

How to Hide MetaData in MLS-Like Secure Group Messaging: Simple, Modular, and Post-Quantum

Keitaro Hashimoto*
Tokyo Institute of Technology & AIST
Tokyo, Japan
hashimoto.k.au@m.titech.ac.jp

Shuichi Katsumata†
AIST & PQShield Ltd.
Tokyo, Japan & Oxford, U.K.
shuichi.katsumata@aist.go.jp

Thomas Prest
PQShield SAS
Paris, France
thomas.prest@pqshield.com

ABSTRACT

Secure group messaging (SGM) protocols allow large groups of users to communicate in a secure and asynchronous manner. In recent years, continuous group key agreements (CGKAs) have provided a powerful abstraction to reason on the security properties we expect from SGM protocols. While robust techniques have been developed to protect the *contents* of conversations in this context, it is in general more challenging to protect *metadata* (e.g. the identity and social relationships of group members), since their knowledge is often needed by the server in order to ensure the proper function of the SGM protocol.

In this work, we provide a simple and generic wrapper protocol that upgrades non-metadata-hiding CGKAs into metadata-hiding CGKAs. Our key insight is to leverage the existence of a unique continuously evolving group secret key shared among the group members. We use this key to perform a group membership authentication protocol that convinces the server in an *anonymous* manner that a user is a legitimate group member. Our technique only uses a standard signature scheme, and thus, the wrapper protocol can be instantiated from a wide range of assumptions, including post-quantum ones. It is also very efficient, as it increases the bandwidth cost of the underlying CGKA operations by at most a factor of two.

To formally prove the security of our protocol, we use the universal composability (UC) framework and model a new ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ capturing the correctness and security guarantee of metadata-hiding CGKA. To capture the above intuition of a “wrapper” protocol, we also define a restricted ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$, which roughly captures a non-metadata-hiding CGKA. We then show that our wrapper protocol UC-realizes $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ in the $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ -hybrid model, which in particular formalizes the intuition that any non-metadata-hiding CGKA can be modularly bootstrapped into metadata-hiding CGKA.

KEYWORDS

secure group messaging; metadata-hiding; messaging layer security; continuous group key agreement; generic construction; post-quantum security

1 INTRODUCTION

A *secure group messaging* (SGM) protocol allows a group of users to asynchronously communicate in an *end-to-end encrypted* (E2EE) fashion. With the advent of Signal, SGM protocols have seen exponential growth in practical use, and currently, IETF is almost

finished standardizing the Messaging Layer Security (MLS) SGM protocol [14] known to offer good scalability properties.

The baseline security notion guaranteed by all SGM protocols is that no adversary, including the server, should be able to read the messages sent among the group members. However, this does not prevent the adversary from collecting *metadata*, such as the identities of the sender and of other group members, which can both be leaked from the exchanged encrypted contents. It has been shown in numerous real-world scenarios [23, 34, 49, 51, 54] that knowledge of metadata alone can cause damaging repercussions, sometimes enough to defeat the purpose of using SGM protocols. These also have negative impacts on the activity and safety of some users, e.g., journalists and activists [47, 48]. Recent media articles [18, 42] report that, in the United States, metadata collection by law enforcement agencies on *secure messaging applications* is a widespread practice, supported by a legal [42] and technical framework [18] that gives it a wide reach.

Metadata in SGM. In SGM protocols, we can informally divide sensitive information into the following three layers:

$$\left\{ \begin{array}{l} \text{1st layer: group secret keys \& messages} \\ \text{2nd layer: static, explicit metadata} \\ \text{3rd layer: dynamic, implicit metadata} \end{array} \right\} =: \text{“metadata”}$$

Securing the 1st layer is the *default* goal of any SGM protocol; exchanging messages in an E2EE fashion is only possible if secure group secret keys¹ are shared among the group. Since the server is not considered an endpoint of the conversation, state-of-the-art SGM protocols aim at protecting the 1st layer from the server.

The 2nd and 3rd layers together constitute the metadata. Since they help the server to ensure the functionality of the SGM protocol, they are often only encrypted using a transport layer encryption protocol (like TLS or Noise [52]) between the server and the participants. In this case, the server has access to this information and may expose it if legally compelled, as discussed earlier.

The 2nd layer captures any *static* metadata that is *explicitly* leaked from the content transmitted over the channel. For instance, the exchanged content may explicitly include the identity of the sender in the clear or the identity of a member being added, as in e.g. vanilla Signal [45] and MLSPlaintext [14]. Static metadata are also defined as a collection of *sender information* and *handshake messages* in the MLS standard draft [14, Sec. 10.1].

The 3rd layer captures any *dynamic* metadata that is *implicitly* leaked from the access pattern between group members and the

*The author was partially supported by JSPS KAKENHI Grant Number 22J13963, Japan.

†The author was partially supported by JSPS KAKENHI Grant Number 22K17892, Japan and JST AIP Acceleration Research JPMJCR22U5, Japan.

Table 1: SGM protocols and the corresponding layers they protect. “✓” (resp. “✗”) indicates that there is a (resp. are no) security proof. “✓” indicates that there is a proof capturing the security of the 2nd & 3rd layers but not of the 1st layer.

Layer	1st	1st & 2nd	1st, 2nd & 3rd
Signal	Vanilla Signal	⊥	Private Groups
Security proof	✗	⊥	✓ [25]
MLS	MLSPplaintext	MLSCiphertext	⊥
Security proof	✓ E.g., [9, 12]	✗	⊥

server via the communication channel. For example, in MLSCiphertext that hides up to the 2nd layer, users connect to the server using a non-anonymous protocol such as TLS. Then, since the user implicitly identifies itself to the server via the channel, the server learns the group member’s identities (and how many times each member accessed) by observing the users accessing the same group identity. In particular, this happens regardless of protecting the 2nd layer metadata. To hide the user identity on the channel, users can use anonymous protocols (e.g., Tor [1, 31]) instead. However, even if an anonymous protocol is used when a user fetches information about a group they belong to, the exact subset of accessed information may be correlated to this user’s identity. This may for example be the case for SGM protocols that arrange group members in complex data structures such as trees. Specifically, hiding the metadata only at the 2nd layer while using anonymous channels is insufficient since similar information may be inferred from the 3rd layer, which incorporates all implicit leakages of dynamic metadata.

In this work, only when all three layers are secured do we say that an SGM protocol is *metadata-hiding*.²

Existing metadata-hiding SGM. Existing SGM protocols and the level of layers they protect are depicted in Tab. 1. Signal recently proposed a metadata-hiding SGM protocol that we call *Private Groups* [3]. This is an extension of *Sealed Sender* [2] — a metadata-hiding *two-user* secure messaging protocol. The main building block of Private Groups is an efficient *MAC-based keyed-verification anonymous credential* (KVAC) [24] that leverages the specific properties of tools from classical group-based cryptography, such as the ElGamal PKE and Schnorr PoK. While there is no formal security proof for Signal’s vanilla SGM, recently Chase, Perrin, and Zaverucha [25] proposed a new security model to capture exactly the metadata layers (2nd & 3rd layers) and provided a partial security proof of Private Groups.

MLS [14] comes in two variants: MLSPplaintext and MLSCiphertext, each corresponding to protocols protecting the 1st and 1st & 2nd layers, respectively. The security of MLSPplaintext has been scrutinized over the past few years [8, 9, 12, 17, 19] and we now have a good understanding of it. However, no formal security proof for MLSCiphertext is known. Moreover, unlike Signal’s Private Groups, constructing any variant of MLS that further hides the *dynamic* metadata is unanswered. Considering that dynamic metadata leaks part of, if not all, static metadata, MLSCiphertext may be leaking more metadata in practice than ideally expected.

²Note that there are other types of metadata we can consider such as access timing [46] and geolocation of users.




Figure 1: (Left) p and c are proposals and commits. (Right) id_1 uploads a proposal p_1 to the server S_v ; id_2 downloads all the stored proposals \bar{p} and uploads a commit c_2 for the next epoch. The single (resp. double, triple) bordered box indicates the 1st (resp. 2nd, 3rd) layer information.

1.1 Goal of This Work

In this work, we focus on *continuous group key agreement* (CGKA) — an abstraction that captures the core protocol underlying the MLS protocol, i.e., TreeKEM [16], and many other MLS-inspired SGM protocols [5, 8–12, 30, 36, 41].³ In brief, CGKA allows an evolving group of users to agree on a continuous sequence of group secret keys. Other than the simple function of sharing a group secret key, CGKA further models the strong notions of *forward secrecy* (FS) and *post-compromise security* (PCS) [7, 28, 29], which allow to greatly limit the scope of a compromise.

In a nutshell, a CGKA works as follows (see also Fig. 1):⁴ A group member may either (a) add a new member, (b) remove a member, and/or (c) update its keys by sending a *proposal* p . In an arbitrary interval, a group member may download the list of proposals $\bar{p} = \{p_i\}_i$ from the server and take them into effect by transmitting a *commit* c — this creates a new epoch where the group state is updated according to \bar{p} . Importantly, a commit also updates the group secret key to achieve PCS.

A proposal p consists of five elements: (i) a string gid identifying the group; (ii) a counter epoch that specifies the current group state; (iii) the identity id_p of the member creating p ; (iv) a string act specifying whether p corresponds to (a), (b), or (c); and (v) other information typically required for authentication. A commit c has a similar structure, where id_c denotes the committer and ct_{key} is a ciphertext encrypting a key used to update the group secret.

As depicted in Fig. 1, key is the 1st layer information, and any other static information included in p and c other than $(\text{gid}, \text{epoch})$ belong to the 2nd layer. Here, $(\text{gid}, \text{epoch})$ needs to be clear so that the receiver can download the appropriate p and c from the server. In the past few years, we have seen several increasingly stronger or different types (e.g., game-based, simulation-based) of security models for CGKA [5, 8–12, 19, 37, 41, 57], however these models only capture security at the 1st layer. Although it is straightforward to construct a CGKA that intuitively secures the 2nd layer once a group secret key is established, it is not clear whether this intuition is correct. Indeed, MLSPplaintext has undergone 13 iterations, and formal security analyses of the 1st layer [8, 12] uncovered some subtle bugs. Thus, our first goal is the following:

- (G1) *Propose a security model capturing the security of the 1st & 2nd layers and prove the security of existing CGKAs.*

As discussed above, securing the 2nd layer alone is insufficient. At first glance, it is tempting to replace the use of TLS for client-server communication with a client-anonymized authenticated channel

³To be accurate, MLS was inspired by *asynchronous ratchet trees* (ART) [30].

⁴We base the explanation on the most recent iterations of TreeKEM (i.e., after version 8 on MLS) following a “propose-and-commit” flow.

(e.g., VPN or an anonymized proxy such as Tor [1, 31]) in order to hide the 3rd layer. Unfortunately, this introduces another issue since, without any authentication on the client side, any adversary who knows (gid, epoch) can upload arbitrary garbage proposals and commits to the server, causing a denial of service (DoS) against the group. It could be possible to rely on the efficient MAC-based KVAC used by Signal’s Private Groups [3], however, their construction is highly limited to a classical, pre-quantum setting, and the security proof is in the generic group model [55]. Considering the modularity of the vanilla Signal and the MLS protocol, having a generic construction that can be efficiently instantiated from versatile assumptions, including but not limited to *post-quantum* assumptions, is highly desirable. Of independent interest, we note that in the face of a compromise or removal of a group member, Private Groups must restart a new group [25]. It remains an interesting problem to construct a protocol that offers any (non-trivial) PCS. This brings us to our second goal:

(G2) *Propose an efficient and generic metadata-hiding CGKA achieving the same level of FS and PCS offered by existing non-metadata-hiding CGKAs.*

Finally, Chase, Perrin, and Zaverucha [25] proposed a security model capturing the 2nd & 3rd layers of Signal. However, it does not capture the 1st layer of security, i.e., confidentiality and integrity of exchanged messages, nor the notion of PCS. Moreover, this model is tailored to the specific construction of Signal’s Private Groups [3] and seems unfit for CGKA. Thus, we arrive at our final goal:

(G3) *Propose a security model for metadata-hiding CGKA.*

1.2 Our Contributions

UC model for the 2nd layer. We propose the first security model of CGKA capturing the security of the 2nd layer, i.e., static metadata. Our security model extends the state-of-the-art universal composability (UC) security model used by Alwen et al. [12] and Hashimoto et al. [36] to analyze TreeKEM version 10 in MLS and Chained CmpPKE, respectively – we denote the ideal functionality as $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$.⁵ The main new ingredients we introduce are *leakage functions* that allow us to formally model the leaked static metadata (e.g., sender’s identity) from the proposals and commits. Similar to the state-of-the-art ideal functionalities, ours captures a strong model where *active adversaries* can tamper with or inject messages, and *malicious insiders* can invite malicious members to the group and arbitrarily fork the group state. With this formalization effort, we answer the first half of (G1), see Sec. 3 for details.

Chained CmpPKE^{ctxt} UC-realizes $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. We prove that a ciphertext variant of Chained CmpPKE by Hashimoto et al. [36], coined as Chained CmpPKE^{ctxt}, UC-realizes the ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$.⁶ Considering the similarity between Chained CmpPKE and TreeKEM, we believe the ciphertext variant of TreeKEM can also be proven secure following a similar proof. The reason why we focused on the former is that it is in some sense a generalization of the latter – it allows members to *selectively download* updates from the server, also known as *filtered* CGKA [11]. This generalization allows obtaining a concretely efficient CGKA even in the *post-quantum* regime,

⁵The subscript “ctxt” is inspired by the protocol name MLSCiphertext

⁶This variation mirrors what MLSCiphertext does to MLSPplaintext, in the sense that it encrypts the static metadata by applying a layer of encryption.

which otherwise could be quite inefficient [36]. This security proof addresses the second half of (G1), see Sec. 3 for details.

UC model for the 3rd layer. We propose the first UC security model of CGKA capturing the security of the 3rd layer, i.e., group access pattern – we denote the ideal functionality as $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$. Any CGKA that UC-realizes $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ is a *metadata-hiding* CGKA. The model captures the fact that a group member performing an upload or download remains anonymous and unlinkable from the server, while also restricting non-group members from accessing the group contents. To formalize the latter property, our ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ captures an *honest-but-curious* server for the first time. All prior models only considered *malicious* servers so it was not possible to define a “correct” behavior of the server, i.e., shutting out non-group members. $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ allows the adversary to corrupt the server, in which case it becomes identical to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ defined above. This answers (G3), see Sec. 5 for details.

A generic and efficient protocol UC-realizing $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$. We provide a simple and generic wrapper protocol W^{mh} that UC-realizes $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ in the $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ -hybrid model. Specifically, given an *arbitrary* CGKA Π_{ctxt} that UC-realizes $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$, W^{mh} in composition with Π_{ctxt} UC-realizes $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$. Unlike Signal’s Private Groups, we do not rely on complex tools such as a MAC-based KVAC, whose known efficient instantiations require classical group-based assumptions. Our key insight is to leverage the unique group secret key shared among the members (which is non-existing in Signal) to perform a proof of membership to the server. The concrete construction of our wrapper protocol only requires a standard signature scheme, which can be efficiently instantiated using either classical or post-quantum assumptions. Our metadata-hiding CGKA inherits all the FS and PCS properties satisfied by the underlying CGKA satisfying $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. For instance, using MLSCiphertext as the underlying CGKA, the upload cost of a key update of our metadata-hiding CGKA can be $O(\log N)$, as opposed to $O(N)$ for Private Groups. This provides a theoretic answer to (G2), see Sec. 4 for details.

Instantiation and efficiency analysis. We provide concrete instantiations of our protocols under either classical or post-quantum assumptions. We then study the bandwidth impact of W^{mh} when applied to Chained CmpPKE^{ctxt}. The impact of W^{mh} is moderate, as it never increases the bandwidth cost of the principal operations (“update”, “add” and “remove” proposals, as well as commit or application messages) by more than a factor of two. In practice, the concrete overhead may be even lower. This illustrates that our notion of metadata-hiding CGKA can be realized at a moderate cost. This covers the efficiency aspects of (G2), see Sec. 6.2 for details.

Statistical leakage from metadata-hiding CGKA. Our security model allows us to prove that a CGKA UC-realizes an ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ *with respect to a specific leakage function*, which defines any inherent metadata leakage that cannot be hidden. What the adversary can learn from this leakage function is another question.

We initiate a discussion on the nature and extent of the information that can be inferred from this leakage. We study several CGKAs [5, 11, 12, 36, 41] and find that all of them leak information through the size of protocol messages (welcome, proposal and/or commit), sometimes in surprising and indirect ways. At this point, we emphasize that the authors of these protocols never claimed

them to be metadata-hiding, so this leakage does not reflect the shortcomings of the designs. We believe that a systematic study of this leakage, as well as proposing countermeasures to provably mitigate it, constitutes a valuable and exciting research direction. For now, this discussion explores the limitations of (G3), see Sec. 7 for details.

We would like to clarify the limitations of this work. First, we only consider the CGKA aspect of SGMs. While it is believed that CGKA captures the essence of SGMs — which is supported by the vast amount of research focusing solely on CGKA [5, 8–12, 19, 36, 41, 57] — to argue a provably secure metadata-hiding SGM, we would need to extend our security model to cover the message exchanging layer as well. It was only recently that a security model that captures the entire SGM at the 1st layer was proposed [9]. Second, our security model and protocol do not prevent an adversary from anonymously registering numerous fake groups on the server. We only prevent an outsider from accessing an existing group. Technically, this seems efficiently solvable using standard anonymous credentials [26] and we leave it as future work to incorporate these into our security model. Finally, metadata outside the scope of our models, such as access timing [46] and device fingerprinting, may circumvent the privacy guarantees provided by a metadata-hiding CGKA.

2 BACKGROUND

We now provide some background on the existing definition of CGKA that models the default security of the 1st layer, i.e., *group secret key*.⁷ In this work, we focus on CGKA defined in the UC framework. With the nice composability property, we will be able to construct a metadata-hiding CGKA in a modular manner.

Due to page limitation, we refer the readers to App. B for some background on the UC framework and a formal definition of the ideal functionality $\mathcal{F}_{\text{CGKA}}$ capturing the UC security of the 1st layer.

2.1 Syntax of CGKA

All the group members internally store the group identifier gid , the current epoch of the group, the current group secret key, and information to identify the current members. Moreover, each member (roughly) stores a public key whose associated secret key is known only to themselves. This is called the *key material*.

Members can upload *proposals* p to the server, which can take on three types: addition or removal of some members, and update of key materials. Members can download any subset of these proposals \vec{p} at an arbitrary interval to create a *commit* c . Any members can then *process* this commit to update their group state to the next epoch according to the content of the proposals that were committed. The group secret key is always updated after a process — if every group member’s key material was uncompromised at that epoch, then this effectively *heals* the group and offers PCS.

In this work, we follow the syntax of Hashimoto et al. [36] (which extends the syntax of Alwen et al. [12]) that models *selective downloading*. This allows the members to only download part of the

commit required to move to the next epoch. Specifically, a commit c is divided into (c_0, \vec{c}) , where c_0 is the member-*independent* commit and $\vec{c} = (\widehat{c}_{id})_{id}$ is the list of member-*dependent* commit. Member id only needs to download (c_0, \widehat{c}_{id}) from the server to advance its epoch. Here, we assume there is a canonical ordering of the group members, and member id requests its index in the list \vec{c} to the server to download $\vec{c}[\text{index}] = \widehat{c}_{id}$. Such selective downloading can bring great efficiency gain — especially in the post-quantum regime where asymmetric ciphertexts tend to be much larger than classical ones (see [36] for further motivations). This formalization was later coined as *filtered CGKA* [11].

Informally, CGKA is defined by the following algorithms, where we assume id is the executing party and omit it from the input.

Group Creation (Create): It initializes a new group state with party id as the only member.⁸

Proposals (Propose, act) $\rightarrow p$: It outputs a proposal p for the *action* act that can take on the value ‘add’- id_t , ‘rem’- id_t , or ‘upd’. The first two actions dictate the adding or removal of id_t . The last updates id_t ’s key material.

