\textbf{know}(c, \ast) \land (c = \text{root}_\ast \lor \textbf{know}(\texttt{Node}[c].\text{pars}, \text{'epoch'})) \land \texttt{Node}[c].\texttt{psk} \neq \text{good}\Big) \lor$

e) // external commits injected by adversary or with bad randomness.
$(\texttt{Node}[c].\texttt{sender_type} = \texttt{new_member} \land \texttt{Node}[c].\text{stat} \neq \text{good} \land \texttt{Node}[c].\texttt{psk} \neq \text{good})$

safe$(c) \iff \neg((\ast, \texttt{true}) \in \texttt{Node}[c].\text{exp} \lor \ast\textbf{can-traverse}(c))$

inj-allowed$(c, \text{spk}) \iff \text{spk} \in \text{Exposed} \land (c = \bot \lor \textbf{know}(c, \text{'epoch'}))$

Figure 7: The safe predicate for the $\mathcal{F}_{\text{ECGKA}}$ and $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ functionalities, adapted from [4] using the color coding from Table 2. The codes highlighted with ● color are only included in $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$.

action, the status $\texttt{Prop}[p].\texttt{stat}$ is set to the randomness flag $\text{RndCor}[\texttt{id}]$. Moreover, $\mathcal{F}_{\text{ECGKA}}$ also stores the input $\texttt{epoch}$ in the helper node $\texttt{ExtProps}[p]$.

Otherwise, the existing node $\texttt{Prop}[p]$ must record $\texttt{id}$ as the origin and the same action $\texttt{act}$. All parent commit nodes of $\texttt{Prop}[p]$ must have the same epoch as the one included in the query input. If the randomness in this query is chosen by the adversary, then $\mathcal{F}_{\text{ECGKA}}$ exposes secret of the invoker $\texttt{id}$'s public key. This captures the fact that the ECDSA signature underlying MLS is vulnerable against randomness leakage attack [5]. In the end, the proposal p is returned.

Interface `ExtCommit`*:* This interface of the form $(\texttt{ExtCommit}, g, \textsf{spk}, \texttt{resync})$ considers the external commit operation. If the aim of this external is to resync, i.e., leave and rejoin the group, indicated by $\texttt{resync} = \texttt{true}$, then the party $\texttt{id}$ must be in the group (i.e., $\textsf{Ptr}[\texttt{id}] \neq \bot$). Otherwise, the party $\texttt{id}$ must be outside the group (i.e., $\textsf{Ptr}[\texttt{id}] = \bot$). Next, $\mathcal{F}_{\text{ECGKA}}$ forwards this query to the adversary and receives an acknowledgment ack, an external commit c, a new groupInfo g', a self-add proposal p_{exi}[3], a self-remove proposal p_{rem}, a group-match commit c_{gm}, the groupInfo g's origin $\textsf{orig}_g$, and a group member list $\textsf{mem}_g$. Then, $\mathcal{F}_{\text{ECGKA}}$ checks whether this query is related to a correct rejoin operation, or acknowledged by the adversary. The input $\textsf{spk}$ must be valid. Afterwards, $\mathcal{F}_{\text{ECGKA}}$ checks whether the groupInfo g maps to an existing commit. If not, then the groupInfo g must be injected by the environment and this external commit query is also considered to be injected. In this case, $\mathcal{F}_{\text{ECGKA}}$ maps the groupInfo g to c_{gm} if its commit node has been created or a new detached root. Within this interface, $\mathcal{F}_{\text{ECGKA}}$ does not create independent proposal nodes for every included proposal. Instead, $\mathcal{F}_{\text{ECGKA}}$ creates "sub-nodes" for the self-add proposal p_{exi} and the self-remove proposal p_{rem} (if available, i.e., $p_{rem} \neq \bot$). This external commit must have a correct member list, i.e, including the invoker $\texttt{id}$ and its $\textsf{spk}$. In the case that the node of the external commit c has not been created, $\mathcal{F}_{\text{ECGKA}}$ creates it, which wraps all above "sub-nodes". Otherwise, $\mathcal{F}_{\text{ECGKA}}$ checks whether the existing node $\texttt{Node}[c]$ is consistent with this external commit and the included proposals. Finally, similar to the regular commit, $\mathcal{F}_{\text{ECGKA}}$ requires that the new groupInfo g' must be non-$\bot$ and point to c. If the consistency and authenticity are not violated, then the commit c and the new groupInfo g' are returned. Note that the external commit only adds the invoker $\texttt{id}$ to the group. This query will not yield any welcome message.

Interface `Process` *(regular commit):* When invoked with a vector of proposals $\vec{p} \neq \bot$, this interface $(\texttt{Process}, c, \vec{p})$ processes a regular commit c and a vector of proposals $\vec{p}$. $\mathcal{F}_{\text{ECGKA}}$ forwards the query to the adversary and processes the regular and external commits differently. If the commit is regular ($\vec{p} \neq \bot$), then $\mathcal{F}_{\text{ECGKA}}$ should receive an acknowledgment ack, a detached root index $\textsf{rootCtr}$, the commit c's origin $\textsf{orig}'$, and its public key $\textsf{spk}'$. $\mathcal{F}_{\text{ECGKA}}$ first checks whether the commit and proposals are correct, or acknowledged by the adversary. Next, $\mathcal{F}_{\text{ECGKA}}$ creates proposal nodes for all injected proposals. If neither the node of the commit c nor the detached root $\textsf{root}_{rt}$ exists, then c must be injected by the adversary. In this case, $\mathcal{F}_{\text{ECGKA}}$ creates the node of c if c will yield a valid member list. Otherwise, either the node of the commit c or the detached root $\textsf{root}_{rt}$ exists. In this case, $\mathcal{F}_{\text{ECGKA}}$ attaches the commit node c to them if the yielding member list and this attachment are valid. If any of the input proposal indicates to remove the invoker $\texttt{id}$ from the group, then the point of $\texttt{id}$ will be set to $\bot$. Otherwise, $\mathcal{F}_{\text{ECGKA}}$ checks whether $\texttt{id}$ is still included in the member list of the node of the commit c, followed by moving the pointer of $\texttt{id}$ to c and mark $\texttt{id}$'s key to $\texttt{true}$ so that it can be challenged. If the consistency and authenticity checks are not violated, then information about the processed commit and its included proposals are returned, i.e., the origins of commit c, the new member list after processing c, and the origins and actions of each proposal in $\vec{p}$.

Interface `Process` *(external commit):* When invoked with $\vec{p} = \bot$, this interface of the form $(\texttt{Process}, c, \vec{p})$ processes an external commit c through ***process-ec**(c), defined in Appendix C. $\mathcal{F}_{\text{ECGKA}}$ should receive an acknowledgment ack, the external commit c's origin $\textsf{orig}'$ and its public key $\textsf{spk}'$, and a vector of proposals $\vec{p}$ that are applied in c. $\mathcal{F}_{\text{ECGKA}}$ first checks whether the commit and proposals are correct, or acknowledged by the adversary. If it holds that $\textsf{Ptr}[\texttt{id}] \neq \bot$, then the invoker $\texttt{id}$ must be a member inside the group. In this case, $\mathcal{F}_{\text{ECGKA}}$ checks whether the node of commit c has been created. If it is not created, then this commit must be injected by the adversary. In this case, $\mathcal{F}_{\text{ECGKA}}$ forwards every proposal $p \in \vec{p}$ to the adversary and receives its action $\textsf{act}$. Next, $\mathcal{F}_{\text{ECGKA}}$ creates the "sub-node" of these proposals and checks whether their actions are valid, i.e., only add and remove proposals are allowed and the added or removed party must be the origin of the commit. Then, $\mathcal{F}_{\text{ECGKA}}$ checks whether the member list for this external commit is valid. If the check passes, $\mathcal{F}_{\text{ECGKA}}$ created the node of the external commit c that wraps above "sub-nodes". If the node of commit c has already been created, $\mathcal{F}_{\text{ECGKA}}$ simply checks

[3] exi is named after the "**Ex**ternal**I**nit" operation in [14, Section 12.1.6], which "is used by new members that want to join a group by using an external commit".

whether the member list for this external commit c is valid and whether the existing commit node is consistent with this external commit c and its included proposals $\vec{p}$. If $\mathsf{Ptr}[\mathtt{id}] = \perp$, then the invoker id must be outside the group. Note that a non-member party can only process external commits that were created by itself. $\mathcal{F}_{\mathrm{ECGKA}}$ simply checks whether the node of commit c has been honestly generated by the invoker id, i.e., not injected by the adversary. Finally, in all cases, $\mathcal{F}_{\mathrm{ECGKA}}$ checks whether id is still included in the member list of the node of the commit c, followed by moving the pointer of id to c and marking id's key to true so that it can be challenged. If the consistency and authenticity checks are not violated, then information about the processed commit c and its included proposals $\vec{p}$ are returned, i.e., the origins of commit c, the new member list after processing c, and the origins and actions of each proposal in $\vec{p}$.

4.3 Comparison between $\mathcal{F}_{\mathrm{CGKA}}$ [4] and our $\mathcal{F}_{\mathrm{ECGKA}}$

Compared to $\mathcal{F}_{\mathrm{CGKA}}$ in [4], the main difference is that our $\mathcal{F}_{\mathrm{ECGKA}}$ captures "External Joins" by including the external operations described in [14, Section 3.3]. External proposals allow a party outside the group to apply for joining the group. Recall that the external proposal node will be appended to all commit nodes of the same epoch. To capture this, $\mathcal{F}_{\mathrm{ECGKA}}$ includes the epoch information into the commit and external proposal nodes. The external commit allows non-members to directly join an "open" groups but also allows group members to resync with the group. To capture this, every commit (regular and external) must be generated together with groupInfo g that can be used by outsiders to generate external commits. Moreover, $\mathcal{F}_{\mathrm{ECGKA}}$ offers different security guarantees in three aspects.

Consistency. The consistency predicate **cons-invariant** in our $\mathcal{F}_{\mathrm{ECGKA}}$ additionally capture the consistency of external proposal nodes, i.e., whether every external proposal has the same epoch as all its parent commit nodes.

Confidentiality. The confidentiality predicate **safe** in our $\mathcal{F}_{\mathrm{ECGKA}}$ has the following three core differences: First, $\mathcal{F}_{\mathrm{ECGKA}}$ encodes the confidentiality requirement for external self-add proposals, similarly to regular add proposals. Second, $\mathcal{F}_{\mathrm{ECGKA}}$ encodes that no rekey path if the proposals included in a commit is either empty or add-only, see condition a) in **heals**. This captures the modifications in MLS since Draft 10 [9]. Third, $\mathcal{F}_{\mathrm{ECGKA}}$ encodes the confidentiality requirement for external commits, identical to normal commits.

Authenticity. The authentication predicate **auth-invariant** in our $\mathcal{F}_{\mathrm{ECGKA}}$ has two core differences: First, we redefine the **inj-allowed** helper function. While **inj-allowed** in $\mathcal{F}_{\mathrm{CGKA}}$ in [4] checks whether the adversary can inject a message on behalf of an id in the group, our **inj-allowed** checks whether an injection can be verified on id's public key spk, as the sender of an external proposal or commit might not be included in any group. Second, our $\mathcal{F}_{\mathrm{ECGKA}}$ encodes the authenticity requirement for the novel external proposals and commit operations, respectively in condition c) and d).

5 RFC 9420 and ETK: Extended Operations TreeKEM

In this section, we introduce our ETK protocol, which is a simplified version of MLS RFC 9420 [14], as an instance of the ECGKA protocol. Our ETK protocol extends ITK [4], which simplifies MLS Draft 12 [15], with the novel external proposal and commit operations as well as some selected updates from Draft 12 to RFC 9420 that have not been formally studied in the literature to the best of our knowledge. In Section 5.1, we present the differences of our ETK with ITK [4], which we recalled in Section 2.2. In Section 5.2, we prove that our ETK protocol securely realizes $\mathcal{F}_{\mathrm{ECGKA}}$. Due to the page limit, we summarize our simplification in Appendix B and provide details on our ETK protocol in Appendix D.

5.1 The ETK Protocol

Similar to [4], we abstract the associated authentication and key services of our ETK as functionalities $\mathcal{F}_{\mathcal{AS}}$ and $\mathcal{F}_{\mathcal{KS}}$, which we depict in Figures 12 and 13 in Appendix C. Compared with ITK [4] that abstracts MLS Draft 12 [15], our ETK incorporates the following updates.

Updated GroupInfo. While ITK outputs groupInfo only included in the welcome messages for parties that have already been invited via regular add proposal, our ETK outputs groupInfo whenever regular or external commits are generated, in order to allows non-member party to join the group by itself. This modification has been included in MLS since Draft 10 [9]. Moreover, while ITK includes both confTransHash and interimTransHash in the groupInfo, our ETK only includes confTransHash and

excludes $\mathsf{interimTransHash}$, since $\mathsf{confTransHash}$ itself is sufficient for the verification on the history of group evolution. This modification has also been included in MLS since Draft 10 [9].

Leaf Nodes and Key Packages. While in ITK the leaf nodes and key packages had the same data structure, in our ETK they have different data structures. More concretely, leaf nodes are now a data structure that can be either contained within a tree, or within a key package. In our ETK, we use a variable `leafNode_source` within leaf node to denote where it is included: key package, commit, or update. This modification has been included in MLS since Draft 13 [11]. Moreover, while ITK encrypts the group secrets in the welcome messages with the public key of the leaf node, our ETK includes a new "init key" in the key packages that is different from the key in the leaf node and uses this "init key" to encrypt the group secrets in the welcome messages. This modification has been included in MLS since Draft 13 [11]. Finally, while ITK refers to every leaf node via an allotLeaf function inputting the hash of their owner's key package, our ETK refers to leaf nodes simply via their leaf index in the tree. Consequently, our allotLeaf function now takes no reference anymore, but assigns the leftmost leaf, and returns the index of this leaf node. This modification is in MLS since Draft 15 [12].

Updated Welcome Messages. While in ITK the `groupInfo` was unencrypted in welcome messages, our ETK encrypts it with the key enclosed in the encrypted group secrets. This modification has been included in MLS since MLS Draft 9 [10]. Moreover, while in ITK outputs every commit along with a single welcome messages for all added parties, our ETK allows the committer to freely decide how many welcome messages to construct, indicated by `wel_type`, and outputs a vector `vec_w` that includes one or more welcome messages, each of which adds either all, or single, or batch parties. This is in MLS since Draft 17 [13].

New Operation: External Self-Add Proposal. Since Draft 15 [12], MLS includes the external self-add proposal, which allows a party to propose to add itself to a group. The proposer `id` must not be in the group. To generate the external self-add proposal, `id` first generates a key package that warps two data structures: a "leaf node" and an "init key", followed by singing the key package using identity ssk. Next, `id` generates the proposal message that includes the intended groupId, the intended epoch, proposal type "`new_member_proposal`", and the signed key package, followed by signing them together. In the end, the signed proposal message is output.

New Operation: External Commit. Since Draft 10 [9], MLS includes the external commit operation that enables an external user to join an "open" group asynchronously without asking any group member to add them. In this case, we have `resync` = `false`. The external commit operation can also be used by an existing group member to replace its prior appearance in the group with a new one. In this case, we have `resync` = `true`.

In both cases, the committer `id` first parses the latest group information and a signature from the input groupInfo object g. The latest group information includes (1) the identifier of the group groupId, (2) the current epoch number epoch, (3) the hash of (the public part) of the binary ratchet tree τ, (4) the transcript hashes confTransHash and interimTransHash, (5) a labeled left-balanced binary ratchet tree τ, (6) a confirmation tag `confTag`, (7) the groupInfo g's sender index `senderIdx`, and (8) a HPKE public key external-pub, the secret key of which is shared by all existing group members. The committer `id` verifies whether the signature can be verified using the sender `senderIdx`'s public key spk encoded in the tree τ. Additionally, `id` checks whether the public ratchet tree is valid.Then, the committer `id` derives the initial secret initSecret for the key schedule using the HPKE export method with external-pub and initializes a new group state for the next epoch. The committer `id` generates an `extInit` proposal that is tailored for external commit to add itself. If `resync` = `true`, then the committer additionally generates a self-removal proposal. Afterwards, similar to the regular commit, the committer `id` adds itself to the ratchet tree at the leftmost available leaf, optionally removes itself from the current group, rekeys path, signs the commit content C using its signature key $\gamma.\mathsf{ssk}$, computes new transcript hash, and runs key schedule. The key schedule derives a confirmation key `confKey` and a joiner secret `joinerSec`. In the end, the committer `id` generates a confirmation tag `confTag` for the new group state, an external commit message ec, and an associated external groupInfo g. A mapping associating the new group state with each pending commit issued by `id` is added to the mapping $\gamma.\mathsf{pendCom}$. The external commit message ec and the new groupInfo g are returned.

5.2 Security Results for ETK and $\mathcal{F}_{\mathrm{ECGKA}}$

Theorem 1. *Assume that* PKE *is* IND-CCA *secure and that* Sig *is* SUF-CMA *secure. The* ETK *protocol securely realizes* $(\mathcal{F}_{\mathcal{AS}}^{\mathcal{IW}}, \mathcal{F}_{\mathcal{KS}}^{\mathcal{IW}}, \mathcal{F}_{\mathrm{ECGKA}})$ *in the* $(\mathcal{F}_{\mathcal{AS}}, \mathcal{F}_{\mathcal{KS}}, \mathcal{G}_{\mathrm{RO}})$*-hybrid model, where* $\mathcal{F}_{\mathrm{ECGKA}}$ *uses the*

predicates **safe** *and* **inj-allowed** *from Figure 7 and calls to* HKDF.Expand*,* HKDF.Extract*, and* MAC *functions are replaced by calls to the global random oracle* $\mathcal{G}_{\mathrm{RO}}$*.*

We provide the full proof of Theorem 1 in Appendix E.

Sketch. Our proof follows [4] and defines four consecutive hybrids games from *H1* to *H4*. We prove that the real world (i.e., *H1*) and the ideal world (i.e., *H4*) are indistinguishable for all environments $\mathcal{Z}$ via intermediate steps.

In hybrid *H1*, a dummy functionality $\mathcal{F}_{dummy}$ forwards all in- and outputs through the simulator $\mathcal{S}$, which executes ETK. Hence, $\mathcal{F}_{dummy}$ encodes no security guarantees and *H1* is identical to the real world.

In hybrid *H2*, $\mathcal{F}_{dummy}$ is replaced by a functionality consisting of $\mathcal{F}_{\mathcal{AS}}^{\mathcal{IW}}$, $\mathcal{F}_{\mathcal{KS}}^{\mathcal{IW}}$ and a modified version of $\mathcal{F}_{\mathrm{ECGKA}}$, where **safe**$(c)$ = `false` and **inj-allowed**$(c, \mathtt{id})$ = `true` for all commits c and parties `id`. The functionality interacts with a trivial simulator that chooses all application secrets according to the protocol. For every input operation, we show that the outputs in ETK (i.e., $\mathcal{F}_{dummy}$) and $\mathcal{F}_{\mathrm{ECGKA}}$ of *H2* is the same. This particularly includes that if *H2* halts due to an **assert** statement, so would ETK. The indistinguishability between *H1* and *H2* proves consistency and correctness of ETK.

The hybrid *H3* is identical to *H2* except that the application secrets in safe epochs are random, i.e., the original **safe** is restored. Moreover, we also create a sequence of sub-hybrid games Hs_i, where the simulator $\mathcal{S}$ sets the secrets for the first i epochs even if they are safe, and all epoch secrets after that are set according to the modified version of $\mathcal{F}_{\mathrm{ECGKA}}$ of *H2*. It is easy to observe that Hs_0 is identical to *H3* and the final sub-hybrid Hs_n for some polynomial n is identical to *H2*. We prove that every two adjacent sub-hybrid games Hs_{i-1} and Hs_i are indistinguishable, and further *H2* and *H3*. The indistinguishability between *H2* and *H3* proves the confidentiality of ETK.

The hybrid *H4* is identical to *H3* except that the messages (proposals and commits) forgery on behalf of uncorrupted parties are excluded, i.e. the original **inj-allowed** is restored. Now, *H4* is identical to the ideal world. We prove that the environment $\mathcal{Z}$ cannot forge any Sig signature or MAC on behalf of any uncorrupted party in *H3* whenever the evaluation of corresponding **inj-allowed** is true except for negligible probability. The indistinguishability between *H3* and *H4* will prove the authenticity of ETK.

□

Remark 1. *Our Theorem 1 requires the underlying signature scheme to be* SUF-CMA*-secure. In particular,* SUF-CMA *is necessary for MLS RFC 9420 to provide commit message integrity against malicious parties, which propagates through the transcripts into the key schedule.*

We give a concrete example for a malicious insider. Suppose the signature scheme is not SUF-CMA *secure, e.g., ECDSA. Consider that Alice sends a commit message c to group members Bob and Charlie. A malicious Bob can create a new commit message c' by mangling the signature within the commit message c without knowing Alice's secret signing key, computing a MAC if present, and send it to Charlie. Although Charlie can successfully process c' and agree on the same content included in the commit with Alice and Bob, their group states diverge, as the signature is included in the transcript hash and the key schedule. Consequently, neither the messages from Charlie nor from Alice can be read by the other. A similar behavior is possible for external committers.*

While this possible behavior with EUF-CMA *schemes seems to runs counter to the design of MLS, we currently do not know of any real-world attack based on this, since an adversary with such capabilities can cause groups to split in other ways, e.g., by dropping messages for some recipients. However this enables a subtle behavior in which participants process the exact same payloads, yet their views and epoch keys diverge. Notably, for non-*SUF-CMA *secure schemes, the so-called* `AuthenticatedContent` *objects in [14] can be mangled, because they include the signature data.*

In practice, RFC 9420 requires implementation of the IETF version of the Ed25519 signature scheme, which is SUF-CMA*-secure [21]. Additionally, RFC 9420 allows for the use of Ed448 and ECDSA variants. ECDSA only satisfies* EUF-CMA*, and it is possible to mangle a signature such that it still verifies as expected [35], thus enabling the behaviors we describe above. In practice, we expect that many implementers will also provide FIPS-compliant secure messengers, which would need to use ECDSA signatures, and hence in these implementations divergence is possible. It seems prudent to either require signature schemes that in fact guarantee* SUF-CMA *(possibly by a wrapper), or reconsider the use in transcripts.*

From a provable security perspective, this shows up as a failure of an assert in $\mathcal{F}_{\mathrm{ECGKA}}$*. However, this behavior should also be captured by* $\mathcal{F}_{\mathrm{CGKA}}$ *in [4], which reduces to* EUF-CMA*. This raises the question of whether this is a bug in the model or proof of [4]. We determined that contrary to the theorem statement, the reduction argument [4, Lemma 3] relies on* SUF-CMA*: the reduction* $\mathcal{A}_{\mathsf{sig}}$ *checks whether a signature-message pair* $(\mathsf{sig}', \mathtt{tbs}')$ *matches any pair appearing in the past signing queries. We confirmed with an author of [4] that the theorem statement is incorrect and should be changed to require* SUF-CMA*.*

6 Stronger Security for Group Re-Synchronization via External Commit

In this section, we formally show that the resumption PSK mechanism of MLS RFC 9420 [14] can be used to improve the security of group re-synchronization. In Section 6.1, we explain the mechanism and the improved security guarantees. In Section 6.2, we introduce our stronger $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ functionality that captures this additional security guarantee. In Section 6.3, we propose our $\mathsf{ETK}^{\text{PSK}}$ protocol. In Section 6.4, we prove that our $\mathsf{ETK}^{\text{PSK}}$ protocol securely realizes $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$.

6.1 Resyncs and Resumption Pre-Shared Keys

A member that cannot process group commits, e.g. because it lost part of its group state, has to re-synchronize with the group. This can be done through a regular remove-then-add P&C mechanism. Alternatively, it can be implemented by the member sending an external commit that removes their old instance while re-adding itself with a new leaf, simply referred to as `resync`. The `resync` is the simpler mechanism, but provides less security: as explained in Section 4.2, the adversary can learn the epoch secret by choosing the external committers randomness. To learn the epoch secret of a regular commit, in addition to choosing the members randomness, the adversary needs to obtain the group secret. Note that a malicious DS - cooperating with an adversary that compromised a members randomness - can manipulate message transmission such that this member falls behind others, and is required to `resync` its state with the group. Improving the security of this `resync` becomes a natural and significant question.

We show that the use of *resumption pre-shared key (PSK)* in MLS improves the security of a `resync`. The resumption PSK mechanism involves the following three steps:

- PSK derivation and storage: As depicted in Figure 2, the key schedule phase derives a resumption PSK along with the other secrets from the epoch secret. The derived PSKs is stored in the local state by every group member for a certain period, determined by the application.
- PSK proposal: A party (insider or outsider) can generate a PSK proposal that specifies a past epoch, the resumption PSK of which will be injected to the next epoch's key schedule phase. The PSK proposals can be committed in both regular and external commits.
- PSK injection: When receiving a commit that includes PSK proposals, every group member retrieves the PSKs of all epochs specified in these proposals from their local state. They then derive the epoch secret by hashing the PSKs with the joiner secret. Note that only members that possess all PSKs are hence able to process such a commit.

When a member performs a `resync` (potentially after being maliciously forced to), this member - in possession of a past PSK of that group - can inject that PSK into the external commit to prove previous group membership. An adversary that controls the member's randomness, but does not know its state and hence PSK, cannot process this commit and is hence locked out.

6.2 The Stronger Functionality $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$

We propose our functionality $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ that extends our $\mathcal{F}_{\text{ECGKA}}$ with the PSK injection, highlighted with orange color (●) in Figures 4, 5 and 7. Below, we explain the additional execution and security guarantees captured by $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$.

Additional Executions. Our $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ includes additional executions in the `Propose`, `Commit`, `ExtCommit`, `Process`, and `Join` interfaces.

Interface `Propose`*:* This interface now allows an additional $\mathsf{act} = \texttt{psk-epoch-use}$, with `epoch` specifying which PSK will be injected, and `use` describing the usage of this PSK. In our work, we only consider the usage "application", which can be used to improve the security of regular or external commits as described above. We leave the analysis of the other usages, "reinit" for group re-initialization and "branch" for creating group branches, for future work. The execution of the PSK proposal $(\texttt{Propose}, \mathsf{act})$ for $\mathsf{act} = \texttt{psk-epoch-use}$ is identical to the one of other proposals.

Interface `Commit`*:* $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ forwards this interface's query to the adversary and retrieves a vector $\vec{c_{psk}}$ that contains all past commits whose PSKs will be injected by this commit's `psk` proposals. $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ runs ***fixate-psk-refs** to ensure that every PSK proposal refers to a valid commit c_{psk} in $\vec{c_{psk}}$. The condition **assert** $\mathsf{res} \neq \bot$ fails when no suitable c_{psk} is found or the epoch does not match. $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$

does not require re-key path for commits containing only `add` and `psk` proposals, hence ***keep-spk** checks if ***only-adds-and-psk**, and the committer's keying material is not forced to be replaced if so.

For every `psk` proposal in the proposals vector, it is required that the committer `id` is in the list of group members who have derived the resumption PSK from the key schedule for the corresponding c_{psk}. Regardless of whether the commit node c or the detached root exists, $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ first checks ***keep-spk** to determine if re-key path is necessary. The secret of this commit node is marked as `bad` if no rekey path is executed or if randomness is leaked. If the commit c includes any proposal that adds new members, the welcome messages must be non-$\perp$. $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ ensures that the messages reference $(c, c_{\vec{psk}})$ and that the vector of welcome messages differs for each c for every $c_{\vec{psk}}$.

Interface `ExtCommit`*:* This interface takes as additional input a vector of epochs $e_{\vec{\text{psk}}}$. The functionality $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ forwards this input $e_{\vec{\text{psk}}}$ to the adversary and receives two additional vectors $p_{\vec{psk}}$ and $c_{\vec{psk}}$. Similar to the `extAdd` and `rem` proposals, $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ does not create independent proposal nodes for each included `psk` proposal. Instead, it zips the `psk` proposals with the corresponding epoch vector $(e_{\vec{\text{psk}}}, p_{\vec{psk}})$ and uses ***create-ext-commit-props** to create sub-nodes for every PSK proposal p_{psk}. Later, $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ sets `ExtCommitProps` via ***fixate-psk-refs** and $c_{\vec{psk}}$, ensuring that every `psk` proposal refers to a valid c_{psk}. The functionality guarantees that for `psk` proposals in `ExtCommitProps`, the committer `id` must be in the list of group members who have derived the resumption PSK from the key schedule.

Interface `Process`*:* For a regular commit ($\vec{p} \neq \perp$), $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ forwards the process query to the adversary and additionally receives $c_{\vec{psk}}$. After verifying correctness and creating proposal nodes for all injected proposals, $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ checks ***fixate-psk-refs** to ensure that every PSK proposal refers to a valid c_{psk} in $c_{\vec{psk}}$. No matter whether the input commit message c is honestly generated or injected by adversary, $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ additionally verifies that at least one PSK for each `psk` proposal is in `id`'s possession. `id` is added to the list of group members who have derived the resumption `psk` of the new epoch.

For an external commit ($\vec{p} = \perp$), after forwarding the process query, $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ receives an additional $c_{\vec{psk}}$ that specifies the commit messages that the included `psk` proposals refer to. If the invoker `id` is a group member and the node of commit c has not been created, after creating the "sub-node" of these proposals, the functionality $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ verifies whether their actions are valid, i.e., either `add`, `rem`, or `psk`. Moreover, the functionality $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ verifies whether every `psk` proposals in `ExtCommitProps` refers to a valid c_{psk}. If the node of commit c has already been created, the $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ additionally verifies whether `psk` proposals in `ExtCommitProps` correctly refer to their valid c_{psk}. `id` is added to the list of group members who have derived the resumption PSK of the new epoch.

Interface `Join`*:* When sending $(\texttt{Join}, \texttt{id}, \texttt{vec_w})$ to the adversary, $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ additionally receives $c_{\vec{psk}}'$. After verifying adversarial inputs, $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ looks for a commit c and its $c_{\vec{psk}}$ that `vec_w` maps to. If no such commit exists, $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ utilizes the adversary-provided $c_{\vec{psk}}'$ to create the node and maps `vec_w` to the newly generated pair $(c, c_{\vec{psk}})$. In all cases, $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ verifies that every `psk` proposal in $\texttt{Node}[c].\mathsf{pro}$ correctly refers to a valid c_{psk}. In the end, `id` is added to the list of group members who have derived the resumption PSK of the new epoch.

Additional Security Guarantees. $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ captures the stronger confidentiality provided by the PSK mechanism in the **safe** predicate in Figure 7 for both regular and external commits.

Regular Commit: Compared to $\mathcal{F}_{\text{ECGKA}}$, there are two differences. First, in the predicate ***secrets-replaced**, in addition to the regular commit with an add-only proposal list, $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ captures the rekey path except for a regular commit that is empty or includes only (regular) add, external self-add, or PSK proposals. This reflects the suggestion in MLS RFC 9420 [14, Section 12.4] that the rekey path can be omitted for a regular commit with add-or-PSK only proposal list.

Second, in condition d) of the helper predicate ***can-traverse**, where $\mathcal{F}_{\text{ECGKA}}$ only takes the leakage of the commit secret (covered by $\textbf{know}(c, *)$) and the initSecret (indicated by $c = \mathsf{root}_* \vee \textbf{know}(\texttt{Node}[c].\mathsf{pars}, \text{'epoch'})$) into account, $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ additionally considers leakage of the `psk-secret` (indicated by $\texttt{Node}[c].\mathsf{psk} \neq$ `good`). In this case, a regular commit c is secure if at least one of the commit-, the init-, and the optional psk-secret is not leaked.

External Commit: In condition e) of the predicate ***can-traverse**, where $\mathcal{F}_{\text{ECGKA}}$ considers an external commit insecure if it involves bad randomness, $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ relaxes the requirement on randomness as long as the included psk-secret is not leaked. This provides the external commit with stronger resilience against randomness leakage attacks by using PSKs from the local states.

6.3 The $\mathsf{ETK}^{\mathrm{PSK}}$ Protocol

Our $\mathsf{ETK}^{\mathrm{PSK}}$ protocol uses the resumption PSK mechanism from RFC 9420 as described in Section 6.1 to extend ETK in Section 5 with the following differences.

New State Variable. Member's group state include an additional `psk_store` that contains the resumption `psk`s and their associated groupId and epoch.

New Operation: PSK Proposal. To perform a PSK proposal with `usage` = "application", the proposer `id` must be in the group. To generate the `psk` proposal, `id` first generates the `pskId` by using groupId, the intended `epoch` and random nonce variables, then signs the proposal and outputs the proposal message p.

Extended Operation: Regular Commit. If the proposal list only includes add-or-PSK proposals, then no re-key path is executed. Once the commit is signed and confTransHash is set, the committer retrieves all PSK values, identified by `pskIds` in the applied `psk` proposals, from its `psk_store`, and derives the `psk-secret`, then hashes both `joinerSec` and `psk-secret` to derive the epoch secrets. If the commit adds any new members, the vector of welcome messages additionally includes `pskIds` to let joiners know which resumption PSKs to use. Then, the `welcome_secret` is derived from both `joinerSec` and `psk-secret` to encrypt groupInfo. The commit message, the welcome vector, and the encrypted groupInfo object are output.

Extended Operation: External Commit. $\mathsf{ETK}^{\mathrm{PSK}}$ inputs an additional vector $\vec{e_{\mathrm{psk}}}$ that identifies the epochs, the PSKs of which are injected in this external commit. For each `epoch` in the vector $\vec{e_{\mathrm{psk}}}$, a `pskId` with `usage` "application" and a `psk` proposal are generated. As in a regular commit execution, the external committer derives a `psk-secret` and injects it into the key schedule phase to derive epoch secrets.

Extended Operation: Process. The extended execution is the same for processing both regular and external commits. As above, if `id` is not removed by the commit, `id` derives the `psk-secret` from all included PSK proposals to inject it into the key schedule phase. In the end, `id` outputs PSK proposals along with the id of the committer and all other proposals.

Extended Operation: Join. The joiner additionally extracts a vector `pskIds` from the decryption of `encGroupSec` from the vector of welcome messages. As above, the joiner derives the `psk-secret` from all local PSK values to compute the welcome secret for recovering the encrypted groupInfo, and to inject in the key schedule for deriving epoch secrets.

6.4 Security Results for $\mathsf{ETK}^{\mathrm{PSK}}$ and $\mathcal{F}_{\mathrm{ECGKA}^{\mathrm{PSK}}}$

Theorem 2. *Assume that* PKE *is* $\mathsf{IND\text{-}CCA}$ *secure and that* Sig *is* $\mathsf{SUF\text{-}CMA}$ *secure. The* $\mathsf{ETK}^{\mathrm{PSK}}$ *protocol securely realizes* $(\mathcal{F}_{\mathcal{AS}}^{\mathcal{IW}}, \mathcal{F}_{\mathcal{KS}}^{\mathcal{IW}}, \mathcal{F}_{\mathrm{ECGKA}^{\mathrm{PSK}}})$ *in the* $(\mathcal{F}_{\mathcal{AS}}, \mathcal{F}_{\mathcal{KS}}, \mathcal{G}_{\mathrm{RO}})$*-hybrid model, where* $\mathcal{F}_{\mathrm{ECGKA}^{\mathrm{PSK}}}$ *uses the predicates* **safe** *and* **inj-allowed** *from Figure 7 and calls to* HKDF.Expand*,* HKDF.Extract*, and* MAC *functions are replaced by calls to the global random oracle* $\mathcal{G}_{\mathrm{RO}}$*.*

Sketch. This proof is similar to the proof of Theorem 1, except for the analysis of the execution related to the resumption `psk`. The $\mathsf{ETK}^{\mathrm{PSK}}$ protocol additionally allows to inject a resumption `psk` from a past epoch to the key derivation schedule at the current epoch by a PSK proposal $p \leftarrow$ (Propose, psk-epoch-use). The `psk`s are stored locally and can be exposed later. Their status `stat` tracks the adversary's application key knowledge. If a member `id` is exposed, the `psk` in its possession are determined by `has_psk` and `stat` of all its `psk` is set to `bad`. $\mathcal{F}_{\mathrm{ECGKA}}$ uses `Prop`[p]`.c-psk` to track the `psk`'s origin commit. Future processors use this to ensure that all members process commits with consistent `psk`s. We additionally show that this approach maintains consistency in $\mathcal{F}_{\mathrm{ECGKA}}$.

Its possible that the processing client does not know the injected `psk`. $\mathcal{F}_{\mathrm{ECGKA}}$ performs the necessary checks for possession for confidentiality to be maintained. Even if the latest epoch state and the secrets included in a commit are corrupted, the environment still cannot distinguish the epoch secret from random as long as the injected PSKs are not leaked. This is because to derive new epoch secret, the PSK values are injected into hashes, which are modeled as random oracles.

