

Continuous Group-Key Agreement: Concurrent Updates without Pruning

Benedikt Auerbach*¹, Miguel Cueto Noval², Boran Erol^{†3}, and Krzysztof Pietrzak²

¹PQShield

benedikt.auerbach@pqshield.com

²ISTA, Klosterneuburg, Austria

{mcuetono, pietrzak}@ista.ac.at

³UCLA, Los Angeles, USA

boranerol03@gmail.com

June 3, 2025

Abstract

Continuous Group Key Agreement (CGKA) is the primitive underlying secure group messaging. It allows a large group of N users to maintain a shared secret key that is frequently rotated by the group members in order to achieve forward secrecy and post compromise security. The group messaging scheme Messaging Layer Security (MLS) standardized by the IETF makes use of a CGKA called TreeKEM which arranges the N group members in a binary tree. Here, each node is associated with a public-key, each user is assigned one of the leaves, and a user knows the corresponding secret keys from their leaf to the root. To update the key material known to them, a user must just replace keys at $\log(N)$ nodes, which requires them to create and upload $\log(N)$ ciphertexts. Such updates must be processed sequentially by all users, which for large groups is impractical. To allow for concurrent updates, TreeKEM uses the “propose and commit” paradigm, where multiple users can concurrently propose to update (by just sampling a fresh leaf key), and a single user can then commit to all proposals at once.

Unfortunately, this process destroys the binary tree structure as the tree gets pruned and some nodes must be “blanked” at the cost of increasing the in-degree of others, which makes the commit operation, as well as, future commits more costly. In the worst case, the update cost (in terms of uploaded ciphertexts) per user can grow from $\log(N)$ to $\Omega(N)$.

In this work we provide two main contributions. First, we show that MLS’ communication complexity is bad not only in the worst case but also if the proposers and committers are chosen at random: even if there’s just one update proposal for every commit the expected cost is already over \sqrt{N} , and it approaches N as this ratio changes towards more proposals.

Our second contribution is a new variant of propose and commit for TreeKEM which for moderate amounts of update proposals per commit provably achieves an update cost of $\Theta(\log(N))$ assuming the proposers and committers are chosen at random.

*Benedikt Auerbach conducted part of this work at ISTA.

†Boran Erol conducted part of this work at ISTA.

Contents

1	Introduction	3
1.1	Our Contributions	4
2	Preliminaries	7
2.1	Continuous Group-Key Agreement	7
2.2	Ratchet Trees	10
2.3	Auxiliary Results	11
3	The Messaging Layer Security Protocol	11
4	The Communication Cost of MLS for Random Sequences of Operations	14
4.1	Considered Scenario	14
4.2	Expected Number of Blanks	15
4.3	Lower Bound on Sent Ciphertexts	17
5	MLS-Cutoff: an Alternative Method for Update Proposals	22
5.1	Protocol Description	23
5.2	Security	23
5.3	Upper Bounds on the Update Cost	25
6	Acknowledgment	27
A	Omitted Formal Security Definition	31
A.1	Public Key Encryption	31
A.2	Digital Signatures	31
A.3	PKI and CGKA	32
B	Omitted Formal Protocol Descriptions	32
B.1	Formal description of MLS	32
B.2	Formal description of MLS-Cutoff	35

1 Introduction

Messaging layer security. In recent years asynchronous secure messaging applications like the WhatsApp, Signal and iMessage have become an important part of every day life with user numbers ranging in the billions. While deployment of these applications has been a great success for practical cryptography, the used protocols do not scale well when used to secure the communication of large groups. For example, Signal limits the group size to 1000 users. Providing efficient secure messaging for large group sizes is the explicit goal of the IETF’s *Messaging Layer Security* (MLS) Standard [BBR⁺23] of 2023. The standard was developed in a joined effort of industry players and the scientific community and aims to enable asynchronous and end-to-end encrypted messaging via an untrusted server, admitting group sizes up to 50.000.

TreeKEM [BBR18], the protocol at MLS core is a so called *continuous group-key agreement* [ACDT20] (CGKA). As a generalization of continuous key agreement [ACD19] to the group setting, it allows a group of users to agree on a shared, evolving secret that intuitively will be used to secure the communication within the group. CGKA is supposed to provide both forward secrecy (FS), meaning that compromising a user does not reveal past messages, and post-compromise security (PCS), i.e., that the group is able to recover from compromise by sampling new key material and exchanging protocol messages. MLS has been taken up with great interest by the community resulting in a vast amount of works analyzing its security [KPPW⁺21, ACJM20, CHK21, BCK21, AJM22, WPBB23, CGWZ25] or proving lower bounds on its communication complexity [BDR20, BDG⁺22, ANPPP23, AAB⁺24]. Further works develop protocol variants making use of an active server [HKP⁺21, AHKM22, AAN⁺22] (no longer simply forwarding protocol messages), being tailored to multiple overlapping groups [AAB⁺21], or enabling resilience against forking of users’ protocol views [AMT23]. Yet another line of works is concerned with the security of alternative, not TreeKEM based, group-messaging protocols [WKHB21, BCG23, CEST24] or analyzes additional properties of CGKA like secure administration [BCV23] or meta data protection [HKP22].

Ratchet trees. While FS can be achieved essentially by means of symmetric cryptography, PCS requires the use of public key cryptography. TreeKEM makes use of so called ratchet trees, that can be seen as a public-key analog of logical key hierarchies [WHA99]. At a high level, a ratchet tree is a balanced binary tree in which each node has an associated public-key encryption (PKE) key-pair. Each group member is assigned to a leaf, the corresponding key-pair can be thought of as the user’s personal key, and the one of the root as the shared group secret. To enable PCS the users will periodically rotate the secret key material known to them by exchanging protocol messages. While all users have a complete view on all public keys, throughout the execution of the protocol TreeKEM enforces the so-called *tree invariant* stating that every user has access to exactly the secret keys of nodes lying on their update path, i.e., the path from their leaf to the root.

To update their key material a user u replaces the keys of nodes on their update path by sampling a fresh key for their leaf, and deriving the remaining keys in a deterministic way from this key. Then every fresh secret key sk is encrypted under the public key of the node’s co-child, i.e., the child not being part of u ’s update path. Now all users whose update path merges into u ’s at the node in question are able to recover sk by decrypting the ciphertext. They can then derive the remaining secret keys up to the root in a deterministic way. Note that in total u has to generate only logarithmically many ciphertexts (in the group size N), making the protocol particularly suited for large groups.

Proposals and Commits. TreeKEM, as described above, requires all updates to be executed in order (i.e., a user can only issue an update after it processed all previous updates), which might not be possible in a large group with frequent updates. To allow for concurrent issuing of protocol messages TreeKEM distinguishes between *proposal messages* and *commit messages*. In the former, a user issuing an update proposal indicates their goal of updating the key material known to them. However, the message only contains a new personal leaf public key and, in particular, does not establish a new shared group secret. Similarly, proposals of adding users to or removing users from the group are handled using proposal messages.

To establish a new group secret and move the group’s state to the next epoch a user can issue a commit

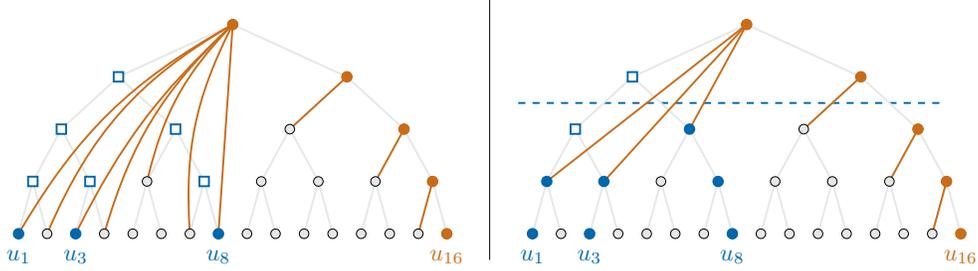


Figure 1: Working principle of update proposals and commits in MLS (left) and MLS-Cutoff (right). In both cases for a fully populated tree users u_1, u_3, u_8 issue an update proposal (depicted in blue) followed by a commit by user u_{16} (depicted in orange). Nodes being blanked are depicted as (empty) squares, and ciphertexts created as part of the commit with orange edges. In the depiction of MLS-Cutoff the blue dashed line marks the cutoff bound.

to a (potentially empty) set of proposal messages. Essentially, this implements all proposals and updates all keys on the committing user u 's update path. Note that to preserve the tree-invariant u cannot refresh the keys on a proposing users' update paths unless they also lie on their own update path (as u must not learn the secret key of such nodes). Moreover, keys on u 's co-path might no longer be considered secure, since the corresponding secret key might lie on the update path of a user issuing a proposal. Thus, TreeKEM handles update proposal by means of *blanking*. I.e., the key-pairs of all (non-leaf) nodes on the update path of proposing users which do not lie on the update path of the committing user are deleted from the ratchet tree. This has the effect of increasing the indegree of some nodes, and thus the commit operation as well as future commits require additional ciphertexts, negatively affecting the protocol's communication complexity. Removals of users from the group is also handled by means of blanking, while adding users employs the unmerged leaves technique. For a depiction of proposals and commits in MLS see Figure 1 (left).

Known results on the communication complexity of MLS. As discussed above, for a group of size N in the best case (in particular if no update proposals are issued at all) the size of a commit message is of order $O(\log(N))$. It is easy to see that in the worst case essentially the whole ratchet tree can end up blank, leading to a linear worst case communication complexity of order $\Omega(N)$. Unfortunately, the latter turns out to be inherent for CGKA schemes built from 'natural' primitives achieving fast PCS [BDR20, BDG⁺22, ANPPP23] (i.e., within a single propose-commit round). Accordingly, while different methods of implementing update proposals have been explored, they either exhibit the same worst-case behavior [KPPW⁺21] or suffer from achieving PCS slower [Wei19, AAN⁺22, AACN⁺24].

While the worst case performance of TreeKEM is a far shot from the advertised $\log(N)$, the known sequences causing it are very artificial. Thus, one might still hope that TreeKEM and in turn MLS perform well for 'natural' random sequences of operations, which would be sufficient for practical applications. However, up to our knowledge there has been very little analysis in this setting so far.

1.1 Our Contributions

The contribution of this work is twofold. We initialize the formal study of MLS's communication complexity for random sequences of operations, which models the lack of coordination between users given the asynchronous nature of CGKA schemes. To this end, we derive lower bounds on the size of its commit messages in an easy to understand model letting randomly chosen users issue update proposals and commits. Our results indicate that, unfortunately, even very small amounts of update proposals per commit operation have a substantial negative impact on MLS's communication complexity.

On the positive side, we propose an alternative way of handling update proposals that provably achieves a logarithmic communication complexity in our model. The changes required to handling proposals are not

too complex, which in our view is a very positive feature, as our proposed method is compatible with the complex mechanisms employed by MLS to ensure authenticity and consistency.

Lower bounds on MLS’s Average Commit Size. Our main goal is to gain an understanding of the impact that blanking nodes in the ratchet tree due to update proposals has on the size of commit messages, for randomly generated sequences of operations. Ideally, one would analyze the communication complexity of MLS using real-world user data. As such data is unfortunately not available, we opt for a simple model of communication. The advantage of this being that it, on the one hand, makes the problem of deriving formal bounds approachable, and, on the other, leads to results that are understandable on an intuitive level.

Our experiment first initializes a group of size N . Then it repeatedly starts issuing update proposals and commits for rounds t as follows.

1. A set U_P consisting of $P \in \mathbb{N}$ group members is sampled uniformly at random.
2. All users in U_P issue an update proposal.
3. A user sampled uniformly at random issues a commit to all update proposals.
4. The size $\text{Cost}(t)$ of the commit is recorded.

To also be able to handle the setting where update proposals are only issued every other round, the experiment is additionally parameterized by $C \in \mathbb{N}$. In the case $C > 1$, the commit to the P proposals is followed by $C - 1$ rounds of commits not implementing any proposals. We are interested in the expectation of $\text{Cost}(t)$ after running the experiment for sufficiently many rounds. Informally, our main result regarding this is the following.

Theorem 4.2 (informal). *Let P, C be constant. Then the expected size of a commit generated according to the sequence described above is of order*

$$\mathbb{E}[\text{Cost}(t)] \geq \Omega(N^{\log_2(1 + \frac{P}{P+C})}).$$

Moreover, if $P = 1 = C$, it holds that

$$\mathbb{E}[\text{Cost}(t)] \geq \log^2(N)/4 - \log(N)/4,$$

which is a more relevant bound for small values of N .

We provide some approximate values of the exponent $e = \log_2(1 + \frac{P}{P+C})$ in the table below. We use $C = 1$ for $P \geq 1$ and for ease of notation write $P = 1/c$ in the case $P = 1, C = c$.

P	1/5	1/2	1	2	5	10	50
e	0.22	0.42	0.58	0.74	0.87	0.93	0.99

Note that already if a single update proposal is issued per commit the MLS’s communication complexity devolves to about \sqrt{N} , a far shot from $\log(N)$. Further, if the number of proposals increases the cost of commits approaches the worst case of a linear commit cost quite fast, with an exponent of 0.99 for $P = 50$.

This means that one either must perform almost all updates sequentially, or commits become almost linear in N . As MLS aims to support groups containing up to $N = 50000$ members, neither of these options is very appealing.

Before turning to our alternative proposal of handling update proposals, we give a bit of intuition on the lower bound’s proof. As the cost of a commit is mainly determined by the blanks in the ratchet tree we first investigate the expected probability of a fixed node v being blank. We notice that the transition of v ’s state between blank and populated forms a Markov chain, allowing us to determine its stationary distribution. We obtain that the probability of v being blank, assuming v is of sufficient depth in the tree, is roughly $P/(P + C)$.

In MLS if a user’s copath contains a node v , then how much v contributes to the communication complexity of a commit by that user depends on whether v is blank. If v is populated, then encrypting to it

only requires one ciphertext. If it is blank, a commit in MLS would include ciphertexts encrypted under the keys of its children and, due to the way MLS works, at least one child must be blank as well meaning that those ciphertexts would have to be encrypted under some of v 's grandchildren and so on. Intuitively, this means that the number of ciphertexts needed to encrypt to a blank node v is $1 + P/(P + C)$ times the number of ciphertexts needed to encrypt to a child of v provided that both v and its child are blank, where 1 corresponds to the child that must be blank and $P/(P + C)$ to the probability that the other child is blank. This argument assumes independence (between the events that the left child of v is blank and the right child of v is blank), which is not justified. However, we show that in expectation the number of ciphertexts does indeed increase by a factor of $1 + P/(P + C)$.

An alternate way of performing update proposals. It is known that if too many users, say, for example, linearly many in the group size, issue concurrent update operations there is no possibility for efficient CGKA. In fact, Bienstock et al. [BDR20] prove that in this case the communication complexity must be linear in the group size. While the lower bound discussed above indicates that even a small number of update proposals per commit operation has a devastating effect on MLS's communication complexity, it leaves open the possibility of an alternative CGKA being more viable in this setting. Recall that update proposals introduce blanks in the tree, as all keys on the updating users' update paths are removed. A natural idea to prevent introducing blanks in the tree would be to not only have the user issuing a commit, but also updating users to re-key their complete update path. In fact this is the approach taken by the continuous group-key agreement scheme CoCoA [AAN⁺22] which in this setting has logarithmic communication complexity. However, proposals can be issued concurrently, meaning users make concurrent modifications to the ratchet tree. This leads to users encrypting new secrets to outdated keys, i.e., keys concurrently being replaced, and, as a second point, is not compatible with the mechanisms used by MLS to ensure integrity of the ratchet tree. As a consequence, CoCoA does achieve PCS slower than MLS, and does not achieve the same strong security notion of insider security.

The second main contribution of this work is a CGKA we denote MLS-Cutoff with the following features.

- For a constant number P of update proposals per commit operation and random operations as described above, the size of its proposal and commit messages is bounded by $O(\log(N))$
- It achieves the same strong security notion as the MLS protocol, namely insider security [AJM22] for the same safety predicate. In a nutshell, the latter says that if all compromised parties issue update proposals that are in turn committed to the group achieves PCS.
- It requires only small modifications to the MLS design. More precisely, it only modifies how the ratchet tree's nodes are blanked or populated by applying update proposals, but keeps other components required, for example, to ensure consistency and authenticity unchanged.

We provide some intuition on our design. Our main observation is that if the number of update proposals is not too large and the updating users are randomly distributed, then updating users' paths merge far up in the tree. Thus, while in MLS-Cutoff update proposals do rekey the issuing user's path, we stop the re-key operation i_{cut} steps before reaching the root, where i_{cut} is the cutoff parameter of the scheme, that is chosen as $\log(\log(N))$. If no update paths collide before the cutoff point, no secrets are encrypted to keys that are concurrently being changed. This on one hand, implies PCS at the same speed as MLS, and on the other is compatible with the tree-integrity mechanisms of tree-hash and parent-hash. In the case that we are unlucky and two updating user's paths collide before the cutoff point, the paths are blanked from (and including) the point where the paths meet. For a depiction on how MLS-Cutoff handles proposals and commits, see Figure 1 (right).

We point out that the worst-case communication complexity of MLS-Cutoff is $\Omega(N)$, as is the case for MLS. This, however, holds for any CGKA achieving fast PCS [BDR20] constructed from standard primitives. We consider MLS-Cutoff to be an attractive alternative to MLS in settings where while update proposals are required their number is not too large. This is a realistic scenario since MLS leaves the update policy to the implementation and, therefore, a rule which simply triggers an update with some probability for users that are online (or enforces it if too much time or communication happened since the last update) would

lead to a fairly random update pattern if there is no user-to-user coordination and thus good performance of MLS-Cutoff. We consider it an interesting question for future work to analyze its performance in more involved settings.

2 Preliminaries

2.1 Continuous Group-Key Agreement

We recall continuous key-agreement schemes (CGKA). As we aim for the security notion of insider security our syntax follows [AJM22].

Public-key infrastructure. We model the public-key infrastructure (PKI) as consisting of two services: an authentication service (AS) and a key service (KS). The former allows parties to register their signing keys and retrieve the signature verification keys of other users, while the latter allows users to register the key packages that are used by other users to add them to a group.

The authentication service stores pairs of users and signing verification keys, (u, svk) , if the key svk has been register under user identity u . Any user u can register a new key and request to receive the secret signing keys that correspond to verification keys registered under them. Moreover, a user can also check if a pair (u, svk) has already been registered and a user can also delete their secrets. We consider the AS as not trusted and therefore in our security model we will give the adversary the possibility of registering keys as well as exposing keys registered by the users (provided that they have not been deleted).

The key service stores pairs of users and key packages, (u, kp) , which are used by other parties to add them to a group. Similarly to the case of AS, users are able to register a new key package as well as retrieving the secret keys sk associated to their key packages and deleting their secret keys. A user can also request a key package for another party. In our model this is done by forwarding the request to the adversary which chooses the package received by the user. This reflects that our security model does not trust the KS. The adversary can also expose the secret keys of a user that have not been deleted.

Syntax. A CGKA scheme allows a group of users to agree on a shared secret evolving over epochs. To this end, the users exchange proposal messages, proposing changes to the group’s state (as for example adding or removing users to or from the group), and commit messages, which implement proposals and advance the group’s state to the next epoch by establishing a new shared secret.

Definition 2.1. *An asynchronous continuous group-key agreement protocol consists of a tuple of algorithms $\text{CGKA.Create}, \text{CGKA.Prop}, \text{CGKA.Com}, \text{CGKA.Proc}, \text{CGKA.Join}, \text{CGKA.Key}$, which can interact with a PKI. Additionally, each user maintains an internal state which we denote by st and is assumed to be an implicit input of all algorithms.*

$\text{CGKA.Create}(\text{svk}_u)$ is run by user u with signing verification key svk_u and creates a group with u as its only member.

$\text{CGKA.Prop}(\text{op}, \text{ad}) \rightarrow \text{pmsg}$ is run by a member who wants to propose an operation op , where $\text{op} \in \{\text{'add'}, \text{'rem'}, \text{'upd'}\}$, ad is some additional data that may depend on the operation. It outputs a proposal message pmsg .

$\text{CGKA.Com}(\text{PMSG}, \text{ad}) \rightarrow (\text{cmsg}, \text{wmsg})$ is run by a member u to commit to a list of proposals $\text{PMSG} = (\text{pmsg}_i)_i$ and it also takes as input some additional data ad . It outputs a commit message cmsg , and a (potentially empty) welcome message wmsg .

$\text{CGKA.Proc}(\text{cmsg}, \text{PMSG}) \rightarrow (\text{PropSemantics}, v)$ is run by a client u to process a commit message cmsg and the corresponding proposal messages PMSG . It outputs some information about the proposals (which we denote PropSemantics) and the party v that generated the commit c .

CGKA.Join(wmsg) \rightarrow (roster, v) is run by user u to process a welcome message and join the respective group. It outputs the set of group members and their signing verification keys (which we refer to as roster) and the party that generated the commit c

CGKA.Key \rightarrow k is run by a member u to obtain the group key associated to their state.

Correctness and security in the UC framework. We follow the work of Alwen et al. ([AJM22] and [ACJM20]). They define a notion of security based on the UC (universal composability) framework, which was first introduced by Canetti in [Can01]. We defer the formal description of the ideal CGKA functionality and the PKI functionalities to Appendix A. In the UC framework one formalizes security of a protocol as the indistinguishability between the ‘real world’ and an ‘ideal world’. We will refer to the distinguisher as the environment.

In the real world the environment \mathcal{Z} gets to choose the inputs (e.g., an instruction to Alice to add Bob) of the machines (e.g., the parties in an execution of a CGKA protocol) of the protocol π and receives their intermediate subroutine-outputs (e.g., a commit generated by Alice). The environment also interacts freely with an adversary \mathcal{A} . \mathcal{A} also interacts with the machines of π and \mathcal{A} has control over the network in the communication between the machines of π . The interaction between \mathcal{A} and π also captures information leakage from the protocol execution (e.g., \mathcal{A} may have the possibility of corrupting the parties).

In the ideal world the protocol π is substituted by an ideal functionality \mathcal{F} and a collection of dummy machines that relay their inputs to \mathcal{F} and the outputs of \mathcal{F} to the corresponding machines. The adversary in the ideal world is denoted by \mathcal{S} because it plays the role of a simulator (of the interaction between \mathcal{A} and \mathcal{Z} in the real world) and we refer to it as an ideal adversary. The adversary \mathcal{S} now communicates with \mathcal{F} and does not interact with the dummy machines. The environment \mathcal{Z} interacts with the dummy machines and the adversary \mathcal{S} just as in the real world.