Commit (Commit, \vec{p}) $\rightarrow (c_0, \vec{c}, \vec{w})$: It commits a vector of proposals \vec{p} and outputs a commit (c_0, \vec{c}) . c_0 is a member-independent commit while $\vec{c} = (\widehat{c}_{id'})_{id'}$ is a list of member-dependent commits, where $|\vec{c}|$ is equal to the current group size. If \vec{p} contains an add proposal, then it outputs a welcome message $\vec{w} = (\widehat{w}_{id_t})_{id_t}$, where id_t denotes the added members.⁹

Process (Process, $c_0, \widehat{c}_{id}, \vec{p}$): It processes a commit (c_0, \widehat{c}_{id}) with the associated proposals \vec{p} , and advances id ’s internal group state to the next epoch.

Join (Join, \widehat{w}_{id}): It allows id to join the group using the welcome message \widehat{w}_{id} . id ’s group state is synced with any member who processes the commit made at the same epoch.

Key (Key) $\rightarrow k$: It outputs the current group secret key k .

2.2 Default UC Security Model of CGKA

All prior works on CGKA capture the rough security notion that the group secret key should remain hidden from the adversary. We follow the state-of-the-art UC security model of [12, 36]. In this model, the malicious server is modeled as an *active adversary* that can tamper with or inject messages. Moreover, a rogue group member is modeled as a *malicious insider* that can invite other malicious members (e.g., server) to the group and arbitrarily fork the group state. The UC security model captures the fact that even in face of such strong adversaries, the group secret key remains indistinguishable from random under certain conditions. Below, we explain the high-level description of the ideal functionality $\mathcal{F}_{\text{CGKA}}$.

History graphs and the safe predicate. The core concept underlying the definition is the so-called *history graph* [9, 12] maintained by $\mathcal{F}_{\text{CGKA}}$. A history graph is a symbolic representation of the group’s evolution, where each node on a graph roughly corresponds to a group state at a particular epoch.¹⁰ It tracks all the generated commits and proposals, and group members’ positions

⁸Following prior definitions [10, 12, 36], we assume Group Creation is run only once.

⁹Although we can also structure \vec{w} to have a party independent w_0 and dependent part, we chose not to do so since it leads to a less secure metadata-hiding protocol. Roughly, by looking at w_0 , the server can infer who will be added to the same group.

¹⁰The formal definition is made with more care since in case an adversary forks the group state, two different nodes with the same epoch can be created.

⁷As mentioned in Sec. 1.2, CGKA only captures the security of the group secret key k . The message, also included in the 1st layer, is handled by a different protocol. Roughly, the messages are sent through E2EE using the established k .

on the history graph. To define the security of the group secret key, \mathcal{F}_{CGKA} is parameterized by a predicate **safe**, which takes the history graph and a node as input, and assigns a random group secret key $k \leftarrow \mathcal{K}$ to a node where **safe** is true. This formalization allows modeling FS and PCS naturally. For example, an adversary \mathcal{A} can corrupt a member at some epoch, thus making **safe** false at epoch. If the member is *healed* at a later epoch' $>$ epoch, then **safe** becomes true again at epoch'. Importantly, **safe** is a scheme-specific predicate that can be defined arbitrarily to capture different levels of FS and PCS.

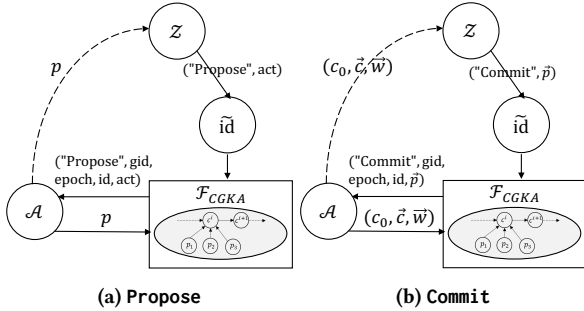


Figure 2: Group operations in the UC security model. \mathcal{Z} invokes the dummy party \tilde{id} , and \mathcal{F}_{CGKA} informs the adversary \mathcal{A} of this invocation. \mathcal{A} simulates the corresponding proposal or commit and sends it to \mathcal{F}_{CGKA} . W.l.o.g., we assume \mathcal{A} sends the same information to \mathcal{Z} , denoted by a dashed line. The shaded region denotes the history graph maintained by \mathcal{F}_{CGKA} .

Example: Ideal propose and commit. We depict the ideal Propose and Commit functions in Fig. 2, where \mathcal{Z} denotes the environment. For example, \mathcal{Z} can invoke the (dummy) party \tilde{id} to execute a commit on an arbitrary list of proposals \vec{p} . \mathcal{F}_{CGKA} then provides to \mathcal{A} the proposals \vec{p} along with all the static metadata (gid, epoch, id) included in a commit, where id denotes the identity of the committer. \mathcal{A} interprets the proposals \vec{p} and simulates the commit (c_0, \vec{c}) , where it further simulates the welcome message \vec{w} if \vec{p} includes an add proposal. The commit and welcome messages are sent to \mathcal{F}_{CGKA} and are registered in the history graph as a new node indicating that a new epoch has been created. Here, if \vec{p} was generated by a group member, then \mathcal{F}_{CGKA} mandates correctness by checking if the commit output by \mathcal{A} was consistent with the actions included in \vec{p} . Otherwise, in case some $p \in \vec{p}$ was not generated by a group member, then this implies that the server or some rogue insider injected a *malicious* proposal.

3 HIDING STATIC METADATA IN CGKA

As a first step, we propose a UC security model capturing the security of the 1st & 2nd layers (i.e., group secret keys and *static* metadata) by defining a new ideal functionality $\mathcal{F}_{CGKA}^{\text{ctxt}}$. Similarly to the **safe** predicate used in \mathcal{F}_{CGKA} to control the levels of FS and PCS, $\mathcal{F}_{CGKA}^{\text{ctxt}}$ comes with five *leakage functions* allowing to control the amount of static metadata leaked from create group, proposal, commit, process, and join. By defining the leakage functions to leak all the static metadata, then $\mathcal{F}_{CGKA}^{\text{ctxt}}$ (essentially) recovers the prior ideal functionality \mathcal{F}_{CGKA} .

We then prove that a ciphertext variant of Hashimoto et al.'s Chained CmpKE [36] satisfies $\mathcal{F}_{CGKA}^{\text{ctxt}}$, where the leakage functions are defined to only leak the minimal static metadata matching our intuition.

3.1 UC Security Model for Static Metadata

$\mathcal{F}_{CGKA}^{\text{ctxt}}$ has the same user interface (or syntax) as \mathcal{F}_{CGKA} . The only difference is how the internals of the ideal functionalities are defined. The full details on $\mathcal{F}_{CGKA}^{\text{ctxt}}$ is provided in App. B. Below, we explain the two main points at which $\mathcal{F}_{CGKA}^{\text{ctxt}}$ differs from \mathcal{F}_{CGKA} .

Modeling static metadata leakage. Recall how \mathcal{F}_{CGKA} defined the ideal proposal function (see Fig. 2). When a party id is invoked on (Propose, act) from the environment \mathcal{Z} , \mathcal{F}_{CGKA} informs the adversary \mathcal{A} with (Propose, gid, epoch, id, act). \mathcal{A} then simulates a proposal p . This models the fact that p in the real world is allowed to leak information on (gid, epoch), the party id creating p , and the type of proposal act $\in \{\text{'add'-id}_t, \text{'rem'-id}_t, \text{'upd'}\}$ included in p .

We control the amount of such leakage from p by a leakage function ***leak-prop**. Informally, ***leak-prop** takes as input the identity id of the member creating the proposal and the current epoch. In case **safe** is false at epoch, ***leak-prop** outputs all the static metadata since the group secret key is compromised at that epoch. Otherwise, it only outputs the static metadata the CGKA is allowed to leak. For example, to model MLSCiphertext, we define ***leak-prop**(id, epoch) to only output (gid, epoch, |id|, |act|) when **safe** is true, where |act| leaks the size of the action included in p . In case every member identity id_t is encoded in the same bit-length, then this implies for instance that p does not leak who created the p and who was added and removed.

Similarly to ***leak-prop**, we define four more leakage functions ***leak-create**, ***leak-proc**, ***leak-com**, and ***leak-wel**. For instance, the last two controls the amount of static metadata leaked from commits and welcome messages, respectively. Unlike proposals, the static metadata that is inherently leaked from commits and welcome messages differs between CGKAs. Thus, care is required when formally defining them. For instance, in MLSCiphertext, a welcome message for member id_t includes a hash of id_t 's key package kp_{id_t} (i.e., key materials used by id_t). To model this fact, ***leak-wel** must also output kp_{id_t} to the adversary (see App. C.2 for more discussion). Another subtlety is that the size of a commit in MLSCiphertext is dictated by how many blank nodes exist in the tree. In Sec. 7, we discuss the real-world consequences of these inherent leakages of static metadata.

Using semantics for nodes in history graphs. In this work, we update the prior definition of history graphs to use the *semantics* of a transcript to identify the nodes in a history graph. That is, we identify a node by a counter that informally counts the number of group operations leading to the node. In previous works [12, 36] (which did not capture the security of the static metadata), each node was identified by the *value* of the *non-encrypted* member-independent commit c_0 . Since c_0 uniquely defined an epoch and group state, it made intuitive sense to identify each node by c_0 .

Unfortunately, this intuition is lost when we try to further secure the static metadata. This is because c_0 is now an *encryption* of the actual commit content. Namely, there can be two distinct c_0 and c'_0 that decrypt to the same commit content. While we would like to

assign different members processing c_0 and c'_0 to the same node in the history graph, it is not clear which c_0 and c'_0 to use to identify the node. Such an issue disappears by using the semantics since the group operation defined inside c_0 and c'_0 are identical.

Independently, Alwen et al. [11] also uses the semantics to define nodes in their security model. This was crucial to capture a more advanced form of (non-metadata-hiding) CGKA coined as *server-aided* CGKA.

Restricted adversary due to commitment problem. In the UC framework, it is typical to restrict the adversary \mathcal{A} from performing corruptions that would cause the so-called *commitment problem*. Informally, this is a type of attack where \mathcal{A} can adaptively choose to corrupt some states *after* being provided with some challenges with respect to the state. While this attack is prohibited by default in any natural game-based definition (e.g., the adversary cannot obtain a secret key after being provided the challenge ciphertext), we need to make this restriction explicit in the UC-based definition.

Compared to prior works in the UC framework [10, 12, 36], the description of the restriction we require is more strict. Previously, when a new node was created due to a commit, the predicate **safe** was undefined for that node. Roughly, this is because the group secret key inside that node was never explicitly used during the real protocol and \mathcal{A} was given the freedom to adaptively decide whether to corrupt that node at an arbitrary moment of the security game. However, in the static metadata-hiding setting, this freedom of the adversary needs to be restricted. This is because when a new node is created, the group secret key inside this node is explicitly used to encrypt the static metadata of the commit content. When **safe** is true at this node, the ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ (roughly) wants to assign a random ciphertext to model the fact that the commit does not leak any information. Since \mathcal{A} is explicitly given the ciphertext, it cannot later decide to corrupt the node (i.e., **safe** remains true), as otherwise, it could trivially distinguish a valid ciphertext from a random one.

3.2 Proof of Static Metadata-Hiding CGKA

Hashimoto et al. [36] proposed Chained CmpKE and showed that it UC-realizes $\mathcal{F}_{\text{CGKA}}$. We show that the ciphertext variant of Chained CmpKE, which we call Chained CmpKE^{ctxt}, UC-realizes $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. The construction itself is a variation mirroring what MLSCiphertext does to MLSPplaintext. Thus, the technical novelty of this section is the first formal security proof that models static metadata-hiding. Considering that the overall structure of Chained CmpKE is similar to TreeKEM, except that the former further supports selective downloading, we expect our proof techniques to naturally translate to the ciphertext variant of TreeKEM. We leave it as an important future work to formally validate this. Below, we provide the core insights of our security proof. The full detail is provided in App. C.

Recycling previous proofs. As explained earlier, the ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ almost degenerates to $\mathcal{F}_{\text{CGKA}}$ when the leakage functions are defined to leak all static metadata. Our proof takes advantage of this fact. Recall that to prove security in the UC framework, we informally need to construct a simulator \mathcal{S} that simulates the real world adversary \mathcal{A} . We thus try to construct a simulator $\mathcal{S}^{\text{ctxt}}$ for $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ that internally executes \mathcal{S} for $\mathcal{F}_{\text{CGKA}}$. Using the description of \mathcal{S} provided in Hashimoto et al. [36], the hope is to

recycle all the proofs provided by them (which spanned 30 pages!). At a high level, during the hybrids where the leakage functions do not hide any static metadata, $\mathcal{S}^{\text{ctxt}}$ can use knowledge of the group encryption key to encrypt whatever non-encrypted proposals and commits \mathcal{S} outputs, and decrypt whatever \mathcal{S} inputs to perform a process. This allows us to recycle all the proofs that ignore the security of the static metadata. We then gradually modify the definition of the leakage functions to arrive at our ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. In the final hybrid, $\mathcal{S}^{\text{ctxt}}$ no longer knows the group encryption key for those epochs where **safe** is true, and $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ (roughly) takes care of generating random encryption of a proposal and commit.

While the high-level idea of recycling the proof of Hashimoto et al. [36] sounds straightforward, it turns out to be easier said than done. The non-triviality comes from the fact that the history graph created by $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ while using the null-leakage function is *not* identical to those created by $\mathcal{F}_{\text{CGKA}}$ – it is only *almost* identical. As explained earlier, our new history graph uses the semantics to identify the nodes and captures settings unique to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. Since all the security arguments boil down to how the predicate **safe** is defined over the history graph constructed during the security proof, it is not clear how to relate the proof of Hashimoto et al. to ours. To this end, we define the notion of *isomorphisms* of history graphs so that a security proof translates within the same class of history graphs. At a high level, two history graphs are isomorphism if the symbolic representation of the group evolution is identical. We prove that the two history graphs created in $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ and $\mathcal{F}_{\text{CGKA}}$ are isomorphic.

Polynomial security loss. We like to point out that our security proof only admits a *polynomial* security loss. This is in contrast to Hashimoto et al. [36] that admitted an *exponential* security loss.¹¹

The main reason for this disparity is because Chained CmpKE^{ctxt} captures a larger part of a CGKA compared to Chained CmpKE. As explained earlier, Chained CmpKE does not use the group secret key to encrypt proposals and commits, while Chained CmpKE^{ctxt} does. Effectively, in the former, an adversary was able to *adaptively* corrupt the group secret key while not trivially winning the security game. However, once the group secret key is used in a higher-level protocol as in Chained CmpKE^{ctxt}, this is no longer the case. An adversary capable of corrupting the group secret key *after* seeing the ciphertext can trivially win the security game.¹² Since the exponential security loss in Hashimoto et al. [36] was due to the adaptivity of the adversary, we can remove this loss by naturally restricting such adversaries in our UC-security model.

REMARK 1 (ADVERSARY-CONTROLLED RANDOMNESS). *In case of corruption, [12, 36] allow the adversary \mathcal{A} in $\mathcal{F}_{\text{CGKA}}$ to set the randomness to be used by the corrupted member id by altering the value $\text{RandCorr}[\text{id}]$. While our ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ models such strong adversaries, we restrict the adversary in our security proof to never alter the randomness for simplicity and better readability.*

¹¹We note that they were able to achieve a polynomial security loss relying on a stronger form of multi-recipient PKE secure against adaptive corruptions.

¹²We note that this problem may theoretically be solved using *non-committing* encryptions [22], but we did not consider them as a viable option as it will add a noticeable overhead in efficiency.

Recent definitions [9, 11] also intentionally disregard this type of attack.¹³ We leave it as future work to incorporate adversary-controlled randomness into the proof.

4 CONSTRUCTING METADATA-HIDING CGKA

We now construct a simple and modular CGKA construction that enables us to secure the 3rd layer (i.e., *dynamic* metadata). This results in the first *metadata-hiding* CGKA. Technically, we construct a wrapper protocol W^{mh} specifically taking care of the 3rd layer security in the $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ -hybrid model. Specifically, W^{mh} can be wrapped around *any* CGKA Π_{ctxt} that UC-realizes $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ and bootstraps Π_{ctxt} into a metadata-hiding CGKA.

The formal security proof of our metadata-hiding CGKA is deferred to Sec. 5, where we propose a new UC security model capturing the 3rd layer. We refer the readers to App. E for the full detail on the construction of W^{mh} . Below, we provide an overview of our protocol.

4.1 Goal of the Wrapper Protocol W^{mh}

To claim that CGKA secures the 3rd layer, our goal is to informally construct a wrapper protocol W^{mh} with the following properties:

- (1) **Anonymous upload.** A group member id can anonymously upload a proposal p or a commit (c_0, \vec{c}) to the server.
- (2) **Anonymous download.** A group member id can anonymously download a commit (c_0, \vec{c}) from the server.
- (3) **Unlinkability.** A member performing multiple uploads and downloads remains unlinkable from the server.
- (4) **Group authentication.** Only members of the group — excluding the *honest* server — can upload to and download from the server.

Here, unlinkability (Item (3)) is a strictly stronger notion compared to anonymity (Item (1)). Even if the member remains anonymous, some protocol may allow the server to link whether two uploads came from the same member. We also note that Item (4) is only relevant when the server is *honest* — a malicious server can always allow a non-member to perform an upload or a download on behalf of the group.

A wrapper protocol W^{mh} satisfying the above conditions is sufficient to bootstrap any CGKA that UC-realizes $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ into a metadata-hiding CGKA. However, if the underlying CGKA supports *selective downloading*, such as Chained CmPKE, Item (2) fails to provide the same efficiency offered by the underlying selective downloading CGKA. We thus strengthen Item (2) as follows.

- (2⁺) **Anonymous selective download.** A group member id can anonymously and *selectively* download a partial commit $(c_0, \widehat{c}_{\text{id}})$ from the server.

Specifically, the wrapper protocol W^{mh} allows the download cost to remain independent of the group member size in case a selective downloading CGKA is internally used.

Finally, even if the group secret key becomes compromised, we want all the above properties to hold again once the group state is healed. Namely, we seek a protocol with the following property.

- (5) **Compromise Resilience.** W^{mh} inherits the FS+PCS guarantees offered by the underlying CGKA UC-realizing $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$.

Insufficiency of client-anonymized authenticated channel.

One natural idea is to use a client-anonymized authenticated channel such as a VPN or an anonymized proxy like Tor [1, 31]. While such a channel solves Item (1) (and Item (2)), this alone cannot solve Items (2⁺) to (5). For instance, when selective downloading is performed, member id needs to specify its index in the group to retrieve the partial commit $(c_0, \widehat{c}_{\text{id}})$. Even if index may not directly leak id , the second time id performs a download on the same index, it will break linkability (Item (3)). Moreover, since the client-side is unauthenticated, it does not prevent external adversaries to perform a denial of service (DoS) attack on the group by uploading garbage contents, thus contradicting Item (4). Recall here that the server can no longer explicitly check if the uploaded contents come from genuine group members since the static metadata including the identity of the uploading member is hidden.

In summary, a client-anonymized authenticated channel alone is not enough to hide dynamic metadata.

4.2 High Level Description of W^{mh}

We provide an overview of our wrapper protocol W^{mh} . The main idea is to use the unique group secret key k exchanged among the group to perform an efficient proof of membership to the server. To make the presentation simple, we deliberately provide an informal description of our protocol.

Below, we assume all parties have access to the ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. Moreover, we assume the party communicates with the server via a client-anonymized authenticated channel, except when a new member is retrieving a welcome message from the server.

Fig. 3a: Group registration. Assume party id_0 wishes to create a group of three members $(\text{id}_0, \text{id}_1, \text{id}_2)$. id_0 first registers a new *empty* group to the server. Informally, id_0 invokes $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ on input (Create) and initializes a new group identifier gid and a group secret key k_0 for epoch = 0. It then *deterministically* creates a group specific signature key $(\text{gvk}_0, \text{gsk}_0) \leftarrow \text{KeyGen}(1^\kappa; \text{PRF}(k_0, \text{'auth'}))$ from the group secret key k_0 . We call this verification key gvk_0 as the *group statement* for epoch = 0. Party id_0 then uploads the pair $(\text{gid}, \text{gvk}_0)$ to the server.¹⁴ Finally, the server creates a new database for the group gid .