□

References

[1] Joël Alwen, Sandro Coretti, Yevgeniy Dodis, and Yiannis Tselekounis. “Modular Design of Secure Group Messaging Protocols and the Security of MLS”. In: *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*. CCS ’21: 2021 ACM SIGSAC Conference on Computer and Communications Security. Nov. 2021.

[2] Joël Alwen, Sandro Coretti, Yevgeniy Dodis, and Yiannis Tselekounis. “Security Analysis and Improvements for the IETF MLS Standard for Group Messaging”. In: *Advances in Cryptology – CRYPTO 2020*. Cham: Springer International Publishing, 2020.

[3] Joël Alwen, Sandro Coretti, Daniel Jost, and Marta Mularczyk. “Continuous Group Key Agreement with Active Security”. In: *Theory of Cryptography*. Vol. 12551. Cham: Springer International Publishing, 2020.

[4] Joël Alwen, Daniel Jost, and Marta Mularczyk. “On the Insider Security of MLS”. In: *Advances in Cryptology – CRYPTO 2022*. Cham: Springer Nature Switzerland, 2022.

[5] Diego F Aranha, Felipe Rodrigues Novaes, Akira Takahashi, Mehdi Tibouchi, and Yuval Yarom. “LadderLeak: Breaking ECDSA with less than one bit of nonce leakage”. In: *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*. 2020.

[6] Richard Barnes. *How Messaging Layer Security Enables Scalable End-to-End Security in Webex*. WebEx Blog. URL: https://blog.webex.com/hybrid-work/scalable-end-to-end-security-in-webex/ (visited on 02/13/2025).

[7] Richard Barnes, Benjamin Beurdouche, Jon Millican, Emad Omara, Katriel Cohn-Gordon, and Raphael Robert. *The Messaging Layer Security (MLS) Protocol*. Internet-Draft draft-ietf-mls-protocol-07. Work in Progress. Internet Engineering Task Force.

[8] Richard Barnes, Benjamin Beurdouche, Jon Millican, Emad Omara, Katriel Cohn-Gordon, and Raphael Robert. *The Messaging Layer Security (MLS) Protocol*. Internet-Draft draft-ietf-mls-protocol-08. Work in Progress. Internet Engineering Task Force, Nov. 2019.

[9] Richard Barnes, Benjamin Beurdouche, Jon Millican, Emad Omara, Katriel Cohn-Gordon, and Raphael Robert. *The Messaging Layer Security (MLS) Protocol*. Internet-Draft draft-ietf-mls-protocol-10. Work in Progress. Internet Engineering Task Force, Oct. 2020.

[10] Richard Barnes, Benjamin Beurdouche, Jon Millican, Emad Omara, Katriel Cohn-Gordon, and Raphael Robert. *The Messaging Layer Security (MLS) Protocol*. Internet-Draft draft-ietf-mls-protocol-09. Work in Progress. Internet Engineering Task Force, Mar. 2020.

[11] Richard Barnes, Benjamin Beurdouche, Raphael Robert, Jon Millican, Emad Omara, and Katriel Cohn-Gordon. *The Messaging Layer Security (MLS) Protocol*. Internet-Draft draft-ietf-mls-protocol-13. Work in Progress. Internet Engineering Task Force, Mar. 2022.

[12] Richard Barnes, Benjamin Beurdouche, Raphael Robert, Jon Millican, Emad Omara, and Katriel Cohn-Gordon. *The Messaging Layer Security (MLS) Protocol*. Internet-Draft draft-ietf-mls-protocol-15. Work in Progress. Internet Engineering Task Force, June 2022.

[13] Richard Barnes, Benjamin Beurdouche, Raphael Robert, Jon Millican, Emad Omara, and Katriel Cohn-Gordon. *The Messaging Layer Security (MLS) Protocol*. Internet-Draft draft-ietf-mls-protocol-17. Work in Progress. Internet Engineering Task Force, Dec. 2022.

[14] Richard Barnes, Benjamin Beurdouche, Raphael Robert, Jon Millican, Emad Omara, and Katriel Cohn-Gordon. *The Messaging Layer Security (MLS) Protocol*. 9420. https://www.rfc-editor.org/info/rfc9420. July 2023.

[15] Richard Barnes, Benjamin Beurdouche, Raphael Robert, Jon Millican, Emad Omara, and Katriel Cohn-Gordon. *The Messaging Layer Security (MLS) Protocol (Draft 12)*. Internet Draft draft-ietf-mls-protocol-12. Internet Engineering Task Force, Oct. 11, 2021.

[16] Richard Barnes, Jon Millican, Emad Omara, Katriel Cohn-Gordon, and Raphael Robert. *The Messaging Layer Security (MLS) Protocol*. Internet-Draft draft-barnes-mls-protocol-00. Work in Progress. Internet Engineering Task Force.

[17] Richard Barnes, Jon Millican, Emad Omara, Katriel Cohn-Gordon, and Raphael Robert. *The Messaging Layer Security (MLS) Protocol*. Internet-Draft draft-ietf-mls-protocol-02. Work in Progress. Internet Engineering Task Force, Oct. 2018.

[18] Mihir Bellare and Phillip Rogaway. “Random Oracles are Practical: A Paradigm for Designing Efficient Protocols”. In: *Proceedings of the 1st ACM Conference on Computer and Communications Security*. CCS ’93. Fairfax, Virginia, USA: Association for Computing Machinery, 1993.

[19] Karthikeyan Bhargavan, Benjamin Beurdouche, and Prasad Naldurg. “Formal Models and Verified Protocols for Group Messaging: Attacks and Proofs for IETF MLS”. In: *Doctoral dissertation, Inria Paris* (2019).

[20] Karthikeyan Bhargavan, Benjamin Beurdouche, and Prasad Naldurg. *Formal Models and Verified Protocols for Group Messaging: Attacks and Proofs for IETF MLS*. [Research Report] Inria Paris, ffhal-02425229f https://inria.hal.science/hal-02425229v1/file/mls-treekem.pdf. 2019.

[21] Jacqueline Brendel, Cas Cremers, Dennis Jackson, and Mang Zhao. *The Provable Security of Ed25519: Theory and Practice*. Cryptology ePrint Archive, Paper 2020/823. 2020.

[22] Chris Brzuska, Eric Cornelissen, and Konrad Kohbrok. “Security Analysis of the MLS Key Derivation”. In: *2022 IEEE Symposium on Security and Privacy (SP)*. 2022 IEEE Symposium on Security and Privacy (SP). May 2022.

[23] R. Canetti. “Universally Composable Security: A New Paradigm for Cryptographic Protocols”. In: *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*. Proceedings 42nd IEEE Symposium on Foundations of Computer Science. 2001.

[24] Céline Chevalier, Guirec Lebrun, Ange Martinelli, and Abdul Rahman Taleb. “Quarantined-TreeKEM: A Continuous Group Key Agreement for MLS, Secure in Presence of Inactive Users”. In: *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*. CCS ’24. Salt Lake City, UT, USA: Association for Computing Machinery, 2024.

[25] Cisco. *Cisco/MLSPP*. https://github.com/cisco/mlspp. Cisco Systems. (Visited on 02/13/2025).

[26] Katriel Cohn-Gordon, Cas Cremers, and Luke Garratt. “On Post-compromise Security”. In: *2016 IEEE 29th Computer Security Foundations Symposium (CSF)*. 2016 IEEE 29th Computer Security Foundations Symposium (CSF). June 2016.

[27] 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: *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. CCS ’18: 2018 ACM SIGSAC Conference on Computer and Communications Security. Oct. 15, 2018.

[28] Cas Cremers, Britta Hale, and Konrad Kohbrok. “The Complexities of Healing in Secure Group Messaging: Why Cross-Group Effects Matter”. In: *Proceedings of the 30th USENIX Security Symposium*. 2021.

[29] Cas Cremers, Charlie Jacomme, and Philip Lukert. “Subterm-Based Proof Techniques for Improving the Automation and Scope of Security Protocol Analysis”. In: *2023 IEEE 36th Computer Security Foundations Symposium (CSF)*. 2023 IEEE 36th Computer Security Foundations Symposium (CSF). Dubrovnik, Croatia: IEEE, July 2023.

[30] Internet Engineering Task Force. *MLS RFC 9420 Datatracker*. https://datatracker.ietf.org/doc/rfc9420/. 2025.

[31] Matthew Hodgson and Chathi. *A Giant Leap Forwards for Encryption with MLS*. [matrix]. July 18, 2023. URL: https://matrix.org/blog/2023/07/a-giant-leap-with-mls/ (visited on 02/13/2025).

[32] Karen Klein, Guillermo Pascual-Perez, Michael Walter, Chethan Kamath, Margarita Capretto, Miguel Cueto, Ilia Markov, Michelle Yeo, Joël Alwen, and Krzysztof Pietrzak. “Keep the Dirt: Tainted TreeKEM, Adaptively and Actively Secure Continuous Group Key Agreement”. In: *2021 IEEE Symposium on Security and Privacy (SP)*. 2021.

[33] AWS Labs. *Awslabs/MLS-RS: An Implementation of Messaging Layer Security (RFC)*. GitHub. 2023. URL: https://github.com/awslabs/mls-rs/tree/main (visited on 02/13/2025).

[34] Rohan Mahy. *More Instant Messaging Interoperability (MIMI) Message Content*. Internet Draft draft-ietf-mimi-content-04. https://datatracker.ietf.org/doc/draft-ietf-mimi-content. Internet Engineering Task Force, June 10, 2024.

[35] Jacques Stern, David Pointcheval, John Malone-Lee, and Nigel P. Smart. “Flaws in Applying Proof Methodologies to Signature Schemes”. In: *Advances in Cryptology - CRYPTO 2002, 22nd Annual International Cryptology Conference, Santa Barbara, California, USA, August 18-22, 2002, Proceedings*. Vol. 2442. Lecture Notes in Computer Science. Springer, 2002.

[36] Théophile Wallez, Jonathan, Benjamin Beurdouche, and Karthikeyan Bhargavan. “TreeSync: Authenticated Group Management for Messaging Layer Security”. In: *32nd USENIX Security Symposium*. (USENIX Security 23). 2023.

$\mathsf{Expt}^{\mathsf{IND\text{-}CCA}}_{\mathsf{PKE}}(\mathcal{A})$:	$\mathbf{Dec}(c)$:
1 $\mathsf{b} \xleftarrow{\$} \{0,1\}$	7 $m \leftarrow \mathsf{PKE.Dec}(\mathsf{sk}, c)$
2 $(\mathsf{pk}, \mathsf{sk}) \xleftarrow{\$} \mathsf{PKE.Kg}()$	8 **return** m
3 $(m_0, m_1) \xleftarrow{\$} \mathcal{A}^{\mathbf{Dec}}(\mathsf{pk})$	
4 $c \xleftarrow{\$} \mathsf{PKE.Enc} \mathsf{pk}, m_{\mathsf{b}}$	
5 $\mathsf{b}' \xleftarrow{\$} \mathcal{A}(c)$	
6 **return** $[\![\mathsf{b} = \mathsf{b}']\!]$	

Figure 8: The IND-CCA security game for a PKE scheme.

A Preliminaries

A.1 Public Key Encryption

Definition 2 (Public Key Encryption)**.** *A* public key encryption *scheme over message space $\mathcal{M}$ is a triple of algorithms, where $\mathsf{PKE} = \{\mathsf{PKE.Kg}, \mathsf{PKE.Enc}, \mathsf{PKE.Dec}\}$ as defined below.*

Key Generation $(\mathsf{pk}, \mathsf{sk}) \xleftarrow{\$} \mathsf{PKE.Kg}()$ *inputs the public parameters* pp *and outputs a public encryption key and a private decryption key pair* $(\mathsf{pk}, \mathsf{sk})$*.*

Encryption $c \xleftarrow{\$} \mathsf{PKE.Enc}(\mathsf{pk}, m)$ *inputs an encryption key* pk *and a message $m \in \mathcal{M}$ and outputs a ciphertext c.*

Decryption $m \leftarrow \mathsf{PKE.Dec}(\mathsf{sk}, c)$ *inputs a private decryption key* sk *and a ciphertext c. It deterministically outputs a message m.*

We call a public key encryption scheme $\mathsf{PKE} = (\mathsf{PKE.Kg}, \mathsf{PKE.Enc}, \mathsf{PKE.Dec})$ *correct* if for all $(\mathsf{pk}, \mathsf{sk}) \xleftarrow{\$} \mathsf{PKE.Kg}()$ and all messages $m \in \mathcal{M}$ it holds that

$$m = \mathsf{PKE.Dec}(\mathsf{sk}, \mathsf{PKE.Enc}(\mathsf{pk}, m)).$$

Definition 3. *We say a public key encryption scheme* $\mathsf{PKE} = (\mathsf{PKE.Kg}, \mathsf{PKE.Enc}, \mathsf{PKE.Dec})$ *is ϵ-*IND-CCA *secure, if the advantage of all adversaries $\mathcal{A}$ that win* $\mathsf{Expt}^{\mathsf{IND\text{-}CCA}}_{\mathsf{PKE}}$ *experiment in Figure 8 is bounded by*

$$|\Pr[\mathsf{Expt}^{\mathsf{IND\text{-}CCA}}_{\mathsf{PKE}}(\mathcal{A}) = 1] - \frac{1}{2}| \leq \epsilon$$

The Generalized Selective Decryption *(GSD) Security.* The GSD game constructs a hypergraph of vertices, where, in the context of ETK, each vertex is referred to as a node u, each storing a seed s_u. For each node, a public-private key pair can be derived using the seed as randomness, or the seed can be used as a symmetric key. Edges connect the nodes to represent dependencies between seeds. Briefly, the GSD game includes the following oracles and functions:

- $\mathbf{Enc}(u, v)$: An *encryption edge* is created upon calling the $\mathbf{Enc}(u, v)$ oracle, encrypting the seed of v under the public key of u.
- $\mathbf{Hash}(u, v, \mathsf{lbl})$: A *key derivation edge* is created when calling $\mathbf{Hash}(u, v, \mathsf{lbl})$, hashing the seed s_u with the label lbl to derive the seed s_v
- $\mathbf{Join\text{-}Hash}(u, u', v, \mathsf{lbl})$: Calling $\mathbf{Join\text{-}Hash}(u, u', v, \mathsf{lbl})$ hashes both s_u and $s_{u'}$ together to derive the seed s_v.
- $\mathbf{Dec}$: Analogous to the IND-CCA game, this oracle decrypts the ciphertext using the private key of u, provided that the seed of u is not valid or u is not a sink node.
- $\mathbf{Corr}$: Outputs the seed of a node and records it in the Corr set.
- $*\mathsf{get\text{-}pk}(u)$: In the GSD experiment, the $*\mathsf{get\text{-}pk}(u)$ function allows $\mathcal{A}$ to obtain the corresponding public key of vertex u by calling the oracle $\mathbf{Enc}(u, 0)$.

The $\mathsf{gsd\text{-}exp}(u)$ function indicates whether the seed s_u of the node u of commit c is exposed. The core component of the proof shows that $\mathsf{gsd\text{-}exp}(u)$ is analogous to predicate $\mathbf{safe}(c)$. For further details, we refer to [4]. The full GSD game is provided in Figure 9. Figure 25 provides a visualization of the GSD graph for a regular commit where Figure 27 visualizes the GSD graph of an external commit.

Game $\text{GSD}_{\mathsf{PKE},\mathcal{A}}$

The game is parameterized by the number of vertices N, the security parameter κ and a hash function Hash.

```
(V, E) ← ([N], ∅) // GSD graph
Corr, Ctxt ← ∅ // corrupted vertices, ciphertexts
s_u, pk_u, sk_u ← ⊥ for each u ∈ [N] // keys for vertex u
u ← ⊥ // challenge vertex
b ←$ {0,1}
s' ←$ {0,1}^κ
b' ← A_PKE^{Enc,Dec,Corr,Chal,Hash,Join-Hash}
if (V, E) acyclic ∧ u is a sink ∧ ¬gsd-exposed(u) then
    return b = b'
else return false
```

Oracle Chal(u)

```
req u = ⊥
u ← u
if b = 0 then return s_u
else return s'
```

Oracle Hash(u, v, lbl)

```
req s_v = ⊥ ∧ (u, *, h-lbl) ∉ E // hash is deterministic
gen-key-if-nec(u)
s_v ← Hash(s_u, lbl)
gen-key-if-nec(v)
E +← (u, v, h-lbl)
return pk_u
```

Oracle Corr(u)

```
req s_u ≠ ⊥
Corr +← u
return s_u
```

Oracle Join-Hash(u, u', v, lbl)

```
req s_v = ⊥ ∧ ((u, u', lbl), *, h-lbl) ∉ E
gen-key-if-nec(u); gen-key-if-nec(u')
s_v ← Hash(s_u, s_u', lbl)
gen-key-if-nec(v)
E +← ((u, u'), v, h-lbl)
return (pk_u, pk_u')
```

Oracle Enc(u, v)

```
gen-key-if-nec(u); gen-key-if-nec(v)
E +← (u, v, e)
c ← PKE.Enc(pk_u, s_v)
Ctxt +← (u, c)
return (pk_u, c)
```

Oracle Dec(u, c)

```
req s_u ≠ ⊥ ∧ u not a sink
req (u, c) ∉ Ctxt
return PKE.Dec(sk_u, c)
```

```
gen-key-if-nec(u)
  if s_u = ⊥ then s_u ←$ {0,1}^κ
  (pk_u, sk_u) ← PKE.Kg(Hash(s_u, node))
  // in ETK and ETK^PSK, the label "node" is used for key generation

gsd-exposed(u)
  return u ∈ Corr
    ∨ ∃(v, u, *) ∈ E : gsd-exposed(v)
    ∨ ∃((v, v'), u, *) ∈ E : gsd-exposed(v) ∧ gsd-exposed(v')
```

Figure 9: The GSD game, modified to explain ETK and $\mathsf{ETK}^{\text{PSK}}$ executions. Taken from [4].

Definition 4. *We say a public key encryption scheme* $\mathsf{PKE} = (\mathsf{PKE.Kg}, \mathsf{PKE.Enc}, \mathsf{PKE.Dec})$ *is* ϵ-GSD *secure, if the advantage of all adversaries* $\mathcal{A}$ *that win* $\text{GSD}_{\mathsf{PKE},\mathcal{A}}$ *game in Figure 9 is bounded by*

$$|\Pr[\text{GSD}_{\mathsf{PKE},\mathcal{A}} = 1] - \frac{1}{2}| \leq \epsilon$$

Theorem 3 ([4])**.** *Let* PKE *denotes a public key encryption scheme. If* PKE *is* IND-CCA *secure, then* PKE *is* GSD *secure.*

A.2 Digital Signature

Definition 5 (Digital Signature)**.** *A* digital signature scheme *over message space* $\mathcal{M}$ *is a triple of algorithms, where* $\mathsf{DS} = (\mathsf{DS.KGen}, \mathsf{DS.Sign}, \mathsf{DS.Vrfy})$ *as defined below.*

Key Generation $(pk, sk) \xleftarrow{\$} \mathsf{DS.KGen}()$ *outputs a public verification and private signing key pair* (pk, sk)*.*

Signing $\sigma \xleftarrow{\$} \mathsf{DS.Sign}(sk, m)$ *inputs a secret key sk as well as the message* $m \in \mathcal{M}$ *to be signed and outputs a signature* σ*.*

Verification $\texttt{true}/\texttt{false} \leftarrow \mathsf{DS.Vrfy}(pk, \sigma, m)$ *inputs a public verification key* pk*, a signature* σ *and message* $m \in \mathcal{M}$*. It deterministically outputs* `true` *if* σ *is a valid signature over* m *wrt.* pk*, else it outputs* `false`*.*

We call a signature scheme $\mathsf{DS} = (\mathsf{DS.KGen}, \mathsf{DS.Sign}, \mathsf{DS.Vrfy})$ *correct* if for all $(pk, sk) \xleftarrow{\$} \mathsf{DS.KGen}()$ and all messages $m \in \mathcal{M}$, it always holds that

$$\texttt{true} = \mathsf{DS.Vrfy}(pk, \mathsf{DS.Sign}(sk, m), m).$$

Definition 6 (EUF-CMA and SUF-CMA security)**.** *Let* $\mathsf{DS} = (\mathsf{DS.KGen}, \mathsf{DS.Sign}, \mathsf{DS.Vrfy})$ *be a digital signature scheme with message space* $\mathcal{M}$ *We say that* DS *is* ϵ-EUF-CMA*-secure if the adversary of all adversaries that win in the* $\mathsf{Expt}_{\mathsf{DS}}^{\mathsf{EUF\text{-}CMA}}$ *experiment is bounded by*

$\underline{\mathsf{Expt}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{DS}}(\mathcal{A})}$:
1 $\mathcal{L}_{\mathsf{Sign}} \leftarrow \varnothing$
2 $(\mathsf{pk}, \mathsf{sk}) \xleftarrow{\$} \mathsf{DS.KGen}(\mathsf{pp})$
3 $(m*, \sigma*) \xleftarrow{\$} \mathcal{A}^{\mathbf{Sign}}(\mathsf{pk})$
4 **return** $[\![\mathsf{DS.Vrfy}(\mathsf{pk}, \sigma*, m*) \wedge m* \notin \mathcal{L}_{\mathsf{Sign}}]\!]$

$\underline{\mathbf{Sign}(m)}$:
5 $\sigma \xleftarrow{\$} \mathsf{DS.Sign}(\mathsf{sk}, m)$
6 $\mathcal{L}_{\mathsf{Sign}} \leftarrow \mathcal{L}_{\mathsf{Sign}} \cup \{m\}$
7 **return** σ

$\underline{\mathsf{Expt}^{\mathsf{SUF\text{-}CMA}}_{\mathsf{DS}}(\mathcal{A})}$:
1 $\mathcal{L}_{\mathsf{Sign}} \leftarrow \varnothing$
2 $(\mathsf{pk}, \mathsf{sk}) \xleftarrow{\$} \mathsf{DS.KGen}(\mathsf{pp})$
3 $(m*, \sigma*) \xleftarrow{\$} \mathcal{A}^{\mathbf{Sign}}(\mathsf{pk})$
4 **return** $[\![\mathsf{DS.Vrfy}(\mathsf{pk}, \sigma*, m*) \wedge (m*, \sigma*) \notin \mathcal{L}_{\mathsf{Sign}}]\!]$

$\underline{\mathbf{Sign}(m)}$:
5 $\sigma \xleftarrow{\$} \mathsf{DS.Sign}(\mathsf{sk}, m)$
6 $\mathcal{L}_{\mathsf{Sign}} \leftarrow \mathcal{L}_{\mathsf{Sign}} \cup \{(m, \sigma)\}$
7 **return** σ

Figure 10: The EUF-CMA and SUF-CMA security games for a DS scheme.

$\underline{\mathsf{Expt}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{MAC}}(\mathcal{A})}$:
1 $\mathcal{L}_{\mathsf{MAC}} \leftarrow \varnothing$
2 $(k \xleftarrow{\$} \mathsf{KGen}(1^{\lambda})$
3 $(m*, t*) \xleftarrow{\$} \mathcal{A}^{\mathbf{MAC}}(.)$
4 **return** $[\![\mathsf{Vrfy}(k, m*, t*) \wedge m* \notin \mathcal{L}_{\mathsf{MAC}}]\!]$

$\underline{\mathbf{MAC}(m)}$:
5 $(m, t) \xleftarrow{\$} \mathsf{MAC}(k, m)$
6 $\mathcal{L}_{\mathsf{MAC}} \leftarrow \mathcal{L}_{\mathsf{MAC}} \cup \{m, t\}$
7 **return** t

Figure 11: The EUF-CMA security game for a MAC scheme.

$$\Pr[\mathsf{Expt}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{DS}}(\mathcal{A}) = 1] \leq \epsilon.$$

Analogously, we say that DS *is* ϵ-SUF-CMA-*secure if the adversary of all adversaries that win in the* $\mathsf{Expt}^{\mathsf{SUF\text{-}CMA}}_{\mathsf{DS}}$ *experiment is bounded by*

$$\Pr[\mathsf{Expt}^{\mathsf{SUF\text{-}CMA}}_{\mathsf{DS}}(\mathcal{A}) = 1] \leq \epsilon$$

, where the $\mathsf{Expt}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{DS}}$ *and* $\mathsf{Expt}^{\mathsf{SUF\text{-}CMA}}_{\mathsf{DS}}$ *experiments are defined in Figure 10.*

A.3 Message Authentication Code

Definition 7 (Message Authentication Code)**.** *A* message authentication code scheme *over message space* $\mathcal{M}$ *and key space* $\mathcal{K}$ *is a tuple of algorithms, where* $\mathsf{MAC} = (\mathsf{MAC.Sign}, \mathsf{MAC.Vrfy})$ *as defined below.*

Signing $t \leftarrow \mathsf{MAC.Sign}(k, m)$ *inputs a key* $k \in \mathcal{K}$*, a message* $m \in \mathcal{M}$ *and outputs a tag* t*.*

Verification `true`/`false` $\leftarrow \mathsf{MAC.Vrfy}(k, m, t)$ *inputs a key* $k \in \mathcal{K}$*, a message* $m \in \mathcal{M}$ *and a tag* t*. It deterministically outputs* `true` *if* t *is a valid tag of* m *w.r.t the key* k*, else it outputs* `false`*.*

We say a $\mathsf{MAC} = (\mathsf{MAC.Sign}, \mathsf{MAC.Vrfy})$ is *correct* if for every key $k \xleftarrow{\$} \mathcal{K}$ and every $m \in \mathcal{M}$, it holds that

$$\mathtt{true} = \mathsf{MAC.Vrfy}(k, m, \mathsf{MAC.Sign}(k, m)).$$

Definition 8 (EUF-CMA security of MAC)**.** *We say a message authentication code* $\mathsf{MAC} = (\mathsf{MAC.Sign}, \mathsf{MAC.Vrfy})$ *is* ϵ-EUF-CMA *secure, if the advantage of all adversaries that break* $\mathsf{Expt}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{MAC}}$ *experiment in Figure 11 is bounded by*

$$\Pr[\mathsf{Expt}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{MAC}}(\mathcal{A}) = 1] \leq \epsilon.$$

A.4 Random Oracle

The Random Oracle Model (ROM), introduced by Bellare and Rogaway [18], conceptualizes an idealized random function. Our work models hash function as an idealized random function, i.e., the random oracle. More concretely, for an adversary aiming to compute the hash function H on input x, it must query the random oracle with x. This process shapes adversarial behavior by granting the ability to choose adversarial inputs to the hash function, which is referred to as extractability. For each query, the random oracle returns uniformly random answers drawn from the predefined range of the hash function. Regardless of how many times the random oracle is queried, the function H always produces the same output.

B Simplifications

Although our ETK and $\mathsf{ETK}^{\mathrm{PSK}}$ protocols are by far, to the best of our knowledge, the closest to the MLS RFC 9420 standard, our study still has the following simplifications:

- **External Proposals:** Our ETK includes the external self-add proposal `extAdd`, while omitting other external proposals, e.g., external removal and external group_context_extension.
- **PSK Usages:** Our $\mathsf{ETK}^{\mathrm{PSK}}$ reuses the resumption PSK mechanisms underlying MLS RFC 9420 standard to improve the security of regular and external commits. However, our $\mathsf{ETK}^{\mathrm{PSK}}$ omits the original applications of the resumption PSK mechanism, e.g., group re-initialization and branch operations.
- **Exporter Secrets:** In the key schedule of MLS RFC 9420, an exporter secret, which can be used by an application to derive new secrets for use outside of MLS, is derived together with other epoch secrets. Since our work focuses only on the evolution of the continuous group key agreement, rather than its external application outside MLS, we omit the exporter secret and leave it as future work.
- **External PSK:** MLS RFC 9420 employs two flavors of pre-shared keys: The resumption PSK and external PSK. The resumption key is derived together with epoch secrets in the key schedule and covered by our $\mathsf{ETK}^{\mathrm{PSK}}$. The external PSK is distributed among group members over some out-of-band channels. For simplicity, our work omits the external PSK mechanism.
- **Private Handshake Messages:** In MLS RFC 9420, the handshake messages, i.e., a message carrying an MLS Proposal or Commit object, can be either public or private. In our work, we consider public handshake message, which means that the handshake messages are signed by its sender and authenticated as coming from a member of the group in a particular epoch, but not encrypted. Note that the private handshake messages are signed by its sender, authenticated as coming from a member of the group in a particular epoch, and encrypted so that it is confidential to the members of the group in that epoch. Intuitively, the security achieved by using public handshake messages can also be achieved by using private handshake messages. For simplicity, we omit the private handshake messages and leave its detailed security analysis as future work.
- **Ciphersuites and Protocol Version:** MLS RFC 9420 introduces several available ciphersuits and their extensions in [14, Section 5.1, 13.1, and 17.1]. Moreover, every group member should specify the MLS protocol version and ciphersuit this member supports in its key packages. For simplicity, we omit this in the key package and simply assume that the MLS protocol version and ciphersuit every party supports are publicly known.
- **Cryptographic Primitive Contexts:** MLS RFC 9420 uses cryptographic primitives often with some labels, which are constant and indicates the computation context. For simplicity, our work continues to use primitives such as Hash, Enc, Sig, and KDF without explicity introducing additional context and labels.
- **Key Expiration:** Our work follows [4] and omits the consideration of expiration of key packages or certificates.
- **Order of Proposal Applications:** Our work follows [4] and omits the formalization of the order that a list of proposal should apply, see [14, Section 12.3]. For simplicity, our ETK and $\mathsf{ETK}^{\mathrm{PSK}}$ protocols enforce a list of proposals to be applied in the order as each proposal appears in the list.
- **Membership Tag:** Our work follows [4] and omits the computation of membership tag on commits that do not include removal proposal.
- **Proposal Reference by Value:** [4] only contains proposals by reference, meaning that a proposal needs to be sent separately from a commit, and is included only via a ProposalRef. External commits require proposals by value in the commit, as the sender is not yet part of the group. For simplicity, we restrict the use of proposals by value to external commits.
- **PSK Life Time:** MLS RFC 9420 standard suggests that "the application SHOULD specify an upper limit on the number of past epochs for which the resumption PSK may be stored" [14, Section 8.6]. In our work, we do not model the life time of resumption PSKs, and simply assume that PSKs are never removed once generated. Consequently, all PSKs in possession of the member are exposed upon state corruption.

Additionally, upon a compromise, $\mathcal{F}_{\mathrm{ECGKA}}$ marks *all past* key packages as exposed and all commits that add such a key package. $\mathcal{F}_{\mathrm{ECGKA}}$ introduces the new init key to key packages, which enhances

forward secrecy in practice: While the leaf node key can be compromised in one of the early epochs after adding, the init key can be deleted immediately after a join. However, We keep the simplification of [4] and also expose all past key packages upon compromise. $\mathcal{F}_{\text{ECGKA}}$ can hence not capture this increased forward secrecy.

C Details on $\mathcal{F}_{\text{ECGKA}}$ and $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ Functionalities

Here we describe the interfaces not described in the body of the paper. Furthermore, we give the formal description of the Authentication Service and Keypackage Service functionalities in Figures 12 and 13. We define the helper functions in Figures 14 to 16. As before, we mark changes and additions from the model of [4] using colors according to Table 2. For details on the security model of $\mathcal{F}_{\text{CGKA}}$, we refer the reader to [4].

Interface `Commit`*:* This interface of the form $(\texttt{Commit}, \vec{p}, \textsf{spk}, \textsf{force-rekey}, \texttt{wel_type})$ considers the regular commit operation and is invoked only if the invoker `id` is in the group. $\mathcal{F}_{\text{ECGKA}}$ forward this query to the adversary and receives an acknowledgment ack, a commit message c, welcome messages `vec_w`, groupInfo g, and an index rt of a detached root. Similar to the interface `Propose`, $\mathcal{F}_{\text{ECGKA}}$ first checks whether the commit is correct, or injected by the adversary. Next, $\mathcal{F}_{\text{ECGKA}}$ creates proposal nodes for all injected proposals. Then, $\mathcal{F}_{\text{ECGKA}}$ checks whether the public key spk of the invoker `id` is valid. If neither the node of the commit c nor the detached root $\textsf{root}_{rt}$ exists, then $\mathcal{F}_{\text{ECGKA}}$ creates the node of c. The secret of this commit node if marked to `Bad` if no rekey-path is executed or the randomness is leaked. Otherwise, either the node of the commit c or the detached root $\textsf{root}_{rt}$ exists. In this case, $\mathcal{F}_{\text{ECGKA}}$ attaches the commit node c to them if the attachment is valid. If this commit c includes any proposal that adds new members, the welcome messages must not be $\perp$ and point to c. The groupInfo g must not be $\perp$ and point to c. If the consistency and authenticity checks are not violated, then the commit c, the welcome messages `vec_w`, and groupInfo g are returned.

Interface `Propose` *(Regular):* This interface of the form $(\texttt{Propose}, \textsf{act})$, $\textsf{act} \in \{\textsf{up-spk}, \textsf{add-}\texttt{id}_t, \textsf{rem-}\texttt{id}_t\}$ considers three conventional proposal operations: (1) update public key spk, (2) add a party $\texttt{id}_t$, and (3) remove a party $\texttt{id}_t$. The invoker `id` must be within the group and $\mathcal{F}_{\text{ECGKA}}$ forwards this query to the adversary $\mathcal{A}$. Upon receiving the proposal p, public key $\textsf{spk}_t$, and an acknowledgment ack, $\mathcal{F}_{\text{ECGKA}}$ checks whether the proposal is correct. Otherwise, it must be maliciously injected by the adversary (i.e., $ack = \texttt{true}$). If the node of proposal p has not been created, then $\mathcal{F}_{\text{ECGKA}}$ creates it using the information in this query (i.e., the query input and the reply from the adversary). Otherwise, $\mathcal{F}_{\text{ECGKA}}$ checks whether the information in this query is consistent with the existing node of p. If the randomness in this query is chosen by the adversary, then $\mathcal{F}_{\text{ECGKA}}$ exposes the secret key corresponding to the invoker `id`'s public key. This models the case in which ETK is used with a signature scheme (such as ECDSA) that is vulnerable to a randomness leakage attack [5]. In the end, the proposal p is returned.

Interface `Join`*:* The interface of the form $(\texttt{Join}, \texttt{vec_w})$ joins a group via a vector of welcome messages `vec_w`. $\mathcal{F}_{\text{ECGKA}}$ forwards this query to the adversary and receives an acknowledgment ack, a commit c', a groupInfo g, an origin $\textsf{orig}'$, and a member list $\textsf{mem}'$. $\mathcal{F}_{\text{ECGKA}}$ checks whether the adversary has acknowledged this query and looks for a commit $c = \texttt{Wel}[\texttt{vec_w}]$ that this vector of welcome messages maps to. In the case that such a commit c does not exist, $\mathcal{F}_{\text{ECGKA}}$ creates and maps `vec_w` to it by using the input provided by the adversary: either simply use c' as c if the node of c' exists, or a new detached node for c with the origin $\textsf{orig}'$, the member list $\textsf{mem}'$, and status `adv`. In all cases, $\mathcal{F}_{\text{ECGKA}}$ moves the pointer of `id` to this commit c that either previously existed or was newly created and marks `id`'s key to `true` so that it can be challenged. The groupInfo g must be non-$\perp$ and point to c. If `id` is included in the member list of this commit c and the consistency and authenticity checks are not violated, then the member list and origin of this commit c are returned.

Interface `Key`*:* This interface sets or challenges the application secret of the current node of the invoker `id`. $\mathcal{F}_{\text{ECGKA}}$ checks whether `id` is a group member and whether the current application secret can be challenged (i.e., $\textsf{HasKey}[\texttt{id}] = \texttt{true}$). If either check fails, then $\mathcal{F}_{\text{ECGKA}}$ exits. Otherwise, if the current application secret of the node c that `id` points to has not been set, $\mathcal{F}_{\text{ECGKA}}$ further checks whether the corresponding safe predicate $\textbf{safe}(c)$ holds. If $\textbf{safe}(c) = \texttt{true}$, then the application secret is sampled uniformly at random and marked as challenged. Otherwise, the application secret is chosen by the adversary and marked as not challenged. In both cases, $\textsf{HasKey}[\texttt{id}]$ is set to `false`, indicating that the application secret of every epoch can be challenged only once, and the new application secret is returned.