In the UC framework one says that a protocol π securely realizes a functionality \mathcal{F} if for all efficient adversaries \mathcal{A} there exists an efficient ideal adversary \mathcal{S} such that the output distribution of any efficient environment \mathcal{Z} after interacting with π and \mathcal{A} is indistinguishable from the output distribution of \mathcal{Z} after interacting with \mathcal{F} and \mathcal{S} . This can be understood as representing the idea that an adversary \mathcal{A} attacking the protocol π does not learn more than the ideal adversary (sometimes called simulator) \mathcal{S} learns from interacting with the ideal functionality.

In the work of Alwen et al. [AJM22], whose security notion we use, one has to consider some additional notions since parties have access to a PKI and certain primitives are modeled as random oracles. We refer to these as setup functionalities. This is done by considering the so called hybrid models (first introduced in [CDPW07]). If \mathcal{G} is a setup functionality, then the environment \mathcal{Z} (as well as π , \mathcal{A} and \mathcal{S}) is allowed to interact with \mathcal{G} . And one says that the a protocol π securely realizes a functionality \mathcal{F} in the \mathcal{G} -hybrid model if the real world and ideal world are indistinguishable even if \mathcal{Z} is allowed to interact with \mathcal{G} .

Ideal functionality for CGKA. The definition of the ideal functionality $\mathcal{F}_{\text{CGKA}}$ already takes correctness into account. In order to do this certain checks are included in $\mathcal{F}_{\text{CGKA}}$ and the environment rather than the adversary is given control over the network. Now we focus our attention on defining security for CGKA following [AJM22]. They consider two more constrains on the quantification of the notion of security described in the previous paragraph. We restrict ourselves to admissible environments, i.e., those that can only corrupt parties at certain times. This is defined as part of the ideal functionality which specifies some conditions that cannot be violated by the environment except with negligible probability. The second condition consists in not considering all possible adversaries but instead only corruption preserving adversaries, i.e., adversaries that trigger a corruption if and only if prompted by the environment.¹ This is done to prevent the ideal adversary from triggering a disallowed corruption since this would make security trivial in the sense that all environments would be disqualified.

In order to define the ideal functionality $\mathcal{F}_{\text{CGKA}}$ we consider the history graph which keeps track of the evolution of the group. Intuitively, the nodes of the graph correspond to the operations done by the users

¹Recall that an arbitrary adversary can choose when to corrupt a party without this being determined by the environment.

and the edges capture the order in which they are issued. The history graph [ACDT21] is an acyclic directed graph which consists of two kinds of nodes, proposal nodes and commit nodes. The edges of the graph capture the notion of how new epochs are derived. Additionally, each node is associated with a collection of attributes ($\text{Node}[c]$ for commit nodes c and $\text{Prop}[p]$ for proposal nodes p) related to the operations of a CGKA protocol. We list these attributes in the following table (divided into three groups: those that all nodes have, those that only proposal nodes have and those that only commit nodes have).

.orig	the party that triggers the creation of the node
.par	the parent commit node
.stat	good for secure nodes, bad for nodes created with adversarially chosen randomness (the adversary knows the secrets associated to this node) or adv for nodes that correspond to messages injected by the adversary
$\text{Prop}[p].\text{act}$	indicates the kind of action and some additional information (upd, svk) if it is an update performed by a party with key svk ($\text{add}, v, \text{svk}_v$) if party v with key svk_v is added (rem, v) if party v is being removed
$\text{Node}[c].\text{pro}$	the list of proposals associated to commit c
$\text{Node}[c].\text{mem}$	list of pairs of members and their signing verification keys
$\text{Node}[c].\text{chall}$	true if the corresponding application secret is chosen at random and false if chosen adversarially
$\text{Node}[c].\text{exp}$	pairs ($u, \text{true/false}$) of corrupted parties u and a boolean value expressing whether the application secret also leaked to adversary

The ideal functionality just keeps track of the evolution of the history graph. It starts building the history graph with just one node root_0 when the group is created. Additionally $\mathcal{F}_{\text{CGKA}}$ stores a value called pointer $\text{Ptr}[u]$ for each party u that maps u to the last commit it processed. When a party u makes a proposal p , $\mathcal{F}_{\text{CGKA}}$ creates a proposal node $\text{Prop}[p]$ whose parent in the history graph is $\text{Node}[\text{Ptr}[u]]$ and u is the origin of the proposal $\text{Prop}[p].\text{orig}$. The functionality also stores information about the kind of operation associated to p in $\text{Prop}[p].\text{act}$. If u makes a commit c to some proposals P , $\mathcal{F}_{\text{CGKA}}$ creates a commit node $\text{Node}[c]$ whose parent in the history graph is $\text{Node}[\text{Ptr}[u]]$, $\text{Node}[c].\text{pro}$ is set to P and $\text{Node}[c].\text{orig}$ is set to u . And if a party v processes c , the functionality just moves v 's pointer to c .

As mentioned before, the ideal functionality imposes certain restrictions on the behaviour of the environment. This is done by including two predicates, namely, $\text{safe}(c)$ and $\text{inj-allowed}(c, u)$ that regulate when an adversary can corrupt a user and when it can inject a commit on behalf of a user, respectively.

In order to guarantee that the only application secrets that can be learned by an adversary in the real world are the ones trivially derived from secrets of exposed users, we let the simulator choose these ‘unsafe’ secrets in the ideal world while the remaining ones are chosen at random by the ideal functionality. Thus, predicate $\text{safe}(c)$ evaluates to **false** if and only if at least one of the following is true:

- (a) $\text{Node}[c].\text{exp} \neq \emptyset$ and the application secret has been leaked to the adversary, or
- (b) the adversary can process c using the secrets of exposed users and the exposed signing keys (in particular this captures the secret keys associated to key packages generated with exposed signing keys or corrupted randomness) of its ancestors.

Analogously, predicate $\text{inj-allowed}(c, u)$ guarantees that in the real world authenticity is only broken when the adversary can trivially forge messages on behalf of a user by restricting when injections are possible in the ideal world. Predicate $\text{inj-allowed}(c, u)$ evaluates to **true** if and only if u 's signing key in $\text{Node}[c]$ has been exposed (i.e., $\text{Node}[c].\text{mem}[u] \in \text{Exposed}$) and at least one of the following two holds:

- (a) $\text{Node}[c].\text{exp} \neq \emptyset$,
- (b) the adversary can process c using the secrets of exposed users and the exposed signing keys (in particular this captures the secret keys associated to key packages generated with exposed signing keys or corrupted randomness) of its ancestors.

The reader may find the precise definition of the predicates $\text{safe}(c)$ and $\text{inj-allowed}(c, u)$ in Figure 17.

2.2 Ratchet Trees

We introduce some basic terminology. We use the word *tree* to refer to a directed acyclic graph $\mathbb{T} = (\mathbb{V}, \mathbb{E})$ with a root, i.e., a node r such that for every node v there exists a unique directed path between r and v . Said path is denoted by $\text{path}(v)$. As in the MLS specification we use the convention that $\text{path} = (v_1, \dots, v_\ell = r)$ is ordered ascending to the root, i.e., in reverse direction of the edges, and does not include v .

Let u and v be two different nodes of some tree $\mathbb{T} = (\mathbb{V}, \mathbb{E})$. We say that u is a parent of v if there exists an edge $(u, v) \in \mathbb{E}$. We say that u is a child of v if and only if v is a parent of u . We denote the set of children by $\text{child}(v)$. We say that u is an ancestor of v if there exists a directed path from u to v . The set of ancestors is denoted by $\text{ancs}(u)$. We say that v is a descendant of u if and only if u is an ancestor of v . We say that u and v are siblings if they are children of the same node. We say that a node is a leaf if it has no children. Thus we have three different kinds of nodes; the root, the leaves and the remaining elements in V which we call intermediate nodes.

In this work we consider complete binary trees with $N = 2^n$ many nodes. Thus each non-leaf node v has exactly two children which we refer to as the left child, denoted by $\text{left}(v)$, and the right child, $\text{right}(v)$. Every node v other than the root has one parent, denoted by $\text{par}(v)$, and a sibling, $\text{sib}(v)$. The depth of a node v is the length of the path from the root to v and we denote it by $\text{depth}(v)$.

A notion closely related to how MLS re-keys users' paths is that of co-path of a node. The co-path of a node v is defined as the set of children of the nodes in v 's path that do not belong to v 's path, namely,

$$\text{co-path}(v) = \cup_{u \in \text{path}(v)} \{w \mid w \in \{\text{left}(u), \text{right}(u)\} \setminus (\text{path}(v) \cup \{v\})\}.$$

Node states. CGKA schemes like MLS make use of a *ratchet tree*, i.e., a complete binary tree $\mathbb{T} = (\mathbb{V}, \mathbb{E})$ with each group member being associated to one of the tree's leaves. Every node $v \in \mathbb{E}$ has as associated data a so called public and private state. Its public state contains

- a PKE public key $v.\text{pk}$,
- a list $v.\text{unm}$ of so called unmerged leaves,
- a parent-hash value $v.\text{pHash}$,
- a signature verification key $v.\text{svk}$ (only if the node is a leaf), and
- a credential $v.\text{cred}$ (only if the node is a leaf).

The private state of the node is $v.\text{sk}$ the secret key associated to $v.\text{pk}$. MLS, as well as the variant defined in this work, enforce the so called tree-invariant stating that $v.\text{sk}$ is known by exactly the user whose leaves have v as an ancestor. Further, a node can be *blank*, meaning that all associated values are empty. We use function $\text{blank}(v)$ to blank a node, i.e., to set all its associated data to \perp .

Resolution and filtered path. The existence of blank nodes leads to the notion of resolution $\text{Res}(v)$ of a node v which, intuitively, corresponds to the smallest covering of the set of descendants of v by non-blank nodes. We define it formally as follows;

- If v is a populated node, its resolution is $\text{Res}(v) = \{v\}$.
- If v is a blank leaf, its resolution is $\text{Res}(v) = \emptyset$.
- If v is a blank internal node, its resolution is defined recursively as

$$\text{Res}(v) = \cup_{u \in \text{child}(v)} \text{Res}(u).$$

If X is a set of nodes, we use the notation $\text{Res}(X)$ to denote $\cup_{v \in X} \text{Res}(v)$.

A leaf v 's filtered path fil-path is the update path after removing all nodes the co-child of which has an empty resolution (and unmerged leafs set). I.e., if $\text{path}(v) = (v_1, \dots, v_\ell = r)$ then using the convention $v_0 = v$ we define

$$\text{fil-path}(v) = (v_i, i \in \{1, \dots, \ell\} \mid (\text{Res}(\text{sib}(v_{i-1})) \cup \text{sib}(v_{i-1}).\text{unm}) \neq \emptyset).$$

Finally, we sometimes have to identify the node in which two leaves' filtered paths meet. Consider leaves $u, v \in \mathcal{V}$ of the same depth with filtered paths $(u_1, \dots, u_\ell = r) \leftarrow \text{fil-path}(u)$ and $(v_1, \dots, v_{\ell'} = r) \leftarrow \text{fil-path}(v)$, respectively. We define their least common ancestor $\text{lca}(u, v)$ as the node u_i in $\text{path}(u)$ where i is minimal such that $u_i \in \text{path}(v)$. Counting from the root we define the index ind-lca of their least common ancestor as

$$\text{ind-lca}(u, v) = \max(a \mid u_{\ell-a} = v_{\ell'-a}, a \in 0, \dots, \ell - 1).$$

2.3 Auxiliary Results

Lemma 2.2 (Bernoulli's Inequality). *For all $x \geq -1$ and $r \in \mathbb{R} \setminus (0, 1)$ it holds that*

$$(1 + x)^r \geq 1 + rx.$$

3 The Messaging Layer Security Protocol

We defer the formal description of the MLS protocol² [BBR⁺23] to Appendix B.1. In this section we provide an overview on its working principle. As the goal of this work is to investigate the protocol's communication complexity which is largely determined by the distribution of blank nodes in the underlying ratchet tree, our exposition mainly focuses on the mechanisms leading to the addition of blanks to the tree, or removal of blanks, respectively. For ease of exposition in this section we treat the key-packages kp (consisting of the corresponding user's PKE pk , signature verification key svk , and leaf credential cred) assigned to leaves in the ratchet tree as public key pk , and, similarly, assign the tree's root a key pair which is not the case in the full protocol (see Figure 19).

Protocol state. The cryptographic state st each user of MLS keeps track of is conceptually divided as follows.

The underlying ratchet tree T . In this simplified exposition we identify users u with their associated leaf, simply writing v_u . As discussed in Section 2 every user u has a complete view of the tree's public part, but only has access to the secret keys $v_u.\text{sk}$ of nodes that lie on their update path v_u . Sometimes when operating on ratchet trees we require the identity of the user running the algorithm, which we denote by $u\text{-own}$.

The key schedule collects a number of secrets used for different purposes. These include, for example, the shared group secret referred to as the application secret, the initialization secret that is combined with the randomness obtained from re-keying a path to advance the key schedule to the next epoch, and a membership key used in combination with the user's signature key to authenticate protocol messages.

Further, the protocol employs consistency mechanism ensuring that all users have the same view of the protocol's state. This includes the transcript hash that essentially encodes the complete history of the group including all operations processed so far, as well as so-called tree-hash and parent-hash values that ensure consistency of the ratchet tree. We will discuss the latter two in a bit more detail at the end of this section.

Generating proposal and commit messages. Users issue proposal messages to add or remove user to or from the group, respectively, as well as to update the personal key pair located at their leaf. In more detail, proposal messages

$$\text{pmsg} = (\text{'type'}, \text{ad})$$

consist of **'type'** indicating the type of proposal, as well as some associated data. In case of an update proposal **'upd'** the latter contains a new public key to be included at the issuing user's leaf, in case of

²Formally, in this work we consider the CGKA underlying MLS. Thus, our exposition ignores the secret tree scheduling keys used to encrypt and authenticate payload messages.

an addition ‘add’ the associated data $\text{ad} = \text{pk}$ is the added party’s personal key, and finally in case of a removal ‘rem’ it indicates the removed user’s leaf. In case of an update proposal the issuing user stores the secret key sk corresponding to pk as a pending update in a list pendUpd .

To prevent unauthorized parties from issuing proposals pmsg has to be *framed*. Essentially, this means the proposal message and group context are signed by the issuing user as well as MACed using a key only available to group members. When processing proposal messages users unframe the message by verifying the authenticity of pmsg before recovering $((\text{‘type’}, \text{ad}), u\text{-snd})$. Here $u\text{-snd}$ denotes the user issuing the message.

A user u can commit to a (potentially empty) list PMSG of proposals, implementing the proposed changes, advancing the group to the next epoch, and establishing a new application secret. This is done by executing the following steps.

1. Verify every proposal $\text{pmsg} \in \text{PMSG}$.
2. Create a copy T' of ratchet tree $\text{st}.T$ and
 - (a) call $\text{apply-props}(\text{PMSG})$ to apply the proposed changes to T' ,
 - (b) call $\text{rekey-path}(v_{u\text{-own}}, 0)$ to rekey $u\text{-own}$ ’s update path in T' and generate an UpdPath object used to communicate the changes to the other group members.
3. Compute the new key schedule based on T' .
4. Authenticate UpdPath and bind it to PMSG .
5. Construct a welcome message if necessary.
6. Store the resulting state st' (containing T') as a pending commit and return the corresponding, framed, commit message cmsg , which, in particular, contains UpdPath .

The largest component of cmsg is UpdPath which itself is comprised of public keys and ciphertexts. The number of ciphertexts that have to be generated depends on the blanks present in T' .

To process a commit message issued by u users essentially follow the same sequence of operations. More precisely, they recover T' by first applying the proposals PMSG to their current view of the tree, and then recovering the new keys on the committing user’s update path from UpdPath . They update the key schedule and verify that all changes were properly authenticated. If the latter is the case, the new key schedule and ratchet tree replace the ones stored in the user’s state. If the processing user is the one issuing the commit operation, they can simply update their state to the one stored as a pending commit.

Applying proposals and path rekeying. When issuing (or processing) a commit operation, blanks are introduced in the ratchet tree by calling apply-props in step 2a and removed from it when calling rekey-path (or applying the re-key operation) in step 2b. For an overview of the two functions see Figures 2 and 3.

apply-props on input a copy T' of the current ratchet tree and list PMSG of proposals applies changes to the ratchet tree as follows. For $\text{pmsg} \in \text{PMSG}$ of type ‘upd’ the issuing user $u\text{-snd}$ ’s leaf key is replaced by the one contained in pmsg . Then $u\text{-snd}$ ’s update path is blanked. If pmsg proposes removing user $u\text{-rem}$, leaf $v_{u\text{-rem}}$ as well as the removed user’s update path are blanked. If the removal results in the right half of the tree not having any populated leaves its size is halved using function truncate . If pmsg proposes adding user $u\text{-add}$ to the group, they get assigned the leftmost unpopulated leaf in the tree, doubling the tree size in the process (by adding blank nodes) if no such leaf exists. The leaf is assigned the public key contained in ad , and added to the set of unmerged leaves for every node in its update path.

To replace the keys on their path the committing user u calls rekey-path . It takes as input T' , u ’s leaf v_u , and, in our presentation, an additional index i_{cut} that allows stopping the rekeying i_{cut} steps early. While MLS uses $i_{\text{cut}} = 0$, looking ahead, it will be used in proposed alternative method of issuing update proposals. Denoting $v_0 = v_u$ and $(v_1, \dots, v_\ell) = \text{fil-path}(v_0)$ the function samples seed s_0 and iteratively generates seed s_i using a key-derivation function, in our presentation written as hash function H_1 . It then generates key pairs $(\text{pk}_i, \text{sk}_i) \leftarrow \text{Pke.Gen}(H_2(s_i))$, H_2 being another hash function, to replace the keys associated to v_i .

```

Algorithm apply-props( $T'$ , PMSG)
00 for pmsg  $\in$  PMSG:
01   parse (('type', ad),  $u$ -snd)  $\leftarrow$  pmsg
02   if ('type', ad) = ('upd', pk):
03     assign-kp( $T'$ ,  $v_{u\text{-snd}}$ , pk)     $\llcorner$  replace  $u$ -snd's leaf pk
04     for  $v \in \text{fil-path}(v_{u\text{-snd}})$      $\llcorner$  blank  $u$ -snd's path
05       blank( $v$ )
06     if  $u$ -own =  $u$ -snd                 $\llcorner$  own update proposal
07       fetch sk  $\leftarrow$  pendUpd(pmsg)
08        $v_{u\text{-own}}.sk \leftarrow$  sk
09   else if ('type', ad) = ('rem',  $u$ -rem):
10     blank-leaf( $T'$ ,  $v_{u\text{-rem}}$ )         $\llcorner$  blank  $u$ -rem's leaf
11     for  $v \in \text{path}(v_{u\text{-rem}})$          $\llcorner$  blank  $u$ -rem's path
12       blank( $v$ )
13      $T' \leftarrow \text{truncate}(T')$      $\llcorner$  truncate tree if necessary
14   else if ('type', ad) = ('add', ( $u$ -add, pk)):
15     add-leaf( $T'$ ,  $u$ -add, pk)       $\llcorner$  assign  $u$ -add a leaf
16     for  $v \in \text{path}(v_{u\text{-add}})$      $\llcorner$  set  $v_{u\text{-add}}$  as unmerged
17        $v.unm \leftarrow \cup \{v_{u\text{-add}}\}$ 
18 return  $T'$ 

```

Figure 2: High-level overview on function `apply-props` for MLS. The function's formal description is in Figure 23.

In order to enable the other group members to recover the secret keys they should have access to, for all $j \in \{1, \dots, \ell\}$ the co-child v_{sib} with respect to v_u 's update path is computed. Then the seed s_j is encrypted under all public keys associated to nodes contained in $\text{Res}(v_{\text{sib}})$, as well as the ones of unmerged leaves resulting in a collection C_j of ciphertexts. The update path object consists of

$$\text{UpdPath} \leftarrow (\text{pk}_0, (\text{pk}_1, C_1), \dots, (\text{pk}_\ell, C_\ell)).$$

A user processing a commit message issued by u -snd replaces the public keys on u -snd's path by the ones contained in `UpdPath`. Then, by computing where their own and u -snd's update paths meet (say for index j) they are able to determine node v_{sib} relevant to them, and using the secret key in $\text{Res}(v_{\text{sib}}) \cup v_{\text{sib}}.unm$ they have access to decrypt one of the ciphertexts contained in C_j to recover s_j . In turn they are able to recover the secret keys sk_j, \dots, sk_ℓ that they should have access to according to the tree invariant by applying H_1 , H_2 , and Pke.Gen .

Consistency mechanisms w.r.t the ratchet tree. MLS employs tree-hashes and parent-hashes (see Figure 21) to ensure consistent views of the tree and provide improved security guarantees for users joining an existing group. The former essentially is a commitment to the entire public state of the ratchet tree, and is computed as a Merkle commitment. I.e., a node's tree hash corresponds to a hash of the node's public values and the tree-hashes of its children.

The parent hash captures information on how the tree's nodes were populated during the protocol execution. Parent hash values are computed as part of a user's commit, along the user's update path, however in reverse order, i.e., from root to leaf. Essentially, a node's parent hash attests to its parent's public key, as well as the tree hash of the sub-tree rooted at its co-child. It plays an important role in preventing malicious insiders from being able to successfully invite users to malformed groups.

```

Algorithm rekey-path( $v_u, i_{\text{cut}}$ )
00  $v_0 \leftarrow v_u$ 
01  $s_0 \leftarrow_{\$} \{0, 1\}^\lambda$  \| generate leaf seed
02  $(v_0.\text{pk}, v_0.\text{sk}) \leftarrow \text{Pke.Gen}(\text{H}_2(s_0))$  \| sample leaf key pair
03  $(v_1, \dots, v_\ell) \leftarrow \text{fil-path}(v_u)$ 
04 for  $j = 1, \dots, \ell - i_{\text{cut}}$ :
05    $s_j \leftarrow \text{H}_1(s_{j-1})$  \| generate seed
06    $(v_j.\text{pk}, v_j.\text{sk}) \leftarrow \text{Pke.Gen}(\text{H}_2(s_j))$  \| sample key pair
07    $C_j \leftarrow ()$ 
08    $v_{\text{sib}} \leftarrow \text{sib}(v_{j-1})$  \| determine co-child
09   for  $w \in \text{Res}(v_{\text{sib}}) \cup v_{\text{sib}}.\text{unm}$ : \| encrypt seed
10      $C_j \leftarrow_{\cup} \text{Pke.Enc}(w.\text{pk}, s_j)$ 
11  $\text{UpdPath} \leftarrow (\text{pk}_0, (\text{pk}_1, C_1), \dots, (\text{pk}_\ell, C_\ell))$ 
12 for  $v \in \text{path}(v_u)$  \| merge leaves
13    $v.\text{unm} \leftarrow \emptyset$ 
14 return  $(\mathcal{T}', \text{UpdPath})$ 

```

Figure 3: High-level overview on function `rekey-path` for MLS. The function’s formal description is in Figure 22.