Fig. 3b: Initial proposal to add members. With the database for gid set up on the server, id_0 next adds id_1 and id_2 to the group. Specifically, id_0 invokes $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ on input (Propose, 'add'- id_i) and generates an add proposal p_i for $i \in [2]$. To upload p_i on the server, id_0 proves that it is a member of the group gid by essentially executing an identification protocol with the server. The server sends a random challenge $ch_i \leftarrow \{0, 1\}^\kappa$ and id_0 creates a signature $\sigma_i \leftarrow \text{Sign}(\text{gsk}_0, ch_i)$. The server verifies that σ_i is a valid signature with respect to the group statement gvk_0 at epoch = 0. If so, it adds p_i to the database. Due to the unforgeability of the signature scheme, no party without the group signing key gsk_0 can impersonate a group member.

¹⁴Our protocol W^{mh} does not prevent a malicious id_0 from registering multiple groups. As explained in Sec. 1.2, one possible way to thwart such a DoS attack would be to use anonymous credentials [26].

¹³Alwen et al.[9] does not allow adversaries to manipulate randomness used for symmetric key encryption.

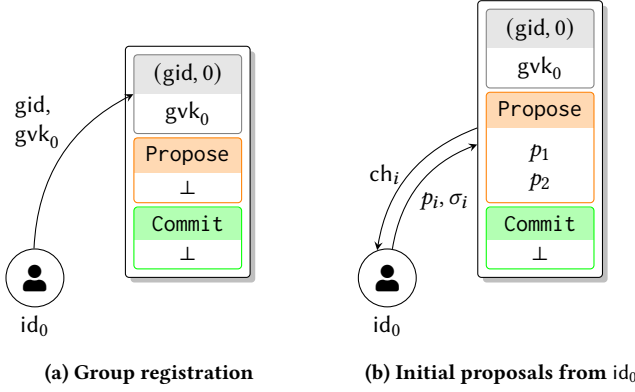


Figure 3: Creation of a group by a party id_0 . The three-part box represents the group state at a given epoch, as stored on the server. The top box stores in the clear the group identifier gid , epoch and group statement gvk_0 . The middle box stores the (encrypted) proposals created during the epoch. The bottom box stores the (encrypted) commit message which concluded this epoch, if it exists.

To prove membership, the server sends a challenge ch_i and id_0 responds with a signature σ_i (Fig. 3b). The contents are exchanged over a client-anonymized authenticated channel.

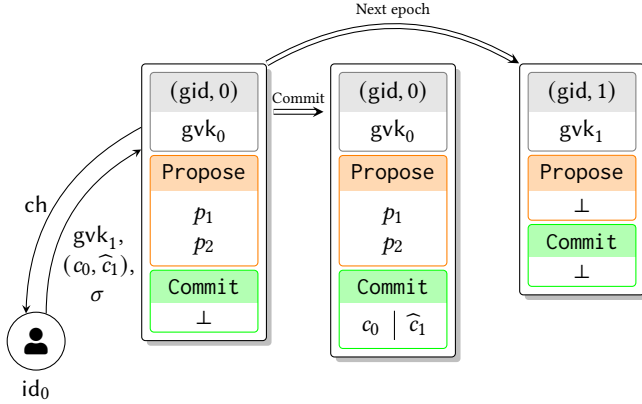


Figure 4: An initial commit made by the group creator id_0 . The server freezes the state of the current epoch = 0, and initializes a new epoch in the database.

Fig. 4: Initial commit to execute initial proposals. To create a new group with members id_1 and id_2 , id_0 must commit the proposals $\vec{p} = (p_i)_{i \in [2]}$. It first invokes $\mathcal{F}_{CGKA}^{ctxt}$ on input $(Commit, \vec{p})$ and generates $(c_0, \vec{c} = (\widehat{c}_i)_{i \in [2]})$ ¹⁵ and (roughly) updates the group secret key k_1 for the next epoch = 1. Similar to the group registration phase, id_0 creates a group statement gvk_1 for epoch = 1. To upload the commit (c_0, \vec{c}) , id_0 performs the same identification protocol as in Fig. 3b to prove that he is indeed a member of the group gid . If the identification protocol succeeds, the server stores

¹⁵Note that id_0 must process the commit (c_0, \widehat{c}_1) to move to the next epoch = 1.

the commit on the database and further creates a new column for the next epoch = 1.

Finally, id_0 uploads the welcome messages \vec{w} to the server. The welcome messages are stored on a party-dependent database and work identically to $\mathcal{F}_{CGKA}^{ctxt}$. In particular, id_1 and id_2 can retrieve \widehat{w}_1 and \widehat{w}_2 , respectively, from the server and execute $\mathcal{F}_{CGKA}^{ctxt}$ on input $(Join, \widehat{w}_i)$ to have the same group state as id_0 . Effectively, they become a member of the group gid at epoch = 1.

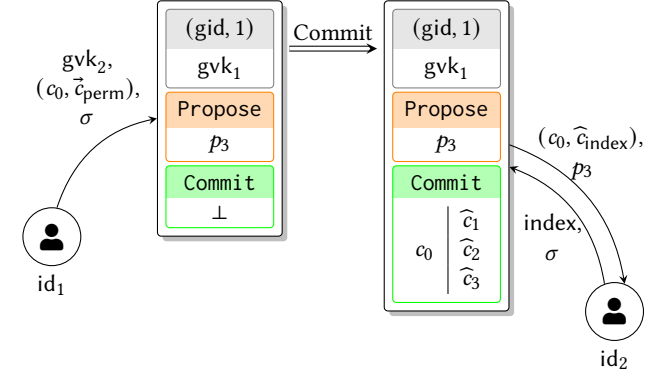


Figure 5: Left: a commit sent by id_1 . Right: a subsequent process made by id_2 . Member-dependent commits $\vec{c}_{perm} := (\widehat{c}_1, \widehat{c}_2, \widehat{c}_3)$ are randomly permuted and id_2 specifies an index to fetch the commit from the server. Challenges sent by the server are omitted for readability.

Commits without selective downloading. Now that id_1 and id_2 joined the group gid at epoch = 1 (i.e., share the same group secret key k_1), they can upload proposals and commits to the database defined with respect to the group statement gvk_1 .

We now explain the structure of a commit message when the group has more than one member. Assume some member made an update proposal p_3 , and id_1 wishes to commit this proposal (left half of Fig. 5). Then, following the same procedure as the initial commit, id_1 invokes $\mathcal{F}_{CGKA}^{ctxt}$ on input $(Commit, p_3)$ and generates $(c_0, \vec{c} = (\widehat{c}_i)_{i \in [3]})$, along with an updated group statement gvk_2 for the next epoch = 2. If selective downloading is not performed, then id_1 can simply upload (c_0, \vec{c}) to the server. The server then initiates a new column for epoch = 2 and the other members can anonymously download the entire commit from the server (by performing the aforementioned identification protocol).

Issues with naive selective downloading. Unfortunately, if selective downloading is naively applied, the above method leaks the access pattern of group members.

Recall that when selective downloading is performed in $\mathcal{F}_{CGKA}^{ctxt}$, a member sends an index and receives the corresponding member-dependent commit \widehat{c}_{index} from the server. When a member selectively downloads commit messages relative to two distinct epochs, they send the same index for both epochs. The server can infer that both requests were made by the same party, contradicting Item (3).

Even worse, suppose that the group secret key, and thus, the group member list is compromised. Although the group secret key may heal after PCS of $\mathcal{F}_{CGKA}^{ctxt}$ kicks in, the access pattern will never

heal; the server who learned the member-index correspondence can permanently break *anonymity* when selective downloading is performed by simply looking at the same index used at every epoch, contradicting Items (2⁺) and (5).

Commits with *oblivious* selective downloading. While the problem exposed above is known to be solvable using relatively complex tools such as private information retrieval (PIR) [27], we provide a much simpler solution using a pseudorandom permutation PRP by taking advantage of the fact that selective downloading is performed by each group member *once per epoch*. Continuing with our above example, when id_1 generates a commit $(c_0, \vec{c} = (\widehat{c}_i)_{i \in [3]})$, it further *deterministically* generates a PRP key $permKey \leftarrow PRF(k_1, 'perm')$ which defines a permutation over the group size, which is $[3] = \{1, 2, 3\}$ for our example. It then creates a permuted member dependent commit \vec{c}_{perm} so that \widehat{c}_i is placed at entry $PRP(permKey, i) \in [3]$. Finally, id_1 uploads $(gsk_2, (c_0, \vec{c}_{perm}))$ to the server by performing the identification protocol using gsk_1 .

When id_2 performs a selective download to retrieve the proposal and (partial) commit from the server, it computes its permuted index = $PRP(permKey, 2)$, where id_2 generates an identical $permKey$ as id_1 . It performs an identification protocol using gsk_1 , sends index, and retrieves $(c_0, \widehat{c}_{index})$ and the proposal p_3 from the server. This is illustrated in the right half of Fig. 5. id_1 can then invoke $\mathcal{F}_{CGKA}^{ctxt}$ on input $(Process, (c_0, \widehat{c}_{index}), p_3)$ and move to the next epoch = 2.

Observe that a member never performs a selective download more than once per epoch. This is the main reason why a PRP suffices — the access would have been linkable if selective downloading was performed more than twice per epoch using the same PRP key. Moreover, since the group secret key is updated at each epoch, the PRP key is also updated, thus satisfying FS and PCS (Item (5)).

REMARK 2 (NON-INTERACTIVE MEMBERSHIP IDENTIFICATION). *We provided a challenge-response type interactive identification protocol to prove group membership. By allowing the server to perform additional checks on the database and further reasonably weakening the security guarantee (i.e., Item (4) is guaranteed only for uploads), we are able to make the protocol completely non-interactive. At a high level, the party simply needs to sign the proposal or commit (rather than a challenge message) to upload, and perform no membership identification to download. The full detail is provided in App. E.2.*

5 FORMAL MODEL FOR METADATA-HIDING CGKA

We define a UC security model capturing the security of the entire 1st, 2nd & 3rd layers (i.e., group secret keys, static and *dynamic* metadata) by defining a new ideal functionality \mathcal{F}_{CGKA}^{mh} . Any CGKA UC-realizing \mathcal{F}_{CGKA}^{mh} is provably a *metadata-hiding* CGKA.

Reusing most of the description of $\mathcal{F}_{CGKA}^{ctxt}$ handling the 1st and 2nd layers, the description of \mathcal{F}_{CGKA}^{mh} can focus mainly on the 3rd layer. Our model succinctly captures all the properties explained in Sec. 4.1, Items (1) to (5). We then show that the wrapper protocol W^{mh} presented in the previous section UC-realizes \mathcal{F}_{CGKA}^{mh} in the $\mathcal{F}_{CGKA}^{ctxt}$ -hybrid model. The full details of this section is provided in App. D. Below, we provide an overview of our idea.




Figure 6: (Left) metadata-hiding CGKA Π_{CGKA}^{mh} when the server Sv is honest and (Right) when Sv is malicious. The red dotted box denotes the entire protocol Π_{CGKA}^{mh} , where it is further decomposed as a combination of the wrapper protocol W^{mh} and ideal functionality $\mathcal{F}_{CGKA}^{ctxt}$. The red shaded region denotes that Sv is corrupted and that (\mathcal{A}, Sv) are viewed as a single adversary \mathcal{A}' . In this case, W^{mh} is ignored and Π_{CGKA}^{mh} degenerates to a UC-realization of $\mathcal{F}_{CGKA}^{ctxt}$.

5.1 Modeling an Honest but Curious Server

In previous constructions of CGKA, the server was assumed to be *always* malicious. This is for instance captured in the ideal functionalities \mathcal{F}_{CGKA} and $\mathcal{F}_{CGKA}^{ctxt}$ by observing that the Commit and Process take as input arbitrary proposals and commits — not just those created by the honest group members.

While assuming the server to be always malicious allows to capture a strong level of security against group secret keys and static metadata, this is far too inflexible for our use case. Recall Sec. 4.1, Item (4). To properly model that any non-member cannot upload and download from the server on behalf of the group, we must model an *honest but curious* server — a server that honestly follows the protocol but tries to learn as much metadata as possible.

To this end, we explicitly incorporate a server into our model as depicted in Fig. 6. We allow the server to be in two states: honest¹⁶ or corrupt. When the server is honest, we are able to properly model Sec. 4.1, Item (4). Otherwise, since a malicious server can arbitrarily choose to accept or reject the identification protocol executed by the wrapper protocol W^{mh} , W^{mh} does not provide any meaningful functionality. In particular, our metadata-hiding CGKA Π_{CGKA}^{mh} degenerates to offer the same functionality as $\mathcal{F}_{CGKA}^{ctxt}$.

5.2 UC Security Model for Dynamic Metadata

As already mentioned, the new ideal functionality \mathcal{F}_{CGKA}^{mh} inherits all functionalities offered by $\mathcal{F}_{CGKA}^{ctxt}$. \mathcal{F}_{CGKA}^{mh} comes with seven additional functionalities:

- RegisterGroup,
- PublishProposal, FetchProposals,
- PublishCommit, FetchCommit,
- PublishWelcome, FetchWelcome.

As the name indicates, Publish-* (resp. Fetch-*) is invoked to upload (resp. download) a proposal, commit, or welcome message from

¹⁶We drop “but curious” for simplicity.

the server. These functions are mainly invoked during an execution of `Create`, `Propose`, `Commit`, `Process`, and `Join`. For instance, when `Commit` is invoked, the member first runs `FetchProposals` to retrieve the proposals \vec{p} from the server and invokes $\mathcal{F}_{CGKA}^{ctxt}$ on input (Commit, \vec{p}) . It then uploads the commit and welcome message output by $\mathcal{F}_{CGKA}^{ctxt}$ using `PublishCommit`.

These functions are defined differently depending on (i) whether the calling party is an honest group member or an adversary,¹⁷ and (ii) whether the server is honest or malicious. As explained above, if the server is malicious, then we let the adversary \mathcal{A} (i.e., server) decide whether `Publish*` or `Fetch*` succeeds. In this case, \mathcal{F}_{CGKA}^{mh} becomes functionally identical to $\mathcal{F}_{CGKA}^{ctxt}$, modulo some syntactical difference due to the inclusion of the server into the model.

Otherwise, if the server is honest, then \mathcal{F}_{CGKA}^{mh} captures the correctness and security guarantees. For correctness, if the calling party is a group member and the group statement for the next epoch was honestly generated, then the functionality demands the server to accept the upload or download. On the other hand, for security, if the calling party is the adversary, then we require the server to reject the upload or download as long as the predicate **safe** at that epoch is true. That is, if the group secret key is not compromised, then the adversary should not be able to upload and download on behalf of the group.

Additionally, the database stored by the server (see Figs. 3 to 5) is modeled in \mathcal{F}_{CGKA}^{mh} by two lists `PropDB` and `ComDB`, each maintaining the proposals and commit for group `gid` at epoch. \mathcal{F}_{CGKA}^{mh} also models the permutation-based selective downloading explained in Sec. 4.2 by an ideal (helper) function `*permute-commit`. Finally, due to the already complex nature of the metadata-hiding CGKA, we did not model adversarially controlled randomness in \mathcal{F}_{CGKA}^{mh} . We leave this as an important future work.

5.3 Proof of Dynamic Metadata-Hiding CGKA

We prove that the wrapper protocol W^{mh} UC-realizes the metadata-hiding ideal functionality \mathcal{F}_{CGKA}^{mh} in the $\mathcal{F}_{CGKA}^{ctxt}$ -hybrid model. The full proof is provided in App. E.3.

The proof is relatively simple and modular since we defined W^{mh} in the $\mathcal{F}_{CGKA}^{ctxt}$ -hybrid model. We can in essence ignore all the description of \mathcal{F}_{CGKA}^{mh} that relates to the 1st and 2nd layers' correctness and security since the same checks can be handled by the simulator \mathcal{S} internally simulating $\mathcal{F}_{CGKA}^{ctxt}$. Specifically, our proof only needs to focus on the 3rd layer of correctness and security. The proof is standard and consists of invoking the security of the pseudorandom permutation and signature scheme.

6 INSTANTIATION AND EFFICIENCY

We now discuss how to instantiate W^{mh} and `Chained CmPKEctxt`. We target the so-called "NIST Level I"¹⁸, which informally states that breaking the protocol is no easier than key-recovery on a block cipher with a 128-bit key (e.g. AES128). This provides a meaningful baseline to discuss post-quantum security and ignoring quantum

¹⁷In our security definition (in App. D), we use `Publish*-Adv` and `Fetch*-Adv` to indicate that the calling party is the adversary.

¹⁸[https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/evaluation-criteria/security-\(evaluation-criteria\)](https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/evaluation-criteria/security-(evaluation-criteria))

Table 2: Comparison of selected mPKEs. An mPKE ciphertext has the form $(ct_0, (\widehat{ct}_{id})_{id \in [N]})$ when uploaded, and $(ct_0, \widehat{ct}_{id})$ when downloaded by id. All sizes are in bytes.

mPKE	Reference	ek	ct ₀	\widehat{ct}_{id}
ElGamal-based	[43] + [40]	32	32	32
SIKEp434-based	[36]	330	330	16
Illum512	[36]	768	704	48
Bilbo640	[36]	10240	10240	24

attacks corresponds to a classical security level of 128 bits. Given a cryptographic object x , we note $|x|$ its size in bytes.

6.1 Instantiation

The main cryptographic primitives used in our protocols are: (a) two signature schemes `SIG` and `SIG'`, (b) a multi-recipient PKE mPKE, and finally (c) symmetric primitives: a pseudo-random function PRF, a symmetric encryption scheme SKE and a pseudo-random permutation PRP.

Multi-recipient PKE. For pre-quantum security, one may combine Kurosawa's multi-recipient variant [43] of ElGamal with the transform of [40], whose decomposability property makes it amenable to the selective downloading performed by `Chained CmPKEctxt`. For post-quantum security, one may readily use one of the mPKEs proposed in [36]: `Illum512`, `Bilbo640` or a `SIKEp434`-based mPKE. For completion, we recall their performance profiles in Tab. 2.

Signature schemes. We choose signature schemes that complement our chosen mPKEs nicely, either by having similar performances, being based on similar assumptions, or both:

- The ECDSA standard is based on similar assumptions as of the ElGamal-based mPKE, and it has a similar performance profile as well.
- The NIST PQC finalist `Falcon` [53] complements our `SIKEp434`-based mPKE, as they both have small communication costs.
- The NIST PQC finalist `Dilithium` [44] is based on Module-LWE (plus Module-SIS), just like `Illum512`.
- `SPHINCS+` [38] and `Bilbo640` are both based on assumptions perceived as conservative (hash-based assumptions and unstructured LWE), and they both have comparatively larger communication costs than other schemes.

For simplicity, we consider that `SIG` and `SIG'` use the same scheme, but using distinct schemes may lead to interesting trade-offs.

Pseudo-random permutation. We require a PRP in order to permute the set of $[N]$ group members. There are at least two viable approaches:

- *Shuffling.* One may use a shuffling algorithm whose randomness is provided by passing `permKey` into a PRF. If cache attacks on group members' devices are not a concern, the Fisher-Yates shuffle is a good choice since it is simple, performs $N - 1$ swaps, and its entropy consumption is optimal. If cache attacks are a concern, one may need to use so-called *oblivious* algorithms. For example, with the Thorp

shuffle [50], each member may compute their permuted index in time $O(\log N)$.

- *Sorting*. This approach assigns to each member id a pseudo-random value $r_{id} = H(\text{permKey}, \text{id})$, then sorts the id's according to the r_{id} 's. The sorting step can be done obliviously in time $O(N \log^2 N)$ using sorting networks, for example, Batchner networks. Assuming H is collision-resistant, this provides a PRP over $[N]$.

While the solutions proposed above are not optimal when used in *permuted-commit-index (each group member needs to shuffle/sort *all* indices before finding their position), we expect them to be significantly less costly than public-key cryptographic operations, even for concretely large groups.

6.2 Efficiency

We now study the impact of W^{mh} on bandwidth efficiency, by applying it to Chained CmPKE^{ctxt}. The results are summarized in Tab. 3, with the overhead of W^{mh} represented in bold red font (+X).

Let us make two observations. First, W^{mh} adds the same overhead to any CGKA protocol realizing the ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ (e.g., our Chained mPKE^{ctxt} and MLSCiphertext¹⁹). Second, as is obvious in Tab. 3, the bandwidth overhead is added only in the *upload* direction, since the additional signatures serve to convince the server that a user is a legitimate group member.