Interface `Expose`*:* The interface of the form $(\texttt{Expose}, \texttt{id})$ enables the adversary to corrupt `id`'s states and all included secrets. If the party `id` is outside the group, i.e., the pointer of `id` equals $\perp$, then $\mathcal{F}_{\text{ECGKA}}$ exits. Otherwise, `id` is a member in the group. In this case, $\mathcal{F}_{\text{ECGKA}}$ marks `id` as a corrupted member in the current commit node that `id` points to. Next, $\mathcal{F}_{\text{ECGKA}}$ updates all status flags in all existing proposal and commit nodes. Then, $\mathcal{F}_{\text{ECGKA}}$ notifies the authentication and key services $\mathcal{F}_{\mathcal{AS}}^{\mathcal{IW}}$ and $\mathcal{F}_{\mathcal{KS}}^{\mathcal{IW}}$ that the current identity and ephemeral keys are exposed. $\mathcal{F}_{\text{ECGKA}}$ marks `id` as a corrupted member in

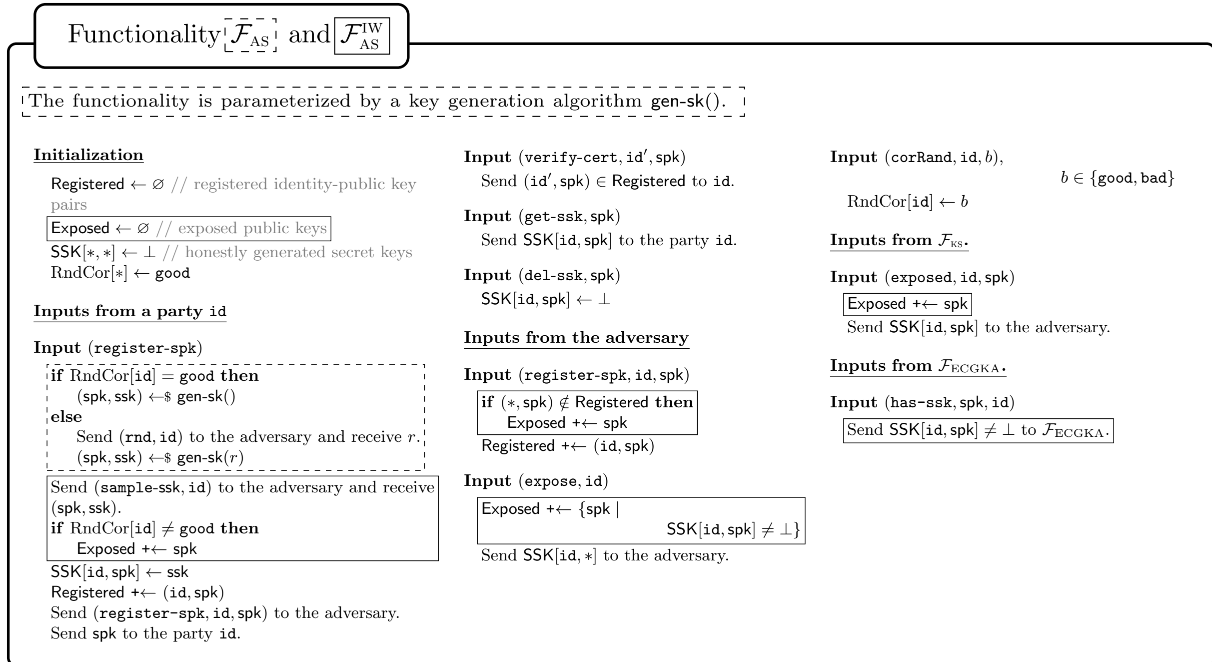

Figure 12: The Authentication functionalities $\mathcal{F}_{\mathrm{AS}}$, and their ideal-world counterparts $\mathcal{F}_{\mathrm{AS}}^{\mathrm{IW}}$. Code marked by solid or dashed boxes are executed only by the respective version, whereas code outside those boxes is shared by both variants. Taken from [4].

all commit nodes that id joins using the corrupted ephemeral keys. Note that this corruption must not infect any commit node that has been challenged. If there exists any commit c that has been challenged[4], i.e., $\mathsf{Node}[c].\mathsf{chall} = \mathtt{true}$, and its safe predicate $\mathbf{safe}(c)$ is changed to false, then all executions in this query will be undone.

Interface CorrRand*:* The interface of the form $(\mathtt{CorrRand}, \mathtt{id}, b)$ enables the adversary to manipulate or recover the randomness of the party id. If $b = \mathtt{true}$, then $\mathcal{F}_{\mathrm{ECGKA}}$ sets RndCor[id] to true, indicating that the adversary now manipulates the randomness of id. If $b = \mathtt{false}$, then $\mathcal{F}_{\mathrm{ECGKA}}$ sets RndCor[id] to false, indicating that the adversary loses control of the random number generator of the party id, i.e., all further values sampled by id are again uniformly at random.

[4]Recall that a commit c can be challenged only if its safe predicate $\mathbf{safe}(c) = \mathtt{true}$.

Functionality $\mathcal{F}_{\text{KS}}$ and $\mathcal{F}_{\text{KS}}^{\text{IW}}$

The functionality is parameterized by a key-package generation algorithm $\mathsf{gen\text{-}kp}(\mathtt{id}, \mathsf{spk}, \mathsf{ssk})$.

Initialization

```
SK[*,*], SPK[*,*], ISK[*,*] ← ⊥ // secret keys and spk's corresponding
to honestly generated keys
RndCor[*] ← good
```

Inputs from a party id

Input (register-kp, spk, ssk)

```
[dashed box]
if RndCor[id] = good then
    (kp, sk, isk) ←$ gen-kp(id, spk, ssk)
    if kp = ⊥ then return
else
    Send (rnd, id) to the adversary and receive r.
    (kp, sk) ← gen-kp(id, spk, ssk; r)
    if kp = ⊥ then return
    Send (exposed, id, spk) to F_AS.
    Send ssk to the adversary.
[solid box]
Send (sample-sk, id, spk, ssk) to the adversary and receive
(kp, sk, isk, ack).
if ¬ack then return
if RndCor[id] ≠ good then
    Send (exposed, id, spk) to F_AS.
[end box]
SK[id, kp] ← sk
SPK[id, kp] ← spk
ISK[id, kp] ← isk
Send (register-pk, id, spk, kp) to the adversary.
Send kp to the party id.
```

Input get-sks

```
Send {(kp, SK[id, kp], ISK[id, kp]) | SK[id, kp] ≠ ⊥} to id.
```

Input (get-kp, id′)

```
Send (get-kp, id, id′) to the adversary and receive kp.
Send kp to id.
```

Input (del-sk, spk)

```
SK[id, kp] ← ⊥
```

Inputs from the adversary

Input (expose, id)

```
Send SK[id, *] to the adversary.
```

Input (corRand, id, b), b ∈ {good, bad}

```
RndCor[id] ← b
```

Figure 13: The Key Service functionalities $\mathcal{F}_{\text{KS}}$, and their ideal-world counterparts $\mathcal{F}_{\text{KS}}^{\text{IW}}$. Code marked by solid or dashed boxes are executed only by the respective version, whereas code outside those boxes is shared by both variants. Reproduced from [4].

Functionality $\mathcal{F}_{\text{ECGKA}}$ and $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$: **Bookkeeping Helpers**

```
// Creating nodes
helper *create-child(c, id, p⃗, mem, stat,
        sender_type = 'member', ExtCommitProps = ⊥)
  if sender_type = 'new_member' then
     req ExtCommitProps ≠ ⊥
  c' = new node with pars ← c, orig ← id, pro ← p⃗,
         mem ← mem, stat ← stat, psk ← stat,
         epoch ← Ptr[id] + 1,
         sender_type ← sender_type,
         ExtCommitProps ← ExtCommitProps, epoch ← Node[c].epoch
  for eP ∈ ExtProps do
     if ExtProps[eP].epoch = epoch then Prop[eP].pars +← c
  return c

helper *create-root(id, mem, stat, epoch,
        sender_type = 'member')
  return new node with pars ← ⊥, orig ← id,
         pro ← ⊥, mem ← mem, stat ← stat,
         epoch ← epoch, mem ← mem,
         sender_type ← sender_type, psk ← stat,

helper *create-prop(c, id, act, stat, c_psk = ⊥)
  return new proposal with pars ← c, orig ← id,
         act ← act, stat ← stat,
         psk ← stat, c_psk ← c_psk

helper *create-ext-prop(id, act, epoch, stat)
  c_par ← {∀c s.t. Node[c].epoch = epoch c}
  return new proposal with pars ← c_par, orig ← id,
         act ← act, stat ← stat, psk ← stat,

helper *fill-props(id, p⃗)
  for p ∈ p⃗ s.t. Prop[p] = ⊥ do
     Send (Proposal, p) to the adversary and receive (orig, act, epoch).
     if epoch ≠ ⊥ then
        Prop[p] ← *create-ext-prop(orig, act, epoch, adv)
     else
        Prop[p] ← *create-prop(Ptr[id], orig, act, adv)

// Does the vector of proposals create an add-only commit?
helper *only-adds-and-psk(p⃗)
  return p⃗ ≠ () ∧ ∀p ∈ p⃗ : Prop[p] ≠ ⊥ ∧ Prop[p].act = (add-∗ ∨psk-∗)

// Does the commit require an update path?
helper *keep-spk(p⃗, force-rekey)
  return *only-adds-and-psk(p⃗) ∧ ¬force-rekey∗

// Output of process and join
helper *output-proc(c)
  (∗, propSem) ← *apply-props(c, Node[c].pro)
  return (Node[c].orig, propSem)

helper *output-join(c)
  return (Node[c].mem, Node[c].orig)

// Is the (new) spk' valid for update or commit?
helper *valid-spk(id, spk')
  spk ← Node[Ptr[id]].mem[id]
  if spk ≠ ⊥ ∧ spk' = spk then return true
  Send (has-ssk, spk', id) to F_AS and receive ack
  return ack

// Generating the group key (secure or insecure)
helper *set-key(c)
  if ¬safe(c) then
     Send (Key, id) to the adversary and receive I.
     Node[c].key ← I
     Node[c].chall ← false
  else
     Node[c].key ←$ 𝓘
     Node[c].chall ← true

// Corruptions
helper *update-stat-after-exp(id)
  for each p s.t. Prop[p] ≠ ⊥ and
       (a) Ptr[id] ∈ Prop[p].pars and
       (b) Prop[p].orig = id and
       (c) Prop[p].act = up
  do Prop[p].stat ← bad
  for each c s.t. Node[c] ≠ ⊥ and
       (a) Node[c].pars = Ptr[id] and
       (b) Node[c].orig = id
  do Node[c].stat ← bad, Node[c].psk ← bad

helper *create-ext-commit-props(parent, act, orig)
  return new proposal with pars ← parent, orig ← orig, act ← act

helper *fixate-psk-refs(props, c⃗_psk)
  pskProps = [prop ∈ props s.t. prop ≠ ⊥ and prop.act = psk-epoch-∗]
  for prop, c_psk ∈ zip(pskProps, c⃗_psk) do
     if prop.c_psk = ⊥ then
        req Node[c_psk] ≠ ⊥
        req Node[c_psk].epoch = epoch
        prop.c_psk = c_psk
  return props
```

Figure 14: Bookkeeping helper functions for $\mathcal{F}_{\text{ECGKA}}$ and $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$.

Functionality $\mathcal{F}_{\text{ECGKA}}$ and $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$: Consistency Helpers

helper *consistent-prop$(p, \text{id}, \text{act})$

// Preexisting p valid for id proposing act?
assert $\text{Prop}[p].\text{orig} = \text{id} \wedge \text{Prop}[p].\text{act} = \text{act}$
$\wedge\ \text{Ptr}[\text{id}] \in \text{Prop}[p].\text{pars}$

helper *consistent-ext-prop$(p, \text{id}, \text{act}, \text{epoch})$

// Preexisting p valid for id proposing act?
assert $\text{Prop}[p].\text{orig} = \text{id} \wedge \text{Prop}[p].\text{act} = \text{act}$
$\wedge\ \forall c \in \text{Prop}[p].\text{pars} : \text{Node}[c].\text{epoch} = \text{epoch}$

helper *valid-successor$(c, \text{id}, \vec{p}, \text{mem})$

// Preexisting node valid for id processing $(c, \vec{p})$?
assert $\text{Node}[c] \neq \perp \wedge \text{Node}[c].\text{mem} = \text{mem}$
$\wedge\ \text{Node}[c].\text{pro} \in \{\perp, \vec{p}\} \wedge\ \text{Node}[c].\text{pars} \in \{\perp, \text{Ptr}[\text{id}]\} \wedge$
$\text{Node}[c].\text{epoch} = \text{Node}[\text{Ptr}[\text{id}]].\text{epoch}$

helper *consistent-comm$(c, \text{id}, \vec{p}, \text{mem})$

// Preexisting c valid for id committing $\vec{p}$?
assert *valid-successor$(c, \text{id}, \vec{p}, \text{mem})$
assert $\text{RndCor}[\text{id}] \neq \text{good} \wedge \text{Node}[c].\text{orig} = \text{id}$

helper *attach$(c, c', \text{id}, \vec{p})$

// Attach detached root c' under new name c as successor of id's node.
assert $c' \neq \text{root}_0$
$\text{Node}[c'].\text{pars} \leftarrow \text{Ptr}[\text{id}]$; $\text{Node}[c'].\text{pro} \leftarrow \vec{p}$; $\text{Node}[c] \leftarrow \text{Node}[c']$; $\text{Node}[c'] \leftarrow \perp$
for $\text{vec_w} : \text{Wel}[\text{vec_w}] = c'$ **do** $\text{Wel}[\text{vec_w}] \leftarrow c$
for $g : \text{groupInfo}[g] = c'$ **do** $\text{groupInfo}[g] \leftarrow c$

helper *consistent-ext-comm$(c, g, \text{id}, \vec{p}, \text{mem}, \text{props})$

// Preexisting c valid for id committing $\vec{p}$?
assert $\text{Node}[c] \neq \perp \wedge \text{Node}[c].\text{mem} = \text{mem} \wedge \text{Node}[c].\text{pro} \in \{\perp, \vec{p}\}$
assert $\text{Node}[c].\text{par} \in \{\perp, \text{groupInfo}[g]\}$
assert $\text{RndCor}[\text{id}] \neq \text{good} \wedge \text{Node}[c].\text{orig} = \text{id}$
assert $\text{Node}[c].\text{ExtCommitProps} = \text{props}$

Functionality $\mathcal{F}_{\text{ECGKA}}$ and $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$: Correctness Helpers

helper *req-correctness$('\text{comm}', \text{id}, \vec{p}, \text{spk}, \text{force-rekey})$

return $\text{Node}[\text{Ptr}[\text{id}]] \neq \perp \wedge$ *members$(\text{id}, [\text{Prop}[p] : p \in \vec{p}], \text{spk})\text{newId}\perp$ // p is valid
$\wedge\ ($*valid-spk$(\text{id}, \text{spk}) \vee ($*keep-spk$(\vec{p}, \text{force-rekey}) \wedge \neg\text{force-rekey}))$
// spk is usable, unless it's an add-only commit

helper *req-correctness$('\text{proc}', \text{id}, c, \vec{p})$

return $\text{Node}[c] \neq \perp \wedge \text{Node}[c].\text{pars} = \text{Ptr}[\text{id}]$
$\wedge\ \text{Node}[c].\text{pro} = \vec{p} \wedge \text{Node}[c].\text{stat} \neq \text{adv}$
$\wedge\ \forall p \in \vec{p} : \text{Prop}[p].\text{stat} \neq \text{adv}$
$\text{Node}[c].\text{s_type} = '\text{member}' \wedge \forall p \in \vec{p}$ with $\text{Prop}[p].\text{act}$
$= \text{psk-}* : \text{Prop}[p].c_{\text{psk}} \neq \perp \wedge \text{id} \in \text{Node}[\text{Prop}[p].c_{\text{psk}}].\text{has_psk}$
$\vee\ \exists p_{rem} \in \vec{p}$ with $\text{Prop}[p_{rem}].\text{act} = \text{rem} - \text{id}$

helper *req-correctness$('\text{ext} - \text{proc}', \text{id}, c, \vec{p})$

return $\text{Node}[c] \neq \perp \wedge \text{Node}[c].\text{pro} = \vec{p} \wedge \text{Node}[c].\text{stat} \neq \text{adv} \wedge \text{Node}[c].\text{s_type} =$
$'\text{new_member}' \wedge \text{Node}[c].\text{ExtCommitProps} \neq \perp$

helper *req-correctness$('\text{prop}', \text{id}, \text{act})$

if $\text{act} = \text{rem-id}_t$ **then**
  return $\text{id}_t \in \text{Node}[\text{Ptr}[\text{id}]].\text{mem}$
else if $\text{act} = \text{up-spk}$ **then**
  return *valid-spk(id, spk)
else if $\text{act} = \text{psk-epoch-use}$ **then**
  return $\text{use} \in \{'\text{reinit}', '\text{application}'\}$
else // Adv can always deliver bad key package
  return false

helper *req-correctness$('\text{ext} - \text{prop}', \text{id}, \text{act})$

return *valid-spk(id, spk)

helper *req-correctness$('\text{ext} - \text{comm}', \text{id}, g,$
$\text{spk}, \text{resync})$

if $\text{resync} \wedge (\text{id}, *) \notin \text{Node}[\text{groupInfo}[g]].\text{mem}$ **then**
  return false
else if $(\text{id}, *) \in \text{Node}[\text{groupInfo}[g]].\text{mem}$ **then**
  return false
return *valid-spk(id, spk)

Figure 15: Consistency and correctness helper functions for $\mathcal{F}_{\text{ECGKA}}$ and $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$.

Functionality $\mathcal{F}_{\text{ECGKA}}$ and $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$: Group State Helpers

```
helper *members(c, id_c, props, spk_c)
  (G, ∗) ← *apply-props(id_c, props, spk_c)
  if (G, ∗) = ⊥ then return ⊥
  else return G

helper *apply-props(c, id_c, p⃗, spk_c)
  // Returns group members G and proposal semantics P resulting from
  applying p⃗ to state Node[c], or ⊥ if p⃗ is invalid.
  req Node[c] ≠ ⊥
  req ∀prop ∈ props : prop ≠ ⊥ ∧ c ∈ prop.pars
  req p⃗ = p⃗_up ++ p⃗_rem ++ p⃗_add ++ p_psk for some p⃗_up, p⃗_rem, p⃗_add, ++ p⃗_psk
              ∨p⃗ = p_exi ++ p_rem ++ p⃗_psk for some p_exi with len(p_exi)=1
  and p_rem with len(p_rem)=1
              with ∀act ∀p ∈ p⃗_act : Node[p].act = act-∗
  G ← Node[c].mem;
  if p_exi ∈ p ∧ ¬p_rem ∈ p then
      req (id_c, ∗) ∉ G
  else
      req (id_c, ∗) ∈ G
      G -← (id_c, ∗);
  G +← (id_c, spk_c)
  L ← {id_c} // set of updated parties
  for p ∈ p⃗_up do
      (id_s, up-spk) ← (prop.orig, prop.act)
      req id_s ∈ G \ L
      G -← (id_s, ∗); G +← (id_s, spk)
      L +← id_s
  for p ∈ p⃗_rem do
      (id_s, rem-id_t) ← (prop.orig, prop.act)
      if p_exi ∉ p then
          req id_s ∈ G ∧ id_t ∈ G \ L
          G -← (id_t, ∗)
      else
          req id_s = id_t = id_c ∧ id_s ∈ G
  for p ∈ p⃗_add do
      (id_s, add-id_t-spk_t) ← (prop.orig, prop.act)
      req id_s ∈ G ∧ id_t ∉ G ∨(id_s = id_t ∧ id_t ∉ G)
      G +← (id_t, spk_t)
  for p ∈ p⃗_psk do
      id_s, psk-epoch-use, c_psk ← (prop.orig, prop.act, prop.c_psk)
      req (∃Node[c] s.t. Node[c].epoch = epoch) ∧ (id ∈ Node[c].has_psk ∨
      ¬∃(id, ∗) ∈ G)
  P ← ((prop.orig, prop.act) : p ∈ p⃗)
  return (G, P)

helper *process-ec(c)
  Send (Process, id, c) to the adversary and
      receive (ack, orig′, spk′, p⃗, c⃗_psk).
  if ¬*req-correctness(′ext − proc′, id, c, p⃗) then
      req ack
  if Ptr[id] ≠ ⊥ then
      if Node[c] = ⊥ then
          for p ∈ p⃗ do
              Send (Proposal, p) to the adversary and receive (act).
              ExtCommitProps[p] ← *create-ext-commit-props(Ptr[id],
              act, orig′)
              req act ∈ {add-∗, rem-∗, psk-∗}
              if act ∈ {add-id_t orrem-id_t} then
                  req id_t = orig′

          try ExtCommitProps = *fixate-psk-refs(ExtCommitProps, c⃗_psk)
          for prop ∈ ExtCommitProps s.t. prop.act = psk-∗ do
              req id ∈ Node[prop.c_psk].has_psk
          mem ← *members(Ptr[id], orig′, ExtCommitProps, spk′)
          assert mem ≠ ⊥
          Node[c] ← *create-child(Ptr[id], orig′, p⃗, mem, adv,
          ′new_member′, ExtCommitProps)
      else
          id_c ← Node[c].orig
          spk_c ← Node[c].mem[id_c]
          ExtCommitProps ← Node[c].ExtCommitProps
          for prop ∈ ExtCommitProps s.t. prop.act = psk-∗ do
              req id ∈ Node[prop.c_psk].has_psk
          mem ← *members(Ptr[id], id_c, ExtCommitProps, spk_c)
          assert mem ≠ ⊥
          *valid-successor(c, id, ExtCommitProps, mem)
  else
      req Node[c] ≠ ⊥
      req Node[c].stat ≠ adv ∧ Node[c].orig = id
  assert id ∈ Node[c].mem
  Ptr[id] ← c
  HasKey[id] ← true
  Node[c].has_psk +← id
  assert cons-invariant ∧ auth-invariant
  return *output-proc(c)
```

Figure 16: Group state helper functions for $\mathcal{F}_{\text{ECGKA}}$ and $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$.

D Details on ETK and $\mathsf{ETK}^{\text{PSK}}$ Protocols

In the following sections, we provide a detailed description of our ETK and $\mathsf{ETK}^{\text{PSK}}$ protocols, highlighting modifications for adapting to external commits, external self-add proposals, and PSK proposals. These changes are color-coded as indicated in Table 2. In this section, we outline the necessary modifications to transition from Draft 12 (partially analyzed in [4]) to the final RFC draft. Additionally, we describe the simplifications and operations that were excluded from ETK and $\mathsf{ETK}^{\text{PSK}}$.

We provide the notation for the public ratchet tree and the protocol state in Tables 3 to 7. We depict the ETK protocol in Figures 17 and 18. We define helper functions in Figures 19 to 24. Changes and additions from the model of [4] are color-coded as indicated in Table 2.

$\tau.\mathsf{root}$	Returns the root.
$\tau.\mathsf{nodes}$	Returns the set of all nodes in the tree.
$v.\mathsf{isroot}$	Returns true iff $v = \tau.\mathsf{root}$.
$v.\mathsf{isleaf}$	Returns true iff v has no children.
$v.\mathsf{par}$	Returns the parent node of v (or $\perp$ if $v.\mathsf{isroot}$).
$v.\mathsf{lchild}$	Returns the left child of v (or $\perp$ if $v.\mathsf{isleaf}$).
$v.\mathsf{rchild}$	Returns the right child of v (or $\perp$ if $v.\mathsf{isleaf}$).
$v.\mathsf{sibling}$	Returns the unique node $u \neq v$ s.t. of $u.\mathsf{par} = v.\mathsf{par}$.
$v.\mathsf{nodeIdx}$	Returns the node index of v.

Table 3: Notation for the public ratchet tree. This object-oriented notation is adapted from [4].

$v.\mathsf{pk}$	The public key of a public-key encryption scheme.
$v.\mathsf{sk}$	The corresponding secret key.
$v.\mathsf{parentHash}$	A hash value binding the node to all of its ancestors.
$v.\mathsf{unmergedLvs}$	The set of leaf indices rooted below v, for which the corresponding party does not know $v.\mathsf{sk}$.
$v.\mathsf{id}$	If $v.\mathsf{isleaf}$: the identity associated with that leaf.
$v.\mathsf{leafIdx}$	If $v.\mathsf{isleaf}$: a unique identifier of the leaf, counted from the leftmost to the rightmost leaf.
$v.\mathsf{spk}$	If $v.\mathsf{isleaf}$: an associated verification key of a signature scheme.
$v.\mathsf{sig}$	If $v.\mathsf{isleaf}$: A signature of the leaf's labels under the signing key corresponding to $v.\mathsf{spk}$.

Table 4: The node labels of the ratchet tree.

Protocol ETK and ETK^{PSK} - Part 1

```
Input (Create, spk)
  req γ = ⊥ ∧ id = id_creator
  γ.groupId, γ.initSecret ←$ {0,1}^κ
  γ.epoch ← 0; γ.interimTransHash ← ε
  γ.certSpks[∗], γ.pendUp[∗], γ.pendCom[∗] ← ⊥
  γ.psk_store ← ()
  γ.τ ← new LBBT_1
  try γ.ssk ← *fetch-ssk-if-nec(γ, spk)
  (ln, sk) ←$ gen-leaf-node(id, spk, ssk, ε, 'commit',
               (γ.groupId, γ.τ.leafIdx))
  γ.τ.leaves[0].leafIdx ← 0
  γ.τ.leaves[0].assignKp(ln)
  γ.τ.leaves[0].sk ← sk

Input (Propose, up-spk)
  req γ ≠ ⊥
  try ssk ← *fetch-ssk-if-nec(γ, spk)
  (ln, sk) ←$ gen-leaf-node(id, spk, ssk, ε, 'update',
               (γ.groupId, γ.τ.leafIdx))
  P ← ('upd', ln); p ← *frameProp(γ, P)
  γ.pendUp[p] ← (ssk, sk)
  return p

Input (Propose, add-id_t)
  req γ ≠ ⊥ ∧ id_t ∉ γ.τ.roster()
  kp_t ← query (get-kp, id_t) to F_KS
  try γ ← *validate-kp(γ, kp_t, id_t, 'key_package', ε)
  P ← ('add', kp_t); p ← *frameProp(γ, P)
  return p

Input (Commit, p⃗, spk, force-rekey, wel_type)
  req γ ≠ ⊥
  γ' ← *init-epoch(γ)
  try (γ', upd, rem, add, psk) ← *apply-props(γ, γ', p⃗)
  req (∗, 'rem'-id) ∉ rem ∧ (id, ∗) ∉ upd
  if force-rekey ∨ p⃗ = () ∨ upd ≠ () ∨ rem ≠ () then
    try (γ', commitSec, updatePath, pathSecs)
                              ← *rekey-path(γ', id, spk)
  else
    commitSec ← 0; updatePath ← ε
    pathSecs[∗] ← ε
  propIDs ← ()
  for p ∈ p⃗ do
    propIDs +← Hash(p)
  C ← (propIDs, updatePath)
  sig ← *signCommit(γ, C)
  γ' ← *set-conf-trans-hash(γ, γ', γ.leafIdx(), C, sig)
  pskIds = *get-pskIds(p⃗)
  try psk_secret ← *derive-psk-secret(pskIds)
  (γ', confKey, joinerSec)
        ← *derive-keys(γ, γ', commitSec, psk-secret)
  confTag ← *conf-tag(γ', confKey)
  if rem ≠ () then
    membTag ← MAC.tag(γ.membKey, C)
  else membTag ← ⊥
  c ← *frameCommit(γ, C, confTag, sig, membTag)
  if add ≠ () then
    (γ', vec_w) ← *welcome-msg(γ', add, joinerSec,
                              pathSecs, confTag, pskIds)
  else vec_w ← ⊥
  g ← *external-GroupInfo(γ', confTag)
  γ' ← *set-interim-trans-hash(γ', confTag)
  γ.pendCom[c] ← (γ', p⃗, upd, rem, add, psk)
  return (c, vec_w, g)

Input (Propose, rem-id_t)
  req γ ≠ ⊥ ∧ id_t ∈ γ.τ.roster()
  leafIdx_t ← γ.τ.leafof(it_t)
  P ← ('rem', leafIdx_t); p ← *frameProp(γ, P)
  return p

Input (Propose, extAdd-id-spk, epoch)
  req γ = ⊥
  try ssk ← query (get-ssk, spk) to F_AS
  (kp, sk, isk) ←$ gen-kp(id, spk, ssk)
  P ← ('add', kp)
  p ← *frameExtProp(groupId, epoch, ssk, P)
  γ.pendAdd[epoch] ← (kp, sk, isk)
  return p

Input (Propose, psk-epoch-use)
  req γ ≠ ⊥
  preSharedKeyId ← *gen-pskId(epoch, use)
  P ← ('psk', preSharedKeyId)
  p ← *frameProp(γ, P)
  return p

Input (Process, c, p⃗)
  req γ ≠ ⊥
  if p⃗ = ⊥ then
    (id_c, (exi, rem)) ← *process-ec(c, p)
    ordered_proposals = exi ++ rem
    return (id_c, ordered_proposals)
  (senderIdx, C, confTag, sig, membTag)
                              ← *unframeCommit(γ, c, sig)
  id_c ← γ.τ.leaves[senderIdx].ID
  if senderIdx = γ.leafIdx() then
    parse (γ', p⃗', upd, rem, add, psk) ← γ.pendCom[c]
    req p⃗ = p⃗'
    γ = γ'
    return (id_c, upd ++ rem ++ add ++ psk)
  parse (propIDs, updatePath) ← C
  req ∀i ∈ [|p⃗|] : Hash(p⃗[i]) = propIDs[i]
  γ' ← *init-epoch(γ)
  try (γ', upd, rem, add, psk) ← *apply-props(γ, γ', p⃗)
  req (∗, id_c) ∉ rem ∧ (id_c, ∗) ∉ upd
  if (∗, 'rem'-id) ∈ rem then
    req MAC.vrf(γ.membKey, c)
    γ ← ⊥
  else
    if updatePath ≠ ε then
      (γ', commitSec)
            ← *apply-rekey(γ', senderIdx, updatePath)
    else
      req p⃗ ≠ () ∧ upd = () ∧ rem = ()
      commitSec ← 0
    γ' ← *set-conf-trans-hash(γ, γ', senderIdx, C, sig)
    pskIds = *get-pskIds(p⃗)
    req psk-secret ← *derive-psk-secret(pskIds)
    (γ', ∗) ← *derive-keys(γ, γ',
                 commitSec, psk-secret)
    req *vrf-conf-tag(γ', confKey, confTag)
    γ ← *set-interim-trans-hash(γ', confTag)
  return (id_c, upd ++ rem ++ add ++ psk)
```

Figure 17: The UC model ETK and ETK^{PSK} as run by party id. The group creator's identity $\mathtt{id}_{creator}$ is encoded a part of the instance's session identifier. As elsewhere we use the color coding from Table 2. Note that code highlighted in orange (●) is only included in ETK^{PSK}, and not in ETK.

Protocol ETK and ETK$^{\mathsf{PSK}}$ - Part 2

```
Input (Join, vec_w)
  req γ = ⊥∨γ.pendAdd[epoch] ≠ ⊥
  parse (encGroupSecs, encGroupInfo, wel_type) ← vec_w
  γ.certSpks[∗], γ.pendUp[∗], γ.pendCom[∗] ← ⊥
  kbs ← query get-sks to F_KS
  kbs ++ γ.pendAdd
  joinerSec, pathSec ← ⊥
  for e ∈ encGroupSecs do
    parse (hash, cipher) ← e
    for (kp, sk, isk) ∈ kbs do
      if hash = Hash(kp) then
        v.sk ← sk
        req v.kp() = kp
        parse (joinerSec, pathSec, pskIds) ← PKE.Dec(isk, cipher)
  req joinerSec ≠ ⊥
  try psk-secret ← *derive-psk-secret(pskIds)
  s ← HKDF.Extract(joinerSec, psk-secret)
  welcome_secret = HKDF.Expand(s, 'welcome')
  welcome_nonce = HKDF.Expand(welcome_secret, 'nonce')
  welcome_key = HKDF.Expand(welcome_secret, 'key')
  groupInfo ← PKE.Dec(welcome_key, welcome_nonce,
             encGroupInfo)
  parse (groupInfoTBS, sig) ← groupInfo
  parse (groupCtxt, γ.τ, confTag, senderIdx) ← groupInfoTBS
  parse (γ.groupId, γ.epoch, γ.treeHash, γ.confTransHash,
         γ.interimTransHash) ← groupCtxt
  req Sig.Vrf(γ.τ.leaves[senderIdx].spk, sig, groupInfoTBS)
  if ∃epoch s.t. v.kp() = γ.pendAdd[epoch] then
    req γ.epoch = epoch
    γ.pendAdd[epoch]epoch⊥
  try γ ← *vrf-tree-state(γ)
  v ← γ.τ.leaves[γ.leafIdx()]
  try γ.ssk ← *fetch-ssk-if-nec(γ, v.spk)
  kbs ← query get-sks to F_KS
  if pathSec ≠ ε then
    v ← γ.τ.lca(γ.leafIdx(), senderIdx)
    while v ≠ ⊥ do
      nodeSec ← HKDF.Expand(pathSec, 'node')
      (sk, v.sk) ← PKE.Kg(nodeSec)
      req v.sk = sk
      pathSec ← HKDF.Expand(pathSec, 'path')
      v ← v.par
  try psk-secret ← *derive-psk-secret(pskIds)
  joiner_psk ← HKDF.Extract(joinerSec, psk-secret)
  (γ, confKey) ← *derive-epoch-keys(γ, joinerSec,
                 psk-secret)
  req *vrf-conf-tag(γ, confKey, confTag)
  return (γ.τ.roster(), γ.τ.leaves[senderIdx].id)

Input Key
  req γ ≠ ⊥
  (k, γ.appSecret) ← (γ.appSecret, ⊥)
  return k
```

```
Input (ExtCommit, g, spk, e⃗_psk, resync)
  if resync then
    req γ ≠ ⊥
  else
    req γ = ⊥
    γ.certSpks[∗], γ.pendUp[∗], γ.pendCom[∗] ← ⊥
  parse (extGroupInfoTBS, sig) ← g
  parse (γ.groupId, γ.epoch, γ.treeHash, γ.confTransHash,
        γ.interimTransHash, γ.τ, γ.τ.confTag, senderIdx,
        external-pub) ← extGroupInfoTBS
  req Sig.Vrf(y.τ.leaves[senderIdx].spk, sig,
        extGroupInfoTBS)
  try γ ← *vrf-tree-state(γ)
  (kem_output, context) ← setupBaseS(external-pub, "")
  γ.initSecret =context.export("MLS 1.0 external
        init secret", KDF.Nh)
  γ' ← *init-epoch(γ)
  p⃗ = ()
  p⃗ ++← ('extInit', kem_output)
  if resync then
    req id ∈ γ.τ.roster()
    p⃗ ++← ('rem', γ.leafInd)
  for epoch ∈ e_psk do
    preSharedKeyId ← *gen-pskId(epoch, 'application')
    p⃗ ++← ('psk', preSharedKeyId)
  try (γ', exi, rem, psk) ← *apply-ec-props(γ, γ', p⃗, id)
  try (γ', commitSec, updatePath, pathSecs) ←
   *rekey-path-upon-join(γ', id, spk)
  C ← (propIDs, updatePath)
  sig ← *signExtCommit(γ, C, ssk)
  γ' ← *set-conf-trans-hash(γ, γ', γ.leafIdx(), C, sig)
  pskIds = *get-pskIds(p⃗)
  try psk-secret ← *derive-psk-secret(pskIds)
  (γ', confKey, joinerSec)
        ← *derive-keys(γ, γ', commitSec, psk-secret,
        kem_output)
  confTag ← *conf-tag(γ', confKey)
  ec ← *frameExtCommit(γ, C, γ'.ssk, sig, confTag)
  g ← *external-GroupInfo(γ', confTag)
  γ' ← *set-interim-trans-hash(γ', confTag)
  γ.pendCom[ec] ← (γ', p⃗, upd, rem, add, psk)
  return (ec, g)
```

Figure 18: The UC model ETK and ETK$^{\mathsf{PSK}}$ as run by party `id`. The group creator's identity $\mathtt{id}_{\mathsf{creator}}$ is encoded as part of the instance's session identifier.

Protocol ETK and ETK$^{\mathsf{PSK}}$: Setup Algorithms

```
Algorithm gen-sk
  (spk, ssk) ← Sig.Kg()
  return(spk, ssk)
Algorithm gen-kp(id, spk, ssk, parentHash = ε, source = ε; r)
  (leafNode, sk) ← gen-leaf-node(id, spk, ssk, parentHash,
              'key_package', source; r)
  (init-key, isk) ← PKE.Kg()
  kpTBS ← (leafNode, initKey)
  sig ← Sig.Sign(ssk, kpTBS)
  kp ← (kpTBS, sig)
  return (kp, sk, isk)
```

```
Algorithm gen-leaf-node(id, spk, ssk,
            parentHash = ε, ln_source, source = ε; r)
  (pk, sk) ← PKE.Kg(r)
  req ln_sourcein{'key_package', 'update', 'commit'}
  if ln_source ≠ 'key_package' then req source ≠ ε
  leafNodeTBS ← (id, pk, spk, parentHash, ln_source, source)
  sig ← Sig.Sign(ssk, leafNodeTBS)
  leafNode ← (id, pk, spk, parentHash, ln_source, sig)
  return (leafNode, sk)
```

Figure 19: The algorithms gen-sk and gen-kp, used by $\mathcal{F}_{\mathsf{AS}}$ and $\mathcal{F}_{\mathsf{KS}}$, respectively.