4 The Communication Cost of MLS for Random Sequences of Operations

We investigate the communication complexity of MLS. Section 4.1 defines a simple experiment used to generate random sequences of operations. We then proceed to compute bounds on the amount of blank nodes the MLS ratchet tree will have on average for such sequences (in Section 4.2) before computing a lower bound on the communication cost of commit operations (in Section 4.3).

4.1 Considered Scenario

As discussed in the introduction MLS exhibits a communication complexity that is logarithmic (in the number of users) in the best case, but can devolve to being linear in the worst case. In fact, the latter must be the case for any protocol that is built from standard primitives and achieves PCS as fast as MLS [BDR20, ANPPP23, BDG⁺22]. Beyond this, very little is known about MLS’s performance in practice.

We initiate the study of MLS’s communication cost under randomized sequences of operations. While ideally, one would evaluate the performance for sequences of operations obtained from real-world user data, such data is unfortunately, up to our knowledge, not publicly available. As a consequence, in this work we take a pragmatic approach and consider sequences determined by a simple experiment. Essentially, we generate a fixed number of update proposals per commit operation, the users performing the operations being sampled uniformly at random. While we do not claim this realistically reflects user behavior, on the one hand, it makes the problem of computing lower bounds approachable, and, on the other, provides simple intuition on the impact that blanking has that is not obscured by technical details of a more complicated model.

We define experiment `Exp-Prop-Com` in Figure 4. It is parameterized by the group size $N = 2^n \in \mathbb{N}$, integers P, C , and the number of rounds t . Essentially, it initializes a group of N users and then for t rounds generates update-proposal and commit operations as follows. First, in every round P users U_P are chosen uniformly at random from the group $[N]$ and issue an update proposal. This is followed by a uniformly random user U_C issuing a commit, which leads to the group moving to the next epoch and ends the round. To allow for the more general scenario that proposals are only issued every couple epochs, we allow for a number of epochs in between sending update proposal in which a commit, but no proposals, is issued. This is done by instead of sampling a single committer to end the round, to sample C committers uniformly at random. To analyze the communication cost of operation sequences generated in this way we record the size

Experiment	Exp-Prop-Com(n, P, C, t)	Round(U_P, U_C)
00	$N \leftarrow 2^n$	16 $\text{cost} \leftarrow ()$
01	$\text{CGKA.Create}_{u_1}(\text{svk})$ $\quad \backslash \backslash$ initialize group	17 $\text{PMSG} \leftarrow \emptyset$
02	$\text{PMSG} \leftarrow \emptyset$	18 for $k \in U_P$:
03	for u in $\{u_2, \dots, u_N\}$:	19 $\text{pmsg} \leftarrow \text{CGKA.Prop}_{u_k}(\text{'upd'}, \perp)$
04	$\text{pmsg} \leftarrow \text{CGKA.Prop}_{u_1}(\text{'add'}, u)$	20 $\text{PMSG} \leftarrow \cup \text{pmsg}$
05	$\text{PMSG} \leftarrow \cup \text{pmsg}$	21 for $k \in U_C$:
06	$(\text{cmsg}, \text{wmsg}) \leftarrow \text{CGKA.Com}_{u_1}(\text{PMSG})$	22 $(\text{cmsg}, \cdot) \leftarrow \text{CGKA.Com}_{u_k}(\text{PMSG})$
07	$\text{CGKA.Proc}_{u_1}(\text{cmsg}, \text{PMSG})$	23 $\text{cost} \leftarrow \cup \text{cmsg} $
08	for u in $\{u_2, \dots, u_N\}$:	24 for $m \in [N]$:
09	$\text{CGKA.Join}_u(\text{wmsg})$	25 $\text{CGKA.Proc}_{u_m}(\text{cmsg}, \text{PMSG})$
10	initialize $\text{Cost} = ()$	26 $\text{PMSG} \leftarrow \emptyset$
11	for i from 1 to t : $\quad \backslash \backslash$ issue rounds	27 return cost
12	$U_P \leftarrow_{\S} \{S \subseteq [N] : U_P = P\}$	
13	$U_C \leftarrow_{\S} \{S \subseteq [N] : U_C = C\}$	
14	$\text{Cost} \leftarrow \cup \text{Round}(U_P, U_C)$	
15	return Cost	

Figure 4: Experiment Exp-Prop-Com with respect to group size $N = 2^n$ and numbers of proposals P and commits C executed over t rounds. We use subscripts u to denote that an algorithm is executed by user u , and omit algorithm outputs not relevant to the experiment.

of each commit message generated during the experiment in an array Cost .

Our goal is to compute a lower bound on the expected value of $\text{Cost}(t)$ after running the experiment for sufficiently many rounds. Recall that the amount of MACs, signatures, and public keys contained in commit messages is at most logarithmic in N . Thus, $\text{Cost}(t)$ is dominated by the number of ciphertexts that have to be generated when calling `rekey-path` (see Figure 2) to communicate the new secret keys on the committer u -snd's path to the other users. Accordingly, we can bound the communication cost from below by

$$\text{Cost}(t) \geq \sum_{v_i \in \text{co-path}(v_{u\text{-snd}})} |\text{Res}(v_i)|,$$

the sum of resolution sizes on the committer's co-path, and in our analysis it suffices to derive a bound for this expression.

4.2 Expected Number of Blanks

Recall that the resolution of a node is determined by which of its descendants are blank. In this section we compute the expected probability of nodes at a certain depth being blank in the ratchet tree generated by experiment Exp-Prop-Com.

Expressing blanks as a Markov chain. Consider the sequence $(T_i)_{i=1}^t$ of ratchet trees generated by experiment Exp-Prop-Com, where T_i is the tree all users store as part of their state after the i th call to round function `Round`. For a fixed internal node v we can define a random variable $(B_i(v))_{i=0}^t$ taking values in $\{0, 1\}$ where 0 corresponds to being a populated node and 1 to v being blank. We write B_i instead of $B_i(v)$ when there is no ambiguity. We observe that $(B_i)_{i=0}^t$ forms a Markov chain. Indeed, v transitioning from blank to populated or vice versa only depends on the choice of the user U_P and U_C issuing proposals or commits in the current round. Since the chain is clearly irreducible, it has a unique stationary distribution. Denoting by $p_{b \rightarrow b'}$ the probability of the node switching from state $B_i = b$ to $B_{i+1} = b'$, the transition matrix of $(B_i)_{i=0}^t$ with respect to state vector $(0, 1)^\top$ is given by

$$M = \begin{pmatrix} p_{0 \rightarrow 0} & p_{1 \rightarrow 0} \\ p_{0 \rightarrow 1} & p_{1 \rightarrow 1} \end{pmatrix}.$$

The eigenvalues of M are given by $p_{0 \rightarrow 0} - p_{1 \rightarrow 0}$ and 1 and an eigenvector of one-norm 1 with eigenvalue 1, corresponding to the stationary distribution, is given by

$$\left(\frac{p_{1 \rightarrow 0}}{p_{1 \rightarrow 0} + p_{0 \rightarrow 1}}, \frac{p_{0 \rightarrow 1}}{p_{1 \rightarrow 0} + p_{0 \rightarrow 1}} \right)^\top. \quad (1)$$

Analyzing the stationary distribution of $(B_i)_{i=0}^t$. Consider the state B_i of node v at the end of round i . If $B_i = 1$ (i.e., v is blank at the end of round i), then after the proposals U_P and commits U_C of round $i + 1$ have happened B_{i+1} will remain blank exactly if none of the users in U_C have v as an ancestor. Further, if $B_i = 0$ (i.e., v is a populated node at the end of round i), then B_{i+1} will remain non-blank unless at least one of the users in U_P has v as an ancestor while none of the users in U_C do. Thus, if we denote

$$p_P := \Pr[v \in \text{ancs}(U_P)] \quad \text{and} \quad p_C := \Pr[v \in \text{ancs}(U_C)]$$

we have that $p_{1 \rightarrow 0} = p_C$ and by independence of the choice of proposer and committer that $p_{0 \rightarrow 1} = (1 - p_C)p_P$. By Equation 1 we obtain that the stationary distribution of a node being blank is given by

$$\text{blanks}_{\text{MLS}} := \left(\frac{p_C}{p_C + p_P(1 - p_C)}, \frac{p_P(1 - p_C)}{p_C + (1 - p_C)p_P} \right)^\top \in [0, 1]^2. \quad (2)$$

We now compute the probability of v being an ancestor of a proposing or committing leaf, respectively. Assume v is at depth $\text{depth}(v) = d$. Then \mathbb{T} has $N2^{-d}$ many leaves that are in $\text{ancs}(v)$. Therefore there are $N(1 - 2^{-d})$ many leaves that are not in $\text{ancs}(v)$ which implies that

$$p_P = \Pr[v \in \text{ancs}(U_P)] = 1 - \Pr[v \notin \text{ancs}(U_P)] = 1 - \frac{\binom{N(1-2^{-d})}{P}}{\binom{N}{P}}.$$

Analogously,

$$p_C = \Pr[v \in \text{ancs}(U_C)] = 1 - \Pr[v \notin \text{ancs}(U_C)] = 1 - \frac{\binom{N(1-2^{-d})}{C}}{\binom{N}{C}}.$$

Using Equation 2 we have shown the following.

Lemma 4.1. *For an execution of experiment Exp-Prop-Com and internal node v of depth $\text{depth}(v) = d$ let $(B_i)_i = (B_i(v))_i$ be defined as above. Then the probability that v is blank in the stationary distribution $B = \text{blanks}_{\text{MLS}}$ is*

$$\mathbb{E}[B(v) = 1] = \frac{\left(1 - \frac{\binom{N(1-2^{-d})}{P}}{\binom{N}{P}}\right) \frac{\binom{N(1-2^{-d})}{C}}{\binom{N}{C}}}{1 - \frac{\binom{N(1-2^{-d})}{C}}{\binom{N}{C}} \frac{\binom{N(1-2^{-d})}{P}}{\binom{N}{P}}}. \quad (3)$$

Before turning to computing a lower bound on the expected commit size in experiment Exp-Prop-Com we look at some simple cases:

When $P = 1 = C$, we have that $p_P = 2^{-d} = p_C$. Therefore,

$$\mathbb{E}[B(v) = 1] = (1 - 2^{-d}) / (2 - 2^{-d}) \approx 1/2 \text{ for } d \text{ large enough.}$$

If $P + C = O(1)$,

$$\begin{aligned} \mathbb{E}[B(v) = 1] &\approx \frac{(1 - (1 - 2^{-d})^P)(1 - 2^{-d})^C}{1 - (1 - 2^{-d})^P(1 - 2^{-d})^C} \\ &\approx \frac{(1 - (1 - P2^{-d}))(1 - C2^{-d})}{1 - (1 - (P + C)2^{-d})} \\ &\approx \frac{P}{P + C} \text{ for } d \text{ large enough.} \end{aligned} \quad (4)$$

In both examples we see a constant probability that a node is blank which, intuitively, means that a commit should have a high upload cost. We formalize this in the next section.

4.3 Lower Bound on Sent Ciphertexts

In this section we prove a lower bound on MLS's communication cost for sequences of operations generated by Exp-Prop-Com. As it turns out, already a constant number of proposals per commit leads to commits being exponential in $n = \log(N)$, a far shot from the best case communication complexity. More precisely we will prove the following.

Theorem 4.2. *If $P+C = O(1)$ and $t > 2N^2 \log(N)$, then there exists a constant $\varepsilon > 0$ such that the expected cost of a commit after Experiment Exp-Prop-Com(n, P, C, t) (Figure 4) satisfies the following inequality:*

$$\mathbb{E}[\text{Cost}(t)] \geq \Omega(N^{\log_2(1+(1-\varepsilon)\frac{P}{P+C})}).$$

Before proving this result, we show a weaker result that already gives some of the main ideas used in the proof of Theorem 4.2. We start by studying the size of the resolution of a node v_d at depth $d \in \{1, \dots, \log(N) - 1\}$:

$$\begin{aligned} \mathbb{E}[|\text{Res}(v_d)|] &\geq \mathbb{E}[|\text{Res}(v_d)| \mid B(v_d) = 0] \Pr[B(v_d) = 0] \\ &\quad + \mathbb{E}[|\text{Res}(v_d)| \mid B(v_d) = 1] \Pr[B(v_d) = 1] \\ &= \Pr[B(v_d) = 0] \\ &\quad + \mathbb{E}[|\text{Res}(\text{right}(v_d))| + |\text{Res}(\text{left}(v_d))| \mid B(v_d) = 1] \Pr[B(v_d) = 1] \\ &= \Pr[B(v_d) = 0] \\ &\quad + 2\mathbb{E}[|\text{Res}(\text{left}(v_d))| \mid B(v_d) = 1] \Pr[B(v_d) = 1]. \end{aligned}$$

One can show that

$$\mathbb{E}[|\text{Res}(\text{left}(v_d))| \mid B(v_d) = 1] \geq \mathbb{E}[|\text{Res}(\text{left}(v_d))|]. \quad (5)$$

This implies that for a node v_d at depth $d \in \{1, \dots, \log(N) - 1\}$,

$$\mathbb{E}[|\text{Res}(v_d)|] \geq \Pr[B(v_d) = 0] + 2\mathbb{E}[|\text{Res}(\text{left}(v_d))|] \Pr[B(v_d) = 1].$$

A recursive application of this inequality yields

$$\mathbb{E}[|\text{Res}(v_d)|] \geq 1 + \sum_{i=0}^{\log(N)-d-1} 2^i \prod_{j=0}^i \Pr[B(v_{d+j}) = 1].$$

Now let's assume that $P = 1 = C$. We know that for $d \geq 3$, $\mathbb{E}[B(v) = 1] \approx 1/2$. Therefore we obtain that

$$\begin{aligned} \mathbb{E}[\text{Cost}(t)] &\geq \sum_{d=3}^{\log(N)} (1 + \sum_{i=0}^{\log(N)-d-1} 1/2) \\ &= \log(N) - 2 + \sum_{d=3}^{\log(N)} (\log(N) - d)/2 \\ &= \log(N) - 2 + (\log(N) - 3)(\log(N) - 2)/4 \\ &\approx \log^2(N)/4 - \log(N)/4. \end{aligned}$$

A similar computation when $P + C = O(1)$ and $P > C$ shows $\mathbb{E}[\text{Cost}(t)] = \Omega(N^{\log_2(\frac{2P}{P+C})})$.

The gap between these bounds and the one stated in Theorem 4.2 is due to Equation 5 being non-tight. This is solved by studying how the size of the resolution of a node grows when several of its ancestors are blank.

Proof of Theorem 4.2. In order to prove the theorem we first establish some helpful results. Let's consider the random variable $|\text{Res}(v)|$ that denotes the size of the resolution of a node v . This corresponds to the number of ciphertexts that have to be uploaded when a commit includes a fresh key for v 's parent. The central idea of the lower bound is to relate the expected resolution size of a node to the expected resolution size of its children, and iterate this expectation in order to get a lower bound that's exponential in $\log(N)$. In the proof, we keep conditioning on an increasingly complicated event, and only consider the probability of a node being blank conditioned on this event. This allows us to argue using a local property whereas the resolution size of a node depends on the entire subtree of the node.

We now introduce some notation. Let $\text{Anc-bl}_t(d, v)$ be the event that all ancestors of v up to depth d are blank after the experiment ends. We usually drop the subindex t as we do with the other random variables considered in this section. When the node v is clear from context, we write $\text{Anc-bl}(d)$. Let v_a denote the ancestor of v such that $\text{depth}(v_a) = d$. We use \mathbb{T}_{sub} to denote the full binary subtree with root v_a and leaves the nodes in \mathbb{T} that have v_a as an ascendant and have the same depth as v . Let n' be the number of leaves in \mathbb{T}_{sub} . Notice that $n' = 2^{\text{depth}(v)-d}$.

Lemma 4.3. *Let $X \sim \text{Geo}(q)$ be distributed according to the geometric distribution with $q \in (0, 1)$ and let $Z = X \bmod (P + C)$. Let $i, j \in \{1, 2, \dots, P + C\}$. Let s be an integer multiple of $P + C$, i.e. $s = m(P + C)$ for some $m > 0$. Then,*

$$\Pr[Z = i \mid X \leq s] \geq (1 - q)^{P+C-1} \Pr[Z = j \mid X \leq s].$$

Proof. This is an almost immediate consequence of some elementary computation with finite geometric series. Let $i, j \in \{1, 2, \dots, P + C\}$. Notice that

$$\begin{aligned} \Pr[Z = i \mid X \leq s] &= \frac{1}{1 - (1 - q)^s} \sum_{h=0}^{m-1} (1 - q)^{(P+C)h+i-1} q \\ &= \frac{1}{1 - (1 - q)^s} (1 - q)^{i-1} q \sum_{h=0}^{m-1} (1 - q)^{(P+C)h}. \end{aligned}$$

Now, notice that the only factor that depends on i is $(1 - q)^i$. The result follows from the observation that $(1 - q)^{i-j} \geq (1 - q)^{P+C-1}$. \square

We now introduce the event All-act. Intuitively, All-act simply states that all users in \mathbb{T}_{sub} acted relatively recently. We'll later show that All-act happens with overwhelming probability. More precisely, All-act simply states that if we look at the last $2N^2 \log N$ actions in \mathbb{T} , ordered from the most recent to the least recent, there's a new user from \mathbb{T}_{sub} in every at most $2N \log N$ steps.

In other words, we're conditioning on the following event: Let X_i be the random variable denoting how many actions it takes for the i th user to appear after the $(i - 1)$ previous users already did. All these users must be different. We're assuming that $X_i \leq 2N \log N$ for every $i \in \{1, 2, \dots, n'\}$. More formally, we'll condition on X_i being less than some $s \geq 2N \log N$ that is a multiple of $P + C$, but we omit this since it doesn't affect the proof and introduces unnecessary notation.

Now, we analyze that the probability of v being blank conditioned on all of its ancestors up to depth d being blank and All-act.

Lemma 4.4. *Let $\varepsilon > 0$ and $d \in \mathbb{N}$ such that $\varepsilon \geq 2^{\text{depth}(v)-d+1}(P + C - 1)/N$ where v is a node at depth $\text{depth}(v) > d$. Then,*

$$\Pr[B(v) = 1 \mid \text{All-act}, \text{Anc-bl}(d)] \geq \frac{1}{2} \left(1 + (1 - \varepsilon) \frac{P}{P + C} \right).$$

Proof. We describe an experiment that will help us determine the probabilities of possible states of \mathbb{T}_{sub} .

For $i \in [n']$, let $X_i \sim \text{Geo}(q_i)$ where $q_i = \frac{n'-i+1}{N} \leq 2^{\text{depth}(v)-d}/N$. Intuitively, X_i corresponds to the number of actions that occurred after the action corresponding to X_{i-1} until a "new" member acted on the tree.

Define $Y_i := \sum_{k=1}^i X_k \bmod (P+C)$. By linearity, we have that $Y_i = X_i + Y_{i-1} \bmod (P+C)$. Intuitively, the Y_i 's help us track whether this action was a commit or a proposal. More formally, notice that if $Y_i \leq C$, then the user that acted i th from last committed, and if $C + 1 \leq Y_i \leq P + C$ the user proposed. We also have that the Y_i 's form a Markov chain, i.e. Y_i conditioned on Y_{i-1} is independent of the previous Y_k 's.

Moreover, we sample a uniformly random string of length n' elements in order to assign each leaf node to a position. For example, $\sigma = 312$ ($\sigma(1) = 3, \sigma(2) = 1, \sigma(3) = 2$) would mean that the leftmost user was the first to act, the 'central' user was last to act and the third user acted second. Notice that $\sigma(1)$ tells us the position in which v acted and $\sigma(2)$ tells us the position in which $\text{sib}(v)$ acted. For example, $\sigma(1) = 2$ would imply that the action on node v was second to last on \mathbb{T}_{sub} . Notice that σ cannot repeat elements. Now we study the probability that v is blank conditioned on the events All-act and Anc-bl(d):

$$\begin{aligned} & \Pr[B(v) = 1 \mid \text{All-act, Anc-bl}(d)] \\ &= \Pr[B(v) = 1 \mid \text{All-act, Anc-bl}(d), \sigma(1) < \sigma(2)] \Pr[\sigma(1) < \sigma(2) \mid \text{All-act, Anc-bl}(d)] \\ &+ \Pr[B(v) = 1 \mid \text{All-act, Anc-bl}(d), \sigma(1) > \sigma(2)] \Pr[\sigma(1) > \sigma(2) \mid \text{All-act, Anc-bl}(d)]. \end{aligned}$$

We have that $\Pr[\sigma(1) < \sigma(2) \mid \text{All-act, Anc-bl}(d)] = \frac{1}{2}$, as which node comes later is clearly independent of All-act, Anc-bl(d) due to symmetry. Moreover, if $\sigma(1) < \sigma(2)$, i.e. v acts later than $\text{sib}(v)$, we have that v is blank since $p(v)$ is blank as we're conditioning on Anc-bl(d). Thus, we have that

$$\begin{aligned} & \Pr[B(v) = 1 \mid \text{All-act, Anc-bl}(d)] \\ &= \frac{1}{2}(1 + \Pr[B(v) = 1 \mid \text{All-act, Anc-bl}(d), \sigma(1) > \sigma(2)]). \end{aligned}$$

For the rest of the proof, we'll prove that

$$\Pr[B(v) = 1 \mid \text{All-act, Anc-bl}(d), \sigma(1) > \sigma(2)] \geq (1 - \varepsilon) \frac{P}{P+C}$$

Before we get into the details, let's explain the intuition behind the proof. The proof rests on the observation that if $X_k \sim \text{Geo}(q)$ for small q , $X_k \bmod (P+C)$ is close to a uniform random variable. Moreover, X_k and Y_{k-1} are independent, so $Y_k = X_k + Y_{k-1} \bmod (P+C)$ is also close to being a uniform distribution. Thus, no matter what complicated event we're conditioning on, the probability that $1 + C \leq Y \leq P + C$ will be approximately $\frac{P}{P+C}$.