For a more concrete perspective on the numbers provided in Tab. 3, we study the impact of applying W^{mh} over Chained CmPKE^{ctxt}, when these protocols are instantiated with the four mPKE-signature pairs selected in Sec. 6.1. The results are given in Tab. 4. For our examples, W^{mh} increases the bandwidth cost of Propose-‘upd’, Propose-‘add’, Propose-‘rem’, Commit, Process and application messages by at most 44%, 63%, 100%, 57%, 39% and 100%, respectively.²⁰ We believe this to be a very reasonable overhead if protecting metadata is considered important.

7 LIMITATION OF EFFICIENT METADATA-HIDING CGKA

We conclude this work by discussing some inherent limitations of metadata-hiding CGKAs. Our model and solution hide metadata from messages but allow leaking message size. This may reveal information about the structure and activity of the group. While remaining informal, our discussion highlights that the nature and extent of this information depend on the inner workings of the CGKA, as well as the precise topology of the group at a given time (especially for tree-based CGKAs).

7.1 Chained CmPKE and TreeKEM

We first discuss Chained CmPKE [36], as well as variants of TreeKEM *without server assistance* [8, 9, 12, 14, 41], meaning that the server forwards messages to group members without editing them.

Leakage from proposal messages. When the size of proposal messages depends on their type, these messages leak the proposal

¹⁹Although no formal proof exists, it is believed that MLSCiphertext realizes $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$.

²⁰For simplicity, numbers for Commit and Process in Tab. 4 consider an idealized setting where no proposal was made during the last epoch. In practice, this is not the case and the bandwidth overhead of W^{mh} for Commit and Process will be even lower in percentage.

type (i.e., ‘upd’, ‘add’, ‘rem’). For Chained CmPKE, this is evident from Tab. 4. Even though this can be fixed by padding, we note that at least for Chained CmPKE, commit messages subsequent to a successful ‘add’ (resp. ‘rem’) will increase (resp. decrease) by $|\widehat{\text{ct}}_{id}|$ bytes. For TreeKEM, a less obvious yet similar correlation exists. On the bright side, the size of proposals is independent of the group size and the sender’s identity.

Leakage from commit messages.

Chained CmPKE [36]. The size of an uploaded commit message is affine in the group size N . Thus, the server can infer N from commit messages. Chained CmPKE allows selective downloading but the wrapper W^{mh} hides the index of each party-dependent message by randomizing indices via per-epoch random permutation. Thus, from the server’s point of view, indices look random.

TreeKEM [8, 9, 12, 14]. The number of PKE ciphertexts in a commit message depends on the topology of the ratchet tree and the sender’s position in this tree, therefore the size of a commit message can leak information about both elements. We provide two examples.

First, suppose that the tree is full (i.e., no blank node) but the group size is not a power of two. Since TreeKEM uses left-balanced binary trees, parties assigned to leaves “at the left” of the tree will send longer commit messages than the ones “on the right”. Hence the server can easily partition the set of parties into two groups depending on the length of commit messages.

Second, we consider the tree in Fig. 7, which has some blank nodes. In this tree, the number of cryptographic materials sent by each party is as follows: parties 1 and 2 (resp. 3 and 4, resp. 5 and 6, resp. 7) send 3 encryption keys + 5 ciphertexts (resp. 3 + 4, resp. 3 + 3, resp. 2 + 2 encryption keys and ciphertexts).

Tainted TreeKEM [41]. In this variant of TreeKEM, commit messages contain cryptographic materials related to tainted nodes, nodes managed by a party other than on the direct path. Message size becomes larger if the sender manages more tainted nodes.

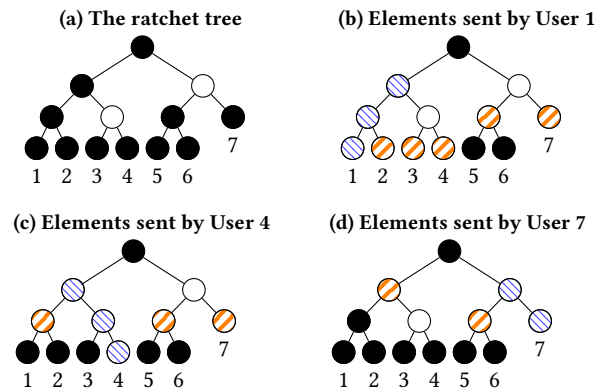


Figure 7: Example of the TreeKEM ratchet tree with blank nodes. The numbers indicate party identities. Figs. 7b to 7d: sending encryption keys (blue circle) and ciphertexts (orange circle) when party 1, 4 or 7 commits, respectively.

Table 3: Bandwidth overhead of our wrapper W^{mh} applied to Chained CmPKE^{ctxt}, for a group of N members in terms of public key cryptography elements. The nominal bandwidth cost of Chained CmPKE^{ctxt} is written in normal font (X). The overhead of W^{mh} is written in bold red (+X). U (resp. A, R) stands for the number of ‘upd’ (resp. ‘add’, ‘rem’) proposals published during the last epoch.

Procedure	Upload					Download				
	ek	ct ₀	ct _{id}	sig	svk	ek	ct ₀	ct _{id}	sig	svk
Propose-‘upd’	1			2 (+1)						
Propose-‘add’	1			1 (+1)	1					
Propose-‘rem’				1 (+1)						
Commit	1	1	N	2 (+2)	(+1)	U+A			2U+A+R	A
Process				(+1)		U+A+1	1	1	2U+A+R+2	A
Applications messages				1 (+1)						

Table 4: Concrete overhead of W^{mh} in terms of public key cryptography elements, when instantiated with: (a) ElGamal-mPKE + ECDSA (E+E), (b) SIKEp434-mPKE + Falcon (S+F), (c) Ilum512 + Dilithium (I+D), (d) Bilbo640 + SPHINCS⁺ (B+S). We assume a group of $N = 256$ members, and numbers for Commit and Process assume no proposal was made during the last epoch. The nominal bandwidth cost of Chained CmPKE^{ctxt} is written in normal font (X), and the overhead of W^{mh} is written in bold red (+X). All sizes are in bytes.

Procedure	E+E		S+F		I+D		B+S	
Propose-‘upd’	160	+64	1 662	+666	5 608	+2 420	44 416	+17 088
Propose-‘add’	128	+64	1 893	+666	4 500	+2 420	27 360	+17 088
Propose-‘rem’	64	+64	666	+666	2 420	+2 420	17 088	+17 088
Commit	8 384	+160	6 088	+2 229	18 600	+6 152	60 800	+34 208
Process	224	+64	2 008	+666	6 360	+2 420	54 680	+17 088
Applications messages	64	+64	666	+666	2 420	+2 420	17 088	+17 088

Leakage from welcome messages. In both protocols, welcome messages leak the receiver’s identity, the key package’s hash, and the group size. Note that the receiver’s identity must be in the clear for the server to deliver welcome messages. Key package hashes are used by the receivers to identify which decryption key is necessary to process the welcome message.

First of all, the server always gets to know when a welcome message was fetched. In addition, if the welcome message did not hide the dynamic metadata (i.e., uploader’s identity), then the server can infer that the party creating the welcome message and the party fetching it belong to the same group. Since groups are created by adding new members through a welcome message, this minor leakage of metadata can be used to trace the entire group member.

However, even if the welcome message hides the dynamic metadata, the server may still be able to link welcome messages to a specific group in some scenarios. For example, assume a party at an insecure epoch_{cur} commits an add proposal and moves to a secure epoch_{next}. Using a client-anonymized authenticated channel to upload the welcome message, the party uploading the welcome message remains anonymous, thus protecting against the above attack. However, since the server gets to see the key package included in the add proposal issued at the insecure epoch_{cur}, the server can link this key package to the key package hash included in the welcome message to infer that the new member joins the group related to epoch_{cur}. In particular, while the key package hash

included in the welcome message is good for efficiency, it may have non-trivial side effects.

A simple way to prevent such information leakage is to remove key package hashes from welcome messages. However, this would require recipients to try decrypting with all their registered decryption keys since they no longer can determine which decryption key can be used to decrypt the received welcome message. Finally, we note that welcome messages will always leak the size of the group to which the party was newly added. This is because a welcome message in essence sends all the current group states to the newly added party.

7.2 Server-Aided Variants of TreeKEM

We discuss recent efficient variations of TreeKEM [5, 11]. These allow the server to perform special tasks, e.g., editing signatures or ratchet trees in addition to delivering messages. They improve efficiency by revealing metadata to the server. Devising metadata-hiding variants of these schemes would likely require devising ways to perform these editing operations in an oblivious way, similarly to the PRP-based solution, we proposed for Chained CmPKE^{ctxt}.

SAIK²¹ [11]. This TreeKEM variant uses reducible signatures, a variant of redactable signatures, to allow parties to selectively download parts of commit messages while still guaranteeing their validity with the same signature that was initially uploaded by the sender. This improves the overall communication cost. However, the server

needs to know the identities of both the sender and receiver in order to reduce the commit message before forwarding it to the receiving party. In SAIK, this reduction is more involved than the selective downloading of Chained CmPKE. Even assuming that the reduction itself can be performed obliviously, the size of the reduced message depends on the positions of both the sender and receiver, so it would still leak information about both parties.

CoCoA²² [5]. This variant allows the server to merge multiple commit messages for the purpose of reducing communication costs. The server keeps the public part of the group state and creates the next group state by merging concurrently issued commit messages. Then, it forwards to each group member the part of the new state they need. This results in $O(\log N)$ upload and download costs, and a $O(\log^2 N)$ total cost. However, the server needs to know the public part of group states, which in particular includes the group member list.

One can say that these schemes improve efficiency by giving more information to the server. In other words, there is a trade-off between efficiency and privacy. We view it as an interesting research direction to construct more efficient (and practical) CGKA protocols while still protecting metadata.

REFERENCES

- [1] [n.d.]. Orbot: Proxy with Tor. <https://guardianproject.info/apps/org.torproject.android/> <https://guardianproject.info/apps/org.torproject.android/>.
- [2] 2018. Technology preview: Sealed sender for Signal. <https://signal.org/blog/sealed-sender/> <https://signal.org/blog/sealed-sender/>.
- [3] 2019. Technology Preview: Signal Private Group System. <https://signal.org/blog/signal-private-group-system/> <https://signal.org/blog/signal-private-group-system/>.
- [4] Ange Albertini, Thai Duong, Shay Gueron, Stefan Kölbl, Atul Luykx, and Sophie Schmieg. 2022. How to Abuse and Fix Authenticated Encryption Without Key Commitment. To Appear at USENIX 2022.
- [5] Joël Alwen, Benedikt Auerbach, Miguel Cueto Noval, Karen Klein, Guillermo Pascual-Perez, Krzysztof Pietrzak, and Michael Walter. 2022. CoCoA: Concurrent Continuous Group Key Agreement. In *EUROCRYPT 2022, Part II (LNCS)*. Springer, Heidelberg, 815–844. https://doi.org/10.1007/978-3-031-07085-3_28
- [6] Joël Alwen, Margarita Capretto, Miguel Cueto, Chethan Kamath, Karen Klein, Iliia Markov, Guillermo Pascual-Perez, Krzysztof Pietrzak, Michael Walter, and Michelle Yeo. 2021. Keep the Dirt: Tainted TreeKEM, Adaptively and Actively Secure Continuous Group Key Agreement. In *2021 IEEE Symposium on Security and Privacy (S&P)*. IEEE Computer Society, Los Alamitos, CA, USA, 596–612. <https://doi.org/10.1109/SP40001.2021.00035>
- [7] Joël Alwen, Sandro Coretti, and Yevgeniy Dodis. 2019. The Double Ratchet: Security Notions, Proofs, and Modularization for the Signal Protocol. In *EUROCRYPT 2019, Part I (LNCS, Vol. 11476)*, Yuval Ishai and Vincent Rijmen (Eds.). Springer, Heidelberg, 129–158. https://doi.org/10.1007/978-3-030-17653-2_5
- [8] Joël Alwen, Sandro Coretti, Yevgeniy Dodis, and Yiannis Tselekounis. 2020. Security Analysis and Improvements for the IETF MLS Standard for Group Messaging. In *CRYPTO 2020, Part I (LNCS, Vol. 12170)*, Daniele Micciancio and Thomas Ristenpart (Eds.). Springer, Heidelberg, 248–277. https://doi.org/10.1007/978-3-030-56784-2_9
- [9] Joël Alwen, Sandro Coretti, Yevgeniy Dodis, and Yiannis Tselekounis. 2021. Modular Design of Secure Group Messaging Protocols and the Security of MLS. In *ACM CCS 2021*, Giovanni Vigna and Elaine Shi (Eds.). ACM Press, 1463–1483. <https://doi.org/10.1145/3460120.3484820>
- [10] Joël Alwen, Sandro Coretti, Daniel Jost, and Marta Mularczyk. 2020. Continuous Group Key Agreement with Active Security. In *TCC 2020, Part II (LNCS, Vol. 12551)*, Rafael Pass and Krzysztof Pietrzak (Eds.). Springer, Heidelberg, 261–290. https://doi.org/10.1007/978-3-030-64378-2_10
- [11] Joël Alwen, Dominik Hartmann, Eike Kiltz, and Marta Mularczyk. 2021. Server-Aided Continuous Group Key Agreement. Cryptology ePrint Archive, Report 2021/1456. <https://eprint.iacr.org/2021/1456>.
- [12] Joël Alwen, Daniel Jost, and Marta Mularczyk. 2022. On the Insider Security of MLS. In *CRYPTO 2022, Part II (LNCS)*. Springer, Heidelberg, 34–68. https://doi.org/10.1007/978-3-031-15979-4_2
- [13] Manuel Barbosa and Pooya Farshim. 2007. Randomness reuse: Extensions and improvements. In *IMA International Conference on Cryptography and Coding*. Springer, 257–276.
- [14] Richard Barnes, Benjamin Beurdouche, Jon Millican, Emad Omara, Katriel Cohn-Gordon, and Raphael Robert. 2022. *The Messaging Layer Security (MLS) Protocol*. Internet-Draft draft-ietf-mls-protocol-13. Internet Engineering Task Force. <https://datatracker.ietf.org/doc/html/draft-ietf-mls-protocol-13> Work in Progress.
- [15] Mihir Bellare and Chanathip Namprempre. 2000. Authenticated Encryption: Relations among notions and analysis of the generic composition paradigm. In *ASIACRYPT 2000 (LNCS, Vol. 1976)*, Tatsuaki Okamoto (Ed.). Springer, Heidelberg, 531–545. https://doi.org/10.1007/3-540-44448-3_41
- [16] Karthikeyan Bhargavan, Richard Barnes, and Eric Rescorla. 2018. *TreeKEM: Asynchronous Decentralized Key Management for Large Dynamic Groups A protocol proposal for Messaging Layer Security (MLS)*. Research Report. Inria Paris. <https://hal.inria.fr/hal-02425247>
- [17] Karthikeyan Bhargavan, Benjamin Beurdouche, and Prasad Naldurg. 2019. *Formal Models and Verified Protocols for Group Messaging: Attacks and Proofs for IETF MLS*. Research Report. Inria Paris. <https://hal.inria.fr/hal-02425229>
- [18] Thomas Brewster. 2022. Meet The Secretive Surveillance Wizards Helping The FBI And ICE Wiretap Facebook And Google Users. Forbes. <https://www.forbes.com/sites/thomasbrewster/2022/02/23/meet-the-secretive-surveillance-wizards-helping-the-fbi-and-ice-wiretap-facebook-and-google-users/>.
- [19] Chris Brzuska, Eric Cornelissen, and Konrad Kohbrok. 2021. Cryptographic Security of the MLS RFC, Draft 11. Cryptology ePrint Archive, Report 2021/137. <https://eprint.iacr.org/2021/137>.
- [20] Ran Canetti. 2001. Universally Composable Security: A New Paradigm for Cryptographic Protocols. In *42nd FOCS*. IEEE Computer Society Press, 136–145. <https://doi.org/10.1109/SFCS.2001.959888>
- [21] Ran Canetti, Yevgeniy Dodis, Rafael Pass, and Shabsi Walfish. 2007. Universally Composable Security with Global Setup. In *TCC 2007 (LNCS, Vol. 4392)*, Salil P. Vadhan (Ed.). Springer, Heidelberg, 61–85. https://doi.org/10.1007/978-3-540-70936-7_4
- [22] Ran Canetti, Uriel Feige, Oded Goldreich, and Moni Naor. 1996. Adaptively Secure Multi-Party Computation. In *28th ACM STOC*. ACM Press, 639–648. <https://doi.org/10.1145/237814.238015>
- [23] Bjorn Carey. 2015. Stanford computer scientists show telephone metadata can reveal surprisingly sensitive personal information. <https://news.stanford.edu/2016/05/16/stanford-computer-scientists-show-telephone-metadata-can-reveal-surprisingly-sensitive-personal-information/>.
- [24] Melissa Chase, Sarah Meiklejohn, and Greg Zaverucha. 2014. Algebraic MACs and Keyed-Verification Anonymous Credentials. In *ACM CCS 2014*, Gail-Joon Ahn, Moti Yung, and Ninghui Li (Eds.). ACM Press, 1205–1216. <https://doi.org/10.1145/2660267.2660328>
- [25] Melissa Chase, Trevor Perrin, and Greg Zaverucha. 2020. The Signal Private Group System and Anonymous Credentials Supporting Efficient Verifiable Encryption. In *ACM CCS 2020*, Jay Ligatti, Xinning Ou, Jonathan Katz, and Giovanni Vigna (Eds.). ACM Press, 1445–1459. <https://doi.org/10.1145/3372297.3417887>
- [26] David Chaum. 1982. Blind Signatures for Untraceable Payments. In *CRYPTO '82*, David Chaum, Ronald L. Rivest, and Alan T. Sherman (Eds.). Plenum Press, New York, USA, 199–203.
- [27] Benny Chor, Oded Goldreich, Eyal Kushilevitz, and Madhu Sudan. 1995. Private Information Retrieval. In *36th FOCS*. IEEE Computer Society Press, 41–50. <https://doi.org/10.1109/SFCS.1995.492461>
- [28] Katriel Cohn-Gordon, Cas Cremers, Benjamin Dowling, Luke Garratt, and Douglas Stebila. 2017. A Formal Security Analysis of the Signal Messaging Protocol. In *2017 IEEE European Symposium on Security and Privacy (EuroS&P)*. 451–466. <https://doi.org/10.1109/EuroSP.2017.27>
- [29] Katriel Cohn-Gordon, Cas Cremers, and Luke Garratt. 2016. On Post-compromise Security. In *2016 IEEE 29th Computer Security Foundations Symposium (CSF)*. 164–178. <https://doi.org/10.1109/CSF.2016.19>
- [30] Katriel Cohn-Gordon, Cas Cremers, Luke Garratt, Jon Millican, and Kevin Milner. 2018. On Ends-to-Ends Encryption: Asynchronous Group Messaging with Strong Security Guarantees. In *ACM CCS 2018*, David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang (Eds.). ACM Press, 1802–1819. <https://doi.org/10.1145/3243734.3243747>
- [31] Roger Dingledine, Nick Mathewson, and Paul F. Syverson. 2004. Tor: The Second-Generation Onion Router. In *USENIX Security 2004*, Matt Blaze (Ed.). USENIX Association, 303–320.
- [32] Yevgeniy Dodis, Paul Grubbs, Thomas Ristenpart, and Joanne Woodage. 2018. Fast Message Franking: From Invisible Salamanders to Encryption. In *CRYPTO 2018, Part I (LNCS, Vol. 10991)*, Hovav Shacham and Alexandra Boldyreva (Eds.). Springer, Heidelberg, 155–186. https://doi.org/10.1007/978-3-319-96884-1_6
- [33] Pooya Farshim, Claudio Orlandi, and Razvan Rosie. 2017. Security of symmetric primitives under incorrect usage of keys. *LACR Transactions on Symmetric*

²¹SAIK stands for Server-Aided Insider-Secure TreeKEM.

²²CoCoA stands for COncurrent Continuous group key Agreement.