τ.clone()	Returns and (independent) copy of τ.
τ.public()	Returns a copy of τ for which all private labels (v.sk) are set to $\bot$.
τ.roster()	Returns the identities of all parties in the tree.
τ.leaves[leafIdx]	Returns the leaf with identifier leafIdx.
τ.leafof(id)	Returns the leaf identifier of the v for which v.id = id.
τ.allotLeaf()	Finds the leaf v with the lowest nodeIdx for which $\neg v$.inuse(), or adds a new leaf v using addLeaf. Returns leafIdx.
τ.directPath(leafIdx)	Returns the direct path, excluding the leaf, as an ordered list from the leaf to root.
τ.coPath(leafIdx)	Returns the co-path to τ.directPath(leafIdx) as an ordered list.
τ.lca(leafIdx$_1$, leafIdx$_2$)	Returns the least common ancestor of the two leafs.
τ.blankPath(leafIdx)	Calls v.blank() on all $v \in \tau$.directPath(leafIdx).
τ.mergeLeaves(leafIdx)	Sets v.unmergedLvs $\leftarrow \varnothing$ for all $v \in \tau$.directPath(leafIdx)
τ.unmergeLeaf(leafIdx)	Sets v.unmergedLvs $+\!\!\leftarrow$ leafIdx for all v returned by τ.directPath(leafIdx)
v.kp()	Returns (v.id, v.pk, v.spk, v.parentHash, v.sig) (undefined if $\neg v$.isleaf).
v.assignKp(kp)	Sets (v.id, v.pk, v.spk, v.parentHash, v.sig) from kp (only allowed if v.isleaf).
v.inuse()	Returns `false` iff all labels except parentHash are $\bot$.
v.blank()	Sets all labels except parentHash to $\bot$.
v.resolution()	Return $\begin{cases} (v) \mathbin{+\!\!+} v.\mathsf{unmergedLvs} & \text{if } v.\mathsf{inuse}() \\ v.\mathsf{lchild.resolution}() \mathbin{+\!\!+} v.\mathsf{lchild.resolution}() & \text{else if } \neg v.\mathsf{isleaf} \\ () & \text{else.} \end{cases}$
v.resolvent(u)	For a descendant u of v, returns the (unique) node in v.resolution() which is an ancestor of u.

Table 5: Helper methods defined on the ratchet tree.

γ.groupId	An identifier of the group.
γ.epoch	The current epoch number.
$\gamma.\tau$	The labeled left-balanced binary tree.
γ.treeHash	A hash of (the public part) of τ.
γ.confTransHash	The confirmed transcript hash.
γ.interimTransHash	The interim transcript hash for the next epoch.
γ.ssk	The current signing key.
γ.certSpks[$*$]	A mapping associating the set of validated signature verification keys to each party id$'$.
γ.pendUp[$*$]	A mapping associating the secret keys for each pending update proposal issued by id.
γ.pendCom[$*$]	A mapping associating the new group state for each pending commit issued by id.
γ.appSecret	The current epoch's ECGKA key.
γ.membKey	The key used to MAC packages.
γ.initSecret	The next epoch's init secret.

Table 6: The protocol state.

γ.leafIdx()	Returns $\gamma.\tau$.leafof(id).leafIdx.
γ.groupCtxt()	Returns (γ.groupId, γ.epoch, γ.treeHash, γ.confTransHash).

Table 7: Helper methods on the protocol state.

Protocol ETK and $\mathsf{ETK}^{\mathsf{PSK}}$: Commit. Process, and Join Helpers - Part 1

```
helper *init-epoch(γ)
  γ' ← γ.clone()
  γ'.epoch ← γ'.epoch + 1
  γ'.pendUp[*], γ'.pendCom[*] ← ⊥
  return γ'

helper *rekey-path(γ', id, spk)
  leafIdx = γ'.leafIdx()
  (updatePathNodes, pathSecs, commitSec,
          leafNodeSec) ← *derive-rekey-values(γ', id, spk, leafIdx)
  try ssk ← *fetch-ssk-if-nec(γ', spk)
  v ← γ'.τ.leaves[γ'.leafIdx()]
  r ← leafNodeSec
  (ln, sk) ← gen-leaf-node(id, spk, ssk, v.parentHash,
          ('commit', γ'.groupId, γ'.τ.leafIdx()); r)
  v.assignKp(ln)
  v.sk ← sk
  γ' ← *set-tree-hash(γ')
  updatePath ← (ln, updatePathNodes)
  return (γ', commitSec, updatePath, pathSecs)

helper *derive-rekey-values(γ', id, spk, leafIdx)
  directPath ← γ'.τ.directPath(γ'.leafIdx())
  coPath ← γ'.τ.coPath(γ'.leafIdx())
  updatePathNodes ← ()
  pathSecs[*] ← ⊥
  leafSec ←$ {0,1}^κ
  leafNodeSec ← HKDF.Expand(leafSec, 'node')
  pathSec ← HKDF.Expand(leafSec, 'path')
  for (v, c) ∈ zip(directPath, coPath) do
    pathSecs[v] ← pathSec
    nodeSec ← HKDF.Expand(pathSec, 'node')
    (v.pk, v.sk) ← PKE.Kg(nodeSec)
    encPathSecs ← ()
    for t ← c.resolution() do
      encPathSecs +← PKE.Enc(t.pk, pathSec)
    updatePathNodes +← (v.pk, encPathSecs)
    pathSec ← HKDF.Expand(pathSec, 'path')
  commitSec ← pathSec
  γ'.τ.mergeLeaves(γ'.leafIdx())
  γ' ← *set-parent-hash(γ', γ'.leafIdx())
  return (updatePathNodes, pathSecs, commitSec,
          leafNodeSec)

helper *rekey-path-upon-join(γ', id, spk)
  try ssk ← *fetch-ssk-if-nec(γ', spk)
  v ← γ'.τ.leaves[γ'.leafIdx()]
  r ← leafNodeSec
  (ln, sk) ← gen-kp(id, spk, ssk, ε, 'commit', (γ'.groupId,
          γ'.τ.leafIdx()); r)
  γ'.τ.leaves[γ'.leafIdx()].assignKp(ln)
  γ'.τ.leaves[γ'.leafIdx()].sk ← sk
  (updatePathNodes, pathSecs, commitSec, leafNodeSec) ←
          *derive-rekey-values(γ', id, spk, leafIdx)
  γ' ← *set-tree-hash(γ')
  updatePath ← (ln, updatePathNodes)
  return (γ', commitSec, updatePath, pathSecs)

helper *apply-rekey(γ', senderIdx, updatePath)
  parse (ln, updatePathNodes) ← updatePath
  (γ', commitSec) ← *derive-apply-rekey-values(γ',
          updatePathNodes, senderIdx)
  commitSec ← pathSec
  v ← γ'.τ.leaves[senderIdx]
  try γ' ← *validate-ln(γ', ln, v.id, 'commit', v.parentHash)
  v.assignKp(kp)
  γ' ← *set-tree-hash(γ')
  return (γ', commitSec)

helper *derive-apply-rekey-values(γ', updatePathNodes,
          senderIdx)
  directPath ← γ'.τ.directPath(senderIdx)
  coPath ← γ'.τ.coPath(senderIdx)
  lca ← γ'.τ.lca(γ'.leafIdx(), senderIdx)
  for (v, c, updatePathNode) ∈ zip(directPath,
                                   coPath, updatePathNodes) do
    parse (v.pk, encPathSecs) ← updatePathNode
    if v = lca then
      r ← c.resolvent(γ'.τ.leaves[γ'.leafIdx()])
      i ← c.resolution().indexofr
      pathSec ← PKE.Dec(r.sk, encPathSecs[i])
    if pathSec ≠ ⊥ then
      nodeSec ← HKDF.Expand(pathSec, 'node')
      (pk, v.sk) ← PKE.Kg(nodeSec)
      req v.pk = pk
      pathSec ← HKDF.Expand(pathSec, )
  commitSec ← pathSec
  γ'.τ.mergeLeaves(senderIdx)
  γ' ← *set-parent-hash(γ', senderIdx)
  return (γ', commitSec)

helper *apply-ec-rekey(γ', id_c, updatePath)
  parse (ln, updatePathNodes) ← updatePath
  try leafIdx ← γ.τ.allotLeaf()
  try γ' ← *validate-ln(γ', ln, id_c, 'commit', ε)
  γ'.τ.leaves[senderIdx].assignKp(ln)
  (γ', commitSec) ← *derive-apply-rekey-values(γ',
          updatePathNodes, senderIdx)
  γ' ← *set-tree-hash(γ')
  return (γ', commitSec)

helper *external-GroupInfo(γ', confTag)
  externalGroupInfoTBS ← (γ'.groupId, γ'.epoch, γ'.treeHash,
          γ'.confTransHash, γ'.interimTransHash, γ'.τ.public(),
          confTag, γ'.leafIdx(), γ'.external_pub)
  sig ← Sig.Sign(γ'.ssk, externalGroupInfoTBS)
  g ← (externalGroupInfoTBS, sig)
  return (γ', g)
```

Figure 20: Helper methods related to the commit and welcome messages.

Protocol ETK and ETK^{PSK} : Commit. Process, and Join Helpers - Part 2

```
helper *truncate-tree(γ')
  v ← γ'.τ.leaves[|γ'.τ.leaves| − 1]
  if ¬v.inuse()
                  ∧ (¬v.par.inuse() ∨ v.par = γ'.τ.root) then
    γ'.τ ← τ.pruneRightmost()
    γ' ← *truncate-tree(γ')
  return γ'
helper *apply-ec-props(γ, γ', p⃗, id_s)
  exi, rem, psk ← ()
  pskIds ← ∅
  for (type, val) ∈ p⃗ do
    if type = 'rem' then
      id_t ← γ.τ.leaves[val].id
      req γ'.τ.leaves[val] ≠ ⊥
      req γ'.τ.leaves[val].inuse() ∧ id_s = id_t
      γ'.τ.leaves[val].blank()
      γ'.τ.leaves[val].blankPath(val)
      γ' ← *truncate-tree(γ')
      rem +← (id_s, 'rem'-id_t)
    else if type = 'psk' then
      parse (∗, psk_data, ∗) ← val
      req γ.psk_store.last((psk_data.groupId,
                psk_data.epoch)) ≠ ⊥
      usage = psk_data.usage
      req usage = 'application'
      psk +← (id_s, 'psk'-epoch-usage)
    else if type = 'exi' then
      exi +← (id_s, val)
    else return ⊥
  return (γ', exi, rem, psk)
helper *gen-pskId(epoch, usage)
  psk_data ← (usage, γ.groupId, epoch)
  nonce ←$ {0, 1}^κ
  return ('resumption', psk_data, nonce)
helper *get-pskIds(p⃗)
  pskIds ← ()
  for p ∈ p⃗ do
    (sender, P) ← *unframeProp(γ, p)
    parse (∗, id) ← P
    pskIds ← id
  return pskIds
helper *derive-psk-secret(pskIds)
  pskSec[0] = 0
  for (i, pskId) ∈ pskIds do
    (∗, pskData, ∗) ← pskId
    psk_[i] = γ.psk_store.last((pskData.groupId,
                pskData.epoch))
    req psk_[i] ≠ ⊥
    psk_extracted[i] = HKDF.Extract(0, psk_[i])
    PSKLabel = (pskId, i, len(pskIds))
    psk_input[i] = HKDF.Expand(psk_extracted[i], 'derivedpsk')
    pskSec[i + 1] = HKDF.Extract(psk_input[i], pskSec[i])
  return pskSec[−1]
```

```
helper *apply-props(γ, γ', p⃗)
  upd, rem, add, psk ← ()
  pskIds ← ∅
  for p ∈ p⃗ do
    if p.sender = 'member' then
      try (senderIdx, P) ← *unframeProp(γ, p)
      id_s ← γ.τ.leaves[senderIdx].id
    else if p.sender = 'new_member_proposal' then
      try (id_s, P) ← *unframeExtProp(γ, p)
    parse (type, val) ← P
    if type = 'upd' then
      req (id_s, ∗) ∉ upd ∧ rem = () ∧ add = ()
      try γ' ← *validate-ln(γ', val, id_s, ε)
      γ'.τ.leaves[senderIdx].assignKp(val)
      γ'.τ.blankPath(senderIdx)
      if senderIdx = γ.leafIdx() then
        parse (ssk, sk) ← γ.pendUp[p]
        γ'.τ.leaves[senderIdx].sk ← sk
        γ'.ssk ← ssk
      spk ← γ'.τ.leaves[senderIdx].spk
      upd +← (id_s, 'upd'-spk)
    else if type = 'rem' then
      id_t ← γ.τ.leaves[val].id
      req γ'.τ.leaves[val] ≠ ⊥
      req γ'.τ.leaves[val].inuse() ∧ (id_t, ∗) ∉ upd ∧ add = ()
      γ'.τ.leaves[val].blank()
      γ'.τ.leaves[val].blankPath(val)
      γ' ← *truncate-tree(γ')
      rem +← (id_s, 'rem'-id_t)
    else if type = 'add' then
      parse (id_t, ∗, spk, ∗, ∗) ← val
      req id_t ∉ γ'.τ.roster()
      try γ' ← *validate-kp(γ', val, id_t, 'key_package'ε)
      try γ'.τ.allotLeaf()
      γ'.τ.leaves[leafIdx()].assignKp(val)
      γ'.τ.unmergeLeaf(leafIdx())
      add +← (id_s, 'add'-id_t-spk)
    else if type = 'psk' then
      req val ∉ PSKIds
      parse(∗, psk_data, ∗) ← val
      req γ.psk_store.last((psk_data.groupID,
                psk_data.epoch)) ≠ ⊥
      usage = psk_data.usage
      psk +← (id_s, 'psk'-epoch-usage)
      PSKIds += val
    else return ⊥
  return (γ', upd, rem, add, psk)
```

Figure 21: The helper methods related to creating and processing the commit and welcome messages.

Protocol ETK and ETK^{PSK} : Commit. Process, and Join Helpers - Part 3

```
helper *welcome-msg(γ, γ', add, joinerSec,
pathSecs, confTag, pskIds, wel_type)
  groupCtxt ← (γ'.groupId, γ'.epoch, γ'.treeHash,
      γ'.confTransHash)
  groupInfoTBS ← (groupCtxt, γ'.τ.public(), confTag, γ'.leafIdx())
  sig ← Sig.Sign(γ'.ssk, groupInfoTBS)
  groupInfo ← (groupInfoTBS, sig)
  try psk-secret ← *derive-psk-secret(pskIds)
  s ← HKDF.Extract(joinerSec, psk-secret)
  welcome_secret = HKDF.Expand(s, 'welcome')
  welcome_nonce = HKDF.Expand(welcome_secret, 'nonce')
  welcome_key = HKDF.Expand(welcome_secret, 'key')
  encGroupInfo ← PKE.Enc(welcome_key, welcome_nonce,
  groupInfo)
  encGroupSecs ← ()
  for (∗, 'add'-id_t-spk_t, ipk-id_t) ∈ add do
    leafIdx_t ← γ'.τ.leafof(id_t)
    v_t ← γ'.τ.leaves[leafIdx_t]
    lca ← γ'.τ.lca(γ'.leafIdx(), leafIdx_t)
    encGroupSec ← PKE.Enc(ipk-id_t, (joinerSec,
    pathSecs[lca], pskIds))
    encGroupSecs +← (Hash(v_t.kp), encGroupSec)
  vec_w ← *parse-wel(encGroupSecs, encGroupInfo, wel_type)
  return (γ', vec_w)

helper *vrf-tree-state(γ')
  req γ'.treeHash = *tree-hash(γ'.τ.root)
  for v ∈ γ'.τ.nodes : v.inuse() ∧ ¬v.isleaf do
    lchild ← v.lchild
    rchild ← *origRChild(v.rchild)
    phr ← *parent-hash-cochild(v, v.rchild)
    phl ← *parent-hash-cochild(v, v.lchild)
    req (lchild.inuse ∧ lchild.parentHash = phr)
        ∨ (rchild.inuse ∧ rchild.parentHash = phl)
  mem ← ∅
  for v ∈ γ'.τ.nodes : v.inuse() ∧ v.isleaf do
    req v.id ∉ mem
    mem +← v.id
    try γ' ← *validate-kp(γ', v.kp(), v.id, v.parentHash)
  return γ'

helper *origRChild(v)
  if v.inuse ∨ v.isleaf then return v
  else return *origRChild(v.lchild)
```

```
helper *parse-wel(encGroupSecs, encGroupInfo, wel_type)
  if wel_type = all then
    vec_w +← (encGroupSecs, encGroupInfo)
  else if wel_type = one then
    for all encGroupSec ∈ encGroupSecs do
      vec_w +← (encGroupSec, encGroupInfo)
  else if wel_type = batch then
    for all encGroupSec stated as batched encGroupSecs do
      vec_w +← (encGroupSec, encGroupInfo)
  return vec_w

helper *process-ec(p⃗)
  (id_c, C, confTag, sig, membTag) ←
          *unframeExtCommit(γ, ec, sig)
  (p⃗, updatePath) ← C
  if id_c = id then
    parse (γ', p⃗, ext − init, rem, psk) ← γ.pendCom[c]
    req p⃗ = p⃗
    return(id_c, ext − init ++ rem ++ psk)
  γ' ← *init-epoch(γ)
  try (γ', ext − init, rem) ← *apply-ec-props(γ, γ', p⃗)
  req len(rem) = 1 ∧ len(ext − init) = 1
  (∗, kem_output) ← parse ext − init
  context = SetupBaseR(kem_output, γ.external_priv, "")
  initSecret = context.export("MLS 1.0 external init secret", KDF.Nh)
  req updatePath ≠ ⊥
  (γ', senderIdx, commitSec) ← *apply-ec-rekey(γ', id_c, updatePath)
  γ' ← *set-conf-trans-hash(γ, γ', senderIdx, C, sig)
  pskIds = *get-pskIds(p⃗)
  req psk-secret ← *derive-psk-secret(pskIds)
  (γ', ∗) ← *derive-keys(γ, confKey, γ', commitSec,
          psk-secret)
  req *vrf-conf-tag(γ', confKey, confTag)
  γ ← *set-interim-trans-hash(γ', confTag)
  return (id_c, upd ++ rem ++ add ++ psk)
```

Figure 22: Helper methods related to commit and welcome messages.

Protocol ETK and ETK^{PSK} : Confirmation-Tag

```
helper *conf-tag(γ', confKey)
  return MAC.tag(confKey, γ'.confTransHash)

helper *vrf-conf-tag(γ', confKey, confTag)
  return MAC.vrf(confKey, confTag, γ'.confTransHash)
```

Protocol ETK and ETK^{PSK} : Tree-Hash

```
helper *set-parent-hash(γ', leafIdx)
  path ← γ'.τ.directPath(leafIdx)
  path ← path.reverse()
  path +← γ'.τ.leaves[leafIdx]
  for v ∈ path do
    if v.isroot then
      v.parentHash ← ε
    else
      v.parentHash ← *parent-hash-cochild(v.par, v.sibling)
  return γ'

helper *parent-hash-cochild(v, u)
  origChildResolution ← u.resolution() \ u.par.unmergedLvs
  return Hash(v.pk, v.parentHash, origChildResolution)

helper *set-tree-hash(γ')
  γ'.treeHash ← *tree-hash(γ'.τ.root)
  return γ'

helper *tree-hash(v)
  if v.isleaf then
    return Hash(v.nodeIdx, v.kp())
  else
    leftHash ← *tree-hash(v.lchild)
    rightHash ← *tree-hash(v.rchild)
    return Hash(v.nodeIdx, v.pk, v.unmergedLvs,
                v.parentHash, leftHash, rightHash)
```

Protocol ETK and ETK^{PSK} : Key-Schedule

```
helper *derive-keys(γ, γ', commitSec, psk-secret,
kem_output)
  if kem_output ≠ ⊥ then
    context = SetupBaseR(kem_output,
              γ.external_priv, "")
    initSecret = context.export("MLS
              1.0externalinitsecret", KDF.Nh)
  else
    initSecret = γ.initSecret
  s ← HKDF.Extract(γ.initSecret, commitSec)
  joinerSec ← HKDF.Expand(s, γ'.groupCtxt())
  joiner_psk ← HKDF.Extract(joinerSec, psk-secret)
  (γ', confKey) ← *derive-epoch-keys(γ', joiner_psk)
  return (γ', confKey, joinerSec)

helper *derive-epoch-keys(γ', joinerSec)
  s ← HKDF.Expand(γ.joinerSec, 'member')
  memberSec ← HKDF.Extract(s, 0)
  e ← HKDF.Expand(memberSec, 'epoch')
  epochSec ← HKDF.Extract(e, γ'.groupCtxt())
  confKey ← HKDF.Expand(epochSec, 'confirm')
  γ'.appSecret ← HKDF.Expand(epochSec, 'app')
  γ'.membKey ← HKDF.Expand(epochSec, 'membership')
  γ'.initSecret ← HKDF.Expand(epochSec, 'init')
  extSecret ← HKDF.Expand(epochSec, 'external')
  γ'.external_priv, γ'.external_pub ← PKE.Kg(extSecret)
  resumptionPsk ←
        HKDF.Expand(epochSec, 'resumption')
  γ'.resumptionPsk ← resumptionPsk
  γ'.psk_store((γ'.groupId, γ'.epoch)) ++ resumptionPsk
  return (γ', confKey)
```

Protocol ETK and ETK^{PSK} : Setup Interaction

```
helper *fetch-ssk-if-nec(γ, spk)
  if γ.τ.leaves[γ.leafIdx()].spk ≠ spk then
    ssk ← query (get-ssk, spk) to F_AS
  else
    ssk ← γ.ssk
  return ssk

helper *validate-kp(γ, kp, id, kp_source, parentHash)
  parse (sig, ln, init-key) ← kp
  req Sig.Vrf(spk, sig, (ln, init-key))
  γ ← *validate-ln(γ, ln, id, kp_source, parentHash)
  return γ

helper *validate-ln(γ, ln, id, ln_source, parentHash)
  parse (id, pk, spk, parentHash', ln_source', source,
         sig) ← ln
  req ln_source = ln_source'
  if spk ∉ γ.certSpks[id] then
    succ ← query (verify-cert, id', spk) to F_AS
    req succ
    γ.certSpks[id] +← spk
  req Sig.Vrf(spk, sig, (id, pk, spk, parentHash))
  return γ
```

Figure 23: Various helper methods for the protocol ETK and ETK^{PSK}.

Protocol ETK and ETK$^{\text{PSK}}$: Message-Framing

```
helper *signCommit(γ, C)
  tbs ← (γ.groupCtxt(), γ.groupId, γ.epoch, γ.leafIdx(),
            'commit', C)
  sig ← Sig.Sign(γ.ssk, tbs)
  return sig

helper *frameCommit(γ, C, confTag, sig, membTag)
  return (γ.groupId, γ.epoch, γ.leafIdx(), 'commit',
            C, confTag, sig, membTag)

helper *unframeCommit(γ, c)
  parse (groupId, epoch, senderIdx, contentType,
            C, confTag, sig, membTag) ← c
  req contentType = 'commit' ∧ groupId = γ.groupId
  reqepoch = γ.epoch
  tbs ← (γ.groupCtxt(), groupId, epoch, senderIdx,
                                  'commit', C)
  req γ.τ.leaves[senderIdx] ≠ ⊥
       ∧ γ.τ.leaves[senderIdx].inuse()
       ∧ Sig.Vrf(γ.τ.leaves[senderIdx].spk, sig, tbs)
  return (senderIdx, C, confTag, sig, membTag)

helper *signExtCommit(γ, C, ssk)
  framedContent ← (γ.groupId, γ.epoch,
            'new_member_commit', 'commit', C)
  sig ← Sig.Sign(γ.ssk, framedContent, γ.groupCtxt())
  return sig

helper *frameExtCommit(γ, C, confTag, sig, ssk)
  framedContent ← (γ.groupId, γ.epoch,
            'new_member_commit', 'commit', C)
  framedContentAuth ← (sig, confTag)
  return (framedContent, framedContentAuth)

helper *unframeExtCommit(γ, ec, sig)
  parse(framedContent, framedContentAuth) ← ec
  parse (γ.groupId, γ.epoch, type, contentType, C) ← framedContent
  req contentType = 'commit' ∧ type = 'new_member_commit'
  reqepoch = γ.epoch ∧ groupId = γ.groupId
  parse (sig, confTag) ← framedContentAuth
  req Sig.Vrf(kp.spk, sig, framedContent)
  return (kp.id, Hash(kp), confTag, sig)

helper *frameProp(γ, P)
  tbs ← (γ.groupCtxt(), γ.groupId, γ.epoch, γ.leafIdx(),
            'proposal', P)
  sig ← Sig.Sign(γ.ssk, tbs)
  tbm ← (tbs, sig)
  membTag ← MAC.tag(γ.membKey, tbm)
  return (γ.groupId, , γ.epoch, γ.leafIdx(), 'proposal', P,
            sig, membTag)

helper *unframeProp(γ, p)
  parse (groupId, epoch, senderIdx, contentType, P,
            sig, membTag) ← p
  req contentType = 'proposal' ∧ groupId = γ.groupId
       ∧ epoch = γ.epoch
  tbs ← (γ.groupCtxt(), groupId, epoch, senderIdx,
            'proposal', P)
  tbm ← (tbs, sig)
  req γ.τ.leaves[senderIdx] ≠ ⊥
       ∧ γ.τ.leaves[senderIdx].inuse()
       ∧ Sig.Vrf(γ.τ.leaves[senderIdx].spk, sig, tbs)
       ∧ MAC.vrf(γ.membKey, membTag, tbm)
  return (senderIdx, P)

helper *frameExtProp(groupId, epoch, ssk, P)
  framedContent ← (groupId, epoch, 'new_member_proposal',
            'proposal', P)
  sig ← Sig.Sign(ssk, framedContent)
  return (framedContent, sig)

helper *unframeExtProp(γ, ep)
  parse (framedContent, sig) ← ep
  parse (groupId, epoch, sender, contentType, P) ← framedContent
  req contentType = 'proposal' ∧ sender =
            'new_member_proposal' ∧ γ.groupId = groupId
       ∧ epoch = γ.epoch
  parse(proposal, kp) ← P
  req proposal = 'add'
  req Sig.Vrf(kp.spk, sig, framedContent)
       ∧ kp.id ∉ roster
  return (kp.id, P)
```

Protocol ETK and ETK$^{\text{PSK}}$: Transcript-Hash

```
helper *set-conf-trans-hash(γ, γ', senderIdx, C, sig)
  commitContent ← (γ.groupId, γ.epoch, senderIdx,
                'commit', C, sig)
  γ'.confTransHash ← Hash(γ.interimTransHash,
                commitContent)
  return γ'

helper *set-interim-trans-hash(γ', confTag)
  γ'.interimTransHash ← Hash(γ'.confTransHash, confTag)
  return γ'
```

Figure 24: The helper methods related to message framing and transcript hashes.

E Proof of Theorem 1

Theorem 1. Assume that PKE is $\mathsf{IND\text{-}CCA}$ secure and that Sig is $\mathsf{SUF\text{-}CMA}$ secure. The ETK protocol securely realizes $(\mathcal{F}_{\mathcal{AS}}^{\mathcal{IW}}, \mathcal{F}_{\mathcal{KS}}^{\mathcal{IW}}, \mathcal{F}_{\mathrm{ECGKA}})$ in the $(\mathcal{F}_{\mathcal{AS}}, \mathcal{F}_{\mathcal{KS}}, \mathcal{G}_{\mathrm{RO}})$-hybrid model, where $\mathcal{F}_{\mathrm{ECGKA}}$ uses the predicates **safe** and **inj-allowed** from Figure 7 and calls to $\mathsf{HKDF.Expand}$, $\mathsf{HKDF.Extract}$, and MAC functions are replaced by calls to the global random oracle $\mathcal{G}_{\mathrm{RO}}$.

Proof. At a high level, our proof has a similar structure to that of [4]. First, we define four consecutive hybrids as follows:

1. **Hybrid 1 (*H1*):** In this first hybrid, a dummy functionality $\mathcal{F}_{dummy}$ forwards all in- and outputs through the simulator $\mathcal{S}_1$. $\mathcal{S}_1$ executes ETK. $\mathcal{F}_{dummy}$ hence encodes no security guarantees.
2. **Hybrid 2 (*H2*):** In the second hybrid, the dummy functionality $\mathcal{F}_{dummy}$ is replaced by a functionality consisting of $\mathcal{F}_{\mathcal{AS}}^{\mathcal{IW}}$, $\mathcal{F}_{\mathcal{KS}}^{\mathcal{IW}}$ and a modified version of $\mathcal{F}_{\mathrm{ECGKA}}$, where the underlying $\textbf{safe}(c) = \texttt{false}$ and $\textbf{inj-allowed}(c, \texttt{id}) = \texttt{true}$ for all commits c and party $\texttt{id}$. The functionality interacts with a trivial simulator that chooses all application secrets according to the protocol. The indistinguishability between *H1* and *H2* will prove the consistency of ETK.
3. **Hybrid 3 (*H3*):** The third hybrid is identical to *H2* except that the application secrets in safe epochs are random, i.e. the original **safe** is restored. The indistinguishability between *H2* and *H3* will prove the confidentiality of ETK.
4. **Hybrid 4 (*H4*):** The fourth hybrid is identical to *H3* except that forgery of messages (for proposals and commits) on behalf of uncorrupted members is excluded, i.e., the original **inj-allowed** predicate is restored, such that functionality exactly matches $\mathcal{F}_{\mathrm{ECGKA}}$. The indistinguishability between *H3* and *H4* will prove the authenticity of ETK.

We respectively prove the indistinguishability between every two adjacent hybrids by Lemma 1 (in Appendix E.1), Lemma 2 (in Appendix E.2), and Lemma 3 (in Appendix E.3). Note that the first hybrid *H1* and the last hybrid *H4* are respectively equivalent to the real world and the ideal world. The proof is concluded. □

E.1 Consistency and Correctness

Intuition. We prove that *H1* (ETK) and *H2* are indistinguishable by proving for every input operation that the outputs in ETK and $\mathcal{F}_{\mathrm{ECGKA}}$ of *H2* are the same. We especially show that if *H2* halts due to an **assert** statement, so would ETK. The **assert** statements in $\mathcal{F}_{\mathrm{ECGKA}}$ of *H2* capture that if one party has processed a proposal, (external) commit, or welcome, all parties that process it in the future end up in the same shared group state. Thus, the indistinguishability between *H1* and *H2* indicates that $\mathcal{F}_{\mathrm{ECGKA}}$ guarantees consistency.

Lemma 1. *Hybrids H1 and H2 are indistinguishable, that is,* ETK *guarantees consistency.*

Proof. We first define the simulator $\mathcal{S}_2$ and examine that (1) $\mathcal{S}_2$ will not trigger any **assert** statement in the $\mathcal{F}_{\mathrm{ECGKA}}$ in hybrid *H2*, and (2) the outputs of $\mathcal{F}_{dummy}$ (i.e., ETK) and the $\mathcal{F}_{\mathrm{ECGKA}}$ in hybrid *H2* are same. Recall that the simulator can monitor the whole history graph. We examine every possible commit node in the history graph. Recall also that every commit node in the history graph includes the application secrets and a $\mathsf{confTransHash}$ value, which hashes the history of commits until c. Since we model the hash function as random oracle, the same $\mathsf{confTransHash}$ values in different commit nodes must indicate that they have the same history of commits leading to them. Thus, every commit node in the history graph must be unique.

A repeating argument in the proof is that a value in ETK matches the attribute that $\mathcal{F}_{\mathrm{ECGKA}}$ stores for consistency. The proof will use the phrasing *the attributes in* ETK *and* $\mathcal{F}_{\mathrm{ECGKA}}$ *correspond directly* for this relationship. The argument is the same in all cases: The attribute value is either set by $\mathcal{F}_{\mathrm{ECGKA}}$ and then passed to the simulator, which returns the identifier encoding exactly these attributes, or the value is read directly out of an (adversarially injected) identifier by the simulator and passed back to $\mathcal{F}_{\mathrm{ECGKA}}$.

We now define $\mathcal{S}_2$ in more detail. Concretely, we define $\mathcal{S}_2$ in several steps: regular internal proposals, external proposals, commits, process, join, and external commit.

Regular (internal) Proposals.

When $\mathcal{F}_{\mathrm{ECGKA}}$ sends $(\texttt{Propose}, \mathsf{act})$, where $\mathsf{act} \in \{\mathsf{up\text{-}spk}, \mathsf{add\text{-}id}_t, \mathsf{rem\text{-}id}_t\}$, then the simulator $\mathcal{S}_2$ executes the ETK protocol to obtain the packet p. We then distinguish two cases:

1. If $\mathsf{act} = \mathsf{add}\text{-}\mathtt{id}_t$, then ETK fetches the key package $\mathtt{kp}_t$ for $\mathtt{id}_t$ from $\mathcal{F}_{\mathcal{KS}}$, which obtains $\mathtt{kp}_t$ by asking the environment $\mathcal{Z}$. Note that the simulator $\mathcal{S}_2$ executes the code of both ETK and $\mathcal{F}_{\mathcal{KS}}$. Thus, $\mathcal{S}_2$ simply uses $\mathtt{kp}_t$ provided by the environment $\mathcal{Z}$. If $p = \bot$, the simulator $\mathcal{S}_2$ sends $ack = \mathtt{false}$ to $\mathcal{F}_{\text{ECGKA}}$. Otherwise, $\mathcal{S}_2$ sends $(p, \mathsf{spk}_t, \mathtt{true})$ to $\mathcal{F}_{\text{ECGKA}}$, where spk_t is taken from $\mathtt{kp}_t$.
2. Otherwise, $\mathsf{act} \in \{\mathsf{up}\text{-}\mathsf{spk}, \mathsf{rem}\text{-}\mathtt{id}_t\}$. If $p = \bot$, the simulator $\mathcal{S}_2$ sends $ack = \mathtt{false}$ to $\mathcal{F}_{\text{ECGKA}}$. Otherwise, $\mathcal{S}_2$ simply sends $(p, \mathsf{spk}_t, \mathtt{true})$ to $\mathcal{F}_{\text{ECGKA}}$, where $\mathsf{spk}_t = \bot$.

Indistinguishability of Output: Both $\mathcal{F}_{\text{ECGKA}}$ and ETK return the same p or $\bot$. $\bot$ is returned if $\mathsf{Ptr}[\mathtt{id}] = \bot$, meaning that id has not yet successfully parsed a welcome or processed its own external commit. In this case, γ in ETK must also be $\bot$.

Assert Statements: There are two **assert** statements.

1. The **assert** in update proposals holds: ***valid-spk** queries has-ssk to $\mathcal{F}_{\mathcal{AS}}^{\mathcal{IW}}$, which means it fails only if the ssk that belongs to spk is not in the secret key store. ETK on the other hand queries get-ssk within a **try** in *fetch-ssk-if-nec, which means the protocol returns $\bot$. In this case, the simulator $\mathcal{S}_2$ has passed $ack = \mathtt{false}$ to $\mathcal{F}_{\text{ECGKA}}$, which also has returned $\bot$. The **assert** statement hence never fails.
2. The **assert** in ***consistent-prop** holds: Note that the proposals in ETK include a membTag, which is computed by using MAC over confTransHash. Recall that we model MAC as a random oracle. Recall also that confTransHash that identifies the history of commits is unique, as we have introduced. Thus, the proposals computed by every party id in node c must be different from those computed in any other node c'. Thus, the **assert** statement in the ***consistent-prop** algorithm holds.

External Proposals. When $\mathcal{F}_{\text{ECGKA}}$ sends (Propose, extAdd-id-spk, epoch) to the simulator $\mathcal{S}_2$, $\mathcal{S}_2$ executes the ETK protocol to obtain the packet p. ETK necessarily fetches the secret key from $\mathcal{F}_{\mathcal{AS}}$. If $p = \bot$, then $\mathcal{S}_2$ sends $ack = \mathtt{false}$ to $\mathcal{F}_{\text{ECGKA}}$. Otherwise, it sends $(p, \mathtt{true})$.

Indistinguishability of Output: Both $\mathcal{F}_{\text{ECGKA}}$ and ETK return the same p or $\bot$. $\bot$ is returned if $\mathsf{Ptr}[\mathtt{id}] \neq \bot$, meaning that id has successfully parsed a welcome or process its own external commit. In this case, γ in ETK must also not be $\bot$ and ETK must also return $\bot$. Otherwise, $\bot$ is returned if the proposal was performed with an invalid spk for id, in which case $\mathcal{S}_2$ has also returned $ack = \mathtt{false}$.