Let $f(y_1, y_2, \dots, y_{n'}, \sigma)$ be a predicate that returns 1 if the assignment is compatible with Anc-bl(d), $\sigma(1) > \sigma(2)$ and 0 otherwise. Let $Z = \Pr[\text{All-act, Anc-bl}(d), \sigma(1) > \sigma(2)]$. For notational simplicity, we'll write \sum instead of

$$\sum_{k=1}^{n'} \sum_{\substack{\sigma \in S_{n'} \\ \sigma(1)=k}} \sum_{y_1=1}^{P+C} \sum_{y_2=1}^{P+C} \dots \sum_{y_{k-1}=1}^{P+C} \sum_{y_{k+1}=1}^{P+C} \dots \sum_{y_{n'}=1}^{P+C}$$

in what follows. Also for notational convenience, we omit the conditioning on All-act, even though it's implicit in every probability we write. Then, we have that $\Pr[B(v) = 1 \mid \text{All-act, Anc-bl}(d), \sigma(1) > \sigma(2)]$ is equal to

$$\frac{1}{Z} \sum_{y_k=1+C}^{P+C} \sum_{y_1=1}^{P+C} f(y_1, y_2, \dots, y_{n'}, \sigma) \frac{1}{(n-1)!} \Pr[Y_1 = y_1] \prod_{l=2}^{n'} \Pr[Y_l = y_l \mid Y_{l-1} = y_{l-1}], \quad (6)$$

where we used the Markov property of the Y_k 's. Also notice that since we're interested in the probability of v being blank, we're only summing up the values in the support of Y_k that start from $C + 1$, so we are ignoring the first C values. Again, for notation simplicity, we introduce a new variable

$$\rho := \Pr[Y_1 = y_1] \prod_{l>1, l \neq k, l \neq k+1} \Pr[Y_l = y_l \mid Y_{l-1} = y_{l-1}].$$

Rearranging Equation 6, we get

$$\begin{aligned} & \frac{1}{Z} \sum \frac{1}{(n-1)!} \rho f(y_1, y_2, \dots, y_{n'}, \sigma) \\ & \cdot \sum_{y_k=1+C}^{C+P} \Pr[Y_{k+1} = y_{k+1} \mid Y_k = y_k] \Pr[Y_k = y_k \mid Y_{k-1} = y_{k-1}] \end{aligned}$$

Let's now argue why it was possible to take $f(y_1, y_2, \dots, y_{n'}, \sigma)$ out of the sum, i.e. why it's independent of $Y_{\sigma(1)}$. If $\sigma(1) < \sigma(2)$, $f(y_1, y_2, \dots, y_{n'}, \sigma) = 0$ no matter what Y_k is. If $\sigma(1) > \sigma(2)$, $\sigma(1)$ has no effect on its ancestors, so Y_k doesn't matter for $\text{Anc-bl}(d)$. Now, notice that if we instead sum all values in the support of Y_k , we have by the law of total probability that

$$\begin{aligned} 1 &= \frac{1}{Z} \sum \frac{1}{(n-1)!} \rho f(y_1, y_2, \dots, y_{n'}, \sigma) \\ & \cdot \sum_{y_k=1}^{P+C} \Pr[Y_{k+1} = y_{k+1} \mid Y_k = y_k] \Pr[Y_k = y_k \mid Y_{k-1} = y_{k-1}] \end{aligned}$$

Moreover, using the Markov property of the Y_i , we have that

$$\begin{aligned} & \Pr[Y_{k+1} = y_{k+1} \mid Y_{k-1} = y_{k-1}] \\ &= \sum_{y_k=1}^{P+C} \Pr[Y_{k+1} = y_{k+1} \mid Y_k = y_k, Y_{k-1} = y_{k-1}] \Pr[Y_k = y_k \mid Y_{k-1} = y_{k-1}] \\ &= \sum_{y_k=1}^{P+C} \Pr[Y_{k+1} = y_{k+1} \mid Y_k = y_k] \Pr[Y_k = y_k \mid Y_{k-1} = y_{k-1}] \end{aligned} \tag{7}$$

Thus, it suffices to show that

$$\begin{aligned} & \sum_{y_k=1+C}^{P+C} \Pr[Y_{k+1} = y_{k+1} \mid Y_k = y_k] \Pr[Y_k = y_k \mid Y_{k-1} = y_{k-1}] \\ & \geq (1 - \varepsilon) \frac{P}{P+C} \Pr[Y_{k+1} = y_{k+1} \mid Y_{k-1} = y_{k-1}]. \end{aligned}$$

Let $t_{max} := \arg \max_t \Pr[Y_{k+1} = y_{k+1} \mid Y_k = t] \Pr[Y_k = t \mid Y_{k-1} = y_{k-1}]$. Notice that

$$\begin{aligned} & \Pr[Y_{k+1} = y_{k+1} \mid Y_k = t_{max}] \Pr[Y_k = t_{max} \mid Y_{k-1} = y_{k-1}] \\ & \geq \frac{1}{P+C} \Pr[Y_{k+1} = y_{k+1} \mid Y_{k-1} = y_{k-1}] \end{aligned} \tag{8}$$

which follows from Equation 7. Then, by Lemma 4.3 (observe that X_k is bounded because of the event All-act), for any $1+C \leq t \leq P+C$, we have that

$$\begin{aligned} & \Pr[Y_{k+1} = y_{k+1} \mid Y_k = t] \Pr[Y_k = t \mid Y_{k-1} = y_{k-1}] \\ & \geq \Pr[Y_{k+1} = y_{k+1} \mid Y_k = t_{max}] \Pr[Y_k = t_{max} \mid Y_{k-1} = y_{k-1}] (1-q)^{2(P+C-1)} \\ & \geq (1-2q(P+C-1)) \frac{1}{P+C} \Pr[Y_{k+1} = y_{k+1} \mid Y_{k-1} = y_{k-1}] \\ & \geq (1-\varepsilon) \frac{1}{P+C} \Pr[Y_{k+1} = y_{k+1} \mid Y_{k-1} = y_{k-1}] \end{aligned}$$

where the second inequality follows from Lemma 2.2 and Equation 8. The last inequality follows from the assumption that $\varepsilon \geq 2^{\text{depth}(v)-d+1}(P+C-1)/N$ and the fact that $q \leq 2^{\text{depth}(v)-d}/N$. Thus,

$$\begin{aligned} & \sum_{t=1+C}^{P+C} \Pr[Y_{k+1} = y_{k+1} \mid Y_k = t] \Pr[Y_k = t \mid Y_{k-1} = y_{k-1}] \\ & \geq (1 - \varepsilon) \frac{P}{P+C} \Pr[Y_{k+1} = y_{k+1} \mid Y_{k-1} = y_{k-1}]. \end{aligned}$$

We thus conclude the proof. \square

We have two more lemmata left to prove before moving onto combining our results.

Lemma 4.5. *Assuming at least $2N^2 \log(N)$ actions took place, $\Pr[\text{All-act}] \geq 1 - \frac{1}{N}$.*

Proof. Notice that for any $i \in \{1, \dots, N\}$,

$$\Pr[X_i > 2N \log N] \leq \left(1 - \frac{1}{N}\right)^{2N \log N} \leq \exp(-2 \log(N)) = N^{-2}.$$

Then, $\Pr[\text{All-act}] \geq 1 - \frac{1}{N}$ by the union bound over all the X_i . \square

Lemma 4.6. *Let $\varepsilon > 0$ and $d \in \mathbb{N}$ such that $\varepsilon \geq 2^{-d}(P+C-1)$ and consider a node v such that $\log(N) - 1 \geq \text{depth}(v) > d$. Then, we have that*

$$\begin{aligned} & \mathbb{E}[|\text{Res}(v)| \mid \text{All-act, Anc-bl}(d)] \\ & \geq (1 + (1 - \varepsilon) \frac{P}{P+C}) \mathbb{E}[|\text{Res}(\text{left}(v))| \mid \text{All-act, Anc-bl}(d, \text{left}(v))]. \end{aligned}$$

Proof. Observe that $\varepsilon \geq 2^{\text{depth}(v)-d+1}(P+C-1)/N$ since $\log(N) - 1 \geq \text{depth}(v)$. By Lemma 4.4, we have that

$$\begin{aligned} & \mathbb{E}[|\text{Res}(v)| \mid \text{All-act, Anc-bl}(d, v)] \\ & \geq \mathbb{E}[|\text{Res}(v)| \mid \text{All-act, Anc-bl}(d, v), B(v) = 1] \\ & \quad \cdot \Pr[B(v) = 1 \mid \text{All-act, Anc-bl}(d, v)] \\ & \geq \mathbb{E}[|\text{Res}(v)| \mid \text{All-act, Anc-bl}(d, v), B(v) = 1] \frac{1}{2} \left(1 + (1 - \varepsilon) \frac{P}{P+C}\right) \\ & = 2 \mathbb{E}[|\text{Res}(\text{left}(v))| \mid \text{All-act, Anc-bl}(d, v), B(v) = 1] \frac{1}{2} \left(1 + (1 - \varepsilon) \frac{P}{P+C}\right) \\ & = \mathbb{E}[|\text{Res}(\text{left}(v))| \mid \text{All-act, Anc-bl}(d, \text{left}(v))] \left(1 + (1 - \varepsilon) \frac{P}{P+C}\right). \end{aligned}$$

\square

We are now ready to prove the section's main theorem.

Proof of Theorem 4.2. Let v be any node at depth $d(v) = d$ and let $\varepsilon > 0$ be a small constant such that $\varepsilon \geq 2^{-d}(P+C-1)$. The fact that such d and ε exist follows from the assumption that $P+C = O(1)$. We show that the expected resolution size of v is $\Omega(N^{\log_2(1+(1-\varepsilon)\frac{P}{P+C})})$. Since any commit will have a node of said depth in its copath, this is sufficient for the proof. We start with the simple observation that

$$\begin{aligned} \mathbb{E}[|\text{Res}(v)|] & \geq \mathbb{E}[|\text{Res}(v)| \mid \text{All-act}, B(v) = 1] \Pr[\text{All-act}, B(v) = 1] \\ & = 2 \mathbb{E}[|\text{Res}(\text{left}(v))| \mid \text{All-act, Anc-bl}(d, \text{left}(v))] \Pr[\text{All-act}, B(v) = 1]. \end{aligned}$$

We use Lemma 4.5 in order to obtain that

$$\begin{aligned} \Pr[\text{All-act}, B(v) = 1] &\geq \Pr[B(v) = 1] - \Pr[\neg\text{All-act}] \\ &\geq \Pr[B(v) = 1] - \frac{1}{N}. \end{aligned}$$

It follows from repeatedly applying Lemmata 4.6 and 4.1 that

$$\begin{aligned} \mathbb{E}[|\text{Res}(v)|] &\geq 2(\Pr[B(v) = 1] - \frac{1}{N})\mathbb{E}[|\text{Res}(\text{left}(v))| \mid \text{All-act}, \text{Anc-bl}(d, \text{left}(v))] \\ &\geq 2(\Pr[B(v) = 1] - \frac{1}{N})(1 + (1 - \varepsilon)\frac{P}{P+C})^{\log_2 N - d(\text{left}(v))} \\ &\geq \Omega\left(\left(\frac{\left(1 - \frac{\binom{N(1-2^{-d})}{P}\right) \binom{N(1-2^{-d})}{C}}{1 - \frac{\binom{N(1-2^{-d})}{C} \binom{N(1-2^{-d})}{P}}\right) - \frac{1}{N}}{1 + (1 - \varepsilon)\frac{P}{P+C}}\right)^{\log_2 N - d}\right). \end{aligned}$$

We make use of Equation 4 and the expression above simplifies to $\mathbb{E}[\text{Cost}(t)] \geq \Omega(N^{\log_2(1+(1-\varepsilon)\frac{P}{P+C})})$, as desired. \square

The proof above does not explicitly state the constant that appears in the lower bound of $\mathbb{E}[\text{Cost}(t)] = \Omega(N^{\log_2(1+(1-\varepsilon)\frac{P}{P+C})})$. However, one can obtain an explicit constant from the proof. For instance if $P = 1 = C$, one can take $\varepsilon = 1/8$ and then for a node v_3 at depth 3 and a node v_4 at depth 4, the proof shows that

$$\begin{aligned} \mathbb{E}[|\text{Res}(v_3)|] &\geq 2 \cdot 1/2 \cdot (1 + (1 - 1/8) \cdot 1/2)^{\log_2 N - 3 - 1} \\ \mathbb{E}[|\text{Res}(v_4)|] &\geq 2 \cdot 1/2 \cdot (1 + (1 - 1/8) \cdot 1/2)^{\log_2 N - 4 - 1}. \end{aligned}$$

This implies that

$$\begin{aligned} \mathbb{E}[\text{Cost}(t)] &\geq \mathbb{E}[|\text{Res}(v_3)|] + \mathbb{E}[|\text{Res}(v_4)|] \\ &\cong 0.39 \cdot N^{0.52}. \end{aligned}$$

For small values of N the previously shown bound of $\log^2(N)/4 - \log(N)/4$ is more relevant.

5 MLS-Cutoff: an Alternative Method for Update Proposals

We introduce MLS-Cutoff a protocol variant of MLS that achieves better communication complexity under random sequences of operations as defined in experiment Exp-Prop-Com. As discussed in the previous section, blanking the path of users issuing an update proposals already for a small constant number of proposals leads to a constant fraction of the tree's node being blank on expectation, which in turn substantially negatively impacts the size of commit messages.

To prevent introducing blanks in the tree, one could have users perform a `rekey-path` operation with every update proposal. In fact, this is the approach taken by the continuous group-key agreement scheme CoCoA [AAN⁺22]. However, taking this approach has two downsides, both due to the fact that proposals can be issued concurrently, meaning users make concurrent modifications to the ratchet tree.

1. If two users u_1, u_2 rekey their path concurrently, u_1 will encrypt a path secret s_j under a public known to u_2 that is supposed to be replaced by u_2 's rekey operation. As a consequence CoCoA suffers from slower PCS. In MLS, if every compromised user issues an update proposal and another user commits to the proposals, then the new application secret will be secure. In CoCoA, on the other hand, it can take up to $\text{depth}(T)$ rounds to recover from compromise.

2. Concurrently rekeying paths prevents users to compute and sign parent hashes, as they cannot anticipate which keys will be placed in the sub-trees rooted at their co-path nodes. (Recall that the parent-hash value of node v is computed with respect to the tree-hash of its sibling $\text{sib}(v)$.) As a consequence, CoCoA only achieves semi-active but neither active [ACJM20] nor insider security [AJM22].

We observe that the issues above only occur from the point on where u_1 and u_2 's update paths meet. Further, if the number of issued update proposals is small, with high probability the users' update paths merge high up in the tree. Our protocol MLS-Cutoff makes use of this observation by having update proposals only re-key the path up to a certain point. Looking ahead, proposals do not affect the $\log(\text{depth}(\mathbb{T}))$ topmost levels of the tree. Still, when applying the proposals it can be the case that a node is affected by more than one rekey operation if the issuing users are situated close enough in the tree. In this case we have to resort to blanking the node. This turns out to be sufficient to prevent the two issues outlined above.

5.1 Protocol Description

In this section we discuss the modifications MLS-Cutoff makes to the MLS protocol and defer its formal description to Appendix B.2. MLS-Cutoff differs in two aspects from MLS, namely how update proposal messages pmsg are generated, and how proposals are applied to the ratchet tree using `apply-props`.

The protocol is parameterized by a cutoff parameter i_{cut} . `CGKA.Prop` works the same as before for removals or additions of users. When issuing an update, however, the user creates a copy \mathbb{T}' of the ratchet tree and calls the path rekeying function (see Figure 3) to obtain

$$(\mathbb{T}', \text{UpdPath}) \leftarrow \text{rekey-path}(v_u, i_{\text{cut}}).$$

The proposal message $\text{pmsg} = (\text{'upd'}, \text{UpdPath})$ is framed and communicated to the rest of the group, and \mathbb{T}' stored as a pending update.

Commit operations work in the same way as in MLS, except for using a modified subroutine `apply-props` in step 2a. For an overview on `apply-props` see Figure 5). The function modifies the copied tree \mathbb{T}' and processes removals and additions in the same way as MLS, i.e., by blanking the leaf and path of removed users and adding added users as unmerged leaves. Update proposals $(\text{'upd'}, \text{UpdPath})$ from issuing user $u\text{-snd}$ are processed by first updating the keys of affected nodes. More precisely, the public keys $\text{pk}_0, \dots, \text{pk}_{\ell-i_{\text{cut}}}$ contained in `UpdPath` are used to replace the ones on $u\text{-snd}$'s update path $(v_1, \dots, v_\ell) = \text{path}(v_0)$, where $v_0 = v_{u\text{-snd}}$. Further, the processing user u recovers all secret keys they should have access to, i.e., the ones from the least common ancestor v_j of v_u and $v_{u\text{-snd}}$ up to $v_{\ell-i_{\text{cut}}}$ (if the v_j is above $v_{\ell-i_{\text{cut}}}$ no secret keys are recovered), and assigns them to the corresponding nodes. This is abstracted as helper function `recover-sks` in Figure 5 that takes as input \mathbb{T}' and ciphertext collection C_j and returns the updated view of \mathbb{T}' with replaced secret keys.

After the keys of nodes affected by pmsg have been updated, `apply-props` blanks nodes if necessary. I.e., all nodes $(v_{\ell-i_{\text{cut}}+1}, \dots, v_\ell)$ on $u\text{-snd}$'s update path after the cutoff point are blanked. Further, nodes affected by at least two update proposals are blanked as well. To this end the function keeps track of a set V_{col} containing all nodes the keys of which have been replaced so far, that is updated at the end of processing pmsg .

After having executed the alternative version of `apply-props` in step 2a the commit algorithm `CGKA.Commit` proceeds in the same way as MLS. In particular, in step 2b the committing user issues a *full* path rekey operation `rekey-path(u, 0)`, which establishes a new root secret.

Commits to a collection of proposals are processed in the analog way to MLS, i.e., the only difference being the use of the modified `apply-props` function.

5.2 Security

MLS-Cutoff achieves the same level of security as MLS. More formally in Appendix B.2 we prove the following theorem

```

Algorithm apply-props( $T'$ , PMSG)
00 if exists  $((\text{'upd'}, \text{ad}), u\text{-own}) = \text{pmsg} \in \text{PMSG}$             $\llcorner$  own proposal
01    $T' \leftarrow \text{pendUpd}(\text{pmsg})$                                 $\llcorner$  recover corresponding tree
02    $V_{\text{col}} \leftarrow \emptyset$                                     $\llcorner$  set collecting potential collisions
03 for  $\text{pmsg} \in \text{PMSG}$ :
04   parse  $((\text{'type'}, \text{ad}), u\text{-snd}) \leftarrow \text{pmsg}$ 
05   if  $(\text{'type'}, \text{ad}) = (\text{'upd'}, \text{UpdPath})$ 
06     if  $u\text{-snd} = u\text{-own}$                                     $\llcorner$  own proposal already handled
07       skip to next pmsg
08      $\llcorner$  rekey nodes
09     parse  $(\text{pk}_0, (\text{pk}_1, C_1), \dots, (\text{pk}_{\ell-i_{\text{cut}}}, C_{\ell-i_{\text{cut}}})) \leftarrow \text{UpdPath}$ 
10      $v_0 \leftarrow v_{u\text{-snd}}$ 
11      $(v_1, \dots, v_\ell) \leftarrow \text{path}(v_0)$ 
12     for  $j = 0, \dots, \ell - i_{\text{cut}}$ 
13        $v_j.\text{pk} \leftarrow \text{pk}_j$ 
14       if  $v_j = \text{lca}(v_{u\text{-own}}, v_{u\text{-snd}})$                   $\llcorner$  least common ancestor
15          $T' \leftarrow \text{recover-sks}(T', C_j)$                   $\llcorner$  recover secret keys
16      $\llcorner$  blank nodes
17     for  $j \in \{1, \dots, \ell - i_{\text{cut}}\}$ 
18       if  $v_j \in V_{\text{col}}$                                       $\llcorner$  blank collisions
19         blank $(v_j)$ 
20     for  $j \in \{\ell - i_{\text{cut}} + 1, \dots, \ell\}$               $\llcorner$  blank remainder of path
21       blank $(v_j)$ 
22      $V_{\text{col}} \leftarrow \{v_1, \dots, v_\ell\}$                   $\llcorner$  update collision set
23      $\llcorner$  process removes, adds as before
24     else if  $(\text{'type'}, \text{ad}) = (\text{'rem'}, u\text{-rem})$ :
25       blank-leaf $(T', v_{u\text{-rem}})$                           $\llcorner$  blank  $u\text{-rem}$ 's leaf
26       for  $v \in \text{path}(v_{u\text{-rem}})$                           $\llcorner$  blank  $u\text{-rem}$ 's path
27         blank $(v)$ 
28        $T' \leftarrow \text{truncate}(T')$                           $\llcorner$  truncate tree if necessary
29     else if  $(\text{'type'}, \text{ad}) = (\text{'add'}, (u\text{-add}, \text{pk}))$ :
30       add-leaf $(T', u\text{-add}, \text{pk})$                           $\llcorner$  assign  $u\text{-rem}$  leaf
31       for  $v \in \text{path}(v_{u\text{-add}})$                           $\llcorner$  set  $v_{u\text{-add}}$  as unmerged
32          $v.\text{unm} \leftarrow \cup \{v_{u\text{-add}}\}$ 
33 return  $T'$ 

```

Figure 5: Applying proposal messages in MLS-Cutoff. The function’s formal description is in Figures 23. We assume PMSG is ordered with update proposals followed by removals, followed by adds.