- Cryptology* (2017), 449–473.
- [34] Ola Flisbäck. 2015. Stalking anyone on Telegram. <https://oflisback.github.io/telegram-stalking/>.
- [35] Paul Grubbs, Jiahui Lu, and Thomas Ristenpart. 2017. Message Franking via Committing Authenticated Encryption. In *CRYPTO 2017, Part III (LNCS, Vol. 10403)*, Jonathan Katz and Hovav Shacham (Eds.). Springer, Heidelberg, 66–97. https://doi.org/10.1007/978-3-319-63697-9_3
- [36] Keitaro Hashimoto, Shuichi Katsumata, Eamonn Postlethwaite, Thomas Prest, and Bas Westerbaan. 2021. A Concrete Treatment of Efficient Continuous Group Key Agreement via Multi-Recipient PKEs. In *ACM CCS 2021*, Giovanni Vigna and Elaine Shi (Eds.). ACM Press, 1441–1462. <https://doi.org/10.1145/3460120.3484817>
- [37] Keitaro Hashimoto, Shuichi Katsumata, Eamonn Postlethwaite, Thomas Prest, and Bas Westerbaan. 2021. A Concrete Treatment of Efficient Continuous Group Key Agreement via Multi-Recipient PKEs. *Cryptology ePrint Archive*, Report 2021/1407. <https://eprint.iacr.org/2021/1407>.
- [38] Andreas Hulsing, Daniel J. Bernstein, Christoph Dobraunig, Maria Eichseder, Scott Fluhrer, Stefan-Lukas Gazdag, Panos Kampanakis, Stefan Kolbl, Tanja Lange, Martin M Lauridsen, Florian Mendel, Ruben Niederhagen, Christian Rechberger, Joost Rijneveld, Peter Schwabe, Jean-Philippe Aumasson, Bas Westerbaan, and Ward Beullens. 2020. *SPHINCS+*. Technical Report. National Institute of Standards and Technology. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-3-submissions>.
- [39] Daniel Jost, Ueli Maurer, and Marta Mularczyk. 2019. A Unified and Composable Take on Ratcheting. In *TCC 2019, Part II (LNCS, Vol. 11892)*, Dennis Hofheinz and Alon Rosen (Eds.). Springer, Heidelberg, 180–210. https://doi.org/10.1007/978-3-030-36033-7_7
- [40] Shuichi Katsumata, Kris Kwiatkowski, Federico Pintore, and Thomas Prest. 2020. Scalable Ciphertext Compression Techniques for Post-quantum KEMs and Their Applications. In *ASIACRYPT 2020, Part I (LNCS, Vol. 12491)*, Shiho Moriai and Huaxiong Wang (Eds.). Springer, Heidelberg, 289–320. https://doi.org/10.1007/978-3-030-64837-4_10
- [41] Karen Klein, Guillermo Pascual-Perez, Michael Walter, Chethan Kamath, Margarita Capretto, Miguel Cueto, Iliia Markov, Michelle Yeo, Joël Alwen, and Krzysztof Pietrzak. 2021. Keep the Dirt: Tainted TreeKEM, Adaptively and Actively Secure Continuous Group Key Agreement. In *2021 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, 268–284. <https://doi.org/10.1109/SP40001.2021.00035>
- [42] Andy Kroll. 2021. FBI Document Says the Feds Can Get Your WhatsApp Data – in Real Time. *Rolling Stone*. <https://www.rollingstone.com/politics/politics-features/whatsapp-imessage-facebook-apple-fbi-privacy-1261816/>.
- [43] Kaoru Kurosawa. 2002. Multi-recipient Public-Key Encryption with Shortened Ciphertext. In *PKC 2002 (LNCS, Vol. 2274)*, David Naccache and Pascal Paillier (Eds.). Springer, Heidelberg, 48–63. https://doi.org/10.1007/3-540-45664-3_4
- [44] Vadim Lyubashevsky, Léo Ducas, Eike Kiltz, Tancrède Lepoint, Peter Schwabe, Gregor Seiler, Damien Stehlé, and Shi Bai. 2020. *CRYSTALS-DILITHIUM*. Technical Report. National Institute of Standards and Technology. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-3-submissions>.
- [45] Moxie Marlinspike and Trevor Perrin. 2016. The double ratchet algorithm. <https://signal.org/docs/specifications/doublerratchet/> <https://signal.org/docs/specifications/doublerratchet/>.
- [46] Ian Martiny, Gabriel Kapthuk, Adam Aviv, Dan Roche, and Eric Wustrow. 2021. Improving Signal’s sealed sender. (2021). To appear at NDSS 2021.
- [47] Susan E. McGregor, Polina Charters, Tobin Holliday, and Franziska Roesner. 2015. Investigating the Computer Security Practices and Needs of Journalists. In *USENIX Security 2015*, Jaeyeon Jung and Thorsten Holz (Eds.). USENIX Association, 399–414.
- [48] Susan E. McGregor, Franziska Roesner, and Kelly Caine. 2016. Individual versus Organizational Computer Security and Privacy Concerns in Journalism. *PopETS* 2016, 4 (Oct. 2016), 418–435. <https://doi.org/10.1515/popets-2016-0048>
- [49] Vaishnavi Krishna Mohan. 2021. WhatsApp’s New Privacy Policy: Collecting Metadata and Its Implications. <https://www.globalviews360.com/articles/whatsapp-s-new-privacy-policy-collecting-metadata-and-its-implications>.
- [50] Ben Morris, Phillip Rogaway, and Till Stegers. 2018. Deterministic Encryption with the Thorp Shuffle. *Journal of Cryptology* 31, 2 (April 2018), 521–536. <https://doi.org/10.1007/s00145-017-9262-z>
- [51] Kurt Opsahl. 2013. Why Metadata Matters. <https://www.eff.org/deeplinks/2013/06/why-metadata-matters>.
- [52] Trevor Perrin. [n.d.]. The Noise Protocol Framework. The Noise Protocol Framework. <http://www.noiseprotocol.org/noise.pdf>.
- [53] Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. 2020. *FALCON*. Technical Report. National Institute of Standards and Technology. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-3-submissions>.
- [54] Charlie Savage. 2013. Court Rejects Appeal Bid by Writer in Leak Case. *The New York Times*. <http://www.nytimes.com/2013/10/16/us/court-rejects-appealbid-by-writer-in-leak-case.html>.
- [55] Victor Shoup. 1997. Lower Bounds for Discrete Logarithms and Related Problems. In *EUROCRYPT’97 (LNCS, Vol. 1233)*, Walter Fumy (Ed.). Springer, Heidelberg, 256–266. https://doi.org/10.1007/3-540-69053-0_18
- [56] Nigel P. Smart. 2005. Efficient Key Encapsulation to Multiple Parties. In *SCN 04 (LNCS, Vol. 3352)*, Carlo Blundo and Stelvio Cimato (Eds.). Springer, Heidelberg, 208–219. https://doi.org/10.1007/978-3-540-30598-9_15
- [57] Matthew Weidner, Martin Kleppmann, Daniel Hugenroth, and Alastair R. Beresford. 2021. Key Agreement for Decentralized Secure Group Messaging with Strong Security Guarantees. In *ACM CCS 2021*, Giovanni Vigna and Elaine Shi (Eds.). ACM Press, 2024–2045. <https://doi.org/10.1145/3460120.3484542>

A DEFINITION OF GENERAL CRYPTOGRAPHIC PRIMITIVES

A.1 Notation

We prepare some notations and keywords used to define the security and construct CGKAs. They are taken almost verbatim from [36, Sec. A.1].

We denote the set of natural numbers (non-negative integers) by \mathbb{N} and the security parameter by $\kappa \in \mathbb{N}$. For an algorithm A , we write $A(\cdot; r)$ to denote that A is run with the explicit randomness r . For $k, n \in \mathbb{N}$ such that $k \leq n$, we write $[k : n]$ to denote the set $\{k, \dots, n\}$. We use the shorthand $[n]$ when $k = 1$. We use $v \leftarrow x$ and $v := x$ to denote assigning the value x to the variable v , and use $v \leftarrow S$ to denote sampling an element v uniformly and randomly from a set S . We denote by $[cond]$ the bit that is 1 if the boolean statement $cond$ is true, and 0 otherwise.

Data structure. If V is a set, we write $V \leftarrow x$ and $V \leftarrow x$ as shorthands for $V \leftarrow V \cup \{x\}$ and $V \leftarrow V \setminus \{x\}$, respectively. For another set W , we write $V \leftarrow W$ and $V \leftarrow W$ as shorthands for $V \leftarrow V \cup W$ and $V \leftarrow V \setminus W$, respectively. For lists (vectors) $x := (x_1, \dots, x_n)$ and $y := (y_1, \dots, y_m)$, we denote the concatenation by $x||y = (x_1, \dots, x_n, y_1, \dots, y_m)$ and use $x \leftarrow v$ as a shorthand for $x \leftarrow x||v$. We further use associative arrays and use $A[i] \leftarrow x$ and $y \leftarrow A[i]$ to assignment and retrieval of element i , respectively. We denote by $A[*] \leftarrow v$ the Initialization of the array to the default value v . For simplicity, we use the wildcard notation when dealing with sets of tuples and multi-argument associative arrays. For example, for an array with domain $\mathcal{I} \times \mathcal{J}$, we write $A[* , j] := \{A[i, j] \mid i \in \mathcal{I}\}$ and for a set $S \subseteq \mathcal{I} \times \mathcal{J}$, we write $(i, *) \in S$ as a shorthand for the condition $\exists j \in \mathcal{J} : (i, j) \in S$.

Keywords. We use the following keywords:

- **req** $cond$ denotes that if the condition $cond$ is false, then the current function unwinds all state changes and immediately returns \perp .
- **parse** $(m_1, \dots, m_n) \leftarrow m$ denotes an attempt to parse a message m as a tuple. If m is not of the correct format, the current function unwinds all state changes and immediately returns \perp .
- **try** $y \leftarrow func(x)$ is a shorthand notation for calling a helper function or subroutine $func$ and executing **req** $y \neq \perp$.
- **assert** $cond$ is only used to describe functionalities. It denotes that if $cond$ is false, then the functionality permanently halts, making the real and ideal worlds trivially distinguishable (this is used to validate inputs of the simulator).

A.2 Decomposable Multi-Recipient Public Key Encryption

A multi-recipient PKE (mPKE) [13, 43, 56] is a type of PKE that allows a party to send the *same* message to a set of recipients of size N , more efficiently than executing N parallel runs of a standard PKE. In this work, we use a *decomposable* mPKE introduced in [40] that mandates a mPKE ciphertext to be decomposable into a recipient-dependent and independent part. [40] showed that many assumptions known to imply PKE (e.g., DDH, LWE, SIDH) can naturally be used to construct an IND-CCA decomposable mPKE in the random oracle model.

Definition A.1 (Decomposable Multi-Recipient Public-Key Encryption). A (single-message) decomposable multi-recipient public-key encryption (mPKE) over a message space \mathcal{M} consists of the following algorithms:

- $\text{mSetup}(1^\kappa) \rightarrow \text{pp}$: On input the security parameter 1^κ , it outputs a public parameter pp .
- $\text{mGen}(\text{pp}) \rightarrow (\text{ek}, \text{dk})$: On input a public parameter pp , it outputs a pair of encryption key and decryption key (ek, dk) .
- $\text{mEnc}(\text{pp}, (\text{ek}_i)_{i \in [N]}, M; r_0, (r_i)_{i \in [N]}) \rightarrow (\text{ct}_0, \vec{\text{ct}} = (\widehat{\text{ct}}_i)_{i \in [N]})$: The (decomposable) encryption algorithm splits into a pair of algorithms $(\text{mEnc}^i, \text{mEnc}^d)$:
 - $\text{mEnc}^i(\text{pp}; r_0) \rightarrow \text{ct}_0$: On input a public parameter pp , it outputs an (encryption key independent) ciphertext ct_0 .
 - $\text{mEnc}^d(\text{pp}, \text{ek}_i, M; r_0, r_i) \rightarrow \widehat{\text{ct}}_i$: On input a public parameter pp , an encryption key ek_i , a message $M \in \mathcal{M}$, it outputs an (encryption key dependent) ciphertext $\widehat{\text{ct}}_i$.
- $\text{mDec}(\text{dk}, \text{ct}_0, \widehat{\text{ct}}_i) \rightarrow M \text{ or } \perp$: On input a decryption key dk and a ciphertext $(\text{ct}_0, \widehat{\text{ct}}_i)$, it outputs either $M \in \mathcal{M}$ or $\perp \notin \mathcal{M}$.

Definition A.2 (Correctness). A mPKE is correct if we have

$$\mathbb{E} \left[\max_{M \in \mathcal{M}} \Pr \left[M = \text{mDec}(\text{dk}, (\text{ct}_0, \widehat{\text{ct}})) : \begin{array}{l} \text{ct}_0 \leftarrow \text{mEnc}^i(\text{pp}), \\ \widehat{\text{ct}} \leftarrow \text{mEnc}^d(\text{pp}, \text{ek}, M) \end{array} \right] \right] \geq 1 - \text{negl}(\kappa),$$

where the expectation is taken over $\text{pp} \leftarrow \text{mSetup}(1^\kappa)$ and $(\text{ek}, \text{dk}) \leftarrow \text{mGen}(\text{pp})$.

Ciphertext-spreadness defines how much entropy a properly generated ciphertext has. For the application of CGKA, we require the party-independent ciphertext to maintain enough min-entropy. Note that [36] defined ciphertext-spreadness with the entire ciphertext but the proof relies on the high min-entropy of the party independent part.

Definition A.3 (Ciphertext-Spreadness of Party Independent Ciphertext). A mPKE has ciphertext-spreadness (with respect to the party independent ciphertext of), if for all $\text{pp} \in \text{mSetup}(1^\kappa)$, we have

$$\mathbb{E}[\max_{\text{ct}_0} \Pr[\text{ct}_0 = \text{mEnc}^i(\text{pp})] \leq \text{negl}(\kappa),$$

where the expectation is taken over $\text{pp} \leftarrow \text{mSetup}(1^\kappa)$ and the encryption randomness.

Definition A.4 (IND-CCA). The security notion is defined by a game illustrated in Fig. 8, where we say the adversary \mathcal{A} wins if the game outputs 1. A mPKE is IND-CCA secure if for all PPT adversaries \mathcal{A} , we have $|\Pr[\mathcal{A} \text{ wins}] - 1/2| \leq \text{negl}(\kappa)$.

GAME IND-CCA

```

1: pp ← mSetup(1κ)
2: foreach i ∈ [N] do
3:   (eki, dki) ← mGen(pp)
4:   (M0, M1) ←  $\mathcal{A}^{\mathcal{D}(\cdot)}$ (pp, (eki)i ∈ [N])
5:   b ←  $\$ \{0, 1\}$ 
6:   (ct0*,  $\vec{\text{ct}}^* := (\widehat{\text{ct}}_i^*)_{i \in [N]}$ ) ← mEnc(pp, (eki)i ∈ [N], Mb)
7:   b' ←  $\mathcal{A}^{\mathcal{D}(\cdot)}$ (pp, (eki)i ∈ [N], ct0*,  $\vec{\text{ct}}^*$ )
8:   return [b = b']

```

Decapsulation Oracle $\mathcal{D}(i, \text{ct}_0, \widehat{\text{ct}}_i)$

```

1: req (ct0,  $\widehat{\text{ct}}_i$ ) ≠ (ct0*,  $\widehat{\text{ct}}_i^*$ )
2: M ← mDec(dki, ct0,  $\widehat{\text{ct}}_i$ )
3: return M

```

Figure 8: IND-CCA security of mPKE.

A.3 Secret Key Encryption

We define secret key encryption SKE.

Definition A.5 (Secret-Key Encryption). Secret-key encryption (SKE) over a key space \mathcal{K} (implicitly parameterized by the security parameter) and message space \mathcal{M} consists of the following two algorithms:

- $\text{Enc}_s(k, M) \rightarrow \text{ct}$: On input a secret key $k \in \mathcal{K}$ and a message $M \in \mathcal{M}$, it outputs a ciphertext ct .
- $\text{Dec}_s(k, \text{ct}) \rightarrow M \text{ or } \perp$: On input a secret key k and a ciphertext ct , it deterministically outputs either $M \in \mathcal{M}$ or $\perp \notin \mathcal{M}$.

Definition A.6 (Correctness). A SKE is correct if $\Pr[\text{Dec}_s(k, \text{Enc}_s(k, M)) = M] \geq 1 - \text{negl}(\kappa)$ holds for all $\kappa \in \mathbb{N}$, $M \in \mathcal{M}$ and $k \in \mathcal{K}$.

We define IND-CPA and IND-CCA security for SKE by the left-or-right version of game-based indistinguishability [15].

Definition A.7 (IND-CPA and IND-CCA for SKE). The security notion is defined by a game illustrated in Fig. 9, where we say the adversary \mathcal{A} wins if the game outputs 1. A SKE is IND-CPA secure (resp. IND-CCA secure) if for all PPT adversaries \mathcal{A} , we have $|\Pr[\mathcal{A} \text{ wins in IND-CPA game (resp. IND-CCA secure)}] - 1/2| \leq \text{negl}(\kappa)$.

We finally define the *key-committing* property [33] which roughly states that it is difficult to find two secret keys that correctly decrypt the same ciphertext (to possibly different messages). As in prior works [4, 32, 33, 35, 36], we define this notion by providing the (non-uniform) adversary oracle access to Enc_s and Dec_s , where we implicitly assume these two algorithms are implemented using an internal hash function modeled as a random oracle.

Definition A.8 (Key-Committing). A SKE has *key-committing* property if for all PPT adversary \mathcal{A} , we have

$$\Pr \left[\begin{array}{l} \text{Dec}_s(k_0, \text{ct}) \neq \perp \\ \wedge \text{Dec}_s(k_1, \text{ct}) \neq \perp \end{array} : (k_0, k_1, \text{ct}) \leftarrow \mathcal{A}^{\text{Enc}_s, \text{Dec}_s}(1^\kappa) \right] \leq \text{negl}(\kappa).$$

<p>GAME IND-CPA</p> <hr/> 1: $k \leftarrow \mathcal{K}$ 2: $b \leftarrow \{0, 1\}$ 3: $b' \leftarrow \mathcal{A}^{\mathcal{LR}(\cdot, \cdot)}(1^\kappa)$ 4: return $[b = b']$	<p>Left-or-Right Oracle $\mathcal{LR}(M_0, M_1)$</p> <hr/> 1: $ct \leftarrow \text{Enc}_s(k, M_b)$ 2: $S \leftarrow ct$ 3: return ct
<p>GAME IND-CCA</p> <hr/> 1: $k \leftarrow \mathcal{K}$ 2: $b \leftarrow \{0, 1\}$ 3: $S \leftarrow \emptyset$ 4: $b' \leftarrow \mathcal{A}^{\mathcal{LR}(\cdot, \cdot), \mathcal{D}(\cdot)}(1^\kappa)$ 5: return $[b = b']$	<p>Decryption Oracle $\mathcal{D}(ct)$</p> <hr/> 1: req $ct \notin S$ 2: $M \leftarrow \text{Dec}_s(k, ct)$ 3: return M

Figure 9: IND-CPA and IND-CCA security of SKE. If the condition following req does not hold, the game terminates by returning a random bit.

A.4 Digital Signatures

We provide the standard notion of digital signatures.

Definition A.9 (Signature Scheme). A signature scheme SIG over a message space \mathcal{M} consists of the following algorithms:

- $\text{Setup}(1^\kappa) \rightarrow pp$: On input the security parameter 1^κ , it outputs a public parameter pp .
- $\text{KeyGen}(pp) \rightarrow (svk, ssk)$: On input a public parameter pp it outputs a pair of verification key and signing key (svk, ssk) .
- $\text{Sign}(pp, ssk, m) \rightarrow sig$: On input a public parameter pp , a signing key ssk and a message m , it outputs a signature sig .
- $\text{Verify}(pp, svk, m, sig) \rightarrow \top/\perp$: On input a public parameter pp , a verification key ssk , a message m and a signature sig , it outputs \top or \perp .