Assert Statements: There are two **assert** statements.

1. The **assert** ***valid-spk** never gets triggered for the same reason as in update proposals: if there was no valid entry in SSK, then ETK will have returned $\bot$. Thus, the simulator $\mathcal{S}_2$ will send $ack = \mathtt{false}$ to $\mathcal{F}_{\text{ECGKA}}$.
2. The second **assert** statement is contained in ***consistent-ext-prop**, which enforces that proposals adding id to an epoch (and hence all nodes of that epoch) with a certain key package are different from those adding id to another epoch or using another key package. This is given by including the epoch in p. It also ensures consistency upon a re-calculation, which holds because attributes in ETK and $\mathcal{F}_{\text{ECGKA}}$ correspond directly.

Commits. The simulator $\mathcal{S}_2$ computes the packets c and the vector of welcome messages vec_w according to ETK. If $c = \bot$, then $\mathcal{S}_2$ sets $ack = \mathtt{false}$. Otherwise, $\mathcal{S}_2$ first checks whether c corresponds to a detached root, indicated by $\mathtt{Node}[c] = \bot$, there exists a vector of welcome messages vec_w such that $\mathtt{Wel}[\mathtt{vec_w}] = \mathsf{root}_{rt}$ and confTransHash in vec_w matches the information in c, sends rt to $\mathcal{F}_{\text{ECGKA}}$ (alongside c and w and c's epoch). Then, $\mathcal{F}_{\text{ECGKA}}$ runs ***fill-props**. For every proposal p in the input $\vec{p}$ without a node, $\mathcal{S}_2$ sets orig and act according to p. This is always doable, as the basic check in ***apply-props** by ETK guarantees that p is well-formed.

Indistinguishability of Output: $\mathcal{F}_{\text{ECGKA}}$ returns c, w, and g, which were generated according to ETK. There are three cases where the Commit algorithm returns $\bot$:

- **req** $\mathsf{Ptr}[\mathtt{id}] \neq \bot$ holds as above.
- **req** ack is false iff ETK failed and returned $\bot$

Assert Statements: The commit algorithm includes the following **assert** statements:

1. **assert** ***valid-spk** succeeds because it is either the committers current spk, in which case `true` is returned, or the spk is updated, in which case `*rekey-path` will be failed upon an invalid spk, the simulator $\mathcal{S}_2$ will return $ack = \texttt{false}$, and $\mathcal{F}_{\text{ECGKA}}$ will return $\perp$ in ***req-correctness**.
2. **assert** mem $\neq \perp$ will succeed because `*apply-props` executed in ***req-correctness** guarantees that at least the committer is part of the group. It is left to show that if `*apply-props` in $\mathcal{F}_{\text{ECGKA}}$ returns $\perp$ due to the proposals not being well-formed, then so would `*apply-props` in ETK and hence the simulator $\mathcal{S}_2$ will return $ack = \texttt{false}$.
   - `Node`$[c] \neq \perp$ with $c =$ Ptr[`id`]: Since Ptr[`id`] is only set after successfully running `Process`, it always points to the commit that produced the secrets available to `id` in ETK. $\mathcal{S}_2$ will return $ack = \texttt{false}$ if Ptr[`id`] was not a valid c. As Ptr[`id`] in $\mathcal{F}_{\text{ECGKA}}$ is only ever set to c with `Node`$[c] \neq \perp$, `Node`[Ptr[`id`]] is hence not $\perp$.
   - $(\texttt{id}_c, *) \in$ `Node`$[c]$.mem: Assuming that `Node`$[c]$.mem correctly reflects the roster of τ after c the committer's membership is checked by ETK when confirming the commit signature. This holds due to the indistinguishability of the output of `Process`, which shows that the changes applied by `*apply-props` are equivalent between $\mathcal{F}_{\text{ECGKA}}$ and ETK.
   - `Prop`$[p] \neq \perp$ is guaranteed by running ***fill-props**.
   - $c \in$ `Prop`$[p]$.par with $c =$ Ptr[`id`], is either given by ***fill-props**, or the confirmation of the membership MAC ETK guarantees that `id` is member of the epoch p is sent to. For external self-add proposals, $\mathcal{F}_{\text{ECGKA}}$ matches them purely based on epoch, and ETK checks that the epoch is correct and parsable for `id`.
   - $p = p_{up} + p_{rem} + p_{add}$: Because of the direct correspondence of `Prop`$[p]$.act and the action of p, ETK performs the same check and $\mathcal{S}_2$ will return $ack = \texttt{false}$. `Prop`$[p]$ accurately reflects the action because either p is created by $\mathcal{S}_2$ according to act, or act is set by $\mathcal{S}_2$ extracting act from p.
   - $\texttt{id}_s \in G \setminus L$ (update) is checked by ETK.
   - $\texttt{id}_s \in G \wedge \texttt{id}_t \in G \setminus L$ (remove) is checked by ETK.
   - $\texttt{id}_s \in G \wedge \texttt{id}_t \notin G$ (internal add) or $\texttt{id}_s = \texttt{id}_t \wedge \texttt{id}_t \notin G$ (external add) is checked during unframing of the proposal and in `*apply-props`.
3. **assert** ***valid-successor**:
   - `Node`$[c] \neq \perp$ is given by the **else** branch.
   - `Node`$[c]$.mem = mem follows from the other requirements: either `Node`$[c]$.par = Ptr[`id`] and `Node`$[c]$.pro $= \vec{p}$, in which case `Node`$[c]$.mem has been set to the same results as the passed parameter mem. Else, `Node`$[c]$.par $= \perp$ and `Node`$[c]$.pro $= \perp$, in which case, c is a detached root but has been returned by $\mathcal{S}_2$ due to matching confTransHash, which includes the tree and hence membership.
   - `Node`$[c]$.pro $\in \{\perp, \vec{p}\}$ holds because p and `Node`$[c]$.pro correspond directly between ETK and $\mathcal{F}_{\text{ECGKA}}$.
   - `Node`$[c]$.par $\in \{\perp,$ Ptr[`id`]$\}$: if `Node`$[c]$.par were neither $\perp$ nor Ptr[`id`], it would be another c'. As shown before $\mathcal{F}_{\text{ECGKA}}$ and ETK place `id` in the same commit. $\mathcal{S}_2$ would have returned $ack = \texttt{false}$, since `id` cannot process c without the epoch's secrets.
4. **assert** `Node`$[c]$.orig = `id`: c as returned by $\mathcal{S}_2$ is unique and includes the committers identity as part of the confTransHash. `id` is encoded in c and `Node`$[c]$.orig correspond directly between ETK and $\mathcal{F}_{\text{ECGKA}}$.
5. **assert** RndCor[`id`] $\neq$ good: If `Node`$[c] \neq \perp$, then $\mathcal{S}_2$ previously returned the same c. $\mathcal{S}_2$ constructs c according to ETK, which uses randomness to generate the new leaf node on an update path. In order to return the same c twice, the member's randomness must have been corrupted by the adversary in the real world. If the adversary in the real world corrupts a member `id`'s randomness, RndCor(`id`, `Bad`) is called in $\mathcal{F}_{\text{ECGKA}}$ and RndCor[`id`] would have been set to bad. If `Node`$[c] = \perp$ and $rt \neq \perp$: then there exists a vector of welcome message `vec_w` whose confTransHash matches that of c. As the hash also includes the new random leaf node, the same argument as above holds.
6. **assert** `vec_w` $\neq \perp$: ETK and also $\mathcal{S}_2$ return `vec_w` if there is an add proposal in $\vec{p}$. If such is the case, the proposal would have received the act = add-$*$ by either `Propose`, or by $\mathcal{S}_2$ during ***fill-props**.
7. **assert** `Wel`[`vec_w`] $\in \{\perp, c\}$ guarantees that the vector of welcome message is different for each c. This is achieved by including the `confTag` over the transcript hash in the Welcome's GroupInfo In Commit, `Wel`[`vec_w`] can previously only have been set to c. The only other option for setting `Wel`[`vec_w`] is in `Join`. There, c is set to the value provided by $\mathcal{S}_2$, which retrieves c as the only fitting commit by comparing the confTransHash, or it is set to a detached node. In the latter case, c is providing the link to the detached node, hence ***attach** would have set `Wel`[`vec_w`] to c.

8. **assert** $g \neq \bot$ succeeds because ETK and $\mathcal{S}_2$ both return g.
9. **assert** $\texttt{groupInfo}[g] \in \{\bot, c\}$: as for Welcomes, the `confTag` in the `groupInfo` guarantees the uniqueness of g for each c. In `Commit` and `ExtCommit`, $\texttt{groupInfo}[g]$ is only ever set to c as retrieved by $\mathcal{S}_2$ by confirming the matching `confTag`.

Process. Process splits the processing of external commits and regular commits based on id handing in the full proposal vector. If `id` hands over the proposal vector, c is treated as a regular commit and the following proceeds:

The simulator $\mathcal{S}_2$ runs ETK protocol and checks whether the receiver would accept the inputs. If the inputs are rejected, then the simulator $\mathcal{S}_2$ sends $ack = \texttt{false}$. Otherwise, the simulator $\mathcal{S}_2$ checks whether c corresponds to a detached root exactly as dealing with `Commit` algorithm above. If c creates a new node, which means that there was no detached root and $\texttt{Node}[c] = \bot$, the simulator $\mathcal{S}_2$ retrieves c's epoch, orig', spk' from c (by checking the committer's key package that was generated in `updatePath`).

In the case of $\vec{p} = \bot$, `Process` treats the commit as an external commit. The simulator $\mathcal{S}_2$ executes ETK (which also treats the commit based as an external commit due to $\vec{p} = \bot$) and sets $ack = \texttt{false}$ if the processing party does not accept the input. Otherwise, $\mathcal{S}_2$ checks if c creates a new node (an existing detached root is impossible in this case, as an external commit does not produce any welcome messages). In this case, $\mathcal{S}_2$ retrieves c's epoch, orig' and spk' from c. In both cases, $\mathcal{S}_2$ also extracts the proposal vector $\vec{p}$ from c. If creating a new node, $\mathcal{S}_2$ retrieves act from the proposal p and returns it.

Indistinguishability of Output: We first explain why the outputs of ETK and $\mathcal{F}_{\text{ECGKA}}$ are the same for regular commits. Recall that a regular commit includes the hash of all proposals and we model Hash function as random oracle. Thus, for every commit c there is only one correct $\vec{p}$ that can be passed by $(\texttt{Process}, c, \vec{p})$. Recall also that the output of `Process` is determined by $\vec{p}$ and the member set mem in c's parent node, both in ETK and ***output-proc** in $\mathcal{F}_{\text{ECGKA}}$. Thus, the outputs of ETK and $\mathcal{F}_{\text{ECGKA}}$ are the same for a regular commit. Note that $\texttt{Node}[c].\mathsf{orig}$ is either determined here by $\mathcal{S}_2$ directly according to ETK or it was set honestly following ETK that was run accordingly. The `Process` returns $\bot$ only in the following cases:

- **req** ack holds as above.

In the case of external commits, the relation between c and p is even more immediate, as $\mathcal{S}_2$ extracts the latter from the former. In this case, ***output-proc** is meanwhile determined by the same parameters as the regular commit, and the argumentation is analogous.The `Process` returns $\bot$ only in the following case:

- **req** ack holds as above.

Assert statements : The `Process` algorithm includes the following **assert** statements for regular commits:

- **assert** $\mathsf{mem} \neq \bot$ means `*apply-props` failed. In case c already existed and `*apply-props` must have been run on the exact same parameters before, especially on $\mathsf{Ptr}[\texttt{id}]$ (or ***req-correctness** would have failed while checking $\texttt{Node}[c].\mathsf{par} = \mathsf{Ptr}[\texttt{id}]$) and passed. Else, ***req-correctness** would have failed. If $\texttt{Node}[c] = \bot$ or c were malformed according to ETK (i.e. a proposal removing the committer), $\mathcal{S}_2$ would return $ack = \texttt{false}$ and `Process` would unwind. This means that c either corresponds to a correct detached root or a correct new node. It is left to show that every single **req** in `*apply-props` is either fulfilled or $\mathcal{S}_2$ would have returned $ack = \texttt{false}$, unwinding the function. This holds for the same reason as for `Commit`.
- **assert** ***valid-successor**: is equivalent to the `Commit`.
- **assert** $\texttt{id} \in \texttt{Node}[c].\mathsf{mem}$: As shown before, $\mathsf{Ptr}[\texttt{id}]$ points to the same c that `id` has most recently processed in ETK. $\mathcal{S}_2$ guarantees that `id` is a member of the previous epoch by either verifying the membership tag or the confirmation tag, else returns $ack = \texttt{false}$, and ***req-correctness** would fail as $\texttt{Node}[c].\mathsf{par}$ is unequal to $\mathsf{Ptr}[\texttt{id}]$. In both cases, ***members** is executed on $\mathsf{Ptr}[\texttt{id}]$ and the same $\vec{p}$, and if no proposal removing id is executed, `id` will be part of $\texttt{Node}[c].\mathsf{mem}$.

In the case of an external commit, `*process-ec`(c) contains the following **assert** statement:

- **assert** $\mathsf{mem} \neq \bot$ for both $\texttt{Node}[c] = \bot$ and $\texttt{Node}[c] \neq \bot$: holds if `*apply-props` runs fully. As ***req-correctness** returned `false`, any malformed c would have lead to an unwinding of `*process-ec`:
  - $\texttt{Node}[c] \neq \bot$ with $c = \mathsf{Ptr}[\texttt{id}]$: given by placement in if statement.

- $(\mathtt{id}_c, *) \in \mathtt{Node}[c].\mathsf{mem}$ with $c = \mathsf{Ptr}[\mathtt{id}]$ (resync): Assuming that $\mathtt{Node}[c].\mathsf{mem}$ correctly reflects the roster of τ after c,the committer's membership is checked by ETK when confirming the commit signature. $(\mathtt{id}_c, *) \notin \mathtt{Node}[c].\mathsf{mem}$ with $c = \mathsf{Ptr}[\mathtt{id}]$ (external commit): same as above.
- $\mathsf{prop} \neq \bot$ is guaranteed by the fact that an external commit is never created as a detached node, and prop is hence always set.
- $c \in \mathsf{prop.par}$ with $c = \mathsf{Ptr}[\mathtt{id}]$ is given by prop being directly included, and ExtCommitProps parents being set according to the including commit. In case that $\mathtt{Node}[c]$ pre-exists, it hence depends on $\mathsf{Ptr}[\mathtt{id}] = \mathtt{Node}[c].\mathsf{par}$. This is the case as o.w. ***req-correctness** would have failed and assuming that $\mathsf{Ptr}[\mathtt{id}]$ points to the commit parsed in ETK, S2 would have returned $ack = \mathtt{false}$.
- $p = p_{exi} + p_{rem}$: (Note: given the only setting in which an external commit node is created is after checking that all included commits are either exi or rem.) $\mathsf{Prop}[p].\mathsf{act}$ directly corresponds to p's action. Since ETK hence performs the same checks $\mathcal{S}_2$ would have returned $ack = \mathtt{false}$.
- $\mathtt{id}_s = \mathtt{id}_t = \mathtt{id}_c$ (remove) is given as ExtCommitProps creates all proposals in such a way that orig' $(= \mathtt{id}_s)$ is equal to the committer $(\mathtt{id}_c)$ without allowing the adversary to set it separately. $\mathtt{id}_t$ being equal is given by the **req** in Process and the way that the proposals are created.

- **assert *valid-successor**: holds for the same reason as in Commit.
- **assert** $\mathtt{id} \in \mathtt{Node}[c].\mathsf{mem}$ is given by $\mathsf{mem} \neq \bot$ and the fact that an external commit does not remove anyone, and adds the committer.

Join. When a party id joins using a vector of welcome message vec_w, the simulator $\mathcal{S}_2$ runs ETK and sets $ack = \mathtt{true}$ if no error occurs. If $ack = \mathtt{false}$ or $\mathtt{Wel}[\mathtt{vec_w}] \neq \bot$, then $\mathcal{S}_2$ simply sends $(ack, \bot, ...)$ to $\mathcal{F}_{\mathrm{ECGKA}}$. Note that if $\mathtt{Wel}[\mathtt{vec_w}] \neq \bot$, $\mathcal{F}_{\mathrm{ECGKA}}$ already knows the information in vec_w and can ignore nay values received from the simulator other than ack. Otherwise, i.e., $ack = \mathtt{true}$ and $\mathtt{Wel}[\mathtt{vec_w}] = \bot$, the simulator $\mathcal{S}_2$ needs to interpret the injected vec_w. Recall that $\mathtt{vec_w} = (\mathtt{encGroupSecs}, \mathtt{encGroupInfo}, \mathtt{wel_type})$ and that $\mathcal{S}_2$ knows all secrets and therefore is able to recover groupInfo from encGroupInfo. In this case, $\mathcal{S}_2$ sends a value c', which either equals an existing node that owns identical confTransHash as the one in groupInfo or $\bot$ if such node does not exist, a value orig' according to groupInfo, a value mem' according to groupInfo, and a value vec_w's epoch to $\mathcal{F}_{\mathrm{ECGKA}}$.

Indistinguishability of Output: $\mathcal{F}_{\mathrm{ECGKA}}$ outputs $\mathtt{Node}[c].\mathsf{mem}$, $\mathtt{Node}[c].\mathsf{orig}$. These are either retrieved from vec_w by $\mathcal{S}_2$ as described above, in which case they are directly equal to the values output by ETK. In the other case, vec_w determines the c from which this info is taken. Note that if a party id can successfully process vec_w, then the corresponding commit c must be uniquely determined for id. This is because that the groupInfo decrypted from vec_w includes a confirmation tag that MAC the hash the history of commits and we model MAC as random oracle. For the commit c, either $\mathcal{F}_{\mathrm{ECGKA}}$ sets orig directly and $\mathcal{S}_2$ supplies the corresponding c, or $\mathcal{S}_2$ derives orig' from c and returns it to $\mathcal{F}_{\mathrm{ECGKA}}$.

The Join returns $\bot$ only in the following case:

- **req** ack: holds as above.

Assert Statements: The Join algorithm includes the following **assert** statements:

- **assert** $\mathtt{id} \in \mathtt{Node}[c].\mathsf{mem}$ never triggers, as ETK implicitly performs a check on the tree membership of id by calling $\gamma.\tau.\mathsf{leaves}[\gamma.\mathsf{leafId}()]$ and $\mathcal{S}_2$ would hence send $ack = \mathtt{false}$. As shown for the indistinguishability, $\mathtt{Node}[c].\mathsf{mem}$ contains the members of the tree, and hence also id.
- **assert** $g \neq \bot$ succeeds because ETK and hence $\mathcal{S}_2$ returns g.
- **assert** $\mathtt{groupInfo}[g] \in \{\bot, c\}$: like for Welcomes, the confTransHash in the groupInfo guarantees the uniqueness of g created by $\mathcal{S}_2$ for a c. In Commit and ExtCommit, $\mathtt{groupInfo}[g]$ is only ever set to c as retrieved by $\mathcal{S}_2$ by confirming the matching confTag.

Consistency: We still need to argue that every joiner id that executes Join with a vector of welcome message vec_w and every group member $\mathtt{id}'$ that executes Process with a commit c' must end up in the same state, where c' is uniquely determined by vec_w for id as we have explained above. Recall that both id and $\mathtt{id}'$ verify the confirmation tag, for which they derive the key from the joiner secret combined with the group context. This guarantees (assuming that MAC is modeled as random oracle) that they agree on the context, which in particular includes the tree hash and the confirmed transcript hash. The tree hash binds the whole ratchet tree, including its structure, spk of all members, and all public keys. In particular,

this implies agreement on the member set mem'. Agreement on the transcript hash implies agreement on the history, including the last committer orig'. The agreement is maintained in descendants of c', since parties agree on the ratchet tree in c'.

External Commit. When an external party sends `ExtCommit`, $\mathcal{S}_2$ constructs the external commit using g according to ETK and sets $ack = \texttt{false}$ if $c = \bot$. Else, it generates all required parameters for $\mathcal{F}_{\text{ECGKA}}$. It generates g, $\vec{p}_{exi}$, and $\vec{p}_{rem}$ according to ETK. These are direct or inlined proposal messages for the commit. It then checks if g is known to $\mathcal{F}_{\text{ECGKA}}$ ($g \neq \bot$). If it is not, but there is a matching commit c_{gmatch} that has the same
$confirmedTH$ as g, $\mathcal{S}_2$ returns this. If there is no such c_{gmatch}, $\mathcal{S}_2$ extracts the sender orig_g and the members mem_g from g.

Indistinguishability of Output: $\mathcal{F}_{\text{ECGKA}}$ returns c and g' generated according to the protocol. The `ExtCommit` returns $\bot$ only in the following cases:

- **req** $\mathsf{Ptr}[\mathtt{id}] \neq \bot$: hold as above.
- **req** ack: holds as above.

Assert Statements: The `ExtCommit` algorithm includes the following **assert** statements:

- **assert** ***valid-spk**$(\mathtt{id}, \mathsf{spk})$: holds because otherwise, `*rekey-path-upon-join` would have failed in ETK and ***req-correctness** in $\mathcal{F}_{\text{ECGKA}}$, and upon $\mathcal{S}_2$ returning $ack = \texttt{false}$, the method unwinds.
- **assert** $\mathsf{mem} \neq \bot \wedge (\mathtt{id}, \mathsf{spk}) \in \mathsf{mem}$: is given as `*apply-props` executed on any $\vec{p}$ contains at least the committer (in this case `id`) when run successfully. It runs successfully because
  - the requirements on `props` of `*apply-props` are fulfilled through the construction in `ExtCommit`.
  - **req** $(\mathtt{id}_c, *) \in G$ (resync) or $(\mathtt{id}_c, *) \notin G$ (external commit): is equally covered by ***req-correctness** in combination with ETK execution.
  - **req** $\mathtt{id}_s = \mathtt{id}_t = \mathtt{id}_c$ and $\mathtt{id}_s \in G$ fulfilled through construction in `ExtCommit` and the previous **req**.
- ***consistent-ext-comm**() holds for a similar reason as for proposals and commits. An external commit does not include a membership tag, but it does include the confirmation tag, which is a MAC over the GroupContext's confirmed transcript hash that provides the same uniqueness guarantee of a commit c. $\mathtt{Node}[c].\mathsf{pro}$ and $\mathtt{Node}[c].\mathsf{par}$ will hence only be set to the proposals included in c and the parent of c included in the $\mathsf{confTransHash}$, or else to $\bot$ in case of a detached root. As a consequence of these two properties, $\mathtt{Node}[c].\mathsf{mem}$ will always be set to the same membership set handed to ***consistent-ext-comm**.

  As `ExtCommitProps` is directly constructed from $\vec{p}$, as a consequence, all its parameters are also equivalent.
- **assert** $g \neq \bot$ succeeds because ETK and hence $\mathcal{S}_2$ returns g.
- **assert** $\mathtt{groupInfo}[g] \in \{\bot, c\}$: like for Welcomes, the `confTag` in the `groupInfo` guarantees the uniqueness of g created by $\mathcal{S}_2$ for a c. In `Commit` and `ExtCommit`, $\mathtt{groupInfo}[g]$ is only ever set to c as retrieved by $\mathcal{S}_2$ by confirming the matching `confTag`.

Consistency Invariant. In the end, we explain that the consistent invariant **cons-invariant** always holds.

1. $\mathtt{Node}[c].\mathsf{pro} \neq \bot$: holds for all regular child nodes. When attaching a detached node with $\mathtt{Node}[c].\mathsf{pro} = \bot$ to a parent node, the proposals are set to $\vec{p}$. `*apply-props` in $\mathcal{F}_{\text{ECGKA}}$ checks whether the commit $c \in \mathtt{Prop}[p].\mathsf{par}$. In every case a new commit node is created or attached to a proposal list (in `Join` and `Process`), the **assert** mem would have hence failed if that is not the case.
2. $\forall \mathtt{id}$ s.t. $\mathsf{Ptr}[\mathtt{id}] \neq \bot$: $\mathtt{id} \in \mathtt{Node}[\mathsf{Ptr}[\mathtt{id}]].\mathsf{mem}$: holds because in both locations where $\mathsf{Ptr}[\mathtt{id}]$ is set to c (`Process` and `Join`) an explicit check is performed that $\mathtt{id} \in \mathtt{Node}[\mathsf{Ptr}[\mathtt{id}]].\mathsf{mem}$.
3. The graph contains no cycles: $\mathcal{F}_{\text{ECGKA}}$ will only attach a detached node c' to an existing graph node c if $\mathcal{S}_2$ provides the corresponding root. $\mathcal{S}_2$ will only do so if the $\mathsf{confTransHash}$ match. Since this hash contains all previous commits, a circular graph cannot exist.
4. $\forall ep$ s.t. $\mathtt{ExtProps}[ep] \neq \bot : (\forall c \in \mathtt{Prop}[ep].\mathsf{par} : \mathtt{Node}[c].\mathsf{epoch} = \mathtt{ExtProps}[ep].\mathsf{epoch})$: holds because the parents of an external proposal ep are only set if the epochs match.

$\square$

E.2 Confidentiality

Intuition. We prove that *H2* and *H3* are indistinguishable by demonstrating that for all commits c for which the predicate **safe**(c) in Figure 7 holds, the group keys are indistinguishable from random. The indistinguishability between *H2* and *H3* indicates the confidentiality of ETK.

To this end, we define a sequence of sub-hybrid Hs_i, where i indicates the index of epochs. In the i-th sub-hybrid Hs_i, the simulator sets the first i epoch secrets, and all epoch secrets after that are set according to $\mathcal{F}_{\text{ECGKA}}$, i.e., randomly if **safe** = `true` and otherwise by the simulator. Note that the environment $\mathcal{Z}$ is executed in polynomial time and therefore can trigger only polynomial epochs of $\mathcal{F}_{\text{ECGKA}}$. Thus, the number of the sequence of sub-hybrid games Hs_i must be polynomial. It is easy to observe that in the final sub-hybrid all epoch secrets are set according to *H2* (and therefore is equivalent to *H2*), and that in the first sub-hybrid Hs_0 all are set according to *H3* (and therefore is equivalent to *H3*). Thus, the indistinguishability between *H2* and *H3* can be reduced to the indistinguishability between every two adjacent sub-hybrids Hs_{i-1} and Hs_i for the sequence of i.

We prove the indistinguishability between every two adjacent sub-hybrids Hs_{i-1} and Hs_i following the strategy in [4], i.e., by introducing the GSD game. Recall in Theorem 3 that the GSD security of any PKE can be reduced to its IND-CCA security. Thus, our proof methodology in this section is to reduce the confidentiality of ETK to the GSD security of the underlying PKE. Since they only behave differently in epoch i, and only if **safe**(c_i) = `true`, we show that if the environment $\mathcal{Z}$ can distinguish between a random key and a simulator chosen key for c_i, then there exists a reduction $\mathcal{B}$ that can win the GSD game (and hence break the IND-CCA security) of the underlying PKE by invoking $\mathcal{Z}$. We need to argue that whenever the environment $\mathcal{Z}$ wins, which can happen only under the condition that **safe**(c) = `true` (implying a valid challenge for $\mathcal{Z}$), the reduction $\mathcal{B}$ can also win, which happens only under the condition that $\neg$gsd-exp(c, u_a) (implying a valid challenge for $\mathcal{B}$ breaking GSD). To this end, we prove that gsd-exp$(c) \implies \neg$**safe**(c) by reasoning through all possible cases of gsd-exp(c).

Important to note that the reduction $\mathcal{B}$ stores one of three values for each key: (1) if the key is unknown to $\mathcal{Z}$ and $\mathcal{B}$ (and hence only known to GSD game), $\mathcal{B}$ stores (gsd, u), with u indicating the GSD node that contains the key value, (2) if the key is known to both $\mathcal{B}$ and $\mathcal{Z}$, $\mathcal{B}$ stores the actual key value, and (3) if the key is known to $\mathcal{Z}$, but not to $\mathcal{B}$ (or GSD game), $\mathcal{B}$ sets it to $\perp$. For bookkeeping, $\mathcal{B}$ maintains a counter u_{ctr} to index new GSD nodes at every point in the game.

Lemma 2. *If* PKE *is* GSD *secure, then Hybrids H2 and H3 are indistinguishable, that is,* ETK *guarantees confidentiality.*

Proof. We show the indistinguishability by hybrid games. More concretely, for any $i > 1$, we define two adjacent sub-hybrids Hs_{i-1} and Hs_i that behave differently only in epoch i: Hs_{i-1} sets the epoch secret in epoch i according to $\mathcal{F}_{\text{ECGKA}}$ (depending on whether **safe**(c) holds), while Hs_i sets is according to the simulator $\mathcal{S}$. Further, we reduce the indistinguishability between every two adjacent hybrid games to the GSD security of the underlying PKE. More concretely, we show that if the environment $\mathcal{Z}$ can distinguish between a random key and a simulator chosen *application* key for c with **safe**(c) = `true`, the reduction $\mathcal{B}$ (i.e., the adversary that aims to break GSD game) can use the environment $\mathcal{Z}$ to win the GSD game with gsd-exp$((c, u))$ = `false` for some associated secret u.

The individual steps of the confidentiality part are organized as follows:

Proof part 2.b): Confidentiality

Show indistinguishability:

I) when not allowing any injections or corruption of randomness by reasoning through all possible cases of gsd-exp$((c, u_{\text{app}}))$ (Appendix E.2):

   (a) the corruption of (c, u_{app})
   (b) the corruption of (c, u_{joi})
   (c) if all three gsd-exp $((c, u_{\text{com}}))$ and gsd-exp ((Node$[c]$.par, u_{ini})) are true.

II) when allowing injections for cases (a) - (c) (Appendix E.2).

III) when allowing the corruption of a client's randomness for cases (a) - (c) (Appendix E.2).

We update the proof to include the changes from draft 12 to the final standard. Additionally, We introduce external commits. We mostly note down changes to the content of [4].

Part 2.b.I) No Injection or Corrupted Randomness Allowed

In this part, we consider the simplest case, where neither malicious injection nor randomness corruption is allowed.

In order to prove by reduction that $\mathsf{gsd\text{-}exp}((c,u)) \implies \neg\mathsf{safe}(c)$, we first outline how reduction $\mathcal{B}$ modifies the code of the functionalities and the simulator in the context where no injection and no corrupted randomness are permitted. Following this, we systematically examine each possible cases of $\mathsf{gsd\text{-}exp}((c,u))$, corresponding to the three conditions specified above, and illustrate how our proof captures the requirements of the safe predicate.

As described in the *gsd* experiment explanation; in our model, for {add, extAdd}-only commits, no UpdatePath is executed, and all of id 's secrets remain the same. This is addressed in ***secrets-replaced**$(c, \mathtt{id})$. Throughout the confidentiality proof, we will examine commits in which an updatePath is required as part of ***secrets-replaced**$(c, \mathtt{id})$, as well as those in which no updatePath is necessary. To this end, we also consider the commits including add, extAdd proposals only as well as those that involve any combination of these proposal types.

From draft 12 to RFC: Different from draft 12, the RFC separates leaf nodes and key packages. While add and extAdd proposals contain key packages, update proposals and updatePath only contain fresh leaf nodes. Both share a similar structure. Therefore, it is necessary for us to change the welcome messages, which now use the additional *init key* of the key package to encrypt the group secrets. Additionally, the committer has the option to create an encrypted welcome message, which can be designed for all members collectively, or for each member individually, or grouped for specific batches of members. Moreover, we need to define how reduction $\mathcal{B}$ processes the extAdd proposals. Hence:

- Key Package Registration: When $\mathcal{Z}$ instructs a party id to register a key package in the emulated $\mathcal{F}_{\mathcal{KS}}^{\mathcal{IW}}$, the reduction $\mathcal{B}$ creates a new GSD vertex by executing $\mathsf{pk} \leftarrow *\mathsf{get\text{-}pk}(u_{\mathrm{ctr}})$ and $\mathsf{init\text{-}pk} \leftarrow *\mathsf{get\text{-}pk}(u_{\mathrm{ctr}})$, respectively followed by u_{ctr}++.[5] The reduction $\mathcal{B}$ uses pk and init-pk to generate the public part of the key package kp, and sets the secret keys $\mathsf{SK}[\mathtt{id},\mathtt{kp}]$ and $\mathsf{ISK}[\mathtt{id},\mathtt{kp}]$ to (gsd, u_{ctr}).
- Proposal extAdd: Whenever the protocol requests a key package for id creating the proposal from $\mathcal{F}_{\mathcal{AS}}^{\mathcal{IW}}$, $\mathcal{Z}$ gets to choose it. Accordingly, the reduction $\mathcal{B}$ stores in the new proposal node the ssk taken from the key package chosen by $\mathcal{Z}$.
- Applying Proposals: Additionally, for each extAdd proposal, secret key of id is set using the value stored in proposal node.
- Commits by reference: Compared with [4], there are two differences. On one hand, following RFC version, the *rekey-path operation is *not* executed if the proposal list in a regular commit satisfies both of the following conditions: (1) it is non-empty, and (2) it contains only add, extAdd proposals.

  The reduction $\mathcal{B}$ emulates the *derive-keys as follows:

  1. *Add vertices to the GSD graph:* The reduction $\mathcal{B}$ generates vertices u_{app}, u_{mem}, u_{conf}, u_{ini}, u_{ext} by querying **Hash** oracle with inputs $(u_{\mathrm{lbl}}, \mathsf{lbl})$ for $\mathsf{lbl} \in \{\mathsf{app}, \mathsf{mem}, \mathsf{conf}, \mathsf{ini}, \mathsf{ext}\}$.
  2. *Create the welcome message:* The welcome message is a component of the welcome vector, consisting of three parts. The final part, wel_type, indicates whether the message contains a single welcome message or multiple welcome messages. In the first part, for each new member $\mathtt{id}_t$, the encryptions of joinerSec and $\mathtt{id}_t$'s pathSec under the init-key in $\mathtt{id}_t$'s key package (obtained from KS by the party adding $\mathtt{id}_t$). Let u_i be the GSD vertex corresponding to the pathSec sent to $\mathtt{id}_t$. If $\mathtt{id}_t$'s init-key is of the form (gsd, u)[6], $\mathcal{A}$ obtains the encryptions by creating encryption edges from u to u_{joi} and from u to u_i. Otherwise, it corrupts u_i, u_{joi} encrypts the values itself. The second part of the welcome message contains the encryption of groupInfo, under the welcome_secret which is simply the combination of joinerSec and the vector of welcome messages. To compute the second part of the vector of welcome messages, the reduction $\mathcal{B}$ creates vertex u_{wel} by querying **Hash** oracle with input $(u_{\mathrm{wel}}, \mathsf{wel})$. The reduction $\mathcal{B}$ gets the respective secrets by corrupting corresponding vertices.
  3. *Create* groupInfo*:* Additionally, the reduction $\mathcal{B}$ creates the groupInfo: It runs $*\mathsf{get\text{-}pk}(u_{\mathrm{ext}})$ to retrieve the external public key and includes it in the groupInfo. Note that at this point, $\mathcal{B}$ does not need to know the external private key.

[5] Recall that we use $\mathsf{pk} \leftarrow *\mathsf{get\text{-}pk}(u)$ to denote that $\mathcal{A}$ obtains the public key pk for a vertex u by calling the oracle $\mathbf{Enc}(u, 0)$ (the special vertex 0 is only used here), following [4].

[6] Following [4], a secret key of the form (gsd, u) means that it is unknown to the environment $\mathcal{Z}$.

- Expose: The private init keys are generated along with the key package secrets, which $\mathcal{A}$ is also required to corrupt.

As the basis for the proof, note that, considering the key derivation illustrated in Figure 2, $\mathsf{gsd\text{-}exp}((c, u_{\mathrm{app}}))$ can only hold true in one of the following cases:

2.b.I.a) The reduction $\mathcal{B}$ corrupts the vertex (c, u_{app}). This happens if and only if it corrupts the state of a party involved in commit c. This case is identical to the case in [4] and will immediately implies $\neg$**safe**(c).

2.b.I.b) The reduction $\mathcal{B}$ corrupts the vertice (c, u_{joi}) to ultimately gain access to the application secret. This happens if and only if the reduction $\mathcal{B}$ computes a vector of welcome messages for $\mathtt{id}_t$ using an exposed key bundle. The mechanism remains as previously described.