Theorem 5.1. *If Pke is IND-CCA2 secure, Sig is SEUF-CMA secure and calls to H_1, H_2, H and Mac are replaced by calls to a random oracle \mathcal{G} , then MLS-Cutoff securely realizes $(\mathcal{F}_{\text{AS}}^1, \mathcal{F}_{\text{KS}}^1, \mathcal{F}_{\text{CGKA}})$ in the $(\mathcal{F}_{\text{AS}}, \mathcal{F}_{\text{KS}}, \mathcal{G})$ -hybrid model.*

We provide some intuition on why the result holds. As MLS-Cutoff closely resembles MLS we are able to largely rely of the security proof of [AJM22]. The recent work of Cremers et al. [CGWZ25] shows that the signature scheme used must be SEUF-CMA secure rather than just EUF-CMA secure as originally stated in [AJM22], though their security proof already made use of SEUF-CMA security implicitly. We note that while in MLS-Cutoff update proposals lead to modifications to ratchet tree nodes beyond users’ personal leaves, it preserves the following property. For an update issuing user u denote their leaf by v_0 , consider their update path (v_1, \dots, v_ℓ) , and assume that the keys of nodes v_1, \dots, v_k have been replaced after processing a commit to the update proposal and v_{k+1}, \dots, v_ℓ have been blanked. Then all sub-trees rooted at the nodes $\text{sib}(v_0), \dots, \text{sib}(v_k)$ must be in the same state as they were when the update proposal was issued. Indeed, would this not be the case for one of the nodes, then it would have caused a collision of update proposals and been blanked. This has the consequence that, on the one hand, keys being added to the tree cannot be encrypted under a public key that was concurrently being replaced, and, on the other hand, that the added nodes are compatible with the parent hash mechanism.

5.3 Upper Bounds on the Update Cost

We compute an upper bound on the communication complexity of MLS-Cutoff for sequences of operations as produces in experiment `Exp-Prop-Com`. We focus on the case $C = 1$, as the performance for $C \geq 1$ is even better.

Expected number of blanks in MLS-Cutoff's ratchet tree. Let $N = 2^n$, $t \in \mathbb{N}$, $C = 1$, and let $P \in \mathbb{N}$ be constant. Consider an execution of `Exp-Prop-Com`(n, P, C, t) with respect to CGKA MLS-Cutoff. As in Section 4.2 consider the sequence $(\mathbb{T}_i)_{i=1}^t$ of ratchet trees generated after the i th application of `Exp-Prop-Com`'s round function. For a fixed internal node v we can define a random variable $(B_i(v))_{i=0}^t$ taking values in $\{0, 1\}$ where 0 corresponds to being a populated node and 1 to v being blank. We write B_i instead of $B_i(v)$ when there is no ambiguity. We note that all users have filtered paths of length n .

Now consider a node v below the cutoff bound i_{cut} , meaning that v 's dept $d = \text{depth}(v) \geq i_{\text{cut}}$, which implies that v is re-keyed if one of its descendant leaves issues an update proposal. We observe that, again, $(B_i)_i$ forms a Markov chain, since the node switching from blank or vice versa only depends on the question, whether it is affected by update proposals and/or commits.

More precisely v moves from blank to populated if it is affected by the commit or by exactly one update proposal, and from populated to blank if it is not affected by the commit and at least two update proposals.

$$\begin{aligned} p_{1 \rightarrow 0} &= \Pr[v \in \text{ancs}(U_C) \vee |\{u \in U_P : v \in \text{ancs}(v_u)\}| = 1] \\ p_{0 \rightarrow 1} &= \Pr[v \notin \text{ancs}(U_C) \wedge |\{u \in U_P : v \in \text{ancs}(v_u)\}| \geq 2]. \end{aligned}$$

Thus, we have

$$\begin{aligned} p_{1 \rightarrow 0} &= \Pr[v \in \text{ancs}(U_C)] + \Pr[|\{u \in U_P : v \in \text{ancs}(v_u)\}| = 1] \\ &\quad - \Pr[v \in \text{ancs}(U_C) \wedge |\{u \in U_P : v \in \text{ancs}(v_u)\}| = 1] \\ &= 1 - \frac{\binom{N(1-2^{-d})}{C}}{\binom{N}{C}} + P(1-2^{-d})^{P-1}2^{-d} - \left(1 - \frac{\binom{N(1-2^{-d})}{C}}{\binom{N}{C}}\right)P(1-2^{-d})^{P-1}2^{-d} \\ &= 1 + (P(1-2^{-d})^{P-1}2^{-d} - 1) \frac{\binom{N(1-2^{-d})}{C}}{\binom{N}{C}}, \end{aligned}$$

and

$$\begin{aligned} p_{0 \rightarrow 1} &= \Pr[v \notin \text{ancs}(U_C)] \Pr[|\{u \in U_P : v \in \text{ancs}(v_u)\}| \geq 2] \\ &= \frac{\binom{N(1-2^{-d})}{C}}{\binom{N}{C}} (1 - (1-2^{-d})^P - P2^{-d}(1-2^{-d})^{P-1}). \end{aligned}$$

Observe that for $C = 1$ and d sufficiently large this simplifies to

$$p_{1 \rightarrow 0} \approx 1 - (1 - 2^{-d}) = 2^{-d}$$

and

$$p_{0 \rightarrow 1} \approx 1 - (1 - 2^{-d})^P - P2^{-d}(1 - 2^{-d})^{P-1} \approx P(P - 1)2^{-2d}$$

By Equation 1 we obtain that the probability of v being blank in the stationary distribution (when $d = \text{depth}(v) \geq i_{\text{cut}}$) is given by

$$\mathbb{E}[B(v) = 1] = \frac{p_{0 \rightarrow 1}}{p_{1 \rightarrow 0} + p_{0 \rightarrow 1}} \approx \frac{P^2 2^{-2d}}{2^{-d} + P^2 2^{-2d}} \approx \frac{P^2 2^{-d}}{1 + P^2 2^{-d}} \in O\left(\frac{P^2}{2^d}\right). \quad (9)$$

As a consequence, the ratchet tree exhibits on expectation a constant amount of about P^2 blank nodes on level d .

Now consider a node v above the cutoff bound i_{cut} , i.e., $d = \text{depth}(v) < i_{\text{cut}}$. This implies that v changes from populated to blank if and only if (at least) one of its descendant leaves issues an update proposal and none of them issue a commit, whereas v changes from populated to blank if and only if at least one of them issues a commit. Using similar arguments to previous ones, we obtain that

$$p_{1 \rightarrow 0} = \Pr[v \in \text{ancs}(U_C)] = 1 - \frac{\binom{N(1-2^{-d})}{C}}{\binom{N}{C}}$$

and

$$p_{0 \rightarrow 1} = \Pr[v \notin \text{ancs}(U_C) \wedge v \in \text{ancs}(U_P)] = \frac{\binom{N(1-2^{-d})}{C}}{\binom{N}{C}} \left(1 - \frac{\binom{N(1-2^{-d})}{P}}{\binom{N}{P}} \right).$$

If $C = 1$ this simplifies to

$$\begin{aligned} p_{1 \rightarrow 0} &= 2^{-d}, \\ p_{0 \rightarrow 1} &\cong (1 - 2^{-d})(1 - (1 - 2^{-d})^P) \cong (1 - 2^{-d})P2^{-d}. \end{aligned}$$

By Equation 1 we obtain that the probability of v being blank in the stationary distribution (when $d = \text{depth}(v) < i_{\text{cut}}$) is given by

$$\mathbb{E}[B(v) = 1] = \frac{p_{0 \rightarrow 1}}{p_{1 \rightarrow 0} + p_{0 \rightarrow 1}} \cong \frac{(1 - 2^{-d})P2^{-d}}{2^{-d} + (1 - 2^{-d})P2^{-d}} \in \left[0, \frac{P}{1 + P} \right]. \quad (10)$$

Bounding $|\text{Res}(v)|$. The size of the resolution of a node v can be bounded by the number of blank nodes in the subtree rooted at v plus one. If $C = 1$, we use this simple fact and Equations 9 and 10 in order to obtain that for nodes of $d = \text{depth}(v) \geq i_{\text{cut}}$

$$\begin{aligned} \mathbb{E}[|\text{Res}(v)|] &\leq 1 + \sum_{d'=\log(N)}^{d+1} 2^{d'-d} \frac{P(P-1)2^{-d'}}{1 + P(P-1)2^{-d'}} \\ &\leq 1 + P(P-1)2^{-d}(\log(N) - d) \end{aligned}$$

and for nodes of $d = \text{depth}(v) < i_{\text{cut}}$

$$\begin{aligned} \mathbb{E}[|\text{Res}(v)|] &\leq 1 + \sum_{d'=\log(N)}^{i_{\text{cut}}} 2^{d'-d} \frac{P(P-1)2^{-d'}}{1 + P(P-1)2^{-d'}} + \sum_{d'=i_{\text{cut}}-1}^{d+1} 2^{d'-d} \frac{P}{1 + P} \\ &\leq 1 + P(P-1)2^{-d}(\log(N) - i_{\text{cut}}) + 2^{-d} \frac{2^d - 2^{i_{\text{cut}}}}{1 - 2} \\ &\leq P(P-1)2^{-d}(\log(N) - i_{\text{cut}}) + 2^{i_{\text{cut}}-d} \end{aligned}$$

Bounding $|\text{pmsg}|$ and $|\text{cmsg}|$. To compute a bound on the communication complexity of MLS-Cutoff for random sequences of operations, recall that by definition of $\text{rekey-path}(\mathbb{T}', \text{id}, i)$, if the $v_0 = v_{\text{id}}$ and $(v_1, \dots, v_\ell) = \text{path}(v_0)$, the amount of ciphertexts is given by

$$|\text{CTXT}| = \sum_{j=0}^{\ell - i_{\text{cut}} - 1} |\text{Res}(\text{sib}(v_j))|.$$

We make use the previous bounds on the size of the resolution of a node and obtain that

$$\begin{aligned}
\mathbb{E}[|\text{CTXT}|] &= \sum_{d=1}^{\log(N)} \mathbb{E}[|\text{Res}(v_d)|] \\
&\leq \sum_{d=1}^{i_{\text{cut}}-1} (P(P-1)2^{-d}(\log(N) - i_{\text{cut}}) + 2^{i_{\text{cut}}-d}) \\
&\quad + \sum_{d=i_{\text{cut}}}^{\log(N)} (1 + P(P-1)2^{-d}(\log(N) - d)) \\
&\leq P(P-1)(\log(N) - i_{\text{cut}}) + 2^{i_{\text{cut}}} + \log(N) - i_{\text{cut}} + 1.
\end{aligned}$$

If we set $i_{\text{cut}} = \log(n) = \log(\log(N))$, we obtain that

$$\mathbb{E}[|\text{CTXT}|] \leq \log(N) + (P(P-1) + 1) \log(N/\log(N)) + 1.$$

We have shown the following.

Lemma 5.2. *Let $C = 1$, and let $P \in \mathbb{N}$ be constant. Let t be sufficiently large. For $N = 2^n$ set $i_{\text{cut}} = \log(n)$. Then, on expectation, proposal messages `pmsg` and commit messages `cmsg` output by MLS-Cutoff for sequences of operations as defined by `Exp-Prop-Com` satisfy*

$$\mathbb{E}[|\text{pmsg}|] \leq \mathbb{E}[|\text{cmsg}|] \in O(\log(N)).$$

We point out that the amount of proposals per commit needs to be small compared to the group size N in order to guarantee a logarithmic communication complexity. Observe that Lemma 5.2 gives a bound that is asymptotic in N while P is assumed to be constant. If P becomes too large compared to the group size, the probability of blanks being introduced in the ratchet tree grows. In the extreme case of all users issuing update proposals, for example, the whole tree is blanked just as in MLS. We point out, that a linear communication complexity is, however, inherent for this type of ‘massive’ concurrency (Compare, e.g., the lower bounds of [BDR20]). As a second point, experiment `Exp-Prop-Com` does not remove users from the group, which would introduce blanks as well. However, observe that these blanks would be removed at a faster pace in MLS-Cutoff compared to MLS due to the additional path rekeying in update proposals.

Summing up, we see MLS-Cutoff as an attractive alternative to how update proposals are handled in MLS, that provides the same security properties, but has an improved communication complexity for moderate numbers of concurrent update proposals. We consider it an interesting open question to evaluate its performance in more realistic models.

6 Acknowledgment

We are very grateful to Guillermo Pascual Perez for the helpful conversations on MLS, its different versions and its consistency mechanisms.

References

- [AAB⁺21] Joël Alwen, Benedikt Auerbach, Mirza Ahad Baig, Miguel Cueto Noval, Karen Klein, Guillermo Pascual-Perez, Krzysztof Pietrzak, and Michael Walter. Grafting key trees: Efficient key management for overlapping groups. In Kobbi Nissim and Brent Waters, editors, *TCC 2021, Part III*, volume 13044 of *LNCS*, pages 222–253. Springer, Cham, November 2021.
- [AAB⁺24] Michael Anastos, Benedikt Auerbach, Mirza Ahad Baig, Miguel Cueto Noval, Matthew Kwan, Guillermo Pascual-Perez, and Krzysztof Pietrzak. The cost of maintaining keys in dynamic groups with applications to multicast encryption and group messaging. In Elette Boyle and Mohammad Mahmoody, editors, *TCC 2024, Part I*, volume 15364 of *LNCS*, pages 413–443. Springer, Cham, December 2024.
- [AACN⁺24] Joël Alwen, Benedikt Auerbach, Miguel Cueto Noval, Karen Klein, Guillermo Pascual-Perez, and Krzysztof Pietrzak. Decaf: Decentralizable cgka with fast healing. In Clemente Galdi and Duong Hieu Phan, editors, *Security and Cryptography for Networks*, pages 294–313, Cham, 2024. Springer Nature Switzerland.
- [AAN⁺22] Joël Alwen, Benedikt Auerbach, Miguel Cueto Noval, Karen Klein, Guillermo Pascual-Perez, Krzysztof Pietrzak, and Michael Walter. CoCoA: Concurrent continuous group key agreement. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 815–844. Springer, Cham, May / June 2022.
- [ACD19] Joël Alwen, Sandro Coretti, and Yevgeniy Dodis. The double ratchet: Security notions, proofs, and modularization for the Signal protocol. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 129–158. Springer, Cham, May 2019.
- [ACDT20] Joël Alwen, Sandro Coretti, Yevgeniy Dodis, and Yiannis Tselekounis. Security analysis and improvements for the IETF MLS standard for group messaging. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part I*, volume 12170 of *LNCS*, pages 248–277. Springer, Cham, August 2020.
- [ACDT21] Joël Alwen, Sandro Coretti, Yevgeniy Dodis, and Yiannis Tselekounis. Modular design of secure group messaging protocols and the security of MLS. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 1463–1483. ACM Press, November 2021.
- [ACJM20] Joël Alwen, Sandro Coretti, Daniel Jost, and Marta Mularczyk. Continuous group key agreement with active security. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part II*, volume 12551 of *LNCS*, pages 261–290. Springer, Cham, November 2020.
- [AHKM22] Joël Alwen, Dominik Hartmann, Eike Kiltz, and Marta Mularczyk. Server-aided continuous group key agreement. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 69–82. ACM Press, November 2022.
- [AJM22] Joël Alwen, Daniel Jost, and Marta Mularczyk. On the insider security of MLS. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 34–68. Springer, Cham, August 2022.
- [AMT23] Joël Alwen, Marta Mularczyk, and Yiannis Tselekounis. Fork-resilient continuous group key agreement. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part IV*, volume 14084 of *LNCS*, pages 396–429. Springer, Cham, August 2023.
- [ANPPP23] Benedikt Auerbach, Miguel Cueto Noval, Guillermo Pascual-Perez, and Krzysztof Pietrzak. On the cost of post-compromise security in concurrent continuous group-key agreement. In Guy N. Rothblum and Hoeteck Wee, editors, *TCC 2023, Part III*, volume 14371 of *LNCS*, pages 271–300. Springer, Cham, November / December 2023.

- [BBR18] Karthikeyan Bhargavan, Richard Barnes, and Eric Rescorla. TreeKEM: Asynchronous Decentralized Key Management for Large Dynamic Groups. <https://mailarchive.ietf.org/arch/attach/mls/pdf1XUH6o.pdf>, May 2018.
- [BBR⁺23] Richard Barnes, Benjamin Beurdouche, Raphael Robert, Jon Millican, Emad Omara, and Katriel Cohn-Gordon. The Messaging Layer Security (MLS) Protocol. RFC 9420, July 2023.
- [BCG23] David Balbás, Daniel Collins, and Phillip Gajland. WhatsApp with sender keys? Analysis, improvements and security proofs. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part V*, volume 14442 of *LNCS*, pages 307–341. Springer, Singapore, December 2023.
- [BCK21] Chris Brzuska, Eric Cornelissen, and Konrad Kohbrok. Cryptographic security of the MLS RFC, draft 11. Cryptology ePrint Archive, Report 2021/137, 2021.
- [BCV23] David Balbás, Daniel Collins, and Serge Vaudenay. Cryptographic administration for secure group messaging. In Joseph A. Calandrino and Carmela Troncoso, editors, *USENIX Security 2023*, pages 1253–1270. USENIX Association, August 2023.
- [BDG⁺22] Alexander Bienstock, Yevgeniy Dodis, Sanjam Garg, Garrison Grogan, Mohammad Hajiabadi, and Paul Rösler. On the worst-case inefficiency of CGKA. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022, Part II*, volume 13748 of *LNCS*, pages 213–243. Springer, Cham, November 2022.
- [BDR20] Alexander Bienstock, Yevgeniy Dodis, and Paul Rösler. On the price of concurrency in group ratcheting protocols. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part II*, volume 12551 of *LNCS*, pages 198–228. Springer, Cham, November 2020.
- [Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145. IEEE Computer Society Press, October 2001.
- [CDPW07] Ran Canetti, Yevgeniy Dodis, Rafael Pass, and Shabsi Walfish. Universally composable security with global setup. In Salil P. Vadhan, editor, *TCC 2007*, volume 4392 of *LNCS*, pages 61–85. Springer, Berlin, Heidelberg, February 2007.
- [CEST24] Kelong Cong, Karim Eldefrawy, Nigel P. Smart, and Ben Terner. The key lattice framework for concurrent group messaging. In Christina Pöpper and Lejla Batina, editors, *ACNS 24 International Conference on Applied Cryptography and Network Security, Part II*, volume 14584 of *LNCS*, pages 133–162. Springer, Cham, March 2024.
- [CGWZ25] Cas Cremers, Esra Günsay, Vera Wesselkamp, and Mang Zhao. ETK: External-operations TreeKEM and the security of MLS in RFC 9420. Cryptology ePrint Archive, Paper 2025/229, 2025.
- [CHK21] Cas Cremers, Britta Hale, and Konrad Kohbrok. The complexities of healing in secure group messaging: Why cross-group effects matter. In Michael Bailey and Rachel Greenstadt, editors, *USENIX Security 2021*, pages 1847–1864. USENIX Association, August 2021.
- [HKP⁺21] Keitaro Hashimoto, Shuichi Katsumata, Eamonn Postlethwaite, Thomas Prest, and Bas Westerbaan. A concrete treatment of efficient continuous group key agreement via multi-recipient PKEs. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 1441–1462. ACM Press, November 2021.
- [HKP22] Keitaro Hashimoto, Shuichi Katsumata, and Thomas Prest. How to hide metadata in mls-like secure group messaging: Simple, modular, and post-quantum. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, CCS '22*, page 1399–1412, New York, NY, USA, 2022. Association for Computing Machinery.

- [KPPW⁺21] K. Klein, G. Pascual-Perez, M. Walter, C. Kamath, M. Capretto, M. Cueto, I. Markov, M. Yeo, J. Alwen, and K. Pietrzak. Keep the Dirt: Tainted TreeKEM, Adaptively and Actively Secure Continuous Group Key Agreement. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 268–284, Los Alamitos, CA, USA, may 2021. IEEE Computer Society.
- [Wei19] Matthew A. Weidner. Group Messaging for Secure Asynchronous Collaboration. Master’s thesis, University of Cambridge, June 2019.
- [WHA99] D. Wallner, E. Harder, and R. Agee. Key management for multicast: Issues and architectures. Request for Comments: 2627, Internet Engineering Task Force, 1999.
- [WKHB21] Matthew Weidner, Martin Kleppmann, Daniel Hugenroth, and Alastair R. Beresford. Key agreement for decentralized secure group messaging with strong security guarantees. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 2024–2045. ACM Press, November 2021.
- [WPBB23] Théophile Wallez, Jonathan Protzenko, Benjamin Beurdouche, and Karthikeyan Bhargavan. TreeSync: Authenticated group management for messaging layer security. In Joseph A. Castellino and Carmela Troncoso, editors, *USENIX Security 2023*, pages 1217–1233. USENIX Association, August 2023.

Supplementary material

A Omitted Formal Security Definition

A.1 Public Key Encryption

Definition A.1. A public-key encryption (PKE) scheme is a tuple of efficient algorithms $\text{Pke} = (\text{Pke.Gen}, \text{Pke.Enc}, \text{Pke.Dec})$ formed by

- a probabilistic key generation algorithm $(\text{pk}, \text{sk}) \leftarrow \text{Pke.Gen}()$ that outputs a pair formed by a public key and a secret key,
- a probabilistic encryption algorithm that on input a public key pk and a message m outputs a ciphertext $c \leftarrow \text{Pke.Enc}(\text{pk}, m)$,
- a deterministic decryption algorithm that on input a ciphertext c and a secret key sk outputs a message $m \leftarrow \text{Pke.Dec}(\text{sk}, c)$ or an error value \perp .

We say it is correct if for all $(\text{pk}, \text{sk}) \leftarrow \text{Pke.Gen}()$ and all messages m it holds that

$$\text{Pke.Dec}(\text{sk}, \text{Pke.Enc}(\text{pk}, m)) = m.$$

Now we proceed to define security. For an adversary \mathcal{A} and a PKE scheme Pke we consider the following experiment $\text{Exp}_{\text{Pke}}^{\mathcal{A}}$:

1. sample $b \leftarrow_{\text{s}} \{0, 1\}$,
2. sample $(\text{pk}, \text{sk}) \leftarrow_{\text{s}} \text{Pke.Gen}()$ and give pk to \mathcal{A} ,
3. \mathcal{A} is allowed to make decryption queries in which it sends a ciphertext c and receives $\text{Pke.Dec}(\text{sk}, c)$,
4. \mathcal{A} sends a pair of messages (m_0, m_1) and receives a challenge ciphertext $c^* \leftarrow \text{Pke.Enc}(\text{pk}, m_b)$
5. \mathcal{A} is allowed to make decryption queries in which it sends a ciphertext c with the condition that $c \neq c^*$ and receives $\text{Pke.Dec}(\text{sk}, c)$,
6. \mathcal{A} outputs a bit b' ,
7. the output of the experiment $\text{Exp}_{\text{PKE}}^{\mathcal{A}}$ is 1 if $b = b'$ and 0 otherwise.