Definition A.10 (Correctness). A signature scheme SIG is correct if for all $\kappa \in \mathbb{N}$, all messages $m \in \mathcal{M}$ and all $pp \in \text{Setup}(1^\kappa)$,

$$\Pr \left[\text{Verify}(pp, svk, m, sig) = \top : \begin{array}{l} (svk, ssk) \leftarrow \text{KeyGen}(pp); \\ sig \leftarrow \text{Sign}(pp, ssk, m) \end{array} \right] \geq 1 - \text{negl}(\kappa).$$

Definition A.11 (sEUF-CMA). A signature scheme is strongly EUF-CMA (sEUF-CMA) secure if for all PPT adversary \mathcal{A} , we have

$$\Pr \left[\begin{array}{l} \text{Verify}(pp, svk, m^*, sig^*) = \top \\ \wedge (m^*, sig^*) \notin L \end{array} : \begin{array}{l} pp \leftarrow \text{Setup}(1^\kappa); \\ (svk, ssk) \leftarrow \text{KeyGen}(pp); \\ (m^*, sig^*) \leftarrow \mathcal{A}^{\mathcal{S}(\cdot)}(pp, svk) \end{array} \right] \leq \text{negl}(\kappa),$$

where $\mathcal{S}(\cdot)$ is defined as $\text{Sign}(ssk, \cdot)$, and L is the set of pairs of messages and signatures generated by the signing oracle. A signature scheme is EUF-CMA secure if we only restrict $(m^*, *) \notin L$ above.

A.5 Message Authentication Codes

We define message authentication codes (MAC).

Definition A.12 (MAC). A (deterministic) message authentication code MAC over a key space \mathcal{K} and a message space \mathcal{M} consists of the following algorithms:

- $\text{TagGen}(k, m) \rightarrow \text{tag}$: On input a key $k \in \mathcal{K}$ and a message $m \in \mathcal{M}$, it (deterministically) outputs a tag tag .
- $\text{TagVerify}(k, m, \text{tag}) \rightarrow \perp/\top$: On input a key k , a message m and a tag tag , it (deterministically) outputs \top or \perp .

Definition A.13 (Correctness). A MAC is correct if for all keys $k \in \mathcal{K}$ and all messages $m \in \mathcal{M}$,

$$\Pr [\text{TagVerify}(k, m, \text{TagGen}(k, m)) = \top] = 1.$$

Following previous works [12, 36], we use a MAC based on a hash function modeled as a random oracle. For instance, we can use HMAC, also used by MLS. Since the security of the MAC is implicitly invoked during in the *generalized selective decryption* (GSD) security game [6, 10, 12, 36], we do not explicitly define the notion of unforgeability.

A.6 HKDF

HKDF is a key derivation function (KDF) based on HMAC. It consists of the two algorithms HKDF.Extract and HKDF.Expand. The extraction algorithm $k \leftarrow \text{HKDF.Extract}(s_0, s_1)$ outputs a uniformly random key k if either s_0 or s_1 has high min-entropy. The expansion algorithm $k_{|l|} \leftarrow \text{HKDF.Expand}(k, |l|)$, on input a key k and a public label $|l|$, outputs a uniformly random key $k_{|l|}$ for $|l|$. Both HKDF.Extract and HKDF.Expand were modeled as a random oracle to prove security of Chained CmpKE [36].

A.7 Pseudorandom Function

Let $F : \mathcal{K} \times \mathcal{D} \rightarrow \mathcal{R}$ be a function family with key space \mathcal{K} , domain \mathcal{D} and finite range \mathcal{R} . We define a pseudorandom function as follows.

Definition A.14 (Pseudorandom Function). We say F is a pseudorandom function (PRF) if for all PPT adversary \mathcal{A} , we have

$$\left| \Pr \left[\begin{array}{l} b = b' : b \leftarrow \{0, 1\}; K \leftarrow \mathcal{K}; RF \leftarrow \mathcal{RF}; \\ b' \leftarrow \mathcal{A}^{\mathcal{F}(\cdot)}(1^\kappa) \end{array} \right] - \frac{1}{2} \right| \leq \text{negl}(\kappa),$$

where \mathcal{RF} is a set of all functions with domain \mathcal{D} and range \mathcal{R} , and $\mathcal{F}(\cdot)$ is defined as $F(K, \cdot)$ if $b = 0$, and $RF(\cdot)$ otherwise.

A.8 Pseudorandom Permutation

Let $\phi : \mathcal{K} \times \mathcal{R} \rightarrow \mathcal{R}$ be a function family of one-to-one functions from \mathcal{R} to \mathcal{R} with key space \mathcal{K} . We define a pseudorandom permutation as follows.

Definition A.15 (Pseudorandom Permutation). We say ϕ is a pseudorandom permutation (PRP) if for all PPT adversary \mathcal{A} , we have

$$\left| \Pr \left[\begin{array}{l} b = b' : b \leftarrow \{0, 1\}; K \leftarrow \mathcal{K}; RP \leftarrow \mathcal{RP}; \\ b' \leftarrow \mathcal{A}^{\mathcal{P}(\cdot)}(1^\kappa) \end{array} \right] - \frac{1}{2} \right| \leq \text{negl}(\kappa),$$

where \mathcal{RP} is the set of all permutations over \mathcal{R} , and $\mathcal{P}(\cdot)$ is defined as $\phi(K, \cdot)$ if $b = 0$, and $RP(\cdot)$ otherwise.

B STATIC METADATA-HIDING CGKA: DEFINITION

In this section, we propose a UC security model capturing the security of the 1st & 2nd layers (i.e., group secret keys and *static* metadata) by defining a new ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. This is an extension of the ideal functionality $\mathcal{F}_{\text{CGKA}}$ [10, 12, 36] that captures the security of the 1st layer.

B.1 Background

B.1.1 Universal Composable Security. We briefly recall the UC framework. We refer to [20, 21] for the full descriptions. The UC security is formalized by the indistinguishability of real and ideal protocols. In the real protocol, parties execute a protocol Π , where an *adversary* \mathcal{A} may corrupt some of the parties. In the ideal protocol, the parties are replaced by *dummy* parties that interact with an *ideal functionality* \mathcal{F} , where a *simulator* \mathcal{S} may corrupt some of the dummy parties. The dummy parties are defined to be the identity function that simply outputs whatever is fed as input. In addition, there is another entity called the *environment* \mathcal{Z} that tries to distinguish the two protocols. In the real (resp. ideal) protocol, \mathcal{Z} can interact arbitrary with \mathcal{A} (resp. \mathcal{S}), and it can also invoke any non-corrupted parties (resp. dummy parties) to honestly run the protocol Π (resp. the ideal functionality \mathcal{F}), where the output is always reported back to \mathcal{Z} . The goal of UC-security is then, given any adversary \mathcal{A} , to construct a simulator \mathcal{S} such that any environment \mathcal{Z} cannot distinguish between the real and ideal protocols. We say the real protocol Π UC-realizes the ideal functionality \mathcal{F} if such \mathcal{S} can be constructed. Put differently, whatever \mathcal{A} can learn from the real protocol Π can be simulated using the information provided by the ideal functionality \mathcal{F} , which is secure by definition.

We often construct a protocol in a setting where a copy of an ideal functionality \mathcal{G} is available. We call such a model \mathcal{G} -hybrid model. If a real protocol Π UC-realizes \mathcal{F} while providing access to an ideal functionality \mathcal{G} , then we say Π UC-realizes \mathcal{F} in the \mathcal{G} -hybrid model. For a real protocol Π in the \mathcal{G} -hybrid model, and a real protocol Π' that realizes \mathcal{G} (in the standard model), we can naturally define a composed protocol $\Pi^{\Pi'}$ in the standard model, in which calls for \mathcal{G} from Π are answered by Π' instead of \mathcal{G} . Canetti [20] proved a *universal composition theorem* stating that, if Π UC-realizes \mathcal{F} in the \mathcal{G} -hybrid model and Π' UC-realizes \mathcal{G} , then $\Pi^{\Pi'}$ UC realizes \mathcal{F} .

The Corruption Model. We define what we mean by “corrupting” a party. In this work, we use the corruption model of continuous state leakage (transient passive corruptions) and adversarially chosen randomness of [10]. This is a standard in CGKA literature, but this is non-standard for typically UC-security. This corruption model allows the adversary to repeatedly corrupt parties by sending them two types of corruption messages: (1) a message `Expose` causes the party to send its current state to the adversary (once), (2) a message `(CorrRand, b)` sets the party’s `rand-corrupted` flag to b . If b is set, the party’s randomness-sampling algorithm is replaced by the adversary providing the coins instead. Ideal functionalities are activated upon corruptions and can adjust their behavior accordingly.

Restricted Environments and Adversaries. To avoid the so-called commitment problem, caused by adaptive corruptions in simulation-based frameworks, we restrict the environment (and thus the adversary) not to corrupt parties at certain times. This roughly corresponds to ruling out “trivial attacks” in game-based definitions, e.g., the adversary cannot compromise the secret key after being provided with the challenge ciphertext. In simulation-based frameworks, such attacks are no longer trivial, but security against them requires relatively strong and inefficient cryptographic tools, e.g., non-committing encryption, and is not achieved by most protocols. We follow prior works [10, 12, 36, 39] and consider a weakened variant of UC-security that only quantifies over a restricted set of so-called admissible environments that do not exhibit the commitment problem. Whether an environment is admissible or not is defined by the ideal functionality \mathcal{F} with statements of the form **restrict cond** and an environment is called admissible (for \mathcal{F}), if it has negligible probability of violating any such *cond* when interacting with \mathcal{F} .

B.1.2 PKI functionality. As in [12, 36], we define our CGKA in a hybrid model where parties can access an ideal functionality that models an (untrusted) PKI. In the real protocol, the parties can interact with the Authentication Service (AS) and Key Service (KS) PKI functionalities. For instance, the environment can instruct the AS (via the party’s protocol) to register a new key for a party. As a result, the AS generates a new key pair for the party and hands the public key to the environment, making the secret key available to the party’s protocol upon request. We note that the adversary can register arbitrary signature keys for any party to capture an insider adversary.

Authentication Service (AS). The authentication service (AS) certifies the ownership of a signature key. The AS is formalized by the functionality \mathcal{F}_{AS} defined in Fig. 10. The definition is identical to that used in [37]. \mathcal{F}_{AS} allows parties to register fresh signature key pairs via `register-svk` query and to check whether a verification key `svk` is registered by a party `id` via the `verify-cert` query. On registration, the new key pair for a party `id` is generated by \mathcal{F}_{AS} using a `genSSK` algorithm (whose concrete specification depends on the CGKA). If `id`’s current randomness source is corrupted (i.e., `RandCorr[id] = ‘bad’`), \mathcal{F}_{AS} asks the adversary to provide the randomness. After registration, `id` receives the new verification key `svk`. Also, parties can retrieve their signing keys via `get-ssk` query and delete registered signing keys via `del-ssk` query. The adversary can register arbitrary verification keys in the name of any party. When a party is corrupted, all signing keys except for the deleted ones are leaked to the adversary. Security is modeled by the ideal-world variant of \mathcal{F}_{AS} , called $\mathcal{F}_{\text{AS}}^{\text{IW}}$. It marks leaked signing keys by storing them in the `ExposedSvk` array (see boxes in Fig. 10). \mathcal{F}_{AS} allows the Key Service functionality \mathcal{F}_{KS} (see below) to signal that a certain `ssk` is leaked. \mathcal{F}_{KS} sends this signal when the signature key is leaked due to a compromise of a key package. Finally, $\mathcal{F}_{\text{AS}}^{\text{IW}}$ always leaks all registered signing keys to the simulator.

Key Service (KS). The Key Service (KS) allows parties to upload one-time key packages used to add them to groups while they are offline. The KS is formalized by the functionality \mathcal{F}_{KS} defined in Fig. 11. The functionality is identical to that used in [37] except that KS checks the validity of maliciously registered key packages

and parties can check whether the key package is registered to KS. In our definition, when the adversary registers a key package to \mathcal{F}_{KS} , \mathcal{F}_{KS} checks whether the registering key packages are valid by the `*validate-kp` function. Thus, \mathcal{F}_{KS} ensures the registered key packages are valid in the sense that the `*validate-kp` function returns `true`. In addition, parties also check the key package is registered to \mathcal{F}_{KS} via a `has-kp` query. This allows parties to check the validity of a key package via the `has-kp` query since \mathcal{F}_{KS} ensures the validity of registered key packages. We introduce these functions to make the syntax of an add and update proposal to look more similar. When a party `id` adds a party `idt`, it first fetches an `idt`'s key package from KS and invokes a CGKA protocol (or \mathcal{F}_{CGKA}) on input `(Propose, 'add'-idt-kpt)`. The party (and \mathcal{F}_{CGKA}) can check the validity of `kpt` through the `has-kp` query to \mathcal{F}_{KS} . This is syntactically similar to an update proposal where, the updated signing key `svk` is validated via the `verify-cert` query.

Other functionalities are identical to that used in [37]. Similar to \mathcal{F}_{AS} , parties can register key packages via the `register-kp` query. Upon receiving the `register-kp` query, \mathcal{F}_{AS} generates a new key package using a `genKP(id, svk, ssk)` algorithm (whose concrete specification depends on the CGKA), which takes a party's identity `id` and a signature key pair `(svk, ssk)` and outputs a key package and the corresponding decryption key. Parties can request another party's key package via `get-kp` query. The returned key package is specified by the adversary. This reflects that the adversary can maliciously inject key packages that were not registered by honest parties. Finally, the ideal-world KS functionality \mathcal{F}_{KS}^{IW} always leaks all decryption keys to the simulator.

B.1.3 History Graph. We use a so-called *history graph* [9, 10, 12, 36] to define the ideal functionality $\mathcal{F}_{CGKA}^{ctxt}$.

Overview. A history graph is a labeled directed graph that acts as a symbolic representation of a group's evolution. It has two types of nodes: commit and proposal nodes, representing all executed commit and propose operations, respectively. Each party is uniquely assigned to a commit node indicating that a party is in a group of members that processed the commit assigned to that specific commit node. The nodes' labels, furthermore, keep track of all the additional information relevant for defining security. For instance, proposal nodes have a label that stores the proposed action, and commit nodes to have labels that store the epoch's application secret and the set of parties corrupted in the given epoch. Security of the application secrets is then formalized by the functionality of choosing a random and independent key for each commit node whenever security is guaranteed; otherwise, the simulator gets to choose the key. Whether security is guaranteed in a given node, is determined via an explicit safe predicate on the node and the history graph. In addition to the secrecy of the keys, the functionality also formalizes authenticity by appropriately disallowing injections.

Formal Definition. As explained in Sec. 3.2, we deviate from the definition of prior history graphs used to define CGKA in the UC framework. Each node in the history graph is identified by node pointers: `prop-id` $\in \mathbb{N}$ for proposal nodes and `node-id` $\in \{0\} \cup \mathbb{N}$ for commit nodes. In contrast, prior works [9, 11] used concrete (non-encrypted) proposals and commits to identify each node. This formalization was well-defined in prior works since each proposal and commit identified a unique group operation in the real protocol.

However, when considering *static metadata-hiding*, two distinct (encrypted) proposals or commits may encrypt the same group operation, in which case, we would like to assign these distinct proposals and commits to the same node. Otherwise, two parties can be in the same group in the real protocol, while they are included in a different commit node in the history graph. Roughly, pointers `prop-id` and `node-id` define the *semantics* of a group operation and allow us to assign semantically equivalent proposals and commits to the same node. The pointer values of a proposal or a commit will be arbitrarily assigned by the simulator, and the ideal functionality checks whether the created history graph maintains consistency, authenticity, and confidentiality.

All nodes in the history graph store the following values:

- `orig`: the identity of the party who created the node, i.e., the message sender.
- `par`: the parent commit node, representing the sender's current epoch.
- `stat` $\in \{\text{'good'}, \text{'bad'}, \text{'adv'}\}$: the flag indicating whether the secrets corresponding to the node are known to the adversary. 'good' means this node is secure, 'bad' means this node is created with adversarial randomness (hence it is well-formed but the adversary knows the secret), and 'adv' means this node is created by the injected message from the adversary.

Proposal nodes further store the following values:

- `act` $\in \{\text{'upd'-kp}, \text{'add'-id_t-kp_t}, \text{'rem'-id_t}\}$: the proposal action. 'upd'-kp means the corresponding party updates its key package to `kp`. 'add'-id_t-kp_t means id_t is added with the key package `kpt`²³.

Commit nodes further store the following values. In this work, history graphs keep `gid`, `epoch`, and new variable `conthide` in addition to the values used in the previous work [12, 37]:

- `gid`: the group identifier.
- `epoch`: the current epoch number.
- `prop`: the ordered list of committed proposals.
- `mem`: the list of a pair of group member's identity and its key package, which is sorted by dictionary order in identities.
- `vcom`: the list of party-dependent commitments associated with this node.
- `key`: the group (application) secret.
- `exp`: the set keeping track of corrupted parties in this node.
- `chall`: the flag indicating whether the group secret is challenged. That is, `chall = true` if a random group key was generated for this node, and `false` if the key was set by the adversary (or not generated).
- `conthide`: the flag indicating whether the static metadata protection is assured at this epoch. That is, `conthide = true` means the messages issued with this epoch's group secret hide metadata, and `false` means the metadata is leaked. This value is initialized when one of proposal/commit/welcome messages is first created at this epoch.

For convenience, we define the following helper function.

- `indexOf(id)`: returns the index of `id` in the list `mem`.

²³The previous models only kept signature keys in `kp`. To capture metadata-hiding property, we need to manage which key packages are being added/updated.

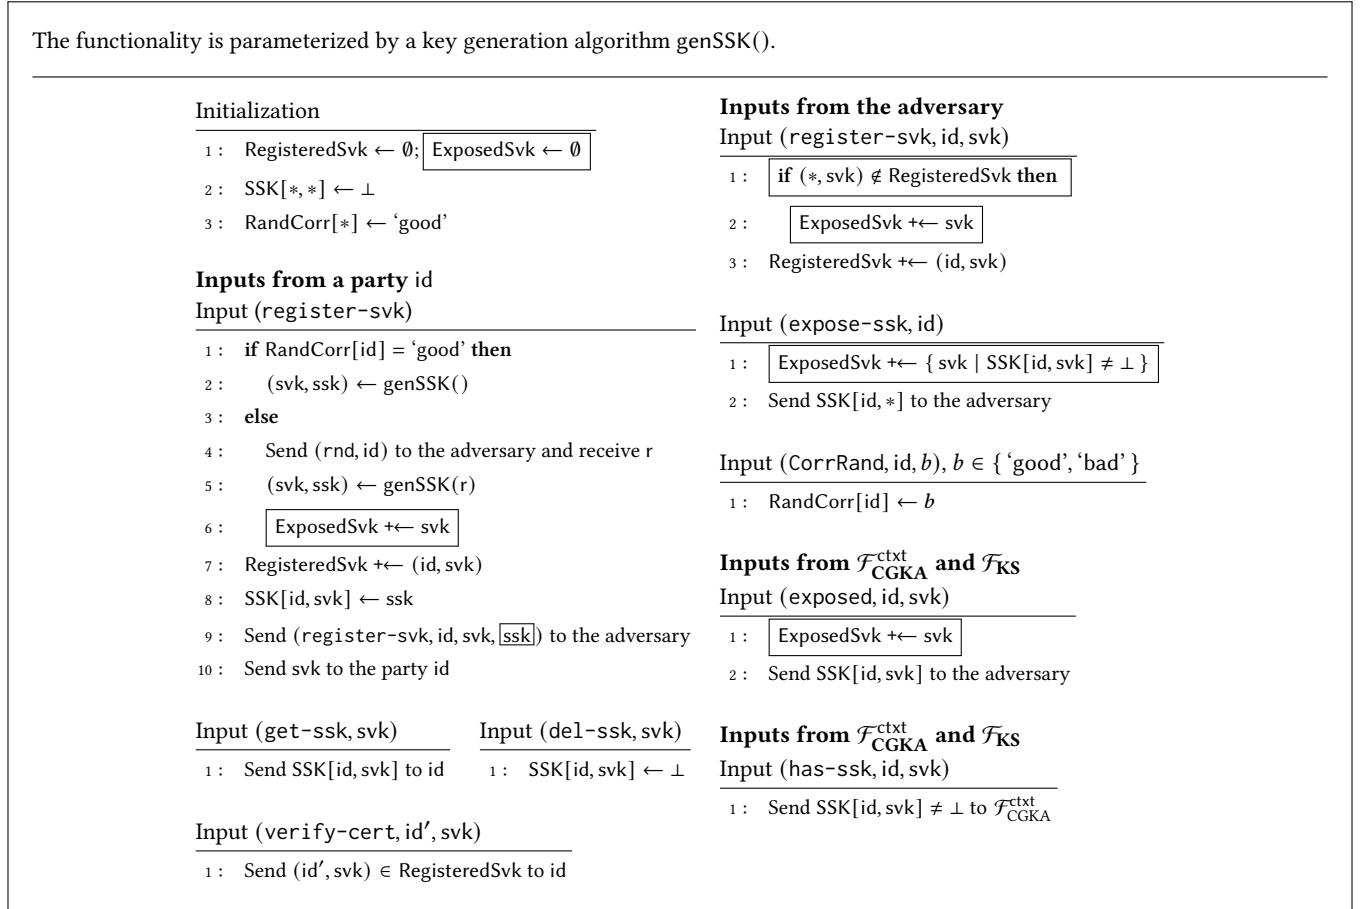


Figure 10: The ideal authentication service functionality \mathcal{F}_{AS} and its variant $\mathcal{F}_{\text{AS}}^{\text{IW}}$ used during the security proof.