2.b.I.c) This case considers the following two elements to be true: $\mathsf{gsd\text{-}exp}((c, u_{\mathrm{com}}))$ and $\mathsf{gsd\text{-}exp}((\mathtt{Node}[c].\mathsf{par}, u_{\mathrm{ini}}))$. We need to show that $\mathsf{gsd\text{-}exp}((c, u_{\mathrm{com}}))$ implies $\mathsf{know}(c, *)$. This implication follows by condition d) of ***can-traverse**. Here, $\mathsf{gsd\text{-}exp}((c, u_{\mathrm{com}}))$ is only possible in one of the following cases:

1) the reduction $\mathcal{B}$ corrupts path secrets u_i, and use them to compute $\mathsf{gsd\text{-}exp}((c, u_{\mathrm{com}}))$
2) the reduction $\mathcal{B}$ calls **Enc** oracle to encrypt path secrets under relevant keys
3) after the commit processed, the reduction $\mathcal{B}$ exposes a party `id` that holds u_i's secret

Firstly, consider case 3) any action requiring `updatePath` removes u_i's secret from its state. For the actions of `id` requiring `updatePath`, the proof remains the same. For {`add`, `extAdd`}-only commits, since no `updatePath` is executed, all of `id`'s secrets remain the same. Thus, if `id` is corrupted, we have $\mathsf{know}(c, \mathtt{id})$ until an `updatePath` issued.

We complete the proof by showing that for cases 1) and 2) $\mathsf{gsd\text{-}exp}((\mathtt{Node}[c].\mathsf{par}, u_{\mathrm{ini}}))$ implies $\mathsf{know}(c,$ 'epoch'). For these cases the proof follows the same pattern as the encapsulating proof for u_{app} as u_{ini} has the exact same lifetime as its epoch's u_{app} is generated from each epoch secret. The proof's condition c) hence recursively proceeds until at some point, $\mathtt{Node}[c].\mathsf{pars} = \mathsf{root}_x$. At that point, condition d) of ***can-traverse**(c) in **safe**(c) becomes true.

External Commit. The external commits include the following operations:

- Derivation of the init key init-key: MLS uses an asymmetric interaction to derive the init key of the new epoch from the external secret. The reduction $\mathcal{B}$ creates a new GSD node (gsd, u_{ctr}) representing the seed s_{exini} chosen by the committer. Then $\mathcal{B}$ creates an edge encrypting the seed with the external public key by querying **Enc** oracle. This encrypted output is shared with the other group members via the external init proposal `extInit` included in the external commit. The init key init-key is derived from this seed by calling **Hash** on the seed node. Then $\mathcal{B}$ proceeds using the resulting node as the u_{exini} node in a commit.
- Generation of leaf node: same as `*rekey-path` in a regular commit but occurs during the `*rekey-path-upon-join`.
- Other than that, the reduction $\mathcal{B}$ proceeds as during a regular commit, corrupting the GSD vertices u_{mem} and u_{conf} in order to create the new commit and `groupInfo`. Note that there is no welcome following an external commit. The reduction $\mathcal{B}$ also corrupts any u_i (path secret) that is not $(gsd, .)$ meaning it is known to $\mathcal{Z}$ and computes all secrets following from that.

Figure 25 visualizes the GSD game accordingly.

We first show that if **safe**(c) holds on an epoch then $\mathsf{gsd\text{-}exp}((c, u_{\mathrm{app}}))$ is false and the epoch is still a valid challenge. We then show how the reduction $\mathcal{B}$ can win the GSD game by utilizing the $\mathcal{Z}$ that can distinguish this epoch between simulated by functionality and by simulator.

The argument for Case 2.b.I.a) and Case 2.b.I.b) remain the same as in regular commit. For Case 2.b.I.c) if there is an edge between (c, u_{exini}) and (c, u_{joi}), means that c is an external commit. As the `*rekey-path-upon-join` proceeds in the same way as `*rekey-path` in a regular commit, the reasoning for $\mathsf{gsd\text{-}exp}((c, u_{\mathrm{com}}))$ implying $\mathsf{know}(c, *)$ also remains the same. The only change occurs when the combination of $\mathsf{gsd\text{-}exp}((\mathtt{Node}[c].\mathsf{par}, u_{\mathrm{exini}}))$ implies $\mathsf{know}(c,' \mathsf{epoch}')$. Similar to the analysis for regular commits, we consider two cases: i) u_{exini} is compromised when a client `id` is exposed while processing the commit, in which case **safe**(c) does not hold, or ii) the reduction $\mathcal{B}$ calculates u_{exini} itself if the seed s_{exini} and is exposed. The case i) happens if an id's key bundle is exposed during the commit processing, which again implies **safe**(c) is false, while the case ii) happens if it is encrypted under a node whose private key is compromised. The seed s_{exini} is encrypted under the key from the u_{ext} node of epoch $\mathtt{Node}[c].\mathsf{par}$. A

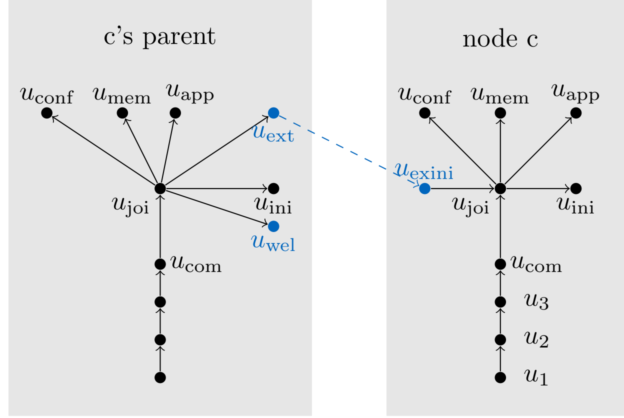

Figure 25: GSD game graph of an external commit, following an image from [4]. Each circle represents a node in the GSD graph. Dashed arrows mark an encryption of the sink using the key of the source, while bold arrows mark a Hash call on the source. The additionally added nodes of the graph are marked in ●. The external secret encrypts the committer-generated seed s_{exini} that is used as the init secret of epoch c.

u_{ext} node is exposed in the exact same cases as a u_{app}, and as in a commit, proof continues recursively until c is a root and hence **safe**(c) is false.

The subgraph G continues to correctly describe all cases in which a node secret stays in the tree: all commits that remove a node (c, u) are either (i) commits by an $\texttt{id}_i$ or (ii) removes of an $\texttt{id}_i$, on the subtree of $\tau.v$. This also includes external commits and resyncs.

Part 2.b.II) Allowing Injections In this part of the proof, we begin by outlining the adversarial modifications to the functionalities and the simulator, incorporating the capability to inject malicious messages.

From draft 12 to RFC

- Injected proposals: Additionally, in case of `extAdd` proposal, similar to `add` proposal, reduction $\mathcal{B}$ creates a new node using the private signing key used from the message.
- Commits injected to process: In the case that $\mathcal{Z}$ makes `id` process an injected commit c' from a party $\texttt{id}_c \neq \texttt{id}$, reduction $\mathcal{B}$ attempts to build new commit node's state. For {`add`, `extAdd` }-only commits, the reduction $\mathcal{B}$ only applies the proposals, as these commits do not require rekey operation. For other commits along with the ciphertexts `ctxt` that `id` would decrypt with keys in the ratchet tree, the reduction $\mathcal{B}$ proceeds as follows:

  i-ii) first two items remain the same.

  iii) if the `ctxt` is copied from an honest commit that generated earlier by $\mathcal{B}$, for which the GSD node associated to both `appSecret` is still a valid challenge, we must show that `id` accepting such a commit c would enable $\mathcal{B}$ to succeed in the GSD game. The trick in this step is to reduce the forgery of the confirmation tag `confTag` in the commit c to the GSD security of the underlying PKE. Let $\mathcal{B}$ challenge the GSD node u and extract the correct seed from $\mathcal{Z}$'s random oracle calls as follows:

  1. Note that $\mathsf{appSecret} = \mathsf{Hash}(\texttt{joinerSec}, .)$ and $\texttt{joinerSec} = \mathsf{Hash}(\mathsf{initSecret}, \mathsf{commitSec}, .)$, where HKDF is modeled as RO. Moreover, note also that `commitSec` must be the same in c and c' for the shared `ctxt` to be accepted.
  2. The reduction $\mathcal{B}$ needs to extract the `joinerSec` of c'. To this end, the reduction $\mathcal{B}$ proceeds as follows: First, recall that the `confTag` is a MAC of `confKey` and `confTransHash`. The only way for $\mathcal{Z}$ to compute a valid tag is to query RO on $(\mathsf{confKey}, \mathsf{confTransHash})$. The reduction $\mathcal{B}$ can search the RO query history and extract `confKey`. Second, recall that `confKey` is derived by applying HKDF to `joinerSec`. The environment $\mathcal{Z}$ needs to query RO on (`joinerSec`) for deriving `confKey`. The reduction $\mathcal{B}$ can then extract `joinerSec` of c' in a similar way.
  3. This case also remains the same.

- Injected vector of welcome messages `vec_w`: with the recent update of RFC, the use of `init-key` to encrypt group secrets results in a significant change to reduction $\mathcal{B}$'s processing order in [4]; however, this does not

have a big impact on the proof. When $\mathcal{Z}$ injects a welcome message $w = (\texttt{encGroupSec}, \texttt{encGroupInfo})$ in the welcome vector vec_w, $\mathcal{B}$ does as follows:

1. Note that wel_type of the corresponding vec_w does not affect the behavior of the reduction. Reduction $\mathcal{B}$ searches for a key package kp such that $\mathsf{SK}[\texttt{id}, \texttt{kp}] \neq \perp$ and $\mathsf{ISK}[\texttt{id}, \texttt{kp}] \neq \perp$ and $\mathsf{H}(\texttt{kp})$ matches an entry $e \in$ encGroupSecs and aborts if no such kp exists.
2. If e is copied from a vector of welcome messages vec_w generated by $\mathcal{B}$ while creating a commit node c and ISK[id, kp] is a GSD node, to win the game, reduction $\mathcal{B}$ proceeds as follows: $\mathcal{B}$ needs to differentiate between the cases, whether the injected welcome message is associated with the commit c that the e stems from. To this end, $\mathcal{B}$ compare encGroupSec. Since e is copied, the joinerSec underlying different welcome messages that include the same e must be the same. Furthermore, an equal encGroupInfo implies an equal confTransHash. In this case, $\mathcal{B}$ moves id to c.

   If the encGroupInfo in the vector of welcome message vec_w is different from the one derived by the reduction $\mathcal{B}$ for c, that means that e was copied but used to create the welcome for a non-existent commit c'. This is not possible as it would allow $\mathcal{B}$ to win the GSD game: In order to construct a encGroupInfo that is decryptable using the secrets stored in the encGroupSecs, the environment $\mathcal{Z}$ must be in possession of the welcome_secret (as $\mathcal{Z}$ cannot copy encGroupInfo from c). In order to derive the welcome_secret, $\mathcal{Z}$ must derive it from the joinerSec using the RO calls. The reduction $\mathcal{B}$ can extract joinerSec by checking queries that have been sent to RO one by one as follows: First, the reduction $\mathcal{B}$ attempts to parse the input of RO into joinerSec candidates. If the parse fails, $\mathcal{B}$ directly omit the following steps and attempts for the next RO input. Second, the reduction $\mathcal{B}$ respective derives welcome_secret and confKey from the joinerSec candidates. Third, the reduction $\mathcal{B}$ derives welcome_key and welcome_nonce from welcome_secret. Fourth, the reduction $\mathcal{B}$ decrypts encGroupInfo to recover groupInfo. Fifth, the reduction $\mathcal{B}$ parses groupInfoTBS from groupInfo and further confTag from groupInfoTBS. Finally, the reduction $\mathcal{B}$ verifies the confTag by using the confKey. If no error occurs, then reduction $\mathcal{B}$ found the correct joinerSec. If any error occurs, then the reduction $\mathcal{B}$ continue to check the next RO query. Note that the environment $\mathcal{Z}$ can send only polynomial queries to RO. The checks that $\mathcal{B}$ needs to execute must also be in polynomial time.

   From joinerSec, the reduction $\mathcal{B}$ can calculate the appSecret for the valid GSD challenge node c and hence win the game. Copying e is hence not possible.
3. Else, if e has not been copied from any vector of welcome messages generated by $\mathcal{B}$, the reduction $\mathcal{B}$ can simply obtain the encrypted joinerSec and pathSec either by using the secret in $\mathsf{ISK}[\texttt{id}, \texttt{kp}]$, if it has been compromised and therefore known by $\mathcal{B}$, or otherwise by querying $\mathcal{B}$'s GSD decryption oracle **Dec**. By using the joinerSec, the reduction $\mathcal{B}$ can decrypt and verify the encGroupInfo in the vector of welcome messages. If the verification passes, then $\mathcal{B}$ further checks whether the decrypted groupInfo indicates an existing node c (by comparing the confTransHash $\in$ groupInfo) or a new node c'. If groupInfo indicates an existing node c, then $\mathcal{B}$ moves id to c. Otherwise, groupInfo indicates a new node c'. In this case, $\mathcal{B}$ creates c' with labels taken from groupInfo and the ratchet tree set to the public part of τ from groupInfo. Then, for any node of τ with a public key for which it has a secret key stored (in another ratchet tree or in ISK and SK), $\mathcal{B}$ copies the secrets into τ. Moreover, $\mathcal{B}$ updates ratchet tree secrets to those derived from pathSec (if any secret key was set to a GSD node, $\mathcal{B}$ uses pathSec to win the game), and computes the epoch secrets from joinerSec. For all other unknown secrets in τ, $\mathcal{B}$ keeps them to be $\perp$.

Injected groupInfo. In case the environment $\mathcal{Z}$ makes id process an injected groupInfo g to create an external commit, the reduction $\mathcal{B}$ proceeds as follows:

First, the reduction $\mathcal{B}$ checks for a commit c that includes the matching confTag, i.e., confTransHash $\in g$. If such commit c exists, then $\mathcal{B}$ further checks whether all other values included in g are equal to another groupInfo g' that was previously generated by $\mathcal{B}$. If so, g indicates an existing node, i.e., environment $\mathcal{Z}$ is re-sending the groupInfo g for c from another client. $\mathcal{B}$ proceeds as for an honest g and appends the external commit to c.

If the confTag corresponds to a commit c and a previously sent g', but any values other than signer and signature are different from that previous g', the reduction $\mathcal{B}$ treats it like any other injected g: if id accepts the message, $\mathcal{B}$ creates the new node with labels taken from groupInfo and the ratchet tree set to the public part of τ from groupInfo. Then, for any node of τ with a public key for which it has a secret key stored (in another ratchet tree or in ISK and SK), $\mathcal{B}$ copies the secrets into τ. Other unknown secrets remain $\perp$.

The reduction $\mathcal{B}$ then proceeds to create the external commit. If the external key is (gsd, u), which qualifies as a valid secret for GSD, $\mathcal{B}$ utilizes the oracle **Enc** to derive the external seed s. Otherwise, $\mathcal{B}$ performs the encryption itself.

This behavior only affects detached trees, as any injected groupInfo that does not correspond to an honest commit immediately creates a new detached root. Consequently, the proof for the main tree remains unaffected, and $\neg$**safe**(c) is addressed by condition b) of ***state-directly-leaks**$(c, \texttt{id})$. For the adaptation of the proof concerning detached trees from Section D.5, refer to the explanation below.

External Commit Injected to Process. The reduction $\mathcal{B}$ constructs the new node including the initSecret as follows: On one hand, if the external-key, which encrypts the shared secret s_{exini} in the extInit proposal of the external commit, is a GSD node, i.e., neither $\mathcal{B}$ nor $\mathcal{Z}$ has knowledge of it, $\mathcal{B}$ uses the **Dec** oracle of GSD to retrieve the shared secret s_{exini}. On the other hand, if the external-key is a value, meaning that both $\mathcal{B}$ and $\mathcal{Z}$ know it, $\mathcal{B}$ directly decrypts s_{exini}. The external-key cannot be $\perp$, as it is not a key that the adversary can inject. From s_{exini}, $\mathcal{B}$ can then derive the new initSecret.

The only case where $\mathcal{B}$ will not be able to decrypt s_{exini} is if the kem_output stems from an honest external commit c' previously generated by $\mathcal{B}$ and copied by $\mathcal{Z}$ and whose appSecret is still a valid challenge. If an id accepts such a copied kem_output, $\mathcal{B}$ can use it to win the GSD-game similarly to a copied path secrets in commits:

1. Since the kem_output is copied, the ext-initSec is identical in both the injected external commit c and the honest external commit c'. At the same time, confTransHash must differ; otherwise, c and c' would be identical.
2. For $\mathcal{Z}$ to compute a valid confTag for the injected commit c, it must query the RO on H(confTransHash, confKey). From this query, $\mathcal{B}$ can extract the confKey. Since confKey is derived from (joinerSec), $\mathcal{Z}$ must also query RO on (joinerSec) to compute a valid tag.
3. For $\mathcal{Z}$ to compute the joinerSec for the injected commit c, it needs to query RO on Hash(ext-initSec, commitSec). From this query, $\mathcal{B}$ can extract the ext-initSec, which is shared with c', by searching the RO query history.
4. $\mathcal{B}$ can then corrupt the commitSec of c', which does not impact the challenge, and compute the joinerSec of c'. Using this joinerSec, the reduction $\mathcal{B}$ calculates the appSecret.

Thus, this case cannot occur.

The reduction $\mathcal{B}$ handles the updatePath of the commit the same way as a regular one, as the argumentation of an injected commit relies on calculating the initSecret of c from the RO calls and corrupting the initSecret of c' directly - regardless of its origin. Note that in the case of an injected external commit, the initSecret always retrieved by $\mathcal{B}$ directly via decryption. This does not affect the argument showing it is impossible to copy updatePath secrets from an honest commit.

The final case to consider is when both s_{exini} and ctxt are copied. If they are copied from different commits, the confTag can be used to break the challenge for both commits in the same way as before. Figure 26 illustrates how $\mathcal{B}$ utilizes $\mathcal{Z}$ to win the GSD game in case of a copied ctxt.

If they are copied from the same commit but the confTransHash differs from the original commit (e.g., when the external commit is injected for a different party), the joinerSec, retrievable via the confKey RO call, are the same between the two commits. The reason behind is that the same ciphertext indicates same path secret and proposals, which further indicate the same commitSec . Combing the fact that s_{exini} is copied from the same commit, which yields the same initSecret, the joinerSec must also be same. Thus, this allows the reduction $\mathcal{B}$ to compute the appSecret to be computed from it.

Proof of Indistinguishability. It remains to show that gsd-exp() $\implies \neg$**safe**(c). As a reminder, this is part **2.b.II** of the main proof. The subproof is structured as follows:

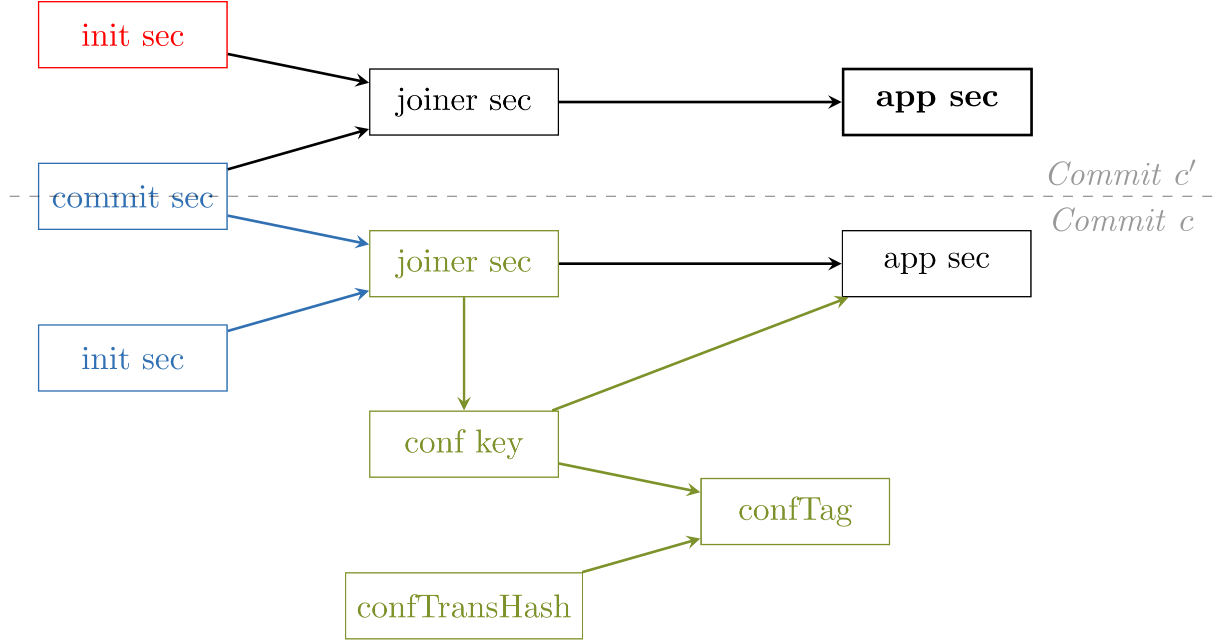

Figure 26: Visualization of $\mathcal{B}$ using the `confTag` to win the GSD game for a commit c' if the environment $\mathcal{Z}$ successfully copies a ciphertext from a commit c' into an injected commit c. The commit secret is shared between both commits. In the first step (in ●), $\mathcal{Z}$ needs to call the random oracle RO to derive the `confTag` from the `confKey` and the confTransHash, which $\mathcal{B}$ now learns from the query. $\mathcal{Z}$ also needs to query the RO to derive `confKey` from `joinerSec`. $\mathcal{B}$ now knows `joinerSec` of c. In order to derive `joinerSec` for commit$c$, $\mathcal{Z}$ also needs to query Hash(`commitSec`, initSecret), which $\mathcal{B}$ now learns (in ●). $\mathcal{B}$ now knows the `commitSec` of commit c', and corrupts the initSecret of commit c' (in ●). $\mathcal{B}$ can now calculate appSecret of c'.

Part 2.b.II): Confidentiality when allowing injections

The reduction $\mathcal{B}$ succeeds in winning the GSD game for commit c' if the environment $\mathcal{Z}$ successfully copies a ciphertext from commit c' into an injected commit c. This can be formulated as follows:

a) Show that if the reduction $\mathcal{B}$ corrupts (c, u_{app}) then $\neg\textbf{safe}(c)$

b) Show that if the reduction $\mathcal{B}$ corrupts (c, u_{joi}) then $\neg\textbf{safe}(c)$

c) Show that if both gsd-exp$((c, u_{\text{com}}))$ and gsd-exp$((\texttt{Node[c].par}, u_{\text{exini}}))$ are true then $\neg\textbf{safe}(c)$. For this, consider all cases where gsd-exp$((c, u_i))$ is true:

  c.1) During an exposure of an `id` that stores u_i

  c.2) When a node $\tau.v$ used in `*rekey-path` is exposed

  c.3) If the secret in $\tau.v$ is $\perp$, which is the case if:

    i) $\mathcal{Z}$ injects a new commit c' on behalf of `id`

    ii) the injected commit c' commits $\tau.v$ as an updated injected leaf

    iii) the injected c' commits $\tau.v$ as an injected key package

The proof is almost the same as in Step 1 (no injections). In the case that c is a regular commit, the proof is extended the same way as in step 2 of the original proof in [4], except for differences are c).

In case c), we still need to prove that gsd-exp$((c, u_{\text{com}})) \implies$ know$(c, .)$, by considering the following three sub-cases, i.e., gsd-exp$((c, u_{\text{com}})) =$ `true` is triggered c.1) during an exposure of an `id` that stores u_i's secret, and c.2) when a node $\tau.v$ used in `*rekey-path` is exposed, and c.3) if the secret in $\tau.v$ is $\perp$.

- c.1) is covered if `id` processed c and has not performed any action *that includes an UpdatePath*, which encompasses any form of external commit but excludes {`add`, `extAdd`}-only commits. Secondly, injecting a node from a previous, honest commit *or external commit* would allow the reduction $\mathcal{B}$ to use the `confTag` to win the GSD game.

- c.2) encompasses the subgraph G that is determined by actions adding and removing the node, *including external commits* and *excluding* {add, extAdd}*-only commits* that do not perform an updatePath.
- c.3) remains the same, as subcase i) $\mathcal{Z}$ injects a new commit c' on behalf of a party id into the subtree, encompasses both scenarios, where the injected c' is a regular or an external commit.

In the case that c is an external commit - if there is an edge between $(c, u_{\texttt{exini}})$ and (c, u_{joi}) - the proof proceeds as follows: gsd-exp(c, u_{app}) is only true in three cases

a) $\mathcal{B}$ corrupts (c, u_{app}) - same as in step 1
b) $\mathcal{B}$ corrupts (c, u_{joi}) - same as in step 1
c) if both gsd-exp$((c, u_{\text{com}}))$ and gsd-exp$((\texttt{Node[c].par}, u_{\text{exini}}))$ are true:
   - For gsd-exp$((c, u_{\text{com}}))$ the argumentation remains the same as for regular commit. Note that an injected external commit c' (case i)) hence implies know(c).
   - In addition, gsd-exp$(((\texttt{Node}[c].\mathsf{par}, u_{\text{exini}})))$ can also be true if injected: When the environment $\mathcal{Z}$ injecting the seed s_{exini} in the extInit proposal, the reduction $\mathcal{B}$ decrypts it and calculates u_{exini} itself. In the case of an injected external commit, condition e) of ***can-traverse** immediately holds, and hence **safe**(c) is false.

In detached trees: Extension D.5 in main reference [4] concerning compromise in detached trees needs to be adapted with only minor changes. The proof shows that if an epoch in a detached tree does not include any exposed signing keys, it also does not include any exposed node keys. The first part shows that if $\tau.v$ has an exposed leaf key, then the signature key must also be leaked. This holds because a leaf key is only compromised if either the owning member is compromised, in which case the signature key spk is also marked as exposed, or the leaf is injected as part of a commit or update, in which case $\mathcal{F}_{\text{ECGKA}}$ also marks spk as exposed. A third, not explicitly named case is the use of an adversarially generated ratchet tree in a welcome or groupInfo. As each leaf node contains a signature using the nodes spk, the adversary must know the spk by corruption to produce the ratchet tree and spk must have been marked as exposed, otherwise $\mathcal{Z}$ could be used to break SUF-CMA.

The second part shows that for any inner node, assuming no node in its subtree is exposed, the leaf key pk is not exposed. The argument stays the same: 1) all signatures verify because at least one party has accepted the commit or welcome *or external commit or the* groupInfo g *to create an external commit.* 2) As the committer of an inner node $\tau.v$, spk cannot have been compromised (by assumption). $\tau.v$ cannot be maliciously generated, as the parentHash of $\tau.v$ is included and signed in the leaf of the committer. The only possible leakage is hence an encryption to an exposed key. However, this is impossible due to the induction hypothesis that states no lower nodes are compromised. The signed parent hash in $\tau.v$ guarantees here that the subtree of $\tau.v$ is the same as when $\tau.v$ was introduced into the tree.

Part 2.b.III) Allowing Bad Randomness *Proposals by Reference and by Value* For the extAdd proposal, $\mathcal{B}$ generates the key package using id's spk. Similarly, by using the randomness provided by $\mathcal{Z}$ creates the proposal message p.

Commit by Reference and by Value. Given randomness provided by $\mathcal{Z}$, $\mathcal{B}$ computes the commit and required secrets as:

1. The reduction $\mathcal{B}$ executes *rekey-path to obtain the path secrets, commitSec, and commitContent. It then generates a valid confTransHash.
2. $\mathcal{B}$ creates a new joinerSec, which is a hash of the current initSecret and the newly computed commitSec, i.e., joinerSec = H (initSecret, commitSec, ...). If initSecret stores a GSD node u, $\mathcal{A}$ queries the H$(u, u_{\text{ctr}},$ commitSec$)$, increments u_{ctr}, and uses the fresh, uncorrupted node as the joinerSec. Otherwise, if initSecret stores a value, $\mathcal{A}$ computes the joinerSec itself.
3. Using the secrets obtained, $\mathcal{B}$ executes the key schedule to compute the confTag. It then computes the commit message c and the vector of welcome messages *welcome-msg accordingly.

Validity of the Challenge. Throughout this study commits are not necessarily calculated using randomness. The reduction $\mathcal{B}$'s actions and the proof remain the same for commits that involve fresh randomness. Recall that the helper predicate ***secrets-replaced** captures the update of the path secret, except for the proposal-list in regular commit c is *non-empty* or only includes regular add or external self-add. Any

commit c'' (except only regular add or external self-add proposals) with ***secrets-replaced**$(c'', \texttt{id})$ would hence replace secrets injected by randomness.

Now we will show with above changes, randomness corruption does not affect existing challenges, where **safe**(c) holds. With corrupted randomness the environment $\mathcal{Z}$ cannot derive the correct `epochSec` because the environment $\mathcal{Z}$ needs to query RO on `joinerSec` to compute `epochSec` = HKDF(`joinerSec`).

We adapt the proof that randomness corruption does not affect the validity of a challenge, and gsd-exp$(c) \implies \neg$**safe**(c) still holds. Due to above mentioned changes, no additional corruption is performed by $\mathcal{B}$ on the `joinerSec`, case b) of gsd-exp holds as before.

Only case that is affected is case c). Here, the proof remains in agreement with [4], showing that gsd-exp$((c, u_{\text{com}})) \implies \neg$know$(c, *)$ also when allowing corruptions. We informed the authors of this bug, and they are in the process of updating the paper accordingly.

Randomness in Detached Trees. In addition to corruption and injection, a leaf or intermediate node $\tau.v$ may also be exposed if it is generated with bad randomness. The proof for detached trees relies on the fact that if any $\tau.v$ is exposed, a member's spk would likewise be revealed, thereby implying $\neg$**safe**(c). However, when committing with corrupted randomness alongside a fresh spk, as described in [4] (corresponding to an `updatePath` in ETK), only the old spk used to sign the commit is leaked, while the fresh spk remains secure. The spk, however, is exposed if it is registered with the AS while the randomness is compromised. The following hence does not contradict **safe**(c) through the condition b) of ***state-directly-leaks**$(c, \texttt{id})$:

- The environment $\mathcal{Z}$ creates a detached root by inviting B to a false group, impersonating A
- B adds C and removes A
- Member B registers spk with $\mathcal{F}_{\mathcal{AS}}^{\mathcal{IW}}$
- The environment $\mathcal{Z}$ corrupts B's randomness
- Member B performs a commit which updates its signature key to spk.

It does, however, contradict **safe**(c) through the condition c) of ***state-directly-leaks**$(c, \texttt{id})$, as condition a) of ***secrets-injected**$(c, \texttt{id})$ holds for commits with corrupted randomness. Since the above scenario requires the node in the detached tree to have healed completely, safe is applicable in the same way as it is in the main tree. After such a heal has occurred, condition b) of ***state-directly-leaks**$(c, \texttt{id})$ is always "overlayed" by other conditions that hold after a compromise.

The proof from [4, Section D.5] that shows that from **safe**(c) follows that no node $\tau.v$ has an exposed public key pk needs to be adapted in the following way: For any leaf node with an exposed pk, either spk must have been exposed, *or the commit c introducing it to the tree was committed with compromised randomness*, both of which contradict **safe**(c). For any intermediate node, since spk is not exposed, so $\tau.v.pk$ is honestly committed (see [4] for details). *Additionally committing with bad randomness would have triggered condition a) of* ***secrets-injected**$(c, \texttt{id})$*, contradicting the assumption that* **safe**$(c) = \texttt{true}$*, and in combination with* ***secrets-replaced**$(c, \texttt{id})$ *would have remained so until removed.* This means the secret key of a node is not chosen by $\mathcal{Z}$.

External Commit from `id`. Behaves similarly to a regular commit: Using the randomness provided by $\mathcal{Z}$, reduction $\mathcal{B}$ computes the external commit and the secrets in new commit node, as follows:

1. $\mathcal{B}$ uses the randomness provided by $\mathcal{Z}$ to execute `*rekey-path-upon-join`, obtaining all path secrets, the `commitSec`, and the intermediate commit packet. Then, it signs the commit packet using the `id` 's spk, along with $\mathcal{Z}$'s randomness, and sets the confTransHash accordingly.
2. A key distinction from a regular commit is that the initSecret for an external commit is determined by the client. As a result, corrupted randomness can expose this initSecret. Reduction $\mathcal{B}$ computes the ext-initSec as follows:
    - If the external secret from the previous epoch stores a GSD node u, $\mathcal{B}$ queries **Enc**(u, u_{ctr}), corrupts u_{ctr}, and sets the external seed s_{exini}, increments u_{ctr}. $\mathcal{B}$ then uses this seed as the initSecret for the external commit.
    - If the external secret from the previous epoch stores a value (and is therefore known to $\mathcal{B}$), $\mathcal{B}$ computes the `joinerSec` directly.
3. Using the `joinerSec` and the confTransHash from previous steps, $\mathcal{B}$ runs the key schedule, computes the `confTag`, and finalizes the commit message c.

Proof for External Commit.

Case a) and case b) of the proof remains the same.

Case c): To show $\mathsf{gsd\text{-}exp}((c, u_{\text{com}})) \implies \neg\mathsf{know}(c, *)$ for external commits, we proceed similar as for regular commits. $\mathsf{gsd\text{-}exp}((c, u_i))$ can also be true if the secret in a ratchet tree node $\tau.v$ used in `*rekey-path-upon-join` stores a seed s generated during an action executed with bad randomness. Consider an external commit c' that inserts s into $\tau.v$. Unlike a regular commit, an external commit injects only the secrets it generates during the commit process. It does not, for instance, include an update proposal from another `id` or a corrupted key package. This injection is covered by condition a) of ***secrets-injected**$(c, \mathtt{id})$. Now consider the case where c itself is the external commit. Since the external committer utilizes the same public ratchet tree as a member performing a regular commit, the proof remains the same to regular commit. □

E.3 Authenticity

Intuition. We prove that *H3* and *H4* are indistinguishable by demonstrating that the environment $\mathcal{Z}$ in *H4* cannot inject any message unless the statement **inj-allowed** is `true`. Note that the evaluation of **inj-allowed** only appears in **auth-invariant**. Thus, we need to prove that **auth-invariant** never triggers. We determine all cases where it would trigger as follows and prove that these events do not occur.

1. We show that if a forged signature event occurs, *H4* could be used to win the SUF-CMA game of the underlying digital signature.
2. We show that if a forged MAC event would occur, *H4* could be used to win the EUF-CMA game of the underlying MAC (and that the advantage is upper bounded by the GSD game).

Note that in ETK, messages are authenticated in two ways: First, public handshake messages include a MAC generated with the membership key derived from the epoch secret, confirming the message is from within the group. Private handshake and application messages on the other hand are encrypted with a group-only key and signed with the sender's leaf signature key, ensuring the sender's identity and message authenticity. Thus, the indistinguishability of *H3* and *H4* indicates that ETK guarantees authenticity.

Lemma 3. *If* Sig *is* SUF-CMA *secure,* MAC *is* EUF-CMA *secure, and* PKE *is* GSD *secure, then Hybrids H3 and H4 are indistinguishable, that is,* ETK *guarantees authenticity.*

Proof. Note that the hybrids *H3* and *H4* are identical unless in *H4* there exists an injected history graph node in $\mathcal{F}_{\text{ECGKA}}$ for which **inj-allowed** requires authenticity. More concretely, the hybrids *H3* and *H4* are identical unless the following happens:

- Event `Bad`: There exists a commit or proposal node with $\mathtt{stat} = \mathtt{adv}$ and **inj-allowed**$(c, \mathtt{id}) = \mathtt{false}$ for its parent c and creator `id`.

Note that **inj-allowed**$(c, \mathtt{id}) = \mathtt{false}$ if `id`'s spk in commit c is not exposed or the epoch key in c is not leaked to the reduction. We then can decompose the event `Bad` into the following two sub-events.

- Event $\mathtt{Bad_{sig}}$: There exists a commit or proposal node with $\mathtt{stat} = \mathtt{adv}$ and $\mathtt{Node}[c].\mathsf{mem} \notin \mathsf{Exposed}$ for its parent c and creator `id`.
- Event $\mathtt{Bad_{MAC}}$: There exists a commit or proposal node with $\mathtt{stat} = \mathtt{adv}$ and $\neg\mathbf{know}(c, '\mathsf{epoch}')$ for its parent c.

The proof is concluded by combining Lemma 4 and Lemma 5, where we respectively prove that the events $\mathtt{Bad_{sig}}$ and $\mathtt{Bad_{MAC}}$ will never happen except for negligible probability. □

Lemma 4. *For any environment $\mathcal{Z}$, there exists a reduction $\mathcal{B}_{\mathtt{sig}}$ that succeeds in the* SUF-CMA *game with a probability that is only polynomially smaller than the probability of $\mathcal{Z}$ triggering* $\mathtt{Bad_{sig}}$.