Definition A.2. We say that a PKE scheme Pke is IND-CCA2 secure if for all efficient adversaries \mathcal{A}

$$|\Pr[\text{Exp}_{\text{Pke}}^{\mathcal{A}} = 1] - \frac{1}{2}|$$

is negligible.

A.2 Digital Signatures

Definition A.3. A digital signature is a tuple of efficient algorithms $\text{Sig} = (\text{GenS}, \text{Sign}, \text{Vrfy})$ formed by

- a probabilistic key generation algorithm $(\text{svk}, \text{ssk}) \leftarrow \text{GenS}()$ that outputs a pair formed by a public verification key and a secret signing key,
- a probabilistic encryption algorithm that on input a secret signing key ssk and a message m outputs a signature $\sigma \leftarrow \text{Sign}(\text{ssk}, m)$,
- a deterministic verification algorithm that on input a message signature pair (m, σ) and a verification key svk outputs a boolean value $b \leftarrow \text{Vrfy}(\text{svk}, m, \sigma)$.

We say it is correct if for all $(\text{svk}, \text{ssk}) \leftarrow \text{GenS}()$ and all messages m it holds that

$$\text{Vrfy}(\text{svk}, \text{Sign}(\text{ssk}, m)) = 1.$$

For an adversary \mathcal{A} and a signature scheme Sig we consider the following experiment $\text{Exp}_{\text{Sig}}^{\mathcal{A}}$:

1. sample $(\text{svk}, \text{ssk}) \leftarrow \text{GenS}()$ and give svk to \mathcal{A} ,
2. \mathcal{A} is allowed to make signing queries in which it sends a message m_i and receives $\sigma_i \leftarrow \text{Sign}(\text{ssk}, m_i)$,
3. \mathcal{A} outputs a message-signature pair (m^*, σ^*) ,
4. the output of $\text{Exp}_{\text{Sig}}^{\mathcal{A}}$ is 1 if for all queries i it holds that $(m^*, \sigma^*) \neq (m_i, \sigma_i)$ and $\text{Vrfy}(\text{svk}, m^*, \sigma^*) = 1$. Otherwise the output is 0.

Definition A.4. We say that a signature scheme Sig is *SEUF-CMA* secure if for all efficient adversaries \mathcal{A}

$$|\Pr[\text{Exp}_{\text{Sig}}^{\mathcal{A}} = 1] - \frac{1}{2}|$$

is negligible.

A.3 PKI and CGKA

For the definition of the PKI functionalities see Figures 6 and 7. The reader may find a pseudo-code description of the ideal functionality $\mathcal{F}_{\text{CGKA}}$ in Figures 8, 9, 10 and 11. The helpers used to describe $\mathcal{F}_{\text{CGKA}}$ can be found in Figures 12, 13, 14, 15 and 16. The reader may find the definition of the predicates $\text{safe}(c)$ and $\text{inj-allowed}(c, u)$ in Figure 17.

B Omitted Formal Protocol Descriptions

B.1 Formal description of MLS

In this section we provide pseudocode for the CGKA underlying MLS. Our protocol description closely follows the one of [AJM22], more precisely the one of protocol ITK, a CGKA introducing fixes to a prior draft of the MLS protocol that were later adapted into the standard. As the focus on this work is on the impact of blanking on the protocol’s communication complexity, makes the same simplifying assumptions. I.e., we focus on the CGKA part of the protocol and do not include its secure messaging layer, assume a fixed protocol version and cipher suite, and omit features like meta data protection and the additional proposal types `PreSharedKey`, `ReInit`, `ExternalInit`, `GroupContextExtensions` and external proposals. As a minor difference we use the parent hash mechanism as defined in the IETF standard [BBR⁺23] that slightly differs from the one used by ITK.

The protocol is defined with respect to public-key encryption scheme $\text{Pke} = (\text{Gen}, \text{Enc}, \text{Dec})$ with key pairs (pk, sk) , signature scheme $\text{Sig} = (\text{GenS}, \text{Sign}, \text{Vrfy})$ with key pairs (svk, ssk) , MAC $\text{Mac} = (\text{Mac.Tag}, \text{Mac.Ver})$, key-derivation function the use of which is modeled as hash functions H_1, H_2 , and hash function H .

Protocol state. We list the content of users’ states, as well as the intuition behind additional values computed throughout the protocol execution in Table 1.

Interaction with KS and AS. Users interact with the key service and authentication service via the following.

`fetch-ssk-if-nec` is used to retrieve (and potentially update) a user’s secret signing key (see Figure 18).

`get-sks` is used to fetch the secret keys of key packages when joining a group (see Figure 7).

`fetch-kp`: fetches a user’s key package from the KS (See Figure 7).

`validate-kp` is used to validate key packages (See Figure 18).

`Gen-kp` generates a key package (Figure 18).

state	
st.groupID	the group identifier
st.ep	the current epoch
st.id	the user’s leaf identifier
st.T	the user’s view of the ratchet tree
st.tHash	tree hash of st.T
st.ssk	the user’s current secret signing key
st.cert-SPks	the signature verification keys associated to each user’s id
st.pendUpd	a list of pending updates
st.pendComm	a list of pending commits
key schedule stored in state	
st.appSec	the current epoch’s CGKA group key
st.interim-tranHash	the interim transcript hash for the following epoch
st.memKey	membership key used to frame protocol messages
st.initSec	the next epoch’s initialization secret
additional values	
comSec	commit secret generated by rekey-path
joinerSec	secret for added group members, derived from initSec, comSec
confKey	key to tag confirmed transcript hash
conf-TransHash	hash encoding operations performed so far

Table 1: Users’ state in MLS.

Ratchet tree state and modifications. The state of nodes in the protocol’s underlying ratchet tree is listed in Table 2. The table further list functions returning particular properties of the tree, as well as functions implementing changes to the ratchet tree.

Protocol algorithms. The group creation, commit, process, join and key-recovery algorithms of MLS can be found in Figure 19. Its proposal algorithm, contrasted against the one of MLS-Cutoff, is in Figure 20. The algorithms make use of helper functionalities `rekey-path`, `apply-rekey` (both in Figure 22), and `apply-props`. The latter is in Figure 23, contrasted against its MLS-Cutoff variant.

Further, the protocol employs several additional helper functionalities. While these play an important role in deriving the key schedule and ensuring consistency and authenticity of protocol messages, they do not play a role in modifications to ratchet tree’s state, which is the main focus of this work. Thus for formal definitions of these functions we refer to [AJM22], providing only an overview on their roles below.

`init-epoch` on input a user’s state `st` returns a copy `st′` of the state with incremented epoch $st′.ep = st.ep + 1$ and empty list of pending updates and commits.

`derive-keys` on input state `st`, pending state `st′` and commit secret `comSec` first computes the joiner secret `joinerSec` from `comSec` and `initSec`. From this it computes confirmation key `confKey`, application secret `appSec`, membership key `memKey`, and updated initialization secret `initSec`. The latter three are stored in `st′`, and the function’s output is $(st′, confKey)$.

`derive-epochKeys` on input pending state `st′` and joiner secret `joinerSec` computes `confKey`, `appSec`, `memKey`, and `initSec` in the same way as `derive-keys` and returns $(st′, confKey, joinerSec)$.

`sign-commit` receives as input the state `st` and value `C` that contains hashes of the applied commits and the `UpdPath` object. It returns a signature σ on the concatenation of `C` with the group context, id, epoch, committer leaf id.

public state of node v	
$v.index$	index of the node
$v.pk$	public PKE key
$v.pHash$	parent hash value
$v.unm$	unmerged leaves
$v.svk$	signature verification key (if v is leaf)
$v.cred$	credential (if v is leaf)
$v.kp$	key package $v.kp = (v.pk, v.svk, v.cred)$ (if v is leaf)
private state of node v	
$v.sk$	secret PKE key
tree state	
$root(T)$	root of the tree
$leaves(T)$	returns the leaves of T
$T.public$	public part of the ratchet tree
v_{id}	leaf associated to leaf identity id
$id(v)$	leaf identity of node v
$path(v)$	direct path from v to the root
$fil-path(v)$	filtered direct path from v_0 to the root
$lca(u, v)$	least common ancestor w.r.t filtered paths of leaves u, v
$ind-lca(u, v)$	index (from root) of the least common ancestor w.r.t filtered paths of leaves u, v
$Res(v)$	resolution of node v
$roster(T)$	returns list of leaf IDs
modifications to tree	
$init-tree$	creates ratchet tree of size 1
$assign-kp(T, id, kp)$	assigns leaf associated to id key package kp
$add-leaf(T, id)$	assign id leftmost free leaf, doubles size of T if necessary
$truncate(T)$	removes right half of T if it contains no populated leaves
$blank-leaf(T, v)$	sets $blank(v)$ for leaf v
$blank-path(T, v, i)$	calls $blank(v_j)$ for $v_j \in path(v)$ with $j \geq i$
$set-as-unmerged(T, v_0)$	sets $v.unm \leftarrow \cup \{v_0\}$ for all $v \in path(v_0)$
$merge-leaves(T, v_0)$	sets $v.unm \leftarrow \emptyset$ for all $v \in path(v_0)$

Table 2: Ratchet tree state and functions modifying the ratchet tree.

`set-int-transHash` on input pending state st' and confirmation tag `confTag` computes the interim transcript hash `interim-tranHash` as a hash of the confirmed transcript hash stored in st' and `confTag`. It stores the resulting value in st' .

`set-conf-trans-hash` on input state st , pending state st' , sender id `id-snd`, C , and signature σ as above computes the confirmed transcript hash `conf-TransHash` as a hash of the latter three values and the interim transcript hash stored in st' . The confirmed transcript hash `conf-TransHash` is stored in st' .

`compute-confTag` on input pending state st' and confirmation key `confKey` returns a MAC `confTag` of the confirmed transcript hash `conf-TransHash` contained in st' .

`verify-confTag` on input pending state st' , confirmation key `confKey`, and confirmation tag `confTag` verifies the tag.

`frame-prop` frames proposal $P = (\text{type}, ad)$ by signing and MACing (under the membership key `memKey`) P together with group context, group id, epoch, leaf id of the sender.

`unframe-prop` on input state st and framed proposal `pmsg`, verifies signature and MAC tag and returns P as well as the sender id.

`frame-comm` on input state st , confirmation tag `confTag`, signature σ (as in `sign-commit`), and membership tag `memTag` frames the commit by appending group id, epoch, and id of the committer.

`unframe-comm` on input state st and framed commit `cmmsg`, verifies the contained signature and MAC tag. If both verify it returns the sender's id, C , the confirmation tag, signature, and the membership tag.

B.2 Formal description of MLS-Cutoff

We provide a formal description of protocol MLS-Cutoff. It is defined with respect to public key encryption scheme `Pke`, Signature scheme `Sign`, message authentication code `Mac`, hash function `H`, and a key-derivation function the use of which is denoted by H_1 and H_2 . The protocol is parameterized by cutoff parameter i_{cut} . The protocol shares most algorithms with MLS differing only in two aspects:

MLS-Cutoff generates proposal messages in a different way. Its proposal algorithm `CGKA.Prop` can be found in Figure 20 (right).

MLS-Cutoff uses a different method of applying proposals to the ratchet tree. I.e., algorithms `CGKA.Com` and `CGKA.Proc` make use of the subroutine `apply-props` of Figure 23 (right).

All remaining aspects are handled as in MLS. I.e., the protocol has the same user state (except for storing the cutoff bound in $st.i_{\text{cut}}$ that remains unchanged throughout the protocol execution), uses `CGKA.Create`, `CGKA.Com`, `CGKA.Proc`, `CGKA.Join`, and `CGKA.Key` as defined in Figure 19, and makes use of the subroutines listed in Appendix B.1 and Figures 22 and 18. Note that when rekeying paths while applying updates MLS-Cutoff calls `apply-rekey` with option `copySt = True`. This means the parent hashes contained in each update proposal are verified with respect to a tree resulting from applying the modifications to $st.T$ instead of the pending tree T' . Thus, the contained key package verifies if the update message was honestly generated, which implies the parent hash was computed with respect to such a tree.

Security of MLS-Cutoff.

Proof of Theorem 5.1. We are able to closely follow the security proof for ITK in [AJM22]. The proof is divided into two parts. First a sequence of hybrids shows that a version of the MLS protocol (MLS*) that has a different parent hash securely realizes $\mathcal{F}_{\text{CGKA}}$ with a slightly different version of the `safe` predicate denoted `safe*`. The second step consists in showing that using the actual parent hash guarantees and the actual `safe` predicate satisfy the following property: if `safe(c) = true`, then no node in the ratchet tree of the epoch that corresponds to c contains an exposed key `pk`.

We argue that both steps can be adapted to the case of MLS-Cutoff. We consider the following hybrids:

- Hybrid \mathcal{H}_0 : Real world.
- Hybrid \mathcal{H}_1 : Instead of the protocol as in the real world, now we consider a dummy functionality $\mathcal{F}_{\text{Dummy}}$ that simply routes its inputs and outputs through a simulator \mathcal{S}_1 that executes the protocol.
- Hybrid \mathcal{H}_2 : $\mathcal{F}_{\text{Dummy}}$ is substituted by the ideal functionalities \mathcal{F}_{AS} , \mathcal{F}_{KS} and $\mathcal{F}_{\text{CGKA}}$ with the only change that $\mathcal{F}_{\text{CGKA}}$ uses as predicates $\text{safe}(c) = \text{false}$ and $\text{inj-allowed}(c, \text{id}) = \text{true}$. The simulator \mathcal{S}_2 still executes the actual protocol.
- Hybrid \mathcal{H}_3 : now $\mathcal{F}_{\text{CGKA}}$ uses the actual predicate safe^* . The simulator \mathcal{S}_3 proceeds as before except that it only sets applications secrets when $\text{safe}^*(c) = \text{false}$.
- Hybrid \mathcal{H}_4 : Ideal world. The difference with the previous hybrid consists in also using the actual predicate inj-allowed . The simulator \mathcal{S}_4 is the same as \mathcal{S}_3 .

Clearly the first two hybrids are indistinguishable. Let's argue that \mathcal{H}_1 and \mathcal{H}_2 are indistinguishable. Despite the fact that in MLS-Cutoff update proposals are created differently, we still frame the proposals in the same way and this guarantees consistency by the same argument as in the proof for MLS.

Concerning \mathcal{H}_2 and \mathcal{H}_3 the only change in the argument that has to be made corresponds to the fact that **rekey-path** is also used in proposals. Therefore the GSD graph also contains nodes and edges corresponding to these new keys just as is the case for commits in [AJM22].

The indistinguishability of hybrids \mathcal{H}_3 and \mathcal{H}_4 follows from the argument made in [AJM22]. Therefore it just remains to show that: if $\text{safe}(c) = \text{true}$, then no node in the ratchet tree of the epoch that corresponds to c contains an exposed key pk . Here we make use of the fact that keys ending up in the ratchet tree after processing commits to update proposals were never encrypted under keys being concurrently updated. In the inductive step we have to consider an additional case, namely, that in which the node v for which $v.\text{pk}$ has been exposed is set as part of proposal that is included in some commit c and the user who generated the commit is a user outsider v 's subtree (in the ratchet tree \mathbb{T} that corresponds to that commit). We denote by id_c the leaf of the party that generated the commit c .

It may be the case that the parent hash of v , pHash_v , cannot be verified by using the parent hash of its parent $v.\text{par}$ in \mathbb{T} as it would be the case when $v.\text{par}$ is blanked. However, all nodes contained in the ratchet tree are parent hash valid, as defined in the MLS standard, which uses a slightly modified parent hash definition compared to ITK.

However the node at which the path of v to the root and the path of id_c to the root merge contains a commitment to the value of pHash_v (by definition of **par-hash-cochild**) which can be verified and this is done by `verify-treeState`. \square

```

Initialization
00 Registered, Exposed  $\leftarrow \emptyset$ , SSK[*,*]  $\leftarrow \perp$ , RndCor[*]  $\leftarrow$  good
Input register-spk from a party  $u$ 
01 if RndCor[ $u$ ] = good then (svk, ssk)  $\leftarrow_s$  GenS()
02 else
03   send rnd to the adversary and receive  $r$ 
04   (svk, ssk)  $\leftarrow$  GenS( $r$ )
05 send (sample-ssk,  $u$ ) to the adversary and receive (svk, ssk)
06 if RndCor[ $u$ ]  $\neq$  good then Exposed  $\leftarrow_{\cup}$  {svk}
07 SSK[ $u$ , svk]  $\leftarrow$  ssk
08 Registered  $\leftarrow_{\cup}$  {( $u$ , svk)}
09 send (register-spk,  $u$ , svk) to the adversary
10 send svk to  $u$ 
Input (verify-cert,  $v$ , svk) from a party  $u$ 
11 Send ( $v$ , svk)  $\in$  Registered to  $u$ 
Input (get-ssk, svk) from a party  $u$ 
12 Send SSK[ $u$ , svk] to  $u$ 
Input (delete-ssk, svk) from a party  $u$ 
13 SSK[ $u$ , svk]  $\leftarrow \perp$ 
Input (register-spk,  $u$ , svk) from the adversary
14 if (*, svk)  $\notin$  Registered then Exposed  $\leftarrow_{\cup}$  {svk}
15 Registered  $\leftarrow_{\cup}$  {( $u$ , svk)}
Input (expose,  $u$ ) from the adversary
16 Exposed  $\leftarrow_{\cup}$  {svk | SSK[ $u$ , svk]  $\neq \perp$ }
17 Send SSK[ $u$ , *] to the adversary
Input (CorrRand,  $u$ ,  $b$ ) from the adversary where  $b \in \{\text{good}, \text{bad}\}$ 
18 RndCor[ $u$ ]  $\leftarrow b$ 
Input (exposed,  $u$ , svk) from  $\mathcal{F}_{KS}$ 
19 Exposed  $\leftarrow_{\cup}$  {svk}
20 Send SSK[ $u$ , svk] to the adversary
Input (has-ssk, svk,  $u$ ) from  $\mathcal{F}_{CGKA}$ 
21 Send SSK[ $u$ , svk]  $\neq \perp$  to  $\mathcal{F}_{CGKA}$ 

```

Figure 6: Functionality \mathcal{F}_{AS} (parametrized by a key generation algorithm for a signature scheme GenS) and the ideal functionality \mathcal{F}_{AS}^1 . Yellow lines (•) are only executed by \mathcal{F}_{AS} while cyan lines (•) are only executed by \mathcal{F}_{AS}^1 .

```

Initialization
00  $SK[*, *], SVK \leftarrow \perp, \text{RndCor}[*] \leftarrow \text{good}$ 
Input (register-kp, svk, ssk) from a party  $u$ 
01 if  $\text{RndCor}[u] = \text{good}$  then  $(kp, sk) \leftarrow_{\text{s}} \text{Gen-kp}()$ 
02   if  $kp \neq \perp$  then return
03 else
04   send rnd to the adversary and receive  $r$ 
05    $(kp, sk) \leftarrow \text{Gen-kp}(r)$ 
06   if  $kp \neq \perp$  then return
07   send (exposed,  $u, svk$ ) to  $\mathcal{F}_{AS}$ 
08   send ssk to the adversary
09 send (sample-sk,  $u, svk, ssk$ ) to the adversary and receive  $(kp, sk, ack)$ 
10 if  $\neg ack$  then return
11 if  $\text{RndCor}[u] \neq \text{good}$  then
12   send (exposed,  $u, svk$ ) to  $\mathcal{F}_{AS}$ 
13  $SK[u, kp] \leftarrow sk, SVK[u, kp] \leftarrow svk$ 
14 send (register-kp,  $u, svk, kp$ ) to the adversary
15 send kp to  $u$ 
Input get-sks from a party  $u$ 
16 Send  $\{(kp, SK[u, kp]) \mid SK[u, kp] \neq \perp\}$  to  $u$ 
Input (fetch-kp,  $v$ ) from a party  $u$ 
17 send (fetch-kp,  $u, v$ ) to the adversary and receive kp
18 sendkp to  $u$ 
Input (delete-sk, svk) from a party  $u$ 
19  $SK[u, kp] \leftarrow \perp$ 
Input (expose,  $u$ ) from the adversary
20 Send  $SK[u, *]$  to the adversary
Input (CorrRand,  $u, b$ ) from the adversary where  $b \in \{\text{good}, \text{bad}\}$ 
21  $\text{RndCor}[u] \leftarrow b$ 

```

Figure 7: Functionality \mathcal{F}_{KS} (parametrized by a key generation algorithm for key packages Gen-kp) and the ideal functionality \mathcal{F}_{KS}^I . Yellow lines (•) are only executed by \mathcal{F}_{KS} while cyan lines (•) are only executed by \mathcal{F}_{KS}^I .

```

Initialization
00 Ptr[*], Node[*], Prop[*], Wel[*]  $\leftarrow \perp$ 
01 RndCor[*]  $\leftarrow$  good
02 HasKey[*]  $\leftarrow$  false
03 rootCtr  $\leftarrow$  0
Input (create, svk) from  $u_{\text{creator}}$   $\parallel$  Used only once
04 req Node[root0] =  $\perp$ 
05 req valid-vsk( $u_{\text{creator}}$ , svk)
06 orig  $\leftarrow$   $u_{\text{creator}}$ 
07 mem  $\leftarrow$  {( $u_{\text{creator}}$ , svk)}
08 Node[root0]  $\leftarrow$  create-root(orig, mem, RndCor[orig])
09 HasKey[ $u_{\text{creator}}$ ]  $\leftarrow$  true
10 Ptr[ $u_{\text{creator}}$ ]  $\leftarrow$  root0