B.2 UC Security Model and $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$

We propose a new ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ capturing the static metadata-hiding property of CGKAs. $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ is based on the prior ideal functionality $\mathcal{F}_{\text{CGKA}}$ [36, 37] that only captured the security of group secret keys. The ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ is formally defined in Figs. 12 to 14, along with several helper functions in Figs. 16 to 20 to aid the readability. By setting the flag $\text{flag}_{\text{contHide}}$ to false and $\text{flag}_{\text{selIDL}}$ to false (resp. true) in the "Initialization," $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ becomes identical to $\mathcal{F}_{\text{CGKA}}$ used in [37] (resp. [36] capturing *selective downloading*).

To specify the ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$, we also need to define the following:

Safety Predicates: **safe**, **sig-inj-allowed**, and **mac-inj-allowed** specify which epoch secrets are secure and when authenticity is guaranteed,

Leakage Functions: ***leak-create**, ***leak-prop**, ***leak-com**, ***leak-wel**, and ***leak-proc** specify information leaked from protocol messages.

The safety predicates and leakage functions are protocol specific. For instance, some CGKA may leak the type of proposal, while others may not. Put differently, a specific CGKA UC-realizes the ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ *with respect to* a particular choice of safety predicates and leakage functions. By tuning the choice, $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ allows modeling a wide variety of CGKAs. A concrete choice of such safety predicates and leakage functions is provided in App. C.2, where we prove UC-security of our CGKA Chained CmpKE^{ctxt}.

Below, we provide an overview of the ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$.

States. Different from the previous functionality in [36], $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ maintains parties' key packages kp instead of signature keys only. Also, it keeps track of party-specific secrets (e.g., CmpKE decryption key) by managing the PendDK array (which stores pending secrets of the key package kp), and the CurrDK array (which stores the current secrets of id). In other words, this modification means that $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ manages key packages as if the key service does. Moreover, $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ switches its functionality according to the flag $\text{flag}_{\text{selIDL}}$ and $\text{flag}_{\text{contHide}}$: $\text{flag}_{\text{selIDL}}$ is set to true if it performs selective downloading, and $\text{flag}_{\text{contHide}}$ is set to true if it offers static metadata-hiding. The other states are identical to $\mathcal{F}_{\text{CGKA}}$

The functionality is parameterized by a key-package generation algorithm $\text{genKP}(\text{id}, \text{svk}, \text{ssk})$ and a key package validation function $\text{*validate-kp}(\text{kp}, \text{id})$.

<p>Initialization</p> <hr/> <p>1: RegisteredKp $\leftarrow \emptyset$</p> <p>2: DK[*,*] $\leftarrow \perp$; SVK[*,*] $\leftarrow \perp$</p> <p>3: RandCorr[*] \leftarrow 'good'</p> <p>Inputs from a party id</p> <p>Input (register-kp, svk, ssk)</p> <hr/> <p>1: if RandCorr[id] = 'good' then</p> <p>2: (kp, dk) \leftarrow genKP(id, svk, ssk)</p> <p>3: if kp = \perp then return</p> <p>4: else</p> <p>5: Send (rnd, id) to the adversary and receive r</p> <p>6: (kp, dk) \leftarrow genKP(id, svk, ssk; r)</p> <p>7: if kp = \perp then return</p> <p>8: Send (exposed, id, svk) to \mathcal{F}_{AS}</p> <p>9: RegisteredKp \leftarrow (id, kp)</p> <p>10: DK[id, kp] \leftarrow dk; SVK[id, kp] \leftarrow svk</p> <p>11: Send (register-kp, id, svk, kp, $\overline{\text{dk}}$) to the adversary</p> <p>12: Send kp to the party id</p> <p>Input (get-dks)</p> <hr/> <p>1: Send { (kp, DK[id, kp]) DK[id, kp] $\neq \perp$ } to id</p>	<p>Input (get-kp, id')</p> <hr/> <p>1: Send (get-kp, id, id') to the adversary and receive kp'</p> <p>2: try *validate-kp(kp, id)</p> <p>3: RegisteredKp \leftarrow (id', kp')</p> <p>4: Send kp' to id</p> <p>Input (del-kp, kp)</p> <hr/> <p>1: DK[id, kp], SVK[id, kp] $\leftarrow \perp$</p> <p>Inputs from id and $\mathcal{F}_{CGKA}^{\text{ctxt}}$</p> <p>Input (has-kp, id, kp)</p> <hr/> <p>1: Send (id, kp) \in RegisteredKp</p> <p>Inputs from the adversary</p> <p>Input (CorrRand, id, b), $b \in \{ \text{'good'}, \text{'bad'} \}$</p> <hr/> <p>1: RandCorr[id] $\leftarrow b$</p> <p>Inputs from the adversary and $\mathcal{F}_{CGKA}^{\text{ctxt}}$</p> <p>Input (exposed, id)</p> <hr/> <p>1: Send DK[id, *] to the adversary</p> <p>2: foreach svk \in SVK[id, *] s.t. svk $\neq \perp$ do</p> <p>3: Send (exposed, id, svk) to \mathcal{F}_{AS}</p>
---	--

Figure 11: The ideal key service functionality \mathcal{F}_{KS} and its variant \mathcal{F}_{KS}^{IW} used during the security proof.

in [36] with some syntactical change. $\mathcal{F}_{CGKA}^{\text{ctxt}}$ maintains the history graph. As explained in App. B.1.3, it identifies proposal nodes by a pointer $\text{prop-id} \in \{0\} \cup \mathbb{N}$ and commit nodes by a pointer $\text{node-id} \in \{0\} \cup \mathbb{N}$. We assume one group is created by an honest party (see Create in Fig. 12). This creates a root (commit) node called the *main root* identified by the pointer $\text{node-id} = 0$. We call the group starting from the main root *main group*. Moreover, other roots may be created without a commit message (e.g., when a party processes an injected welcome message that is not directly related to the main group). Such roots are called *detached root*. $\mathcal{F}_{CGKA}^{\text{ctxt}}$ also stores a pointer $\text{Ptr}[\text{id}]$ for each party id. $\text{Ptr}[\text{id}]$ identifies id's current commit node (i.e., current epoch). If id is not in the group, $\text{Ptr}[\text{id}] = \perp$.

Interface. $\mathcal{F}_{CGKA}^{\text{ctxt}}$ offers interfaces to create a group, create a proposal, commit to a list of proposals, process a commit, join a group, and retrieve the group secret key. All interfaces except create and join are for group members only (i.e., parties for which $\text{Ptr}[\text{id}] \neq \perp$). We explain each interface in more detail below.

Group creation (See Fig. 12) $\mathcal{F}_{CGKA}^{\text{ctxt}}$ allows one main group to be created by a designated party $\text{id}_{\text{creator}}$. Initially, the main group has a single party $\text{id}_{\text{creator}}$, and it can invite additional members by issuing add proposals and committing to them. $\mathcal{F}_{CGKA}^{\text{ctxt}}$ checks the

validity of $\text{id}_{\text{creator}}$'s signature key by *valid-svk . Then, $\mathcal{F}_{CGKA}^{\text{ctxt}}$ informs the adversary \mathcal{S}^{24} of the creation of a new group by sending the message (Create) to \mathcal{S} . (This models the fact that a server knows when a group is created.) If $\text{flag}_{\text{contHide}} = \text{false}$, the adversary also receives the identity and signature key of the group creator. Otherwise, the adversary receives ***leak-create**(id, svk). The adversary returns the new group's identity gid. Then, $\mathcal{F}_{CGKA}^{\text{ctxt}}$ generates the initial key package by the *update-kp function. The *update-kp function generates a new key package by itself if both $\text{flag}_{\text{contHide}}$ and **safe** are true; otherwise asks the key package to the adversary \mathcal{S} . This models, in the ideal metadata-hiding CGKA protocol, an honest party generates a new key package, but it is hidden from the adversary.²⁵ Then, $\mathcal{F}_{CGKA}^{\text{ctxt}}$ initializes the root node. Note that the epoch of an initial group is set to 0.

Creating proposals (See Fig. 12) A party id can be invoked by the environment \mathcal{Z} to create a proposal with a specific action act. $\mathcal{F}_{CGKA}^{\text{ctxt}}$ informs the adversary \mathcal{S} that a proposal message is being created.

²⁴In the UC framework, it is conventional to call \mathcal{S} appearing in the ideal functionality as the "adversary." We use the term "simulator" during the security proof.

²⁵If the flag $\text{flag}_{\text{contHide}}$ is false, the key package is asked to the adversary \mathcal{S} ; this means the key package is known to the adversary because the group creator is known.

\mathcal{S} receives ***leak-prop**(id, act) and returns a flag *ack*, an ideal proposal p and a node pointer prop-id. If $\text{flag}_{\text{contHide}} = \text{false}$, then \mathcal{S} obtains all the information $(\text{Ptr}[\text{id}], \text{id}, \text{act}) = \text{*leak-prop}(\text{id}, \text{act})$ that can be inferred from a non-encrypted proposal. $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ allows \mathcal{S} to send $\text{ack} = \text{false}$ to report that the protocol fails. If $\text{act} = \text{'upd'}$, $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ updates act with the new updated key package returned from ***update-kp** function. The ***update-kp** function generates a new key package by itself if both $\text{flag}_{\text{contHide}}$ and **safe** are true; otherwise asks the key package to the adversary \mathcal{S} . This models, in the ideal metadata-hiding CGKA protocol, an honest party generates a new key package, but the proposal message hides the key package from the adversary.²⁶ If the protocol succeeds, and if no node associated with p exists, $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ creates a new proposal node $\text{Prop}[\text{propCtr}]$ and assigns propCtr to p (setting $\text{PropID}[p] \leftarrow \text{propCtr}$). In certain situations, $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ may not create a new proposal node. For example, id proposes to remove the same party twice in the same epoch. Another example is when the adversary \mathcal{S} controls the party's randomness (via setting $\text{RandCorr} = \text{'bad'}$) and the party proposes to update using the same randomness twice. In these cases, \mathcal{S} can specify to attach the created proposal p to an existing proposal node prop-id. $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ then enforces that the states on the existing proposal node are consistent with the expected one using ***consistent-prop**. $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ marks whether the current epoch is secure or not using ***mark-content-hidden-epoch**. This information is used to determine epochs the adversary is allowed to corrupt (see the **restrict** check run within **Expose** in Fig. 14). Finally, if all check passes, $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ returns the proposal message p to the calling (dummy) party id, which simply relays it to the environment \mathcal{Z} .

Committing to proposals (See Fig. 12) A party id can be invoked by the environment \mathcal{Z} to create a commit with a list of proposals \vec{p} , along with a (possibly fresh) signature verification key svk. $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ informs the adversary \mathcal{S} that a commit message is sent and provides ***leak-com**(id, \vec{p} , svk). If $\text{flag}_{\text{contHide}} = \text{false}$, then \mathcal{S} obtains all the information $(\text{Ptr}[\text{id}], \text{id}, \vec{p}, \text{svk}, \text{mem}) = \text{*leak-com}(\text{id}, \vec{p}, \text{svk})$ that can be inferred from a non-encrypted commit. $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ receives a flag *ack*, a commit node node-id, and a commit message (c_0, \vec{c}) . Here, \vec{c} is a list of party dependent messages $(\vec{c}_d)_{\text{id}}$; if selective downloading is performed (i.e., $\text{flag}_{\text{selDL}}$ is true), then party id only needs to retrieve (c_0, \vec{c}_d) from the server. The adversary \mathcal{S} sets $\text{ack} := \text{false}$ to report that the protocol fails. $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ then obtains the new updated key package via the ***update-kp** function. This models, in the ideal metadata-hiding CGKA protocol, an honest committer generates a new key package, and it is hidden from the adversary. If the commit protocol succeeds, $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ first asks \mathcal{S} to interpret the injected proposals, i.e., proposals where no node has been created, by calling ***fill-prop**. It then computes the new member set resulting from applying \vec{p} to the current member set by calling ***next-members** (which returns \perp if \vec{p} contains invalid proposals).

$\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ then checks the format of \vec{c} specified by \mathcal{S} . If $\text{flag}_{\text{selDL}}$ is true, $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ requires that \vec{c} contains the same number of party-dependent messages as the number of the current members. Else, \vec{c} must be \perp .

²⁶If the flag $\text{flag}_{\text{contHide}}$ is false, the key package is asked to the adversary \mathcal{S} ; this means the key package is known to the adversary.

$\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ then either creates a new commit node or verifies that the existing node is consistent by ***consistent-com**. The adversary \mathcal{S} can specify an existing node-id. This case may happen for example when the adversary makes a party process an injected commit message c_0 and then makes another party commit the same c_0 by controlling its randomness. If the specified node $\text{Node}[\text{node-id}]$ is a detached root, $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ attaches it to id's current node by calling ***attach**. Once the detached root is attached to the main group, the root's tree achieves the same security guarantee as the main tree. Since attaching a detached root changes the topology of the history graph, $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ enforces two invariants: **cons-invariant** enforces the consistency of the graph, and **auth-invariant** enforces the authenticity guarantee.

When add proposals are committed (i.e., $\text{addedMem} \neq \perp$), $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ informs the adversary \mathcal{S} that a welcome message is sent and provides ***leak-wel** on input the id's current epoch $\text{Ptr}[\text{id}]$, the new epoch $\text{Node}[c_0]$ and the receiver's identity id_t . If $\text{flag}_{\text{contHide}} = \text{false}$, then \mathcal{S} obtains all the information that can be inferred from a non-encrypted welcome message. \mathcal{S} returns a simulated welcome message \vec{w} to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. We note that in prior definitions [36, 37], the commit and welcome messages were simultaneously simulated by \mathcal{S} . We consciously divide this process into two. This allows us to model the fact that a welcome message does not necessarily leak information about the group. That is, the server can observe that a party id is invited to some group but will not know which group.²⁷ $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ assigns the welcome message to the commit node created above. Finally, $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ marks whether the current epoch is secure or not using ***mark-content-hidden-epoch** and returns (c_0, \vec{c}, \vec{w}) to the calling (dummy) party id, which simply relays it to the environment \mathcal{Z} .

Processing commits (See Fig. 13) A party id can be invoked by the environment \mathcal{Z} to process a commit message with an associating list of proposals (c_0, \vec{c}, \vec{p}) . We explain the case where selective downloading is performed (i.e., $\text{flag}_{\text{selDL}}$ is true). $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ first checks that \vec{c} is the correct id-dependent message associated with c_0 , and outputs \perp if it is incorrect. (In case $\text{flag}_{\text{selDL}}$ is false, \vec{c} must be \perp .) $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ then calls \mathcal{S} on input ***leak-prop**(id) and (c_0, \vec{c}, \vec{p}) . \mathcal{S} sets $\text{ack} := \text{false}$ to report that the protocol failed. If the process succeeds, $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ first asks the adversary to interpret the injected proposals by calling ***fill-prop**. $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ then either creates a new commit node or verifies that the existing node is consistent. The adversary can specify the existing node-id. If the node corresponding to c_0 does not exist and the adversary does not specify any existing node, $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ checks the validity of \vec{p} and creates a new commit node with the committer identity orig' and its signature key svk' which are \mathcal{S} interprets from (c_0, \vec{c}, \vec{p}) . Note that the new node holds the same group identity as id's current node and the epoch is incremented. If the commit message was assigned a node (i.e., $\text{Node}[c_0] \neq \perp$) or the adversary \mathcal{S} specifies an existing node, $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ checks the validity of the group identity and epoch and enforces that it is a valid successor of id's current node by calling ***valid-successor**. If c_0 is assigned to a detached root, $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ attaches the root to id's current node. If c_0 is not assigned a node, $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ assigns the adversary-specified node-id to c_0 .

²⁷Note that we can capture the situation where a welcome message leaks the group by defining ***leak-wel** to output the node pointer of the commit message.

Finally, if c_0 removes id , $\mathcal{F}_{CGKA}^{ctxt}$ sets $\text{Ptr}[id] = \perp$. Otherwise, $\mathcal{F}_{CGKA}^{ctxt}$ updates id 's secrets if necessary and moves $\text{Ptr}[id]$ to the new commit node. The calling party receives the committer's identity, the semantics of the applied proposals, and the list of (id, svk) -pair.

Joining a group (See Fig. 13) A party id can be invoked by the environment \mathcal{Z} to join a group using the welcome message \widehat{w} . $\mathcal{F}_{CGKA}^{ctxt}$ forwards (id, \widehat{w}) to the adversary \mathcal{S} and receives the interpreted result. As usual, the adversary sets $ack := \text{false}$ to report that the protocol failed. If the process succeeds, $\mathcal{F}_{CGKA}^{ctxt}$ identifies the commit node $\text{node-id} = \text{Wel}[id, \widehat{w}]$ corresponding to \widehat{w} . If this is the first time $\mathcal{F}_{CGKA}^{ctxt}$ sees \widehat{w} , i.e., $\text{Wel}[id, \widehat{w}] = \perp$, \mathcal{S} can specify $\text{node-id}'$. If the commit node for $\text{node-id}'$ does not exist (i.e., $\text{Node}[\text{node-id}'] = \perp$), $\mathcal{F}_{CGKA}^{ctxt}$ creates a new detached root where all the stored values are chosen by \mathcal{S} . Finally, $\mathcal{F}_{CGKA}^{ctxt}$ updates id 's secrets (registered in the key service \mathcal{F}_{KS}) and returns the state of the joining group (the committer's identity, the group identity, the epoch, and the list of (id, kp) -pair) to the calling (dummy) party id .

Group keys (See Fig. 13) Parties can fetch the current group secret via the Key query. The returned group secret k is random if the protocol guarantees its confidentiality (identified by the **safe** predicate). Otherwise, k is set by the adversary. Unlike prior definitions [36, 37], we also prepare two extra group secret keys: the metadata key k_{mh} that can be obtained using Key_{mh} , and the next metadata key k_{mh}' that can be obtained using NextKey_{mh} . While Key_{mh} is defined identically to Key , NextKey_{mh} is a function that allows obtaining the metadata key at the next epoch. These keys are explicitly used to hide the *dynamic* metadata. That is, if we only care about hiding static metadata, these two group secret keys can be safely omitted from the definition.

Corruptions (See Fig. 14) The adversary \mathcal{S} can corrupt a party id using the Expose query. When the adversary \mathcal{S} inputs (Expose, id) to $\mathcal{F}_{CGKA}^{ctxt}$, the ideal functionality records the following leaked information:

- The current group secret keys and id 's key materials (e.g., encryption key and signing key). This is recorded by adding id to the exposed set of id 's current node (cf. Line 2 in (Expose, id) query).
- The key materials created by id during an update or a commit at the current epoch. This is recorded by setting the status of all the child commit nodes created by id (i.e., nodes with $\text{par} = \text{Ptr}[id]$) to 'bad' (cf. $\text{*update-stat-after-exp}$ function).
- The current signature signing key ssk . This is recorded by signaling to \mathcal{F}_{AS} that svk is exposed and sends ssk to the adversary (cf. Line 5 in (Expose, id) query).