Proof. For any environment $\mathcal{Z}$, $\mathcal{B}_{\mathtt{sig}}$ emulates the functionalities and simulator by embedding its challenge spk as one of the public keys honestly created during the experiment. To emulate commits signed under ssk, $\mathcal{B}_{\mathtt{sig}}$ calls **Sign** oracle. When $\mathtt{Bad_{sig}}$ occurs, reduction $\mathcal{B}_{\mathtt{sig}}$ stops the experiment and sends to its challenger to forgery consisting of the sig', $\mathtt{tbs}'$ from the injected node c'.

First assume c' is a commit node. We show that if $\mathtt{Bad_{sig}}$ occurs and $\mathsf{spk} = \mathtt{Node}[c].\mathsf{mem}[\mathtt{id}]$, then $\mathcal{B}_{\mathtt{sig}}$ wins.

1. We know sig' is a valid signature over $\mathtt{tbs}'$. The injected node was created when some party accepted c' which means that it verified sig' under $\mathsf{spk} = \mathtt{Node}[c].\mathsf{mem}$.
2. Simulation differs from the experiment when $\mathsf{spk} = \mathtt{Node}[c].\mathsf{mem}$ is corrupted, however this does not happen as the event guarantees $\mathsf{spk} = \mathtt{Node}[c].\mathsf{mem} \notin \mathsf{Exposed}$.
3. The reduction $\mathcal{B}_{\mathtt{sig}}$ has never queried **Sign** on $\mathtt{tbs}'$ for sig' to its signing oracle. Assuming $(\mathsf{sig}', \mathtt{tbs}')$ is the same as $(\mathsf{sig}^*, \mathtt{tbs}^*)$ queried by $\mathcal{B}_{\mathtt{sig}}$ to the sign oracle for c^*, where c' is injected and c^* is honestly generated by $\mathcal{B}_{\mathtt{sig}}$. By showing our assumption implies $c' = c^*$, we will achieve a contradiction. Now we will divide the proof for two, according to the type of the commit c'.
   - We first consider that c' is a regular commit. c' contains $(\mathsf{groupId}', \mathsf{epoch}', \mathsf{leafIdx}', \mathtt{'Commit'}, C', \mathtt{confTag}', \mathsf{sig}', \mathtt{membTag}')$ and $\mathtt{tbs}' = (\mathsf{groupCtxt}', \mathsf{groupId}', \mathsf{epoch}', \mathtt{senderIdx}', \mathtt{'Commit'}, C')$ and c^* contains the analogous values. Then c' and c^* can only differ on $\mathtt{confTag}' \neq \mathtt{confTag}^*$ or $\mathtt{membTag}' \neq \mathtt{membTag}^*$. Note that $\mathtt{membTag}' = \mathsf{MAC.tag}(\mathsf{membKey}', C')$ and that $\mathtt{membTag}^* = \mathsf{MAC.tag}(\mathsf{membKey}^*, C^*)$. According to our assumption $\mathtt{tbs}' = \mathtt{tbs}^*$ implies $C' = C^*$ and $\mathsf{groupCtxt}' = \mathsf{groupCtxt}^*$. Group context determines the epoch and key schedule. We have $\mathsf{membKey}' = \mathsf{membKey}^*$ and further $\mathtt{confKey}' = \mathtt{confKey}*$. Accordingly $\mathtt{membTag}' = \mathtt{membTag}^*$ and $\mathtt{confTag}' = \mathtt{confTag}^*$, which gives a contradiction.
   - Then we consider that c' is an external commit. c' contains two components: $\mathtt{framedContent} = (\mathsf{groupId}', \mathsf{epoch}', \mathsf{'new_member_commit'}, \mathsf{'commit'}, C', \mathsf{sig}')$ and $\mathtt{framedContentAuth}' = (\mathsf{sig}', \mathtt{confTag}')$ and c^* contains the analogous values. Moreover, $\mathtt{tbs}' = (\mathsf{groupCtxt}', \mathsf{groupId}', \mathsf{epoch}', \mathtt{senderIdx}', \mathtt{'Commit'}, C')$ and $\mathtt{tbs}^*$ contains the analogous values. Similar to regular commit, c' and c^* can only differ on $\mathtt{confTag}' \neq \mathtt{confTag}^*$. By definition it holds that $\mathtt{confTag}' = \mathsf{Hash}(\mathsf{confTransHash}', \mathtt{confKey}')$ and $\mathtt{confTag}^* = \mathsf{Hash}(\mathsf{confTransHash}^*, \mathtt{confKey}^*)$. Since $\mathtt{confTag}$ is derived on the framed content, we have $\mathtt{confTag}' = \mathtt{confTag}^*$. However this gives us contradiction as c' is injected and c^* is honestly generated.

$\square$

Lemma 5. *For any environment $\mathcal{Z}$, there exists a reduction $\mathcal{B}_{\mathtt{MAC}}$ and $\mathcal{B}_{\mathtt{PKE}}$ such that the probability that $\mathcal{Z}$ triggering $\mathtt{Bad}_{\mathtt{MAC}}$ is upper bounded by $p(\epsilon_{\mathsf{MAC}} + \epsilon_{\mathsf{PKE}})$, where p is a polynomial, ϵ_{MAC} is the advantage of $\mathcal{B}_{\mathtt{MAC}}$ against the security of MAC and ϵ_{PKE} is the advantage of $\mathcal{B}_{\mathtt{PKE}}$ against the security of* PKE.

Proof. For an injected packet, $\mathcal{Z}$ can trigger $\mathtt{Bad}_{\mathtt{MAC}}$ in the following cases. First, we consider the case where all receivers verify a MAC tag $\mathtt{confTag}$ but not $\mathtt{membTag}$. The reduction $\mathcal{B}_{\mathtt{MAC}}$ runs $\mathcal{Z}$ which emulates the UC experiment exactly, except that it uses its oracles in the EUF-CMA game of MAC, instead of using the MAC key $\mathsf{membKey}$ in epoch c. If $\mathtt{Bad}_{\mathtt{MAC}}$ occurs for parent node c and injected child c', reduction $\mathcal{B}_{\mathtt{MAC}}$ outputs forgery $(\mathtt{confTag}', \mathtt{tbm}')$ where $\mathtt{tbm}'$ denotes the content to be MACed.

1. The difference between the probability of $\mathtt{Bad}_{\mathtt{MAC}}$ occurring for c in the experiment emulated by the reduction $\mathcal{B}_{\mathtt{MAC}}$'s and in Hybrid *H3* (or *H4*) is upper bounded by the advantage of the reduction $\mathcal{B}_{\mathtt{PKE}}$ in the GSD game.
2. If $\mathtt{Bad}_{\mathtt{MAC}}$ occurs for c, then the reduction $\mathcal{B}_{\mathtt{MAC}}$ wins with $(\mathtt{confTag}', \mathtt{tbm}')$.

For the first item, $\mathtt{Bad}_{\mathtt{MAC}}$ for c implies that $\mathbf{safe}^*(c)$ is true until the event and no party in c is exposed. We can construct $\mathcal{B}_{\mathtt{PKE}}$ implicit authentication, i.e, confidentiality proof. The reduction $\mathcal{B}_{\mathtt{PKE}}$ embeds the challenge in $\mathsf{membKey}$ in c and proceeds the experiment as until $\mathtt{Bad}_{\mathtt{MAC}}$ for c occurs. Thus it has no effect on the probability of the event.

For the second item, we show that for valid tag $(\mathtt{confTag}')$ that is verified by a party accepting the injected c', $(\mathtt{confTag}', \mathtt{tbm}')$ was not queried by the reduction $\mathcal{B}_{\mathtt{MAC}}$ to the MAC oracle. Assume $(\mathtt{confTag}', \mathtt{tbm}') = (\mathtt{confTag}^*, \mathtt{tbm}^*)$, where $(\mathtt{confTag}^*, \mathtt{tbm}^*)$ was queried by the reduction $\mathcal{B}_{\mathtt{MAC}}$ to the oracle when it honestly generated packet c^*. We need to prove that the reduction $\mathcal{B}_{\mathtt{MAC}}$ has never queried $\mathtt{tbm}'$ to its MAC oracle. Note that $\mathcal{B}_{\mathtt{MAC}}$ has only queried other $\mathtt{tbm}^*$ for other node c^* to its MAC oracle. Thus, for any $c' \neq c^*$, it always holds that $\mathtt{tbm}' \neq \mathtt{tbm}^*$. The proof for commits requiring a $\mathtt{membTag}$ can be demonstrated in a similar manner. External commits, on the other hand, do not include a $\mathtt{membTag}$; thus, the proof remains consistent with the proof provided for $\mathtt{confTag}$ above.

Now, consider commits that include only $\mathtt{extAdd}$ proposals. Note that such commits do not contain a framed MAC tags, e.g., $\mathtt{confTag}$ or $\mathtt{membTag}$. Consequently, the authenticity of these commits relies on the first part of **inj-allowed**, specifically the security of the Sig. $\square$

F Proof of Theorem 2

Theorem 2. Assume that PKE is IND-CCA secure and that Sig is SUF-CMA secure. The $\mathsf{ETK}^{\mathrm{PSK}}$ protocol securely realizes $(\mathcal{F}_{\mathcal{AS}}^{\mathcal{IW}}, \mathcal{F}_{\mathcal{KS}}^{\mathcal{IW}}, \mathcal{F}_{\mathrm{ECGKA}^{\mathrm{PSK}}})$ in the $(\mathcal{F}_{\mathcal{AS}}, \mathcal{F}_{\mathcal{KS}}, \mathcal{G}_{\mathrm{RO}})$-hybrid model, where $\mathcal{F}_{\mathrm{ECGKA}^{\mathrm{PSK}}}$ uses the predicates **safe** and **inj-allowed** from Figure 7 and calls to HKDF.Expand, HKDF.Extract, and MAC functions are replaced by calls to the global random oracle $\mathcal{G}_{\mathrm{RO}}$.

Proof. For the proof of $\mathsf{ETK}^{\mathrm{PSK}}$, we retain the description of *H1* − *H4*. Since $\mathsf{ETK}^{\mathrm{PSK}}$ differs from ETK only regarding PSK injection, most of the proof is identical to the one of ETK in Appendix E. Hence, we will focus on stating additional requirements and checks in this section. We again prove the indistinguishability between every two adjacent hybrids by Lemma 6 (in Appendix F.1), Lemma 7 (in Appendix F.2), and Lemma 8 (in Appendix F.3). □

F.1 Consistency and Correctness of $\mathsf{ETK}^{\mathrm{PSK}}$

Intuition. We prove the indistinguishability between *H1* ($\mathsf{ETK}^{\mathrm{PSK}}$) and *H2* by proving for every input operation that the outputs in $\mathsf{ETK}^{\mathrm{PSK}}$ and $\mathcal{F}_{\mathrm{ECGKA}^{\mathrm{PSK}}}$ of *H2* are the same. The **assert** statements in $\mathcal{F}_{\mathrm{ECGKA}^{\mathrm{PSK}}}$ of *H2* capture that if one party has processed a proposal, (external) commit, or welcome, all parties that process it in the future end up in the same shared group state. Thus, the indistinguishability between *H1* and *H2* indicates that $\mathcal{F}_{\mathrm{ECGKA}^{\mathrm{PSK}}}$ guarantees consistency.

Lemma 6. *Hybrids H1 and H2 are indistinguishable, that is,* $\mathsf{ETK}^{\mathrm{PSK}}$ *guarantees consistency.*

Regular (internal) Proposals. We consider two cases when $\mathcal{F}_{\mathrm{ECGKA}^{\mathrm{PSK}}}$ sends (Propose, act), firstly act ∈ {up-spk, add-id$_t$, rem-id$_t$, psk-epoch-use}, or secondly act ∈ {up-spk, rem-id$_t$, psk-epoch-use}. In both cases, assert statements and consecutively the indistinguishability of the statements remain the same.

External Proposals. All the assert statements and the proof remain the same.

Commits. Additionally, for each PSK proposal contained in $\vec{p}$, $\mathcal{S}_2$ determines the commit of the history graph the injected PSK pertains to. This cannot be derived from the protocol directly, hence $\mathcal{S}_2$ needs to maintain a separate mapping from commit to PSK. This is possible since $\mathcal{S}_2$ has access to the whole history graph. $\mathcal{S}_2$ returns these c in the order of the PSK proposals as $c_{\vec{psk}}$. If any pskId appears twice in a commit, $\mathcal{S}_2$ returns ⊥ instead. Then, $\mathcal{F}_{\mathrm{ECGKA}^{\mathrm{PSK}}}$ runs ***fill-props**. For every proposal p in the input $\vec{p}$ without a node, $\mathcal{S}_2$ sets orig and act according to p. This is always possible, as the basic check in ***apply-props** by $\mathsf{ETK}^{\mathrm{PSK}}$ guarantees that p is well-formed.

Indistinguishability of Output: An additional **req** statement may have the Commit algorithm return ⊥:

- **req** id ∈ Node[Prop[p].c_{psk}].has_psk: ***fixate-psk-refs** guarantees that every PSK proposal refers to a valid c_{psk}. This **req** statement will fail if id did not perform a Process of c_{psk} correctly (yet) or the same pskId appears twice. Due to the indistinguishability of Process between ETK and $\mathcal{F}_{\mathrm{ECGKA}}$, the same holds for $\mathsf{ETK}^{\mathrm{PSK}}$, and the PSKs of c_{psk} would have not been added to id's psk_store. Hence $\mathsf{ETK}^{\mathrm{PSK}}$ would have failed in `*derive-psk-secret`.

Assert Statements: The commit algorithm includes the following additional **assert** statements:

1. **assert** res ≠ ⊥: holds if ***fixate-psk-refs**((Prop[p] : $p \in \vec{p}$), $c_{\vec{\mathrm{psk}}}$) runs through. It would fail either because $\mathcal{S}_2$ could not find a suitable c_{psk} or because the epoch does not match. The latter must hold because of the direct correspondence of Node[c].epoch in $\mathcal{F}_{\mathrm{ECGKA}}$ and the epoch of c and the way that $\mathcal{S}_2$ choses c_{psk}. The former fails only if there is no preSharedKeyId for the epoch in the psk_store in which case, due to the direct correspondence of Node[c].has_psk in $\mathcal{F}_{\mathrm{ECGKA}}$ and the commit's existence in the psk_store in $\mathsf{ETK}^{\mathrm{PSK}}$, both ***req-correctness** will fail and *ack* will be false.
2. **assert** ***valid-spk** remains the same.
3. **assert** mem ≠ ⊥ will succeed in the following cases.
   - $p = p_{up} + p_{rem} + p_{add} + p_{psk}$ includes the additional p_{psk}. Because of the direct correspondence of Prop[p].act and the action of p, $\mathsf{ETK}^{\mathrm{PSK}}$ performs the same check and $\mathcal{S}_2$ will return ack = false. Prop[p] accurately reflects the action because either p is created by $\mathcal{S}_2$ according to act, or act is set by $\mathcal{S}_2$ extracting act from p.

- $\texttt{Node}[c].\mathsf{epoch} = \mathsf{epoch} \wedge \texttt{id} \in \texttt{Node}[c].\texttt{has_psk} \vee \nexists(\texttt{id}, *) \in G$ states that at least one PSK for each PSK proposal must be in id's possession. This is checked by ETK.

4. **assert** ***valid-successor** has the additional case:

   - $\texttt{Node}[c].\mathsf{epoch} = \texttt{Node}[\mathsf{Ptr}[\texttt{id}]].\mathsf{epoch}$: if $\texttt{Node}[c]$ is a detached root, the epoch was read out of c' and must match, as it is included in the confTransHash, based on which $\mathcal{S}_2$ chooses the root. If it was a pre-existing node, the confTransHash also matches (else c would not be equal to a previous node), and hence the same holds.

5. **assert** $\texttt{Wel}[\texttt{vec_w}] \in \{\perp, (c, c_{\text{psk}})\}$ guarantees that the vector of welcome message is different for each c. This is achieved by including the confTag over the transcript hash in the Welcome's groupInfo. In Commit, $\texttt{Wel}[\texttt{vec_w}]$ can previously only have been set to c and following the argumentation for Process, the probability for another $\vec{p}$ matching c is low, and hence $\vec{\texttt{psk}}$ is also the same according to the standard hybrid argument. The only other option for setting $\texttt{Wel}[\texttt{vec_w}]$ is in Join. There, c and $\vec{\texttt{psk}}$ are set to the value provided by $\mathcal{S}_2$, which retrieves c as the only fitting commit by comparing the confTransHash, or it is set to a detached node. In the latter case, c is providing the link to the detached node, hence ***attach** would have set $\texttt{Wel}[\texttt{vec_w}]$ to $(c, c_{\vec{\text{psk}}})$.

Process. $\mathsf{ETK}^{\text{PSK}}$ distinguishes the processing of external commits and regular commits by id handing in either an empty vector or the full proposal vector. Both for regular and (for both cases of) external commits, the simulator $\mathcal{S}_2$ retrieves c's epoch, orig', spk' from c. $\mathcal{S}_2$ also extracts the proposal vector $\vec{p}$ from c and the vector $c_{\vec{psk}}$ consisting of the commits of all injected PSKs or $\perp$ if any of the pskId appear twice.

Indistinguishability of Output:

For the regular commits, in addition to cases mentioned in ETK, the Process returns $\perp$ in the following cases:

- **req** $\mathsf{res} \neq \perp$: holds if **req** ***fixate-psk-refs**$(\texttt{Prop}[p] : p \in \vec{p})$ runs through. This is the case for a valid commit, where c_{psk} has been set before. Otherwise, **req** ***req-correctness** would have failed, and ETK would either also have failed or provided the necessary correct commits in $c_{\vec{psk}}$.
- **req** $(\texttt{id} \in \texttt{Node}[c_{\text{psk}}].\texttt{has_psk} \vee \nexists(\texttt{id}, *) \in G)$ holds as for Commit.

For the external commit we consider the following additional cases where Process returns $\perp$:

- **req** $\texttt{ExtCommitProps} \neq \perp$ holds as for Commit: if ***fixate-psk-refs** fails because of an unseen commit, ***req-correctness** would have failed, and ETK would have failed if it did not have the appropriate $c_{\vec{psk}}$.
- **req** $\texttt{id} \in \texttt{Node}[c_{\text{psk}}].\texttt{has_psk}$: due to the direct correspondence between has_psk and the PSK being in id's psk_store in ETK, this fails exactly if ETK fails.

Assert statements : The Process algorithm of $\mathsf{ETK}^{\text{PSK}}$ includes the following **assert** statements for regular commits:

- **assert** $\mathsf{mem} \neq \perp$ means *apply-props failed. If c already existed, then *apply-props must have been run on the exact same parameters before, especially on $\mathsf{Ptr}[\texttt{id}]$ (or ***req-correctness** would have failed while checking $\texttt{Node}[c].\mathsf{par} = \mathsf{Ptr}[\texttt{id}]$) and passed. Else, ***req-correctness** would have failed. If $\texttt{Node}[c] = \perp$ or c were malformed according to ETK (i.e. a proposal removing the committer), $\mathcal{S}_2$ would return $ack = \texttt{false}$ and Process would unwind. This means that c either corresponds to a correct detached root or a correct new node. It is left to show that every single **req** in *apply-props is either fulfilled or $\mathcal{S}_2$ would have returned $ack = \texttt{false}$, unwinding the function. This holds for the same reason as for Commit, the only exception being PSK proposals, since $\texttt{Node}[c].\mathsf{epoch} = \mathsf{epoch} \wedge \texttt{id} \in \texttt{Node}[c].\texttt{has_psk} \vee \nexists(\texttt{id}, *) \in G$ might hold for the committer but not the processing party. In such a case of a valid commit with unprocessable PSK proposals, ***req-correctness** would have failed and $\mathcal{S}_2$ would have returned $ack = \texttt{false}$. The same holds for an injected commit.

In the case of an external commit, *process-ec(c) contains the following **assert** statements:

- **assert** $\mathsf{mem} \neq \perp$ for both $\texttt{Node}[c] = \perp$ and $\texttt{Node}[c] \neq \perp$: holds if *apply-props runs fully. The only additional case that results in the unwinding of *process-ec is as follows:

- $\exists \texttt{Node}[c].\mathsf{epoch} = \mathsf{epoch} \land \texttt{id} \in \texttt{Node}[c].\texttt{has_psk} \lor \nexists(\texttt{id}, *) \in G$ states that at least one PSK for each PSK proposal must be in id's possession. This is guaranteed by the **req** in the line before.

Join. The Join returns $\perp$ only in the following case:

- **req** $\texttt{id} \in \texttt{Node}[c_{\text{psk}}].\texttt{has_psk}$: if $\perp$ were invalid, $\mathcal{S}_2$ would have returned $ack = \texttt{false}$. Due to the direct correspondence of has_psk and the PSK being in id's PSK store, this **req** fails exactly if ETK does.

Assert Statements: The proof for the **assert** statements of the Join algorithm remains the same.

Consistency: The proof showing that every joiner id that executes Join with a vector of welcome message vec_w and every group member id′ that executes Process with a commit c' must end up in the same state remains the same.

External Commit. After $\mathcal{S}_2$ constructs the external commit and generates all required parameters and variables and makes the necessary checks as described in ETK proof, it additionally sends the vector $c_{\vec{psk}}$ consisting of the commits of all injected PSKs or $\perp$ if any of the pskId's appear twice.

Indistinguishability of Output:
$\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ returns c and g' generated according to the protocol. The ExtCommit returns $\perp$ additionally in the following cases:

- If **req** $\texttt{ExtCommitProps} \neq \perp$ fails then because $\mathcal{S}_2$ could not supply suitable c_{psk} or there was the same pskId twice. In this case, $\mathsf{ETK}^{\text{PSK}}$ would also fail.
- **req** $\texttt{id} \in \texttt{Node}[c_{\text{psk}}].\texttt{has_psk}$: Due to the direct correspondence of has_psk and the PSK being in id's psk_store, this **req** fails exactly if $\mathsf{ETK}^{\text{PSK}}$ does.

Assert Statements: The ExtCommit algorithm includes changes only for the following **assert** statements:

- **assert** $\mathsf{mem} \neq \perp \land (\texttt{id}, \mathsf{spk}) \in \mathsf{mem}$: is given as *apply-props executed on any $\vec{p}$ contains at least the committer (in this case id) when run successfully. It runs successfully because:
  - the requirements on props of *apply-props are fulfilled through the construction in ExtCommit.
  - **req** $(\texttt{id}_c, *) \in G$ (resync) or $(\texttt{id}_c, *) \notin G$ (external commit): is equally covered by ***req-correctness** in combination with the $\mathsf{ETK}^{\text{PSK}}$ execution.
  - **req** $\texttt{id}_s = \texttt{id}_t = \texttt{id}_c$ and $\texttt{id}_s \in G$ is fulfilled through the construction in ExtCommit and the previous **req**.
  - $\exists \texttt{Node}[c].\mathsf{epoch} = \mathsf{epoch} \land \texttt{id} \in \texttt{Node}[c].\texttt{has_psk} \lor \nexists(\texttt{id}, *) \in G$ states that at least one PSK for each PSK proposal must be in id's possession. This is guaranteed by the **req** in the line before.

- ***consistent-ext-comm**() holds for a similar reason as for proposals and commits. An external commit does not include a membership tag, but it does include the confirmation tag, a MAC over the GroupContext's confirmed transcript hash, which provides the same uniqueness guarantee of a commit c. $\texttt{Node}[c].\mathsf{pro}$ and $\texttt{Node}[c].\mathsf{par}$ will hence only be set to the proposals included in c and the parent of c included in the confTransHash, or else to $\perp$ in case of a detached root. As a consequence of these two properties, $\texttt{Node}[c].\mathsf{mem}$ will always be set to the same membership set handed to ***consistent-ext-comm**.

  As ExtCommitProps is directly constructed from $\vec{p}$, as a consequence, all its parameters are also equivalent, the only exception being c_{psk}, whose injection, however, is encoded in the confTransHash.

F.2 Confidentiality of $\mathsf{ETK}^{\text{PSK}}$

Lemma 7. *If* PKE *is* GSD *secure, then Hybrids H2 and H3 are indistinguishable, that is,* $\mathsf{ETK}^{\text{PSK}}$ *guarantees confidentiality.*

Proof. We show the indistinguishability by hybrid games. More concretely, for any $i > 1$, we define two adjacent sub-hybrids Hs_{i-1} and Hs_i that behave differently only in epoch i: Hs_{i-1} sets the epoch secret in epoch i according to $\mathcal{F}_{\text{ECGKA}}$ (depending on whether **safe**(c) holds), while Hs_i sets is according to the simulator $\mathcal{S}$. Further, we reduce the indistinguishability between every two adjacent hybrid games to the GSD security of the underlying PKE. More concretely, we show that if the environment $\mathcal{Z}$ can distinguish between a random key and a simulator chosen *application* key for c with **safe**$(c) = \texttt{true}$, the reduction $\mathcal{B}$

(i.e., the adversary that aims to break GSD game) can use the environment $\mathcal{Z}$ to win the GSD game with $\textsf{gsd-exp}((c, u)) = \texttt{false}$ for some associated secret u.

The individual confidentiality proof steps of the $\textsf{ETK}^{\text{PSK}}$ are organized as follows:

Proof part 2.b): Confidentiality

Show indistinguishability:

I) when not allowing any injections or corruption of randomness by reasoning through all possible cases of $\textsf{gsd-exp}((c, u_{\text{app}}))$ (Appendix F.2):

  (a) the corruption of (c, u_{app})
  (b) the corruption of both (c, u_{joi}) and (c, u_{psk}) .
  (c) if all three $\textsf{gsd-exp}$ $((c, u_{\text{com}}))$ and $\textsf{gsd-exp}$ $((\text{Node}[c].\text{par}, u_{\text{ini}}))$ and $\textsf{gsd-exp}$ $((c, u_{\text{psk}}))$ are true.

II) when allowing injections for cases (a) - (c) (Appendix F.2).

III) when allowing the corruption of a client's randomness for cases (a) - (c) (Appendix F.2).

Here we mostly note down the `psk` injection effect on the proof as many proof steps identical to Appendix E.2

Part 2.b.I) No Injection or Corrupted Randomness Allowed

Here in case of neither malicious injection nor randomness corruption is allowed, we prove by reduction that $\textsf{gsd-exp}((c, u)) \implies \neg\textbf{safe}(c)$. Reduction $\mathcal{B}$ has the following modifications on the code of the functionalities and the simulator in the context where no injection and no corrupted randomness are permitted.

Following this, we systematically examine each possible case of $\textsf{gsd-exp}((c, u))$, corresponding to the three conditions specified above, and illustrate how our proof captures the requirements of the safe predicate.

As described in the *gsd* experiment explanation; psk injection affects our model, for {`add`, `extAdd`, `psk`}-only commits, no `updatePath` is executed, and all of `id` 's secrets remain the same. This is addressed in ***secrets-replaced**$(c, \texttt{id})$ with color coding. Different than ETK in Appendix E.2, for no `updatePath` is executed case, we will also consider `psk` proposals. To this end, we also consider the commits including `add`, `extAdd`, `psk` proposals only, as well as those that involve any combination of these proposal types.

From draft 12 to RFC: All the changes related to leaf key separation, welcome messages and init key usage are also valid for $\textsf{ETK}^{\text{PSK}}$. Additionally:

- Proposal `psk`: Whenever $\mathcal{Z}$ chooses an `pskIds`, the reduction $\mathcal{B}$ stores `pskIds` in the new proposal node.
- Applying Proposals: Additionally, for each `psk` proposal with list $\texttt{pskIds} = \{\texttt{pskId}_1, \texttt{pskId}_2, ...\}$ (recall that the reduction has already stored `pskIds` in the corresponding node), the reduction $\mathcal{B}$ recursively queries the **Join-Hash** oracle with input $(\textsf{temp}_i, \texttt{pskId}_i, \textsf{temp}_{i+1}, \textsf{lbl} = \textsf{derive-psk})$ for $i \geq 1$, where $\textsf{temp}_i = 0$. Note that the secret included in the node, which is derived from the last **Join-Hash** query, is indeed the `psk-secret` of the `psk` proposal. We denote this node by u_{psk}.
- Commits by reference: Compared with the ETK case in Appendix E.2, the `*rekey-path` operation is *not* executed if the proposal list in a regular commit satisfies the condition that it contains only `add`, `extAdd`, or `psk` proposals.

  On the other hand, if the commit includes any `psk` proposal, the reduction $\mathcal{B}$ emulates the `*derive-keys` as follows:

  1. *Add vertices to the GSD graph:* Additionally the reduction $\mathcal{B}$ creates a $u_{\text{joi-psk}}$ vertex by querying the **Join-Hash** oracle with input $(u_{\text{joi}}, u_{\text{psk}}, u_{\text{joi-psk}}, \textsf{lbl} = \textsf{joi-psk})$. Afterwards, the reduction $\mathcal{B}$ generates vertices u_{app}, u_{mem}, u_{conf}, u_{ini}, u_{ext}, u_{psk} by querying the **Hash** oracle as in [4], but with inputs $(u_{\text{joi-psk}}, u_{\text{lbl}}, \textsf{lbl})$ for $\textsf{lbl} \in \{\textsf{app}, \textsf{mem}, \textsf{conf}, \textsf{ini}, \textsf{ext}, \textsf{psk}\}$.
  2. *Create the packet:* To complete the set of values needed to compute the commit packet, the reduction $\mathcal{B}$ additionally has to corrupt u_{psk}.
  3. *Create the welcome message:* The welcome message is a component of the welcome vector, consisting of three parts. The final part, `wel_type`, indicates whether the message contains a single welcome

message or multiple welcome messages. In the first part, for each new member $\texttt{id}_t$, the encryptions of `joinerSec` and $\texttt{id}_t$'s `pathSec` together with `pskIds` under the init-key in $\texttt{id}_t$'s key package (obtained from KS by the party adding $\texttt{id}_t$). Let u_i be the GSD vertex corresponding to the `pathSec` sent to $\texttt{id}_t$. If $\texttt{id}_t$'s init-key is of the form (gsd, u)[7], $\mathcal{A}$ obtains the encryptions by creating encryption edges from u to u_{joi} and from u to u_i. Otherwise, it corrupts u_i and u_{joi} encrypts the values itself. The second part of the welcome message contains the encryption of `groupInfo`, under the `welcome_secret` which is simply the combination of `joinerSec`, `psk-secret` (derived from `pskIds`) and the vector of welcome messages. To compute the second part of the vector of welcome messages, the reduction $\mathcal{B}$ creates the vertex u_{wel} by querying **Hash** oracle with input $(u_{\text{joi-psk}}, u_{\text{wel}}, \text{wel})$. The reduction $\mathcal{B}$ gets the respective secrets by corrupting corresponding vertices.

- Expose: Additionally, the state of `id` includes `psk-secret`, as specified by ***update-stat-after-exp**.

As the basis for the proof, considering the key derivation that also considers the `psk` key material illustrated in Figure 2, gsd-exp$((c, u_{\text{app}}))$ can only hold true in one of the following cases:

2.b.I.a) The reduction $\mathcal{B}$ corrupts the vertex (c, u_{app}). This case identical to ETK.

2.b.I.b) The reduction $\mathcal{B}$ corrupts the vertices (c, u_{joi}) and (c, u_{psk}) to ultimately gain access to the application secret. This happens if and only if the reduction $\mathcal{B}$ computes a vector of welcome messages for $\texttt{id}_t$ using an exposed key bundle. Upon exposure of the key bundle, the initial key also becomes corrupted along with u_{psk} (since we are considering only the resumption `psk`'s). So the mechanism remains as previously described.

2.b.I.c) This case considers the following three elements to be true: gsd-exp$((c, u_{\text{com}}))$, gsd-exp$((\texttt{Node}[c].\text{par}, u_{\text{ini}}))$, and gsd-exp$((c, u_{\text{psk}}))$. We need to show that gsd-exp$((c, u_{\text{com}}))$ implies know$(c, *)$. This implication follows by condition d) of ***can-traverse**. Here, gsd-exp$((c, u_{\text{com}}))$ is only possible in one of the following cases that are identical to ETK.

1) the reduction $\mathcal{B}$ corrupts path secrets u_i, and uses them to compute gsd-exp$((c, u_{\text{com}}))$
2) the reduction $\mathcal{B}$ calls the **Enc** oracle to encrypt path secrets under relevant keys
3) after the commit is processed, the reduction $\mathcal{B}$ exposes a party `id` that holds u_i's secret

Additionally the reduction $\mathcal{B}$ needs to assure Node$[c]$.`psk` $\neq$ good. This only possible if the reduction $\mathcal{B}$ corrupts some `id` who owns the `psk` (not necessarily $\texttt{id}_t$ that commits c) or exposes `pskId` that generates the `psk-secret` during the commit.

Considering case 3) any actions of `id` requiring `updatePath` removes u_i's secret from its state, thus, the proof remains the same. For {`add`, `extAdd`, `psk`}-only commits, since no `updatePath` is executed, all of `id`'s secrets remain the same. Thus, if `id` is corrupted, know$(c, \texttt{id})$ is true until an `updatePath` issued. We complete the proof by showing that for cases 1) and 2) the combination of gsd-exp$((\texttt{Node}[c].\text{par}, u_{\text{ini}}))$ and gsd-exp$((c, u_{\text{psk}}))$ implies know$(c, \text{'epoch'})$. For these cases the proof follows the same pattern as the encapsulating proof for u_{app}, as both u_{ini} and u_{psk} have the exact same lifetime as its epoch's u_{app}. This is because the resumption `psk`'s with *application* usage are generated from each epoch secret. The proof's condition c) hence recursively proceeds until at some point, $\texttt{Node}[c].\text{pars} = \text{root}_x$. At that point, condition d) of ***can-traverse**(c) in **safe**(c) becomes true.

External Commit. The proof is identical to that of ETK presented in Appendix E.2, with the exception of a minor modification in Case 2.b.I.c). Figure 27 visualizes the GSD game with the effect of a `psk` injection. There is an edge between (c, u_{exini}) and (c, u_{joi}), meaning that c is an external commit. As the `*rekey-path-upon-join` proceeds in the same way as `*rekey-path` in a regular commit, the reasoning for gsd-exp$((c, u_{\text{com}}))$ implying know$(c, *)$ also remains the same. The only change occurs when the combination of gsd-exp$((\texttt{Node}[c].\text{par}, u_{\text{exini}}))$ and gsd-exp$((c, u_{\text{psk}}))$ implies know$(c, '\text{epoch}')$.

Similar to the analysis for regular commits, we consider two cases: i) u_{exini} is compromised together with u_{psk} when a client `id` is exposed while processing the commit, in which case **safe**(c) does not hold, or ii) the reduction $\mathcal{B}$ calculates u_{exini} and u_{psk} itself if the seed s_{exini} and s_{psk} are exposed. Case i) happens if an `id`'s key bundle is exposed during the commit processing, which again implies **safe**(c) is false, while case ii) happens if it is encrypted under a node whose private key is compromised. The seed s_{exini} is encrypted under the key from the u_{ext} node of epoch $\texttt{Node}[c].\text{par}$. A u_{ext} node is exposed in the exact same cases as a u_{app}, and as in a commit, and the the u_{psk} has the same life time as u_{app}, proof continues recursively until c is a root and hence **safe**(c) is false.

[7]Following [4], a secret key of the form (gsd, u) means that it is unknown to the environment $\mathcal{Z}$.

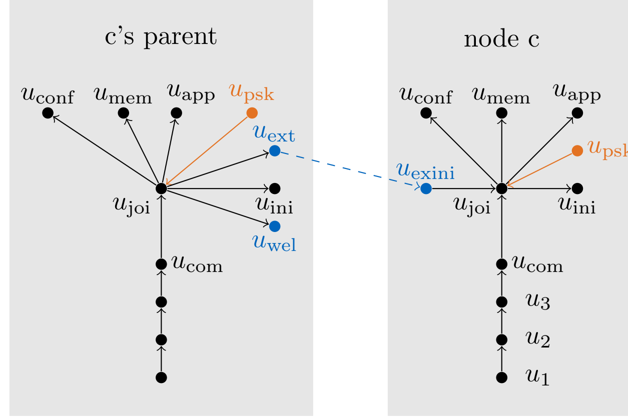

Figure 27: GSD game graph of an external commit, following an image from [4]. Each circle represents a node in the GSD graph. Dashed arrows mark an encryption of the sink using the key of the source, while bold arrows mark a Hash call on the source. The additionally added nodes of the graph are marked in ●. The external secret encrypts the committer-generated seed s_{exini} that is used as the init secret of epoch c.