```

Figure 8: Initialization of $\mathcal{F}_{\text{CGKA}}$.

```

Input (propose, act) from  $u$ 
00 req  $\text{Ptr}[u] \neq \perp$ 
01 send (propose,  $u$ , act) to the adversary and receive  $(p, \text{svk}_v, \text{ack})$ 
02 if  $\neg \text{req-correct-prop}(u, \text{act})$  then req  $\text{ack}$ 
03 if  $\text{act} = (\text{upd}, \text{svk})$  then assert  $\text{valid-vsk}(u, \text{svk})$ 
04 if  $\text{act} = (\text{add}, v)$  then  $\text{act} \leftarrow (\text{add}, v, \text{svk}_v)$ 
05 if  $\text{Prop}[p] = \perp$  then
06    $\text{Prop}[p] \leftarrow \text{create-prop}(\text{Ptr}[u], u, \text{act}, \text{RndCor}[u])$ 
07 else  $\text{consistent-prop}(p, u, \text{act}, \text{RndCor}[u])$ 
08 if  $\text{RndCor}[u] = \text{bad}$  then send (exposed,  $u$ ,  $\text{svk}$ ) to  $\mathcal{F}_{\text{AS}}$ 
09 return  $p$ 
Input (commit,  $P$ ,  $\text{svk}$ ) from  $u$ 
10 req  $\text{Ptr}[u] \neq \perp$ 
11 send (commit,  $u$ ,  $P$ ,  $\text{svk}$ ) to the adversary and receive  $(c, w, \text{ack}, \text{rt})$ 
12 if  $\neg \text{req-correct-commit}(u, P, \text{svk})$  then req  $\text{ack}$ 
13 fill-props( $u$ ,  $P$ )
14 assert  $\text{valid-vsk}(u, \text{svk})$ 
15  $\text{mem} \leftarrow \text{members}(c, u, P, \text{svk})$ 
16 assert  $\text{mem} \neq \perp \wedge (u, \text{svk}) \in \text{mem}$ 
17 if  $\text{Node}[c] = \perp \wedge \text{rt} = \perp$  then
18    $\text{Node}[c] \leftarrow \text{create-child}(\text{Ptr}[u], u, P, \text{mem}, \text{RndCor}[u])$ 
19 else
20   if  $\text{Node}[c] = \perp$  then  $c' \leftarrow \text{root}_{\text{rt}}$ 
21   else  $c' \leftarrow c$ 
22    $\text{consistent-commit}(c', u, P, \text{mem})$ 
23   if  $c \neq c'$  then  $\text{attach}(c, c', u, P)$ 
24 assert  $w \neq \perp$  if and only if  $\exists p \in P: \text{Node}[p].\text{act} = (\text{add}, *)$ 
25 if  $w \neq \perp$  then
26   assert  $\text{Wel}[w] \in \{\perp, c\}$ 
27    $\text{Wel}[w] \leftarrow c$ 
28 assert  $\text{cons-invariant} \wedge \text{auth-invariant}$ 
29 if  $\text{RndCor}[u] = \text{bad}$  then send (exposed,  $u$ ,  $\text{Node}[u].\text{mem}[u]$ ) to  $\mathcal{F}_{\text{AS}}$ 
30 return  $(c, w)$ 

```

Figure 9: Proposals, commits and process of $\mathcal{F}_{\text{CGKA}}$.

```

Input (process,  $c, P$ ) from  $u$ 
00 send (process,  $u, c, P$ ) to the adversary and receive ( $rt, \text{orig}', \text{svk}', \text{ack}$ )
01 if  $\neg \text{req-correct-proc}(u, c, P)$  then req  $\text{ack}$ 
02 fill-props( $u, P$ )
03 if  $\text{Node}[c] = \perp \wedge rt = \perp$  then
04    $\text{mem} \leftarrow \text{members}(\text{Ptr}[u], \text{orig}', P, \text{svk}')$ 
05   assert  $\text{mem} \neq \perp \wedge \text{inj-allowed}(\text{Ptr}[u], u)$ 
06 else
07   if  $\text{Node}[c] = \perp$  then  $c' \leftarrow \text{root}_{rt}$ 
08   else  $c' \leftarrow c$ 
09    $v \leftarrow \text{Node}[c'].\text{orig}$ 
10    $\text{svk}_v \leftarrow \text{Node}[c'].\text{mem}[v]$ 
11    $\text{mem} \leftarrow \text{members}(\text{Ptr}[u], v, P, \text{svk}_v)$ 
12   assert  $\text{mem} \neq \perp$ 
13   valid-succesor( $c', u, P, \text{mem}$ )
14   if  $c \neq c'$  then attach( $c, c', u, P$ )
15 if  $\exists p \in P: \text{Prop}[p].\text{act} = (\text{rem}, u)$  then  $\text{Ptr}[u] \leftarrow \perp$ 
16 else
17   assert  $u \in \text{Node}[c].\text{mem}$ 
18    $\text{Ptr}[u] \leftarrow c$ 
19    $\text{HasKey}[u] \leftarrow \text{true}$ 
20 assert  $\text{cons-invariant} \wedge \text{auth-invariant}$ 
21  $(*, \text{PropSemantics}) \leftarrow \text{apply-props}(c, \text{Node}[c].\text{pro})$ 
22 return ( $\text{Node}[c].\text{orig}, \text{PropSemantics}$ )
Input (join,  $w$ ) from  $u$ 
23 send (join,  $u, w$ ) to the adversary and receive ( $c', \text{orig}', \text{mem}', \text{ack}$ )
24 req  $\text{ack}$ 
25  $c \leftarrow \text{Wel}[w]$ 
26 if  $c = \perp$  then
27   if  $\text{Node}[c'] \neq \perp$  then  $c \leftarrow c'$ 
28   else
29      $\text{rootCtr} \leftarrow \text{rootCtr} + 1$ 
30      $c \leftarrow \text{root}_{\text{rootCtr}}$ 
31      $\text{Node}[c] \leftarrow \text{create-root}(\text{orig}', \text{mem}', \text{adv})$ 
32    $\text{Wel}[w] \leftarrow c$ 
33  $\text{Ptr}[u] \leftarrow c$ 
34  $\text{HasKey}[u] \leftarrow \text{true}$ 
35 assert  $u \in \text{Node}[c].\text{mem} \wedge \text{cons-invariant} \wedge \text{auth-invariant}$ 
36 return ( $\text{Node}[c].\text{mem}, \text{Node}[c].\text{orig}$ )
Input key from  $u$ 
37 req  $\text{Ptr}[u] \neq \perp \wedge \text{HasKey}[u]$ 
38 if  $\text{Node}[u].\text{key} = \perp$  then set-key( $\text{Ptr}[u]$ )
39  $\text{HasKey}[u] \leftarrow \text{false}$ 
40 return  $\text{Node}[\text{Ptr}[u]].\text{key}$ 

```

Figure 10: Process, join and key in $\mathcal{F}_{\text{CGKA}}$.

```

Input (expose, u) from the adversary
00 This input is not allowed if  $\exists c$  such that  $\text{Node}[c].\text{chall} \wedge \neg \text{safe}(c)$ 
01 if  $\text{Ptr}[u] \neq \perp$  then
02    $\text{Node}[\text{Ptr}[u]].\text{exp} \leftarrow_{\cup} \{(u, \text{HasKey}[u])\}$ 
03    $\text{update-stat-after-exp}(u)$ 
04   send (exposed, u, Node[u].mem[u]) to  $\mathcal{F}_{\text{AS}}$ 
05 send get-sks to  $\mathcal{F}_{\text{KS}}$  and receive SK and SVK.
06 for each kp such that  $\text{SK}[u, kp] \neq \perp \wedge \text{SVK}[u, kp] = \text{svk}$  do
07   for each c such that  $\exists p \in \text{Node}[c].\text{pro} : \text{Prop}[p].\text{act} = (\text{add}, u, \text{svk})$  do
08      $\text{Node}[c].\text{exp} \leftarrow_{\cup} \{(u, \text{true})\}$ 
Input (CorrRand, u, b) from the adversary where  $b \in \{\text{good}, \text{bad}\}$ 
09  $\text{RndCor}[u] \leftarrow b$ 

```

Figure 11: Corruptions in $\mathcal{F}_{\text{CGKA}}$.

```

Helper  $\text{create-root}(u, \text{mem}, \text{stat})$ 
00 return new node with  $\text{par} \leftarrow \perp, \text{orig} \leftarrow u, \text{pro} \leftarrow \perp, \text{mem} \leftarrow \text{mem}, \text{stat} \leftarrow \text{stat}$ 
Helper  $\text{create-prop}(c, u, \text{act}, \text{stat})$ 
01 return new proposal node with  $\text{par} \leftarrow c, \text{orig} \leftarrow u, \text{act} \leftarrow \text{act}, \text{stat} \leftarrow \text{stat}$ 
Helper  $\text{create-child}(c, u, P, \text{mem}, \text{stat})$ 
02 return new node with  $\text{par} \leftarrow c, \text{orig} \leftarrow u, \text{pro} \leftarrow P, \text{mem} \leftarrow \text{mem}, \text{stat} \leftarrow \text{stat}$ 
Helper  $\text{fill-props}(u, P)$ 
03 for  $p \in P$  such that  $\text{Prop}[p] = \perp$  do
04   send (propose, p) to the adversary and receive (orig, act)
05    $\text{Prop}[p] \leftarrow \text{create-prop}(\text{Ptr}[u], \text{orig}, \text{act}, \text{adv})$ 

```

Figure 12: Helpers for creating new nodes in $\mathcal{F}_{\text{CGKA}}$.

<pre> Helper valid-vsk(u, svk') 00 $svk \leftarrow \text{Node}[\text{Ptr}[u]].\text{mem}[u]$ 01 if $svk \neq \perp \wedge svk = svk'$ then return true 02 send ($\text{has-ssk}, svk', u$) to \mathcal{F}_{AS} and receive ack 03 return ack Helper set-key(c) 04 if $\neg \text{safe}(c)$ then 05 send (key, u) to the adversary and receive K 06 $\text{Node}[c].\text{key} \leftarrow K, \text{Node}[c].\text{chall} \leftarrow \text{false}$ 07 else $\text{Node}[c].\text{key} \leftarrow_s \mathcal{K}, \text{Node}[c].\text{chall} \leftarrow \text{true}$ Helper update-stat-after-exp(u) 08 for each p such that $\text{Prop}[p] \neq \perp$ and 09 $\text{Prop}[p].\text{par} = \text{Ptr}[u]$ and 10 $\text{Prop}[p].\text{orig} = u$ and 11 $\text{Prop}[p].\text{act} = \text{upd}$ 12 do $\text{Prop}[p].\text{stat} \leftarrow \text{bad}$ 13 for each c such that $\text{Node}[c] \neq \perp$ and 14 $\text{Node}[c].\text{par} = \text{Ptr}[u]$ and 15 $\text{Node}[c].\text{orig} = u$ 16 do $\text{Node}[c].\text{stat} \leftarrow \text{bad}$ Helper members(c, u, P, svk) 17 $(G, *) \leftarrow \text{apply-props}(c, u, P, svk)$ 18 if $(G, *) = \perp$ then return \perp 19 else return G </pre>	<pre> Helper apply-props(c, u, P, svk) 20 $\text{req } \text{Node}[c] \neq \perp \wedge (u, *) \in \text{Node}[c].\text{mem}$ 21 $\text{req } \forall p \in P: \text{Prop}[p] \neq \perp \wedge \text{Prop}[p].\text{par} = c$ 22 $G \leftarrow \text{Node}[c].\text{mem}$ 23 $G \leftarrow G \setminus \{(u, *)\}$ 24 $G \leftarrow_{\cup} \{(u, svk)\}$ 25 $L \leftarrow \{u\}$ 26 for each $p \in P$ such that $\text{Prop}[p].\text{act} = (\text{upd}, *)$ do 27 $(v, (\text{upd}, svk')) \leftarrow (\text{Prop}[p].\text{orig}, \text{Prop}[p].\text{act})$ 28 $\text{req } v \in G \setminus L$ 29 $G \leftarrow G \setminus \{(v, *)\}$ 30 $G \leftarrow_{\cup} \{(v, svk')\}$ 31 $L \leftarrow_{\cup} \{v\}$ 32 for each $p \in P$ such that $\text{Prop}[p].\text{act} = (\text{rem}, *)$ do 33 $(v, (\text{rem}, v')) \leftarrow (\text{Prop}[p].\text{orig}, \text{Prop}[p].\text{act})$ 34 $\text{req } v \in G \wedge v' \in G \setminus L$ 35 $G \leftarrow G \setminus \{(v', *)\}$ 36 for each $p \in P$ such that $\text{Prop}[p].\text{act} = (\text{add}, *)$ do 37 $(v, (\text{add}, v', svk_{v'})) \leftarrow (\text{Prop}[p].\text{orig}, \text{Prop}[p].\text{act})$ 38 $\text{req } v \in G \wedge v' \notin G$ 39 $G \leftarrow_{\cup} \{(v', svk_{v'})\}$ 40 $\text{PropSemantics} \leftarrow ((\text{Prop}[p].\text{orig}, \text{Prop}[p].\text{act}): p \in P)$ 41 return $(G, \text{PropSemantics})$ </pre>
---	---

Figure 13: Other helpers for \mathcal{F}_{CGKA} .

<pre> Helper consistent-prop($p, u, \text{act}, \text{stat}$) 00 assert $\text{Prop}[p].\text{orig} = u \wedge \text{Prop}[p].\text{act} = \text{act} \wedge \text{Prop}[p].\text{par} = \text{Ptr}[u]$ Helper valid-succesor(c, u, P, mem) 01 $\text{Node}[c] \neq \perp \wedge \text{Node}[c].\text{mem} = \text{mem} \wedge \text{Node}[c].\text{pro} \in \{\perp, P\} \wedge \text{Node}[c].\text{par} \in \{\perp, \text{Ptr}[u]\}$ Helper consistent-commit(c, u, P, mem) 02 assert $\text{valid-succesor}(c, u, P, \text{mem})$ 03 assert $\text{RndCor}[u] \neq \text{good} \wedge \text{Node}[c].\text{orig} = u$ Helper attach(c, c', u, P) 04 assert $c' \neq \text{root}_0$ 05 $\text{Node}[c'].\text{par} \leftarrow \text{Ptr}[u], \text{Node}[c'].\text{pro} \leftarrow P, \text{Node}[c] \leftarrow \text{Node}[c'], \text{Node}[c'] \leftarrow \perp$ 06 for each w such that $\text{Wel}[w] = c'$ do $\text{Wel}[c] \leftarrow c$ </pre>
--

Figure 14: Consistency helpers for \mathcal{F}_{CGKA} .

```

Helper req-correct-prop( $u, \text{act}$ )
00 if  $\text{act} = (\text{rem}, v)$  then
01   return  $v \in \text{Node}[\text{Ptr}[u]].\text{mem}$ 
02 else if  $\text{act} = (\text{upd}, \text{svk})$  then
03   return  $\text{valid-vsk}(u, \text{svk})$ 
04 else return false
Helper req-correct-commit( $u, P, \text{svk}$ )
05 return  $\text{apply-props}(u, P, \text{svk}) \neq \perp \wedge \text{valid-vsk}(u, \text{svk})$ 
Helper req-correct-proc( $u, c, P$ )
06 return  $\text{Node}[c] \neq \perp \wedge \text{Node}[c].\text{par} = \text{Ptr}[u] \wedge \text{Node}[c].\text{pro} = P \wedge$ 
07    $\text{Node}[c].\text{stat} \neq \text{adv} \wedge \forall p \in P: \text{Prop}[p].\text{stat} \neq \text{adv}$ 

```

Figure 15: Correctness helpers for $\mathcal{F}_{\text{CGKA}}$.

```

Helper auth-invariant
00 return true iff
01    $\forall c$  if  $\text{Node}[c].\text{stat} = \text{adv}$  then  $\text{inj-allowed}(\text{Node}[c].\text{par}, \text{Node}[c].\text{orig})$  and
02    $\forall p$  if  $\text{Prop}[p].\text{stat} = \text{adv}$  then  $\text{inj-allowed}(\text{Prop}[p].\text{par}, \text{Prop}[p].\text{orig})$ 
Helper cons-invariant
03 return true iff
04    $\forall c$  such that  $\text{Node}[c].\text{par} \neq \perp$  it holds that
05      $\text{Node}[c].\text{pro} \neq \perp$  and  $\forall p \in \text{Node}[c].\text{pro}: \text{Prop}[p].\text{par} = \text{Node}[c].\text{par}$ 
06   and  $\forall u$  such that  $\text{Ptr}[u] \neq \perp: u \in \text{Node}[\text{Ptr}[u]].\text{mem}$ 
07   and the history graph contains no cycles

```

Figure 16: Invariant helpers for $\mathcal{F}_{\text{CGKA}}$.

```

safe(c)
00 safe(c) = true if and only if  $\neg((*, \text{true}) \in \text{Node}[c].\text{exp} \vee \text{can-traverse}(c))$ 
inj-allowed(c, u)
01 safe(c) = true if and only if  $\text{Node}[c].\text{exp} \in \text{Exposed} \wedge \text{know}(c, \text{'ep'})$ 
know(c, u) = true if and only if
02 1) state-directly-leaks(c, u)  $\vee$ 
03 2)  $(\text{Node}[c].\text{par} \neq \perp \wedge \neg \text{secrets-replaced}(c, u) \wedge \text{know}(\text{Node}[c].\text{par}, u)) \vee$ 
04 3)  $\exists c': (\text{Node}[c'].\text{par} = c \wedge \neg \text{secrets-replaced}(c', u) \wedge \text{know}(c', u))$ 
state-directly-leaks(c, u) = true if and only if
05 1)  $(u, *) \in \text{Node}[c].\text{exp} \vee$ 
06 2)  $\exists \text{rootCtr}: (\text{root}_{\text{rootCtr}} \text{ is ancestor of } c \wedge \exists \text{svk} \in \text{Exposed}: (u, \text{svk}) \in \text{Node}[c].\text{mem}) \vee$ 
07 3)  $(u, *) \in \text{Node}[c].\text{mem} \wedge \text{secrets-injected}(c, u)$ 
secrets-injected(c, u) = true if and only if
08 1)  $(\text{Node}[c].\text{orig} = u \wedge \text{Node}[c].\text{stat} \neq \text{good}) \vee$ 
09 2)  $\exists p \in \text{Node}[c].\text{pro}: (\text{Prop}[p].\text{act} = (\text{add}, u, \text{svk}) \wedge \text{svk} \in \text{Exposed}) \vee$ 
10 3)  $\exists p \in \text{Node}[c].\text{pro}: (\text{Prop}[p].\text{act} = (\text{upd}, *) \wedge \text{Prop}[p].\text{orig} = u \wedge \text{Prop}[p].\text{stat} \neq \text{good})$ 
secrets-replaced(c, u) = true if and only if
11  $\text{Node}[c].\text{orig} = u \vee (\exists p \in \text{Node}[c].\text{pro}: \text{Prop}[p].\text{act} \in \{(\text{add}, u), (\text{rem}, u)\} \vee (\text{Prop}[p].\text{act} = (\text{upd}, *) \wedge \text{Prop}[p].\text{orig} = u))$ 
know(c, 'ep') = true if and only if
13  $\text{Node}[c].\text{exp} \neq \emptyset \vee \text{can-traverse}(c)$ 
can-traverse(c) = true if and only if
14 1)  $(\text{Node}[c].\text{par} = \perp \wedge \exists \text{svk} \in \text{Exposed}: (*, \text{svk}) \in \text{Node}[c].\text{mem}) \vee$ 
15 2)  $\exists p \in \text{Node}[c].\text{pro}: (\text{Prop}[p].\text{act} = (\text{add}, u, \text{svk}) \wedge \text{svk} \in \text{Exposed}) \vee$ 
16 3)  $\text{leaf-welcome-key-reuse}(c) \vee$ 
17 4)  $(\text{know}(c, *) \wedge (c = \text{root}_* \vee \text{know}(\text{Node}[c].\text{par}, \text{'ep'})))$ 
leaf-welcome-key-reuse(c) = true if and only if
18  $\exists u \exists p \in \text{Node}[c].\text{pro}: \text{Prop}[p].\text{act} = (\text{add}, u) \wedge$ 
19 1)  $\text{Prop}[p].\text{act} = (\text{add}, u) \wedge$ 
20 2)  $\exists c': c \text{ is ancestor of } c' \wedge (u, *) \in \text{Node}[c'].\text{exp} \wedge$ 
21  $(\nexists c'' \text{ in the path between } c \text{ and } c' \text{ such that } \text{secrets-replaced}(c'', u))$ 

```

Figure 17: Predicates $\text{safe}(c)$ and $\text{inj-allowed}(c, u)$.

<pre> Helper fetch-ssk-if-nec(st, svk) 00 if $v_{st.id}.svk \neq svk$ then 01 send (get-ssk, svk) to \mathcal{F}_{AS} and receive ssk 02 else ssk \leftarrow st.ssk 03 return ssk Helper validate-kp(st, kp, id, pHash) 04 parse $(id', pk, svk, pHash', \sigma) \leftarrow$ kp 05 req $id = id' \wedge pHash = pHash'$ 06 if $svk \notin st.cert-SPks[id]$ then 07 send (verify-cert, id, svk) to \mathcal{F}_{AS} and receive b 08 req b 09 st.cert-SPks[id] $\leftarrow_{\cup} \{svk\}$ 10 req $Vrfy(svk, \sigma, (id', pk, svk, pHash', \sigma))$ 11 return st </pre>	<pre> Helper Gen-kp(id, svk, ssk) 12 $(pk, sk) \leftarrow_{\\$}$ Pke.Gen() 13 $\sigma \leftarrow$ Sign(ssk, (id, pk, svk, pHash)) 14 kp \leftarrow (id, pk, svk, pHash, σ) 15 return (kp, sk) </pre>
---	---

Figure 18: Helper functions for the interaction with KS and AS.

<pre> Algorithm CGKA.Create(svk) 00 require $st = \perp \wedge u = u_{creator}$ 01 $st.T \leftarrow \text{init-tree}$ 02 $st.groupID \leftarrow_{\mathcal{S}} \{0, 1\}^{\kappa}$ 03 $st.appSec \leftarrow_{\mathcal{S}} \{0, 1\}^{\kappa}$ 04 $st.ep \leftarrow 0$ 05 $st.interim-tranHash \leftarrow \varepsilon$ 06 try $st.ssk \leftarrow \text{fetch-ssk-if-nec}(st, svk)$ 07 $(kp, sk) \leftarrow \text{Gen-kp}(svk, ssk)$ 08 $\text{id}(\text{leaves}(T)) \leftarrow H(pk)$ 09 $\text{assign-kp}(T, v_{id}, kp)$ 10 $v_{id}.sk \leftarrow sk$ Algorithm CGKA.Com(st, PMSG) 11 $st' \leftarrow \text{init-epoch}(st)$ 12 $T' \leftarrow st'.T$ 13 try $(st', \text{upd}, \text{rem}, \text{add}) \leftarrow \text{apply-props}(st, PMSG)$ 14 require $(st.id, \cdot, \cdot) \notin \text{rem} \wedge (st.id, \cdot, \cdot) \notin \text{upd}$ 15 try $(st', \text{comSec}, \text{UpdPath}, \text{pathSec})$ $\leftarrow \text{rekey-path}(st', st.id, v_{id}.svk, 0)$ 16 $\text{proplds} \leftarrow ()$ 17 for $pmsg \in PMSG$: 18 $\text{proplds} \leftarrow_{\cup} H(pmsg)$ 19 $C \leftarrow (\text{proplds}, \text{UpdPath})$ 20 $\sigma \leftarrow \text{sign-commit}(st, C)$ 21 $st' \leftarrow \text{set-conf-trans-hash}(st, st', st.id, C, \sigma)$ 22 $(st', \text{confKey}, \text{joinerSec}) \leftarrow \text{derive-keys}(st, st', \text{comSec})$ 23 $\text{confTag} \leftarrow \text{compute-confTag}(st', \text{confKey})$ 24 if $\text{rem} \neq ()$ 25 $\text{memTag} \leftarrow \text{Mac.Tag}(st.\text{memKey}, C)$ 26 else $\text{memTag} \leftarrow \perp$ 27 $\text{cmsg} \leftarrow \text{frame-comm}(st, \text{confTag}, \sigma, \text{memTag})$ 28 if $\text{add} \neq ()$ 29 $(st', \text{wmsg}) \leftarrow \text{comp-wmsg}(st',$ $\text{add}, \text{joinerSec}, \text{pathSec}, \text{confTag})$ 30 else $\text{wmsg} \leftarrow \perp$ 31 $st' \leftarrow \text{set-int-transHash}(st', \text{confTag})$ 32 $st.\text{pendComm}(\text{cmsg}) \leftarrow (st', PMSG, \text{upd}, \text{rem}, \text{add})$ 33 return $(\text{cmsg}, \text{wmsg})$ Algorithm CGKA.Key 34 $(k, st.appSec) \leftarrow (st.appSec, \perp)$ 35 return k </pre>	<pre> Algorithm CGKA.Proc(cmsg, PMSG) 36 $(\text{id-snd}, C, \text{confTag}, \sigma, \text{memTag}) \leftarrow \text{unframe-comm}(st, \text{cmsg})$ 37 if $\text{id-com} = st.id$ 38 $(st', PMSG', \text{upd}, \text{rem}, \text{add}) \leftarrow st.\text{pendComm}(\text{cmsg})$ 39 require $PMSG' = PMSG$ 40 $st \leftarrow st'$ 41 return 42 parse $(\text{proplds}, \text{UpdPath}) \leftarrow C$ 43 require $\text{proplds}(i) = H(PMSG(i))$ for all i 44 $st' \leftarrow \text{init-epoch}(st)$ 45 try $(st', \text{upd}, \text{rem}, \text{add}) \leftarrow \text{apply-props}(st, PMSG)$ 46 require $\text{id-com} \notin \text{rem} \wedge \text{id-com} \notin \text{upd}$ 47 if $st.id \in \text{rem}$ 48 require $\text{Mac.Ver}(st.\text{memKey}, \text{memTag})$ 49 $st \leftarrow \perp$ 50 else 51 $(st', \text{comSec}) \leftarrow \text{apply-rekey}(st, st', \text{id-com}, \text{UpdPath}, 0, \text{False})$ 52 $st' \leftarrow \text{set-conf-trans-hash}(st, st', \text{id-com}, C, \sigma)$ 53 $(st', \text{confKey}, \text{joinerSec}) \leftarrow \text{derive-keys}(st, st', \text{comSec})$ 54 require $\text{verify-confTag}(st', \text{confKey}, \text{confTag})$ 55 $st' \leftarrow \text{set-int-transHash}(st', \text{confTag})$ 56 $st \leftarrow st'$ 57 return $((\text{upd}, \text{rem}, \text{add}), \text{id-snd})$ Algorithm CGKA.Join(wmsg) 58 parse $(\text{encGrpSec}, \text{GrpInfo}) \leftarrow \text{wmsg}$ 59 parse $(\text{GrpInfoTBS}, \sigma) \leftarrow \text{GrpInfo}$ 60 parse $(st.groupID, st.ep, st.tHash, st.conf-TransHash,$ $st.interim-tranHash, st.confTag, \text{id-com}) \leftarrow \text{GrpInfoTBS}$ 61 require $\text{Sig.Vrfy}(v_{\text{id-com}}.svk, \sigma, \text{GrpInfoTBS})$ 62 try $\text{verify-treeState}(st)$ 63 try $\text{fetch-ssk-if-nec}(st, v_{st.id}.svk)$ 64 $\text{kbs} \leftarrow \text{get-sks}$ 65 $(\text{joinerSec}, \text{pathSec}) \leftarrow (\perp, \perp)$ 66 for $e \in \text{encGrpSec}$ 67 parse $(h, c) \leftarrow e$ 68 for $(kp, sk) \in \text{kbs}$ 69 if $h = H(kp)$ 70 $v.sk \leftarrow sk$ 71 require $v_{st.id}.kp = kp$ 72 parse $(\text{joinerSec}, \text{pathSec}) \leftarrow \text{Pke.Dec}(sk, c)$ 73 if $s_j = \text{pathSec} \neq \perp$ 74 $a \leftarrow st.T.\text{ind-lca}(v_{st.id}, v_{\text{id-com}})$ 75 $(v_1, \dots, v_\ell) \leftarrow \text{fil-path}(v)$ 76 for $i \in \{\ell - a, \dots, \ell - 1\}$ 77 $(pk, v_i.sk) \leftarrow \text{Pke.Gen}(H_2(s_i))$ 78 require $pk = v_i.pk$ 79 $s_{i+1} \leftarrow H_1(s_i)$ 80 $(st, \text{confKey}) \leftarrow \text{derive-epochKeys}(st, \text{joinerSec})$ 81 require $\text{verify-confTag}(st, \text{confKey}, \text{confTag})$ 82 return $(\text{roster}(T), st.id)$ </pre>
--	---

Figure 19: Algorithms CGKA.Create, CGKA.Com, CGKA.Proc, CGKA.Join, CGKA.Key shared between MLS and MLS-Cutoff. For subroutines rekey-path, apply-props, comp-wmsg, see Figure 22, 23, 18, respectively.

Algorithm CGKA.Prop(op, ad) (used in MLS)	Algorithm CGKA.Prop(op, ad) (used in MLS-Cutoff)
00 $T \leftarrow st.T$	18 $T \leftarrow st.T$
01 if op = 'upd':	19 if op = 'upd':
02 try st.ssk \leftarrow fetch-ssk-if-nec(st, svk)	20 try st.ssk \leftarrow fetch-ssk-if-nec(st, svk)
03	21 $st' \leftarrow$ init-epoch(st)
04 $(kp, sk) \leftarrow$ Gen-kp(svk, ssk)	22 $(id-own, svk) \leftarrow (st'.id-own, v_{id-own}\text{-svk})$
05 pmsg \leftarrow frame-prop(st, ('upd', kp))	23 try (st', ·, UpdPath, pathSec) \leftarrow rekey-path(st', id-own, svk, st.i _{cut})
06 st.pendUpd(pmsg) \leftarrow (ssk, sk)	24 pmsg \leftarrow frame-prop(st, ('upd', UpdPath))
07 if op = ('add'):	25 st.pendUpd(pmsg) \leftarrow (ssk, v'_{id-own}\text{-sk}, pathSec)
08 parse u \leftarrow ad	26 if op = ('add'):
09 require $u \notin$ roster(T)	27 parse u \leftarrow ad
10 send (fetch-kp, u) to \mathcal{F}_{KS} and get kp_a	28 require $u \notin$ roster(T)
11 try validate-kp(kp_a)	29 send (fetch-kp, u) to \mathcal{F}_{KS} and get kp_a
12 pmsg \leftarrow frame-prop(st, ('add', kp_a))	30 try validate-kp(kp_a)
13 if op = ('rem'):	31 pmsg \leftarrow frame-prop(st, ('add', kp_a))
14 parse u \leftarrow ad	32 if op = ('rem'):
15 require $u \in$ roster(T)	33 parse u \leftarrow ad
16 pmsg \leftarrow frame-prop(st, ('rem', id-rem))	34 require $u \in$ roster(T)
17 return pmsg	35 pmsg \leftarrow frame-prop(st, ('rem', id-rem))
	36 return pmsg

Figure 20: Generating proposals in MLS (left) and protocol variant MLS-Cutoff (right). Differences between the protocols are marked in yellow (•) and cyan (•). In line 25 v'_{id-own} is to be understood as the issuing user's leaf with respect to st' .

<pre> Algorithm set-tHash(st') 00 T' ← st'.T 01 st'.tHash ← tree-hash(T', root(T')) Algorithm tree-hash(T', v) 02 if v is leaf 03 return H(v.index, v.kp) 04 else 05 lHash ← tree-hash(T', left(v)) 06 rHash ← tree-hash(T', right(v)) 07 return H(v.index, v.pk, v.unm, v.pHash, lHash, rHash) Algorithm set-pHash(st', id) 08 T' ← st'.T 09 v₀ ← v_{id} 10 parse (v₁, ..., v_ℓ) ← fil-path(v₀) 11 v_ℓ.pHash ← ε 12 for j = ℓ - 1, ..., 0 13 v_j.pHash ← par-hash-cochild(v_{j+1}, sib(v_j)) 14 return st' Algorithm par-hash-cochild(T', v, u) 15 oriRes ← Res(u) ∪ (u.unm \ v.unm) 16 copy T' to T'' 17 for v'' ∈ v.unm 18 blank-path(T'', v'') 19 for w ∈ path(v'') 20 w.unm ← w.unm \ {v''} 21 origSibTreeHash ← tree-hash(T'', u) 22 return H(v.pk, v.pHash, origSibTreeHash) </pre>	<pre> Algorithm verify-treeState(st') 23 T' ← st'.T 24 require st'.tHash = tree-hash(T', root(T')) 25 for v ∈ T'.V 26 if v not blank ∧ v not leaf 27 l ← par-hash-cochild(v, left(v)) 28 r ← par-hash-cochild(v, right(v)) 29 require (left(v) not blank ∧ left(v).pHash = r) ∨ ∨ (right(v) not blank ∧ right(v).pHash = l) 30 M ← () 31 for v ∈ T'.V 32 if v not blank ∧ v is leaf 33 require id(v) ∉ M 34 M ← ∪ {id(V)} 35 try st' ← validate-kp(st', v.kp, id(v), v.pHash) 36 return st' </pre>
---	--

Figure 21: Helper functions set-tHash, set-pHash, and verify-treeState, used to authenticate the ratchet tree's state.

<p>Algorithm rekey-path(st', id, svk, i)</p> <pre> 00 $T' \leftarrow st'.T$ 01 $UpdPathNodes \leftarrow ()$ 02 $pathSec \leftarrow ()$ 03 $(v_1, \dots, v_\ell) \leftarrow \text{fil-path}(v_{id})$ 04 $s_0 \leftarrow_s \{0, 1\}^\lambda$ 05 for $j = 1, \dots, \ell - i$: 06 $s_j \leftarrow H_1(s_{j-1})$ 07 if $j < \ell$: 08 $(v_j.pk, v_j.sk_j) \leftarrow \text{Gen}(H_2(s_j))$ 09 $C \leftarrow ()$ 10 $v_{sib} \leftarrow \text{sib}(v_{j-1})$ 11 for $w \in \text{Res}(v_{sib}) \cup v_{sib}.unm$: 12 $C \leftarrow_{\cup} c_w := \text{Pke.Enc}(w.pk, s_j)$ 13 $UpdPathNodes \leftarrow_{\cup} (v_j.pk, C)$ 14 $pathSec \leftarrow_{\cup} (s_j)$ 15 if $i = 0$ 16 $comSec \leftarrow s_\ell$ 17 else $comSec = \perp$ 18 $\text{merge-leaves}(T', v_{id})$ 19 $\text{set-pHash}(T', v_{id})$ 20 try $\text{fetch-ssk-if-nec}(T', svk)$ 21 $(kp, sk) \leftarrow \text{Gen-kp}(id, svk, ssk, v.pHash; s_0)$ 22 require $v_{H(kp)} = \perp$ 23 $id(v_{id}) \leftarrow H(kp)$ 24 $\text{assign-kp}(T', H(kp), kp)$ 25 $v_{id}.sk \leftarrow sk$ 26 $\text{set-tHash}(T')$ 27 $UpdPath \leftarrow (UpdPathNodes, kp)$ 28 return $(T', comSec, UpdPath, pathSec)$ <p>Algorithm comp-wmsg($st', add, joinerSec, pathSec, confTag$)</p> <pre> 29 $tbs \leftarrow (st'.groupID, st'.ep, st'.tHash,$ $st'.conf-TransHash, st'.interim-tranHash,$ $st'.T.public, st'.confTag, st'.id)$ 30 $\sigma \leftarrow \text{Sign}(st'.ssk, tbs)$ 31 $GrpInfo \leftarrow (tbs, \sigma)$ 32 $\text{encGrpSec} \leftarrow ()$ 33 for $(\cdot, id-add, svk) \in add$ 34 $a \leftarrow \text{ind-lca}(v_{st'.id}, v_{id-add})$ 35 $c \leftarrow \text{Pke.Enc}(v_{id-add}.pk, (joinerSec, pathSec(\ell - a)))$ 36 $\text{encGrpSec} \leftarrow_{\cup} (H(v_{id-add}.kp), c)$ 37 return $(st', wmsg)$ </pre> </pre>	<p>Algorithm apply-rekey($st, st', id-snd, UpdPath, i, copySt$)</p> <pre> 38 $T' \leftarrow st'.T$ 39 $\text{parse}(UpdPathNodes, kp) \leftarrow UpdPath$ 40 $(v_1, \dots, v_\ell) \leftarrow \text{fil-path}(T', v_{st.id})$ 41 $(v'_1, \dots, v'_{\ell'}) \leftarrow \text{fil-path}(T', v_{id-snd})$ 42 $\text{parse}((pk_1, C_1), \dots, (pk_m, C_m)) \leftarrow UpdPathNodes$ 43 for $i \in \{1, \dots, m\}$ 44 $v'_i.pk \leftarrow pk_i$ 45 $a \leftarrow \text{ind-lca}(v_{st.id}, v_{id-snd})$ 46 if $i > a$: 47 $comSec \leftarrow \perp$ 48 else 49 $\text{parse}(c_w)_w \leftarrow C_{\ell-a}$ 50 if $v'_0 \in v'_{\ell'-a-1}.unm$: 51 $w \leftarrow v'_0$ 52 else: 53 $m \leftarrow \max(m' \in \{0, \dots, \ell' - a\} \mid v'_{m'} \text{ not blank})$ 54 $w \leftarrow v'_m$ 55 $s_{\ell-a} \leftarrow \text{Pke.Dec}(sk_w, c_w)$ 56 for $j = \ell' - a + 1, \dots, \ell' - i$: 57 $s_j \leftarrow H_1(s_{j-1})$ 58 if $j < \ell$ 59 $(pk_j, sk_j) \leftarrow \text{Gen}(H_2(s_j))$ 60 $v'_j.sk \leftarrow sk_j$ 61 require $v'_j.pk = pk_j$ 62 if $i = 0$ 63 $comSec \leftarrow s_\ell$ 64 else $comSec = \perp$ 65 $\text{merge-leaves}(T', v_{id})$ 66 $st' \leftarrow \text{set-pHash}(T', v_{id})$ 67 if $copySt = \text{False}$ 68 try $\text{validate-kp}(st'', kp, id-snd, st''.pHash)$ 69 else 70 $st'' \leftarrow \text{init-epoch}(st)$ 71 $T'' \leftarrow st''.T$ 72 $(v''_1, \dots, v''_{\ell''}) \leftarrow \text{fil-path}(T'', v_{id-snd})$ 73 for $i \in \{1, \dots, m\}$ 74 $v''_i.pk \leftarrow pk_i$ 75 $\text{merge-leaves}(T'', v_{id})$ 76 $st'' \leftarrow \text{set-pHash}(T'', v_{id})$ 77 try $\text{validate-kp}(st'', kp, id-snd, st''.pHash)$ 78 require $v_{H(kp)} = \perp$ 79 $id(v_{id-snd}) \leftarrow H(kp)$ 80 $\text{assign-kp}(T', H(kp), kp)$ 81 $\text{set-tHash}(T')$ 82 return $(st', comSec)$ </pre>
--	---

Figure 22: Helper functions rekey-path and apply-rekey used to rekey users' paths.

```

Algorithm apply-props(st, st', PMSG) (used in MLS)
00 (upd, rem, add)  $\leftarrow$  ( $\emptyset$ ), ( $\emptyset$ ), ( $\emptyset$ )
01 fetch  $T' \leftarrow st'.T$ 
02 for pmsg  $\in$  PMSG:
03   ((‘type’, ad), id-snd)  $\leftarrow$  unframe-prop(st, pmsg)
04   if ‘type’ = ‘upd’:
05     require (id-snd,  $\cdot$ )  $\notin$  upd  $\wedge$  rem = ( $\emptyset$ )  $\wedge$  add = ( $\emptyset$ )

06     try validate-kp(st', ad, id-snd,  $\varepsilon$ )
07     assign-kp( $T'$ ,  $v_{id-snd}$ , ad)
08     if st'.id = id-snd

09     parse (ssk, sk)  $\leftarrow$  pendUpd(pmsg)
10      $v_{id-snd}.sk \leftarrow sk$ 
11     st'.ssk  $\leftarrow$  ssk

12     blank-path( $T'$ ,  $v_{id-snd}$ , 1)

13     svk  $\leftarrow$  leaf( $T'$ , id-snd).svk
14     require  $v_{H(kp)} = \perp$ 
15     id( $v_{id-snd}$ )  $\leftarrow$  H(kp)
16     upd  $\leftarrow_{\cup}$  (id-snd, ‘upd’, svk)
17   else if ‘type’ = ‘rem’:
18     require id-rem := ad  $\neq$  id-snd  $\wedge$   $v_{id-rem} \neq \perp$ 
19     require (ad,  $\cdot$ )  $\notin$  upd  $\wedge$  add = ( $\emptyset$ )
20     blank-leaf( $T'$ ,  $v_{id-rem}$ )
21     blank-path( $T'$ ,  $v_{id-rem}$ , 1)
22      $T' \leftarrow$  truncate( $T'$ )
23     rem  $\leftarrow_{\cup}$  (id-snd, ‘rem’, id-rem)
24   else if ‘type’ = ‘add’:
25     try st'  $\leftarrow$  validate-kp(st', ad,  $u_{new}$ ,  $\varepsilon$ )
26     id-add  $\leftarrow$  H(ad)
27     add-leaf( $T'$ , id-add)
28     set-as-unmerged( $T'$ , id-add)
29     add  $\leftarrow_{\cup}$  (id-snd, ‘add’,  $v_{id-add}.svk$ )
30   else return  $\perp$ 
31 return (st', upd, rem, add)

```

```

Algorithm apply-props(st, st', PMSG) (used in MLS-Cutoff)
32 (upd, rem, add)  $\leftarrow$  ( $\emptyset$ ), ( $\emptyset$ ), ( $\emptyset$ )
33 fetch  $T' \leftarrow st'.T$ 
34 for pmsg  $\in$  PMSG:
35   ((‘type’, ad), id-snd)  $\leftarrow$  unframe-prop(st, pmsg)
36   if ‘type’ = ‘upd’:
37     require (id-snd,  $\cdot$ )  $\notin$  upd  $\wedge$  rem = ( $\emptyset$ )  $\wedge$  add = ( $\emptyset$ )
38     if st'.id  $\neq$  id-snd

39       ( $st'.\cdot$ )  $\leftarrow$  apply-rekey(st, st', id-snd, ad, st'. $i_{cut}$ , True)
40     if st'.id = id-snd
41       parse ((pk1,  $\cdot$ ), ..., (pk $\ell - i_{cut}$ ,  $\cdot$ ), kp)  $\leftarrow$  ad
42       assign-kp( $T'$ ,  $v_{id-snd}$ , kp)
43       parse (ssk, sk, pathSec)  $\leftarrow$  pendUpd(pmsg)
44        $v_{id-snd}.sk \leftarrow sk$ 
45       st'.ssk  $\leftarrow$  ssk
46       parse ( $v_1, \dots, v_\ell$ )  $\leftarrow$  path(id-snd)
47       for  $j = 1, \dots, \ell - st'.i_{cut}$ 
48         ( $v_j.pk, v_j.sk$ )  $\leftarrow$  (pk $j$ , pathSec( $j$ ))
49       blank-path( $T'$ ,  $v_{id-snd}$ ,  $\ell - st'.i_{cut} + 1$ )
50       parse ( $v_1, \dots, v_\ell$ )  $\leftarrow$  path( $v_{id-snd}$ )
51       for  $j = 1, \dots, \ell - st'.i_{cut}$ 
52         if  $v_j \in V_{col}$ 
53           blank( $v_j$ )
54            $V_{col} \leftarrow_{\cup} \{v_j\}$ 
55       svk  $\leftarrow$  leaf( $T'$ , id-snd).svk
56       require  $v_{H(kp)} = \perp$ 
57       id( $v_{id-snd}$ )  $\leftarrow$  H(kp)
58       upd  $\leftarrow_{\cup}$  (id-snd, ‘upd’, svk)
59   else if ‘type’ = ‘rem’:
60     require id-rem := ad  $\neq$  id-snd  $\wedge$   $v_{id-rem} \neq \perp$ 
61     require (ad,  $\cdot$ )  $\notin$  upd  $\wedge$  add = ( $\emptyset$ )
62     blank-leaf( $T'$ ,  $v_{id-rem}$ )
63     blank-path( $T'$ ,  $v_{id-rem}$ , 1)
64      $T' \leftarrow$  truncate( $T'$ )
65     rem  $\leftarrow_{\cup}$  (id-snd, ‘rem’, id-rem)
66   else if ‘type’ = ‘add’:
67     try st'  $\leftarrow$  validate-kp(st', ad,  $u_{new}$ ,  $\varepsilon$ )
68     id-add  $\leftarrow$  H(ad)
69     add-leaf( $T'$ , id-add)
70     set-as-unmerged( $T'$ , id-add)
71     add  $\leftarrow_{\cup}$  (id-snd, ‘add’,  $v_{id-add}.svk$ )
72   else return  $\perp$ 
73 return (st', upd, rem, add)

```

Figure 23: Helper function apply-props for MLS (left) and MLS-Cutoff (right). Differences between the protocols are marked in yellow (•) and cyan (•).