Part 2.b.II) Allowing Injections In the case of allowing the injection of malicious injections, additional to the capabilities we mentioned in Appendix E.2, including `psk` injection, the functionalities and the simulator incorporate the following adversarial modifications:

From draft 12 to RFC

- Injected proposals: Additionally, in the case of `psk` proposal, the reduction $\mathcal{B}$ stores `pskIds` included in the `psk` proposal in a new proposal node.
- Commits injected to process: In the case that $\mathcal{Z}$ makes `id` process an injected commit c' from a party $\texttt{id}_c \neq \texttt{id}$, reduction $\mathcal{B}$ attempts to build new commit node's state. For {`add`, `extAdd`, `psk` }-only commits, the reduction $\mathcal{B}$ only applies the proposals, as these commits do not require rekey operation. For other commits along with the ciphertexts `ctxt` that `id` would decrypt with keys in the ratchet tree, the reduction $\mathcal{B}$ proceeds as follows:

  i-ii) remain the same.

  iii) if the `ctxt` is copied from an honest commit that was generated earlier by $\mathcal{B}$, for which the GSD node associated to both appSecret is still a valid challenge, we must show that `id` accepting such a commit c would enable $\mathcal{B}$ to succeed in the GSD game. The trick in this step is to reduce the forgery of the confirmation tag `confTag` in the commit c to the GSD security of the underlying PKE. Let $\mathcal{B}$ challenge the GSD node u and extract the correct seed from $\mathcal{Z}$'s random oracle calls as follows:

  1. Note that appSecret = Hash(`joinerSec`, `psk-secret`, .) and `joinerSec` = Hash(initSecret, `commitSec`, .), where HKDF is modeled as RO. Moreover, note also that `commitSec` must be the same in c and c' for the shared `ctxt` to be accepted.
  2. The reduction $\mathcal{B}$ needs to extract the `joinerSec` and `psk-secret` of c'. To this end, the reduction $\mathcal{B}$ proceeds as follows: First, recall that the `confTag` is a MAC of `confKey` and confTransHash. The only way for $\mathcal{Z}$ to compute a valid tag is to query RO on (`confKey`, confTransHash). The reduction $\mathcal{B}$ can search the RO query history and extract `confKey`. Second, recall that `confKey` is derived by applying HKDF to `joinerSec` and `psk-secret`. The environment $\mathcal{Z}$ needs to query RO on (`joinerSec`, `psk-secret`) for deriving `confKey`. The reduction $\mathcal{B}$ can then extract `joinerSec` and `psk-secret` of c' in a similar way.
  3. This case also remains the same.

- Injected vector of welcome messages `vec_w`: When $\mathcal{Z}$ injects a welcome message w = (`encGroupSec`, `encGroupInfo`) in the welcome vector `vec_w`, $\mathcal{B}$ does as follows:

1. Similar to proof of ETK, the reduction $\mathcal{B}$ searches for a key package `kp` such that $\mathsf{SK}[\texttt{id},\texttt{kp}] \neq \bot$ and $\mathsf{ISK}[\texttt{id},\texttt{kp}] \neq \bot$ and $\mathsf{H}(\texttt{kp})$ matches an entry $e \in$ `encGroupSecs` and aborts if no such `kp` exists.

2. If e is copied from a vector of welcome messages `vec_w` generated by $\mathcal{B}$ while creating a commit node c and `ISK[id, kp]` is a GSD node, to win the game, reduction $\mathcal{B}$ proceeds as follows: $\mathcal{B}$ needs to distinguish whether the injected welcome message is associated with the commit c that the e stems from. To this end, $\mathcal{B}$ compares the `encGroupSec`. Since e is copied, the `joinerSec` and `pskIds` underlying different welcome messages that include the same e must be the same. Moreover, `psk-secret` must be the same as is it derived from `pskIds`. Furthermore, an equal `encGroupInfo` implies an equal confTransHash. In this case, $\mathcal{B}$ moves `id` to c.

   If the `encGroupInfo` in the vector of welcome message `vec_w` is different from the one derived by the reduction $\mathcal{B}$ for c, that means that e was copied but used to create the welcome for a non-existent commit c'. This is not possible as it would allow $\mathcal{B}$ to win the GSD game: In order to construct a `encGroupInfo` that is decryptable using the secrets stored in the `encGroupSecs`, the environment $\mathcal{Z}$ must be in possession of the `welcome_secret` (as $\mathcal{Z}$ cannot copy `encGroupInfo` from c). In order to derive the `welcome_secret`, $\mathcal{Z}$ must derive it from the `joinerSec` and `psk-secret` using the RO calls. The reduction $\mathcal{B}$ can extract `joinerSec` and `psk-secret` by checking queries that have been sent to RO one by one as follows: First, $\mathcal{B}$ attempts to parse the input of RO into `joinerSec` and `psk-secret` candidates. If the parse fails, $\mathcal{B}$ directly omits the following steps and attempts for the next RO input. Second, the reduction $\mathcal{B}$ derives `welcome_secret` and `confKey` from the `joinerSec` and `psk-secret` candidates, respectively. Third, the reduction $\mathcal{B}$ derives `welcome_key` and `welcome_nonce` from `welcome_secret`. Fourth, the reduction $\mathcal{B}$ decrypts `encGroupInfo` to recover `groupInfo`. Fifth, the reduction $\mathcal{B}$ parses `groupInfoTBS` from `groupInfo` and further `confTag` from `groupInfoTBS`. Finally, the reduction $\mathcal{B}$ verifies the `confTag` by using the `confKey`. If no error occurs, then reduction $\mathcal{B}$ found the correct `joinerSec` and `psk-secret`. If any error occurs, then the reduction $\mathcal{B}$ continues to check the next RO query. Note that the environment $\mathcal{Z}$ can send only polynomial queries to RO. The checks that $\mathcal{B}$ needs to execute must also be in polynomial time.

   From `joinerSec` and `psk-secret`, the reduction $\mathcal{B}$ can calculate the appSecret for the valid GSD challenge node c and hence win the game. Copying e is hence not possible.

3. Else, if e has not been copied from any vector of welcome messages generated by $\mathcal{B}$, the reduction $\mathcal{B}$ can simply obtain the encrypted `joinerSec`, `pskIds` and `pathSec` either by using the secret in ISK[`id`, `kp`], if it has been compromised and therefore known by $\mathcal{B}$, or otherwise by querying $\mathcal{B}$'s GSD decryption oracle **Dec**. Related to `pskIds`, we consider two cases. First, all `psk` in the list `pskIds` have been exposed. In this case, the reduction $\mathcal{B}$ also knows them and therefore can derive `psk-secret` by itself. By using the `joinerSec` and `psk-secret`, the reduction $\mathcal{B}$ can decrypt and verify the `encGroupInfo` in the vector of welcome messages. If the verification passes, then $\mathcal{B}$ further checks whether the decrypted `groupInfo` indicates an existing node c (by comparing the confTransHash $\in$ `groupInfo`) or a new node c'. If `groupInfo` indicates an existing node c, then $\mathcal{B}$ moves `id` to c. Otherwise, `groupInfo` indicates a new node c'. In this case, $\mathcal{B}$ creates c' with labels taken from `groupInfo` and the ratchet tree set to the public part of τ from `groupInfo`. Then, for any node of τ with a public key for which it has a secret key stored (in another ratchet tree or in ISK and SK), $\mathcal{B}$ copies the secrets into τ. Moreover, $\mathcal{B}$ updates the ratchet tree secrets to those derived from `pathSec` (if any secret key was set to a GSD node, $\mathcal{B}$ uses `pathSec` to win the game), and computes the epoch secrets from `joinerSec`. For all other unknown secrets in τ, $\mathcal{B}$ keeps them as $\perp$.

   Second, there exists some PSK with identifier `pskId` $\in$ `pskIds` that is unexposed. In this case, the environment $\mathcal{Z}$ must create `welcome_secret` by querying RO. Similar to above case for "H(`kp`) $\in$ `encGroupSecs` but `encGroupInfo` is forged", we can prove that this case is impossible otherwise $\mathcal{B}$ can easily win the GSD game.

Injected `groupInfo`. The proof is identical to ETK for injected `groupInfo`.

External Commit Injected to Process.

The reduction $\mathcal{B}$ constructs the new node including the initSecret as follows: On one hand, if the external-key, which encrypts the shared secret s_{exini} in the `extInit` proposal of the external commit, is a GSD node, i.e., neither $\mathcal{B}$ nor $\mathcal{Z}$ has knowledge of it, $\mathcal{B}$ uses the **Dec** oracle of GSD to retrieve the shared secret s_{exini}. On the other hand, if the external-key is a value, meaning that both $\mathcal{B}$ and $\mathcal{Z}$ know it, $\mathcal{B}$ directly decrypts s_{exini}. The external-key cannot be $\perp$, as it is not a key that the adversary can inject. From s_{exini}, $\mathcal{B}$ can then derive the new initSecret. The only case where $\mathcal{B}$ will not be able to decrypt s_{exini} is if the `kem_output` stems from an honest external commit c' previously generated by $\mathcal{B}$ and copied by $\mathcal{Z}$ and whose appSecret is still a valid challenge. If an `id` accepts such a copied `kem_output`, $\mathcal{B}$ can use it to win the GSD-game similarly to a copied path secrets in commits:

1. Since the kem_output is copied, the ext-initSec is identical in both the injected external commit c and the honest external commit c'. At the same time, confTransHash must differ; otherwise, c and c' would be identical.
2. For $\mathcal{Z}$ to compute a valid confTag for the injected commit c, it must query the RO on H(confTransHash, confKey). From this query, $\mathcal{B}$ can extract the confKey. Since confKey is derived from (joinerSec, psk-secret), $\mathcal{Z}$ must also query RO on (joinerSec, psk-secret) to compute a valid tag. This allows the reduction $\mathcal{B}$ to extract psk-secret.
3. For $\mathcal{Z}$ to compute the joinerSec for the injected commit c, it needs to query RO on Hash(ext-initSec, commitSec). From this query, $\mathcal{B}$ can extract the ext-initSec, which is shared with c', by searching the RO query history.
4. $\mathcal{B}$ can then corrupt the commitSec of c', which does not impact the challenge, and compute the joinerSec of c'. Using this joinerSec and psk-secret, the reduction $\mathcal{B}$ calculates the appSecret.

Thus, this case cannot occur.

The reduction $\mathcal{B}$ handles the updatePath of the commit the same way as a regular one, as the argumentation of an injected commit relies on calculating the initSecret and psk-secret of c from the RO calls and corrupting the initSecret and psk-secret of c' directly - regardless of its origin. Note that in the case of an injected external commit, both the initSecret and psk-secret are always retrieved by $\mathcal{B}$ directly via decryption. This does not affect the argument showing it is impossible to copy updatePath secrets from an honest commit.

The final case to consider is when both s_{exini} and ctxt are copied. If they are copied from different commits, the confTag can be used to break the challenge for both commits in the same way as before. Figure 28 illustrates how $\mathcal{B}$ utilizes $\mathcal{Z}$ to win the GSD game in case of a copied ctxt.

If they are copied from the same commit but the confTransHash differs from the original commit (e.g., when the external commit is injected for a different party), the joinerSec and psk-secret, retrievable via the confKey RO call, are the same between the two commits. The reason behind is that the same ciphertext indicates same path secret and proposals, which further indicate the same commitSec and psk-secret. Combing the fact that s_{exini} is copied from the same commit, which yields the same initSecret, the joinerSec must also be same. Thus, this allows the reduction $\mathcal{B}$ to compute the appSecret to be computed from it.

Proof of Indistinguishability. We show that gsd-exp() $\implies \neg$**safe**(c). As a reminder, this is part **2.b.II** of the main proof of $\mathsf{ETK}^{\mathrm{PSK}}$. The subproof is structured as follows:

Part 2.b.II): Confidentiality when allowing injections

The reduction $\mathcal{B}$ succeeds in winning the GSD game for commit c' if the environment $\mathcal{Z}$ successfully copies a ciphertext from commit c' into an injected commit c. This can be formulated as follows:

a) Show that if the reduction $\mathcal{B}$ corrupts (c, u_{app}) then $\neg$**safe**(c)

b) Show that if the reduction $\mathcal{B}$ corrupts (c, u_{joi}) and (c, u_{psk}) then $\neg$**safe**(c)

c) Show that if both gsd-exp$((c, u_{\text{com}}))$ and gsd-exp((Node[c].par, u_{exini})) and gsd-exp$((c, u_{\text{psk}}))$ are true then $\neg$**safe**(c). For this, consider all cases where gsd-exp$((c, u_i))$ is true:

- c.1) During an exposure of an id that stores u_i
- c.2) When a node $\tau.v$ used in *rekey-path is exposed
- c.3) If the secret in $\tau.v$ is $\perp$, which is the case if:
  - i) $\mathcal{Z}$ injects a new commit c' on behalf of id
  - ii) the injected commit c' commits $\tau.v$ as an updated injected leaf
  - iii) the injected c' commits $\tau.v$ as an injected key package

In the case that c is a regular commit, the proof is extended the same way as in step 2 of the original proof in [4], except for two differences at b) and c).

In case b), we still need to prove that the corruption of (c, u_{joi}) and (c, u_{psk}) implies $\neg$**safe**(c), where now the node (c, u_{psk}) might be created by the psk-proposal with pskIds injected by the environment $\mathcal{Z}$.

Below, we prove that allowing the environment $\mathcal{Z}$ to inject PSK proposals with pskIds does not affect $\mathcal{Z}$'s advantages by considering the following two additional sub-cases: whether the environment $\mathcal{Z}$ has exposed all of the PSK values of pskId $\in$ pskIds or not, i.e., Node$[c]$.psk = good or bad.

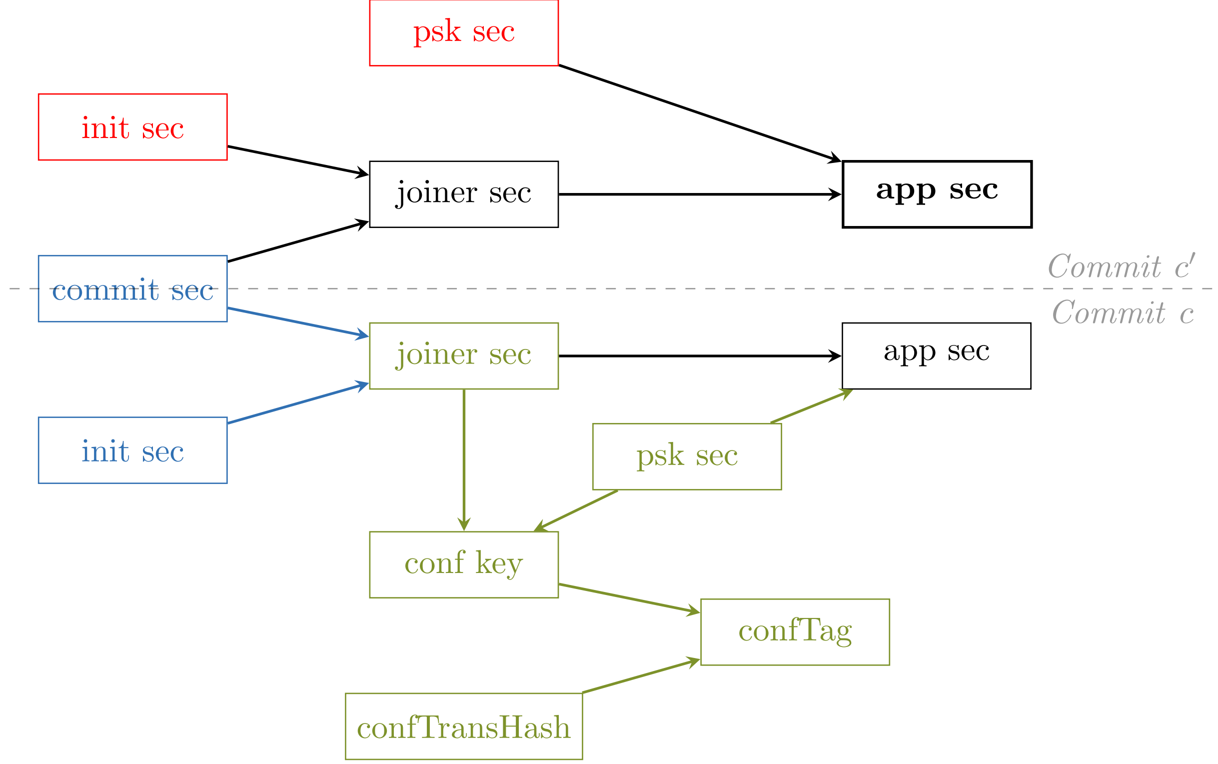

Figure 28: Visualization of $\mathcal{B}$ using the `confTag` to win the GSD game for a commit c' if the environment $\mathcal{Z}$ successfully copies a ciphertext from a commit c' into an injected commit c. The commit secret is shared between both commits. In the first step (in ●), $\mathcal{Z}$ needs to call the random oracle RO to derive the `confTag` from the `confKey` and the confTransHash, which $\mathcal{B}$ now learns from the query. $\mathcal{Z}$ also needs to query the RO to derive `confKey` from `joinerSec` and `psk-secret`. $\mathcal{B}$ now knows `joinerSec` and `psk-secret` of c. In order to derive `joinerSec` for commit$c$, $\mathcal{Z}$ also needs to query Hash(`commitSec`, initSecret), which $\mathcal{B}$ now learns (in ●). $\mathcal{B}$ now knows the `commitSec` of commit c', and corrupts the initSecret and `psk-secret` of commit c' (in ●). $\mathcal{B}$ can now calculate appSecret of c'.

Recall that PSK injection allows the environment $\mathcal{Z}$ to inject the `pskId`, but not the PSK values themselves. The environment cannot manipulate the exact PSK values and further the yielding `psk-secret`. In the first sub-case, if the environment $\mathcal{Z}$ has exposed all PSK values with identifiers `pskId` $\in$ `pskIds`, i.e., `Node`$[c]$`.psk` $=$ `bad`, then the environment $\mathcal{Z}$ can compute `psk-secret`. At the same time, we have gsd-exp$((u, u_{\text{psk}}))$ $=$ `true`. However, in order to derive the correct epoch-secret `epochSec` $=$ HKDF(`joinerSec`, `psk-secret`), the environment $\mathcal{Z}$ still needs to query RO with correct `joinerSec`. Except for correctly guessing `joinerSec` with negligible probability, the only way that $\mathcal{Z}$ can obtain `joinerSec` is to corrupt (c, u_{joi}). Otherwise, in the second sub-case, the environment $\mathcal{Z}$ does not expose all PSK values of the corresponding `pskId` $\in$ `pskIds`, i.e., `Node`$[c]$`.psk` $=$ `good`. In this case, we have gsd-exp$((c, u_{\text{psk}}))$ $=$ `false`. In order to derive the correct epoch-secret `epochSec` $=$ HKDF(`joinerSec`, `psk-secret`), the environment $\mathcal{Z}$ has to query RO with the correct `psk-secret`, which is generated by the recursive invocation of HKDF$(., \texttt{psk}_{\texttt{pskId}})$ for all `pskIds` $\in$ `pskIds`, where $\texttt{psk}_{\texttt{pskId}}$ is the PSK value of corresponding `pskIds`. Thus, the environment $\mathcal{Z}$ has to query RO with correct $\texttt{psk}_{\texttt{pskId}}$ for all `pskId` $\in$ `pskIds`, in particular, for the `pskId` that $\mathcal{Z}$ has not exposed. Thus, this case also happens with negligible probability.

Changes in case c), on the other hand, are identical to the proof for ETK, except for minor modifications in c.1) and c.2). Commits *that include an* `updatePath`, which encompass any form of regular commit but exclude `add`, `extAdd`-only commits, are considered as `add`, `extAdd`, `psk`.

In the case that c is an external commit - if there is an edge between $(c, u_{\texttt{exini}})$ and (c, u_{joi}) - the proof proceeds as follows: gsd-exp(c, u_{app}) is only true in three cases:

a) $\mathcal{B}$ corrupts (c, u_{app}) - same as in step 1

b) $\mathcal{B}$ corrupts (c, u_{joi}) and (c, u_{psk}) - same as in step 1

c) if all three gsd-exp$((c, u_{\text{com}}))$, gsd-exp$((\texttt{Node[c].par}, u_{\text{exini}}))$ and gsd-exp$((c, u_{\text{psk}}))$ are true:

- For $\mathsf{gsd\text{-}exp}((c, u_{\text{com}}))$ the argumentation remains the same as for regular commit. Note that an injected external commit c' (case i)) hence implies $\mathsf{know}(c)$.
- In addition, $\mathsf{gsd\text{-}exp}(((\mathtt{Node}[c].\mathsf{par}, u_{\text{exini}})))$ can also be true if injected: When the environment $\mathcal{Z}$ injecting the seed s_{exini} in the `extInit` proposal, the reduction $\mathcal{B}$ decrypts it and calculates u_{exini} itself. In the case of an injected external commit, condition e) of ***can-traverse** immediately holds, and hence **safe**(c) is false.

In detached trees, the proof remains the same as ETK.

Part 2.b.III) Allowing Bad Randomness

Proposals by Value. In case of a `psk` proposal, the reduction $\mathcal{B}$ computes the proposal message p by using randomness provided by $\mathcal{Z}$ (as 'nonce' value), the current `memberSec`, and `id`'s spk (all of which are known by $\mathcal{B}$), and stores the `pskId` in `pskIds`.

Commit by Reference and by Value. Given randomness provided by $\mathcal{Z}$, $\mathcal{B}$ computes the commit and required secrets as:

1. The reduction $\mathcal{B}$ executes `*rekey-path` to obtain the path secrets, `commitSec`, and commitContent. If a `psk` is required, it computes `pskId` using the provided randomness. It then generates a valid confTransHash.
2. $\mathcal{B}$ creates a new `joinerSec`, which is a hash of the current initSecret and the newly computed `commitSec`, i.e., `joinerSec` = H (initSecret, `commitSec`, ...). If initSecret stores a GSD node u, $\mathcal{A}$ queries the $\mathsf{H}(u, u_{\text{ctr}}, \mathtt{commitSec})$, increments u_{ctr}, and uses the fresh, uncorrupted node as the `joinerSec`. Otherwise, if initSecret stores a value, $\mathcal{A}$ computes the `joinerSec` itself.
3. Additionally, $\mathcal{B}$ is required to compute `psk-secret`. Since $\mathcal{B}$ knows the full `pskIds` list, $\mathcal{B}$ can compute `psk-secret` same as in Step 1. Using the secrets obtained, $\mathcal{B}$ executes the key schedule to compute the `confTag`. It then computes the commit message c and the vector of welcome messages `*welcome-msg` accordingly.

Validity of the Challenge and Randomness in `psk`*.* The reduction $\mathcal{B}$'s actions and the proof remain the same for commits that involve fresh randomness. Recall that the helper predicate ***secrets-replaced** captures the update of the path secret, except for the proposal-list in regular commit c is *non-empty* or only includes regular add, external self-add, or PSK proposals. Any commit c'' (except only regular add, external self-add, or PSK proposals) with ***secrets-replaced**$(c'', \mathtt{id})$ would hence replace secrets injected by randomness.

Now we will show with above changes, randomness corruption does not affect existing challenges, where **safe**(c) holds. With corrupted randomness the environment $\mathcal{Z}$ cannot manipulate the exact values of any psk, and therefore `psk-secret`. Instead, $\mathcal{Z}$ can inject `pskId` by using corrupted "nonce" value, i.e, $\mathtt{Node}[c].\mathtt{psk} = \mathsf{bad}$.

This case is same with part 2 of the proof: although $\mathcal{Z}$ can compute `psk-secret`, $\mathcal{Z}$ cannot derive the correct epochSec unless $\mathsf{gsd\text{-}exp}((u, u_{\text{psk}})) = \mathtt{true}$, because the environment $\mathcal{Z}$ needs to query RO on `joinerSec` and `psk-secret` to compute $\mathtt{epochSec} = \mathsf{HKDF}(\mathtt{joinerSec}, \mathtt{psk\text{-}secret})$. Hence, reduction $\mathcal{B}$ cannot get additional advantage by `psk-secret`.

We adapt the proof that randomness corruption does not affect the validity of a challenge, and $\mathsf{gsd\text{-}exp}(c) \implies \neg\mathbf{safe}(c)$ still holds. Due to above mentioned changes, no additional corruption is performed by $\mathcal{B}$ on the `joinerSec`, case b) of gsd-exp holds as before.

Only case that is affected is case c). Here, the proof remains in agreement with [4], showing that $\mathsf{gsd\text{-}exp}((c, u_{\text{com}})) \implies \neg\mathsf{know}(c, *)$ also when allowing corruptions. This is the same bug we mentioned in ETK proof.

Randomness in Detached Trees. The proof remains identical to the ETK proof, allowing randomness in detached trees.

External Commit from `id`*.* Similar to regular commit, using the randomness provided by $\mathcal{Z}$, reduction $\mathcal{B}$ computes the external commit and the secrets in new commit node, as follows:

1. $\mathcal{B}$ uses the randomness provided by $\mathcal{Z}$ to execute `*rekey-path-upon-join`, obtaining all path secrets, the `commitSec`, and the intermediate commit packet. It computes `pskId` using the same randomness. Then, it signs the commit packet using the `id` 's spk, along with $\mathcal{Z}$'s randomness, and sets the confTransHash accordingly.

2. A key distinction from a regular commit is that the initSecret for an external commit is determined by the client. As a result, corrupted randomness can expose this initSecret. Reduction $\mathcal{B}$ computes the ext-initSec as follows:
   - If the external secret from the previous epoch stores a GSD node u, $\mathcal{B}$ queries $\mathbf{Enc}(u, u_{\text{ctr}})$, corrupts u_{ctr}, and sets the external seed s_{exini}, increments u_{ctr}. $\mathcal{B}$ then uses this seed as the initSecret for the external commit.
   - If the external secret from the previous epoch stores a value (and is therefore known to $\mathcal{B}$), $\mathcal{B}$ computes the `joinerSec` directly.
3. Using the `joinerSec`, `psk-secret`, and the confTransHash from previous steps, $\mathcal{B}$ runs the key schedule, computes the `confTag`, and finalizes the commit message c.

Proof for External Commit. Case a) of the proof remains the same. Case b): The reduction $\mathcal{B}$ now also corrupts (c, u_{joi}) and (c, u_{psk}) when executing an external commit with compromised randomness. Any node that is injected with bad randomness has $\texttt{Node}[c].\mathsf{st} = \texttt{bad}$ and $\texttt{Node}[c].\mathsf{psk} = \texttt{bad}$. Additionally, external commit nodes are marked via the sender type `new_member`. An external commit with bad randomness hence fulfills condition e) of ***can-traverse**(c), and hence **safe**$(c) = \texttt{false}$.

Case c): To show $\mathsf{gsd\text{-}exp}((c, u_{\text{com}})) \implies \neg\mathsf{know}(c, *)$ for external commits, we proceed similar as for regular commits. $\mathsf{gsd\text{-}exp}((c, u_i))$ can also be true if the secret in a ratchet tree node $\tau.v$ used in `*rekey-path-upon-join` stores a seed s generated during an action executed with bad randomness. Consider an external commit c' that inserts s into $\tau.v$. Unlike a regular commit, an external commit injects only the secrets it generates during the commit process. It does not, for instance, include an update proposal from another `id` or a corrupted key package. This injection is covered by condition a) of ***secrets-injected**$(c, \texttt{id})$. Now consider the case where c itself is the external commit. Since the external committer utilizes the same public ratchet tree as a member performing a regular commit, the proof remains the same to regular commit.

□

F.3 Authenticity

We follow the same intuition with ETK to prove authenticity; specifically, we aim to demonstrate that *H3* and *H4* are indistinguishable by showing that **auth-invariant** is never triggered.

Lemma 8. *If* Sig *is* SUF-CMA *secure,* MAC *is* EUF-CMA *secure, and* PKE *is* GSD *secure, then Hybrids H3 and H4 are indistinguishable, that is,* $\mathsf{ETK}^{\text{PSK}}$ *guarantees authenticity.*

Proof. Similarly, the hybrids *H3* and *H4* are identical unless the event `Bad` occurs. We further decompose the event `Bad` into two sub-events: $\texttt{Bad}_{\texttt{sig}}$ and $\texttt{Bad}_{\texttt{MAC}}$. The proof is concluded by combining Lemma 9 and Lemma 5, where we respectively prove that the events $\texttt{Bad}_{\texttt{sig}}$ and $\texttt{Bad}_{\texttt{MAC}}$ will never happen except for negligible probability. For $\mathsf{ETK}^{\text{PSK}}$, the case where $\texttt{Bad}_{\texttt{MAC}}$ is triggered remains identical to that in ETK. Therefore, we focus here on examining the scenario where the event $\texttt{Bad}_{\texttt{sig}}$ occurs.

□

Lemma 9. *For any environment* $\mathcal{Z}$*, there exists a reduction* $\mathcal{B}_{\texttt{sig}}$ *that succeeds in the* SUF-CMA *game with a probability that is only polynomially smaller than the probability of* $\mathcal{Z}$ *triggering* $\texttt{Bad}_{\texttt{sig}}$.

Proof. For any environment $\mathcal{Z}$, $\mathcal{B}_{\texttt{sig}}$ emulates the functionalities and simulator by embedding its challenge spk as one of the public keys honestly created during the experiment. To emulate commits signed under ssk, $\mathcal{B}_{\texttt{sig}}$ calls **Sign** oracle. When $\texttt{Bad}_{\texttt{sig}}$ occurs, reduction $\mathcal{B}_{\texttt{sig}}$ stops the experiment and sends to its challenger to forgery consisting of the sig', $\texttt{tbs}'$ from the injected node c'.

First assume c' is a commit node. We show that if $\texttt{Bad}_{\texttt{sig}}$ occurs and $\mathsf{spk} = \texttt{Node}[c].\mathsf{mem}[\texttt{id}]$, then $\mathcal{B}_{\texttt{sig}}$ wins.

1. We know sig' is a valid signature over $\texttt{tbs}'$. The injected node was created when some party accepted c' which means that it verified sig' under $\mathsf{spk} = \texttt{Node}[c].\mathsf{mem}$.
2. Simulation differs from the experiment when $\mathsf{spk} = \texttt{Node}[c].\mathsf{mem}$ is corrupted, however this does not happen as the event guarantees $\mathsf{spk} = \texttt{Node}[c].\mathsf{mem} \notin \mathsf{Exposed}$.

3. The reduction $\mathcal{B}_{\texttt{sig}}$ has never queried **Sign** on $\texttt{tbs}'$ for sig' to its signing oracle. Assuming $(\mathsf{sig}', \texttt{tbs}')$ is the same as $(\mathsf{sig}^*, \texttt{tbs}^*)$ queried by $\mathcal{B}_{\texttt{sig}}$ to the sign oracle for c^*, where c' is injected and c^* is honestly generated by $\mathcal{B}_{\texttt{sig}}$. By showing our assumption implies $c' = c^*$, we will achieve a contradiction. Now we will divide the proof for two, according to the type of the commit c'.
    - We first consider that c' is a regular commit. c' contains $(\mathsf{groupId}', \mathsf{epoch}', \mathsf{leafIdx}', '\texttt{Commit}', C', \texttt{confTag}', \mathsf{sig}', \texttt{membTag}')$ and $\texttt{tbs}' = (\mathsf{groupCtxt}', \mathsf{groupId}', \mathsf{epoch}', \texttt{senderIdx}', '\texttt{Commit}', C')$ and c^* contains the analogous values. Then c' and c^* can only differ on $\texttt{confTag}' \neq \texttt{confTag}^*$ or $\texttt{membTag}' \neq \texttt{membTag}^*$. Note that $\texttt{membTag}' = \mathsf{MAC.tag}(\mathsf{membKey}', C')$ and that $\texttt{membTag}^* = \mathsf{MAC.tag}(\mathsf{membKey}^*, C^*)$. According to our assumption $\texttt{tbs}' = \texttt{tbs}^*$ implies $C' = C^*$ and $\mathsf{groupCtxt}' = \mathsf{groupCtxt}^*$. As group context and $\texttt{psk}$ (if there exists) determines the epoch and key schedule. We have $\mathsf{membKey}' = \mathsf{membKey}^*$ and $\texttt{psk-secret}' = \texttt{psk-secret}^*$, and further $\texttt{confKey}' = \texttt{confKey}*$. Accordingly $\texttt{membTag}' = \texttt{membTag}^*$ and $\texttt{confTag}' = \texttt{confTag}^*$, which gives a contradiction.
    - Then we consider that c' is an external commit. c' contains two components: $\texttt{framedContent} = (\mathsf{groupId}', \mathsf{epoch}', '\mathsf{new_member_commit}', '\mathsf{commit}', C', \mathsf{sig}')$ and $\texttt{framedContentAuth}' = (\mathsf{sig}', \texttt{confTag}')$ and c^* contains the analogous values. Moreover, $\texttt{tbs}' = (\mathsf{groupCtxt}', \mathsf{groupId}', \mathsf{epoch}', \texttt{senderIdx}', '\texttt{Commit}', C')$ and $\texttt{tbs}^*$ contains the analogous values. Similar to regular commit, c' and c^* can only differ on $\texttt{confTag}' \neq \texttt{confTag}^*$. By definition it holds that $\texttt{confTag}' = \mathsf{Hash}(\mathsf{confTransHash}', \texttt{confKey}')$ and $\texttt{confTag}^* = \mathsf{Hash}(\mathsf{confTransHash}^*, \texttt{confKey}^*)$. Since $\texttt{confTag}$ is derived on the framed content, we have $\texttt{confTag}' = \texttt{confTag}^*$. However this gives us contradiction as c' is injected and c^* is honestly generated.
4. Now consider all proposals nodes except $\texttt{extAdd}$-only commits, p' contains $(\mathsf{groupId}', \mathsf{epoch}', \mathsf{leafIdx}', '\texttt{Propose}', P', \mathsf{sig}', \texttt{membTag}')$ and $\texttt{tbs}' = (\mathsf{groupCtxt}', \mathsf{groupId}', \mathsf{epoch}', \mathsf{leafIdx}()', '\texttt{Propose}', P')$ and c^* contains the analogous values. Note that proposal have $\texttt{membTag}$ but not $\texttt{confTag}$. The proof for this case is analogous to the proof for commits.
5. For the $\texttt{extAdd}$ proposal on the other hand, p' contains $(\texttt{tbs}', \mathsf{sig}')$, where $\texttt{tbs}' = (\mathsf{groupId}', \mathsf{epoch}', \texttt{new_member}', '\texttt{Propose}', P')$ and c^* contains the analogous values. Note that here as there is no $\texttt{membTag}$ or $\texttt{confTag}$, commits c' and c^* that includes only $\texttt{extAdd}$ proposals can only differ on $\texttt{psk-secret}' \neq \texttt{psk-secret}^*$. According to our assumption $\texttt{tbs}' = \texttt{tbs}^*$ implies $C' = C^*$ and $\mathsf{groupCtxt}' = \mathsf{groupCtxt}^*$. As group context and $\texttt{psk}$ determines the epoch and key schedule. So $\texttt{psk-secret}' = \texttt{psk-secret}^*$ must hold. This results in a contradiction.

□

Contents

1 Introduction 1

2 Background 2
- 2.1 The Evolution of the MLS Protocol . . . 2
- 2.2 Formal Security Analysis of MLS Draft 12 . . . 4
- 2.3 Notation . . . 6

3 Extended Continuous Group Key Agreement: Syntax 7

4 Security Model and Functionality 8
- 4.1 History Graph . . . 8
- 4.2 Functionality $\mathcal{F}_{\text{ECGKA}}$ . . . 9
- 4.3 Comparison between $\mathcal{F}_{\text{CGKA}}$ [4] and our $\mathcal{F}_{\text{ECGKA}}$ . . . 16

5 RFC 9420 and ETK: Extended Operations TreeKEM 16
- 5.1 The ETK Protocol . . . 16
- 5.2 Security Results for ETK and $\mathcal{F}_{\text{ECGKA}}$ . . . 17

6 Stronger Security for Group Re-Synchronization via External Commit 19
- 6.1 Resyncs and Resumption Pre-Shared Keys . . . 19
- 6.2 The Stronger Functionality $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ . . . 19
- 6.3 The $\mathsf{ETK}^{\text{PSK}}$ Protocol . . . 21
- 6.4 Security Results for $\mathsf{ETK}^{\text{PSK}}$ and $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ . . . 21

A Preliminaries 24
- A.1 Public Key Encryption . . . 24
- A.2 Digital Signature . . . 25
- A.3 Message Authentication Code . . . 26
- A.4 Random Oracle . . . 26

B Simplifications 27

C Details on $\mathcal{F}_{\text{ECGKA}}$ and $\mathcal{F}_{\text{ECGKA}^{\text{PSK}}}$ Functionalities 29

D Details on ETK and $\mathsf{ETK}^{\text{PSK}}$ Protocols 35

E Proof of Theorem 1 44
- E.1 Consistency and Correctness . . . 44
- E.2 Confidentiality . . . 50
- E.3 Authenticity . . . 59

F Proof of Theorem 2 61
- F.1 Consistency and Correctness of $\mathsf{ETK}^{\text{PSK}}$ . . . 61
- F.2 Confidentiality of $\mathsf{ETK}^{\text{PSK}}$ . . . 63
- F.3 Authenticity . . . 71