Then, the ideal functionality gives id 's current epoch $\text{Ptr}[id]$, the associated information $\text{Node}[\text{Ptr}[id]]$, and id 's current secret keys stored in the CurrDK array. (Note that $\mathcal{F}_{CGKA}^{ctxt}$ manages users' key packages and secret information, see above.) Also, the adversary is allowed to corrupt a non-group member id as well. In such a case, the key packages that id registered to the key service \mathcal{F}_{KS} are leaked. $\mathcal{F}_{CGKA}^{ctxt}$ signals to \mathcal{F}_{KS} that key packages (including the signing key) are exposed and send the corresponding decryption

keys and signing keys to the adversary (cf. Line 7 in (Expose, id) query).

If an adversary is allowed to compromise key materials that can be used to compute a group secret key, which $\mathcal{F}_{CGKA}^{ctxt}$ has already assigned random values, then \mathcal{Z} trivially distinguishes a real protocol from an ideal protocol. For instance, if \mathcal{Z} queries Key for a commit node where the predicate **safe** is true, then $\mathcal{F}_{CGKA}^{ctxt}$ assigns a random value to the group secret key k . Then, if the adversary at some later point compromises a party id via an (Expose, id) and can compute the *real* group secret key k' from the compromised key materials, \mathcal{Z} can distinguish the two protocols by checking if $k = k'$.

To avoid such trivial attacks, we restrict the environment to not be able to corrupt key materials for those commit nodes with $\text{chall} = \text{true}$ or $\text{conthide} = \text{true}$. The former is identical to those used in prior works [36, 37]; if a random group secret key was set, then \mathcal{Z} cannot corrupt a party that allows recovering of the real group secret key. The latter is new to this work. For static metadata-hiding, recall that we must encrypt the proposal and commits. This is modeled in $\mathcal{F}_{CGKA}^{ctxt}$ by requiring the adversary \mathcal{S} to create the (encrypted) proposal and commits *without* knowing the message when the predicate **safe** is true. If \mathcal{Z} were to compromise a party that allows recovering the real group secret key, it can try to decrypt the encrypted proposal or commit to trivially distinguish between a real and ideal protocol. $\text{conthide} = \text{true}$ indicates that a commit node created random encryption and restricts \mathcal{Z} from later corrupting it. Note that this implies that the predicate **safe** for honestly generated commit nodes cannot be switched from true to false once created. This is in sharp contrast to previous definitions since the group secret key was never implicitly used as part of the real protocol. We note that this is not a weakness of our security model but rather a natural consequence of considering CGKAs in a larger system.

Initialization

```

1: Ptr[*], Prop[*], Node[*], Wel[*] ← ⊥
2: propCtr, nodeCtr ← 1
3: / Flag is set to true if selective downloading is performed.
4: flagSelDL = true
5: / Flag is set to true if propose and commit contents are hidden
6: flagContHide = true
7: PropID[*], NodeID[*] ← ⊥
8: PendDK[*], CurrDK[*] ← ⊥
9: RandCorr[*] ← 'good'

```

Inputs from a party $id_{creator}$

Input (Create, svk)

```

1: req Ptr[idcreator] = ⊥
2: Send (Create, *leak-create(idcreator, svk)) to S and receive gid
3: req *valid-svk(idcreator, svk)
4: (kp, dk) ← *update-kp(id, svk)
5: mem ← { (idcreator, kp) }; CurrDK[id] ← dk
6: Node[0] ← *create-root(gid, 0, idcreator, mem, RandCorr[id])
7: Ptr[idcreator] ← 0

```

Inputs from a party id

Input (Propose, act), act ∈ { 'upd'-svk, 'add'-id_t-kp_t, 'rem'-id_t }

```

1: req Ptr[id] ≠ ⊥
2: Send (Propose, *leak-prop(id, act)) to S and receive (ack, prop-id, p)
3: req ack
4: if act = 'upd'-svk then
5:   req *valid-svk(id, svk)
6:   (kp, dk) ← *update-kp(id, svk)
7:   act ← 'upd'-kp
8:   PendDK[kp] ← dk
9: if act = 'add'-idt-kpt then req *valid-kp(idt, kpt)
10: if PropID[p] = ⊥ ∧ prop-id = ⊥ then
11:   Prop[propCtr] ← *create-prop(Ptr[id], id, act, RandCorr[id])
12:   PropID[p] ← propCtr; propCtr++
13: else
14:   if PropID[p] = ⊥ then (prop-id', PropID[p]) ← (prop-id, prop-id)
15:   else prop-id' ← PropID[p]
16:   *consistent-prop(prop-id', id, act)
17: if act = 'upd'-svk ∧ RandCorr[id] = 'bad' then
18:   Send (exposed, id, svk) to FAS
19: / Mark whether generated messages hide contents (static metadata)
20: *mark-content-hidden-epoch(Ptr[id])
21: return p

```

Input (Commit, \vec{p} , svk)

```

1: req Ptr[id] ≠ ⊥
2: Send (Commit, *leak-com(id,  $\vec{p}$ , svk)) to S and receive (ack, node-id, c0,  $\vec{c}$ )
3: req *valid-svk(id, svk)
4: (kpnew, dknew) ← *update-kp(id, svk)
5: PendDK[kpnew] ← dknew
6: req *succeed-com(id,  $\vec{p}$ , kpnew) ∨ ack
7: *fill-prop( $\vec{p}$ )
8: (mem', *) ← *next-members(Ptr[id], id,  $\vec{p}$ , kpnew)
9: assert mem' ≠ ⊥ ∧ (id, kpnew) ∈ mem'
10: / If selective downloading is performed, then member specific  $\vec{c}$  has the same size
    / as the current member. Otherwise,  $\vec{c}$  is ⊥
11: if flagSelDL then
12:   assert | $\vec{c}$ | = |Node[Ptr[id]].mem|
13: else
14:   assert  $\vec{c}$  = ⊥
15: if NodeID[c0] = ⊥ ∧ node-id = ⊥ then
16:   Node[nodeCtr] ← *create-child(Ptr[id], id,  $\vec{p}$ ,  $\vec{c}$ , mem', RandCorr[id])
17:   NodeID[c0] ← nodeCtr; nodeCtr++
18: else
19:   if NodeID[c0] = ⊥ then (node-id', NodeID[c0]) ← (node-id, node-id)
20:   else node-id' ← NodeID[c0]
21:   *consistent-com(node-id', id,  $\vec{p}$ , mem)
22:   if Node[node-id'].par = ⊥ then
23:     *attach(node-id', id,  $\vec{p}$ )
24: / Create welcome message for added members
25: addedMem ← Node[NodeID[c0]].mem \ Node[Ptr[id]].mem
26:  $\vec{w}$  ← ∅
27: foreach (idt, *) ∈ addedMem do
28:   Send (Welcome, *leak-wel(Ptr[id], NodeID[c0], idt)) to S and
29:   receive (ack,  $\vec{w}$ )
30:   req ack
31:   parse (idt, *) ←  $\vec{w}$ 
32:   assert Wel[idt,  $\vec{w}$ ] ∈ { ⊥, NodeID[c0] }
33:   Wel[idt,  $\vec{w}$ ] ← NodeID[c0]
34:    $\vec{w}$  ←  $\vec{w}$  ∪  $\vec{w}$ 
35: assert cons-invariant ∧ auth-invariant
36: if RandCorr[id] = 'bad' then
37:   Send (exposed, id, svk) to FAS
38: / Mark whether generated messages hide contents (static metadata)
39: *mark-content-hidden-epoch(Ptr[id])
40: if  $\vec{w} ≠ ∅$  then *mark-content-hidden-epoch(NodeID[c0])
41: return (c0,  $\vec{c}$ ,  $\vec{w}$ )

```

Figure 12: The ideal static metadata-hiding CGKA functionality $\mathcal{F}_{CGKA}^{ctxt}$: Create, Propose, Commit. The modifications in order for the functionality to keep track of key packages are highlighted in orange.

Input (Process, $c_0, \widehat{c}, \vec{p}$)

```

1: req Ptr[id]  $\neq \perp$ 
2: / If  $c_0$  and  $\vec{p}$  were generated honestly, then use the existing node-id  $\neq \perp$ .
   / Otherwise, the group secret must be exposed and  $\mathcal{S}$  decides where to attach Ptr[id].
3: Send (Process, *leak-proc(id),  $c_0, \widehat{c}, \vec{p}$ ) to  $\mathcal{S}$  and
4:   receive (ack, node-id, orig', kp')
5: req *succeed-proc(id,  $c_0, \widehat{c}, \vec{p}$ )  $\vee$  ack
6: *fill-prop( $\vec{p}$ )
7: / If selective downloading is performed,
   / then only the member-specific commitment  $\widehat{c}$  is accepted when  $c_0$  is honestly generated.
   if flagselDL then
8:   node-id  $\leftarrow$  NodeID[ $c_0$ ]
9:   if node-id  $\neq \perp \wedge$  Node[node-id].stat = 'good' then
10:     indexid  $\leftarrow$  Node[Ptr[id]].indexOf(id)
11:     req  $\widehat{c} =$  Node[node-id].vcom[indexid]
12:   else
13:     assert  $\widehat{c} = \perp$  / Otherwise,  $\widehat{c}$  is  $\perp$ .
14: / If node-id =  $\perp$ , create a new node; Otherwise, check consistency with existing node.
   if NodeID[ $c_0$ ] =  $\perp \wedge$  node-id =  $\perp$  then
15:   try (mem', *)  $\leftarrow$  *next-members(Ptr[id], orig',  $\vec{p}$ , kp')
16:   assert mem'  $\neq \perp$ 
17:   Node[nodeCtr]  $\leftarrow$  *create-child(Ptr[id], orig',  $\vec{p}$ , mem', 'adv')
18:   (NodeID[ $c_0$ ], node-id)  $\leftarrow$  (nodeCtr, nodeCtr)
19:   nodeCtr++
20: else
21:   if NodeID[ $c_0$ ] =  $\perp$  then (node-id', NodeID[ $c_0$ ])  $\leftarrow$  (node-id, node-id)
22:   else node-id'  $\leftarrow$  NodeID[ $c_0$ ]
23: / After processing, require gid to remain the same and epoch to be incremented by 1.
24: assert Node[Ptr[id]].gid = Node[node-id'].gid
25: assert Node[Ptr[id]].epoch = Node[node-id'].epoch + 1
26: idc  $\leftarrow$  Node[node-id'].orig; kpc  $\leftarrow$  Node[node-id'].mem[idc]
27: (mem', *)  $\leftarrow$  *next-members(Ptr[id], idc,  $\vec{p}$ , kpc)
28: assert mem'  $\neq \perp$ 
29: *valid-successor(node-id', idc,  $\vec{p}$ , mem')
30: if Node[node-id'].par =  $\perp$  then *attach(node-id', id,  $\vec{p}$ )
31: / Mark whether generated messages hide contents (static metadata)
32: *mark-content-hidden-epoch(Ptr[id])
33: if  $\exists p \in \vec{p} : \text{Prop}[p].\text{act} = \text{'rem'-id}$  then Ptr[id]  $\leftarrow \perp$ 
34: else
35:   assert (id, *)  $\in$  Node[NodeID[ $c_0$ ]].mem
36: / Fetch new dk if key package is updated
37: if Node[Ptr[id]].mem[id]  $\neq$  Node[NodeID[ $c_0$ ]].mem[id] then
38:   kp  $\leftarrow$  Node[NodeID[ $c_0$ ]].mem[id]
39:   CurrDK[id]  $\leftarrow$  PendDK[kp]
40:   Ptr[id]  $\leftarrow$  NodeID[ $c_0$ ]
41: assert cons-invariant  $\wedge$  auth-invariant
42: return *output-proc(node-id')
```

Input (Join, id, \widehat{w})

```

1: req Ptr[id] =  $\perp$ 
2: / If  $\widehat{w}$  was generated honestly, then use the existing node-id  $\neq \perp$ .
3: Send (Join, id,  $\widehat{w}$ ) to  $\mathcal{S}$  and
   receive (ack, node-id', gid', epoch', orig', mem')
4: req *succeed-wel(id,  $\widehat{w}$ )  $\vee$  ack
5: node-id  $\leftarrow$  Wel[id,  $\widehat{w}$ ]
6: if node-id =  $\perp$  then
7:   if Node[node-id']  $\neq \perp$  then node-id  $\leftarrow$  node-id'
8:   else
9:     Node[nodeCtr]  $\leftarrow$  *create-root(gid', epoch', orig', mem', 'adv')
10:    node-id  $\leftarrow$  nodeCtr
11:    nodeCtr++
12:    Wel[id,  $\widehat{w}$ ]  $\leftarrow$  node-id
13:   kp  $\leftarrow$  Node[node-id].mem[id]
14:   CurrDK[id]  $\leftarrow$  DK[id, kp] / Fetch the registered dk used to join
15:   Ptr[id]  $\leftarrow$  node-id
16:   assert (id, *)  $\in$  Node[node-id].mem
17:   assert cons-invariant  $\wedge$  auth-invariant
18:   return *output-join(node-id)
```

Input (Key)

```

1: req Ptr[id]  $\neq \perp$ 
2: if Node[Ptr[id]].key =  $\perp$  then *set-key(Ptr[id])
3: return Node[Ptr[id]].key
```

Input (Key_{mh})

```

1: req Ptr[id]  $\neq \perp$ 
2: if Node[Ptr[id]].kmh =  $\perp$  then *set-key(Ptr[id])
3: return Node[Ptr[id]].kmh
```

Input (NextKey_{mh}, c_0)

```

1: req Ptr[id]  $\neq \perp \wedge$  NodeID[ $c_0$ ]  $\neq \perp$ 
2: req Node[NodeID[ $c_0$ ]].par = Ptr[id]  $\wedge$  Node[NodeID[ $c_0$ ]].orig = id
3: if Node[NodeID[ $c_0$ ]].kmh =  $\perp$  then *set-key(NodeID[ $c_0$ ])
4: return Node[NodeID[ $c_0$ ]].kmh
```

Figure 13: The ideal static metadata-hiding CGKA functionality $\mathcal{F}_{CGKA}^{\text{ctx}}$: Process, Join, Key, Key_{mh}, and NextKey_{mh}. The last two Key_{mh} and NextKey_{mh} are to be used in a higher layer protocol. The modifications in order for the functionality to keep track of key packages are highlighted in orange

Input (Expose, id)	Input (CorrRand, id, b), $b \in \{ \text{'good'}, \text{'bad'} \}$
1: if $\text{Ptr}[\text{id}] \neq \perp$ then	1: $\text{RandCorr}[\text{id}] \leftarrow b$
2: $\text{Node}[\text{Ptr}[\text{id}]].\text{exp} \leftarrow \text{id}$	
3: $\text{*update-stat-after-exp}(\text{id})$ / Pending secrets are marked as exposed.	
4: $\text{svk} \leftarrow \text{Node}[\text{Ptr}[\text{id}]].\text{mem}[\text{id}].\text{svk}$ / Take svk from id's key package	
5: Send (exposed, id, svk) to \mathcal{F}_{AS}	
6: Send ($\text{Ptr}[\text{id}], \text{Node}[\text{Ptr}[\text{id}]]$) to \mathcal{S} / All information stored in $\text{Node}[\text{Ptr}[\text{id}]]$ is sent to \mathcal{S} .	
7: Send $\text{CurrDK}[\text{id}]$ to \mathcal{S} / id's secret key is sent to \mathcal{S} .	
8: Send (exposed, id) to \mathcal{F}_{KS}	
9: restrict $\forall \text{node-id}$:	
10: if $\text{Node}[\text{node-id}].\text{chall} = \text{true}$ then $\text{safe}(\text{node-id}) = \text{true}$	
11: if $\text{Node}[\text{node-id}].\text{conthide} = \text{true}$ then $\text{safe}(\text{node-id}) = \text{true}$	

Figure 14: The static metadata-hiding CGKA functionality $\mathcal{F}_{CGKA}^{\text{ctx}}$: Corruptions from the adversary \mathcal{S} . The difference between those of \mathcal{F}_{CGKA} [37] is highlighted in yellow. The modifications in order for the functionality to keep track of key packages are highlighted in orange

$\text{*create-root}(\text{gid}, \text{epoch}, \text{id}, \text{mem}, \text{stat})$	$\text{*set-key}(\text{node-id})$
1: return new node with $\text{par} \leftarrow \perp, \text{orig} \leftarrow \text{id}, \text{gid} \leftarrow \text{gid}$ epoch \leftarrow epoch, $\text{prop} \leftarrow \perp, \text{mem} \leftarrow \text{mem}, \text{stat} \leftarrow \text{stat}$.	1: if $\text{safe}(\text{node-id})$ then
	2: $\text{Node}[\text{node-id}].\text{key} \leftarrow \mathcal{K}$
	3: $\text{Node}[\text{node-id}].\text{chall} \leftarrow \text{true}$
	4: else
	5: Send (Key, id) to \mathcal{S} and receive k
	6: $\text{Node}[\text{node-id}].\text{key} \leftarrow k$
	7: $\text{Node}[\text{node-id}].\text{chall} \leftarrow \text{false}$
$\text{*create-child}(\text{node-id}, \text{id}, \vec{p}, \vec{c}, \text{mem}, \text{stat})$	$\text{*mark-content-hidden-epoch}(\text{node-id})$
1: $(\text{gid}, \text{epoch}) \leftarrow (\text{Node}[\text{Ptr}[\text{id}]].\text{gid}, \text{Node}[\text{Ptr}[\text{id}]].\text{epoch})$	1: if $\text{safe}(\text{node-id})$ then
2: / Create a new node with an incremented epoch.	2: $\text{Node}[\text{node-id}].\text{conthide} \leftarrow \text{true}$
3: return new node with $\text{par} \leftarrow \text{node-id}, \text{orig} \leftarrow \text{id}$, epoch \leftarrow epoch + 1, $\text{prop} \leftarrow \vec{p}, \text{vcom} \leftarrow \vec{c}$, mem \leftarrow mem, $\text{stat} \leftarrow \text{stat}$.	3: else
	4: $\text{Node}[\text{node-id}].\text{conthide} \leftarrow \text{false}$
$\text{*create-prop}(\text{node-id}, \text{id}, \text{act}, \text{stat})$	$\text{*update-stat-after-exp}(\text{id})$
1: return new node with $\text{par} \leftarrow \text{node-id}, \text{orig} \leftarrow \text{id}$, act \leftarrow act, $\text{stat} \leftarrow \text{stat}$.	1: foreach prop-id s.t. $\text{Prop}[\text{prop-id}] \neq \perp$
	2: $\wedge \text{Prop}[\text{prop-id}].\text{par} = \text{Ptr}[\text{id}]$
	3: $\wedge \text{Prop}[\text{prop-id}].\text{orig} = \text{id}$
	4: $\wedge \text{Prop}[\text{prop-id}].\text{act} = \text{'upd'}$ * do
	5: $\text{Prop}[\text{prop-id}].\text{stat} \leftarrow \text{'bad'}$
	6: foreach node-id s.t. $\text{Node}[\text{node-id}] \neq \perp$
	7: $\wedge \text{Node}[\text{node-id}].\text{par} = \text{Ptr}[\text{id}]$
	8: $\wedge \text{Node}[\text{node-id}].\text{orig} = \text{id}$ do
	9: $\text{Node}[\text{node-id}].\text{stat} \leftarrow \text{'bad'}$
$\text{*fill-prop}(\vec{p})$	
1: foreach $p \in \vec{p}$ s.t. $\text{PropID}[p] = \perp$ do	
2: Send (Propose, $\text{Ptr}[\text{id}], p$) to \mathcal{S} and receive ($\text{prop-id}, \text{orig}, \text{act}$)	
3: if $\text{prop-id} = \perp$ then	
4: $\text{Prop}[\text{propCtr}] \leftarrow \text{*create-prop}(\text{Ptr}[\text{id}], \text{orig}, \text{act}, \text{'adv'})$	
5: $\text{PropID}[p] \leftarrow \text{propCtr}$	
6: $\text{propCtr}++$	
7: / If $\text{flag}_{\text{contHide}} = \text{true}$ and p hides the same content, then \mathcal{S} outputs $\text{prop-id} \neq \perp$. / In this case, check consistency with the exiting node.	
8: else	
9: $\text{*consistent-prop}(\text{prop-id}, \text{orig}, \text{act})$	
10: $\text{PropID}[p] \leftarrow \text{prop-id}$	

Figure 15: The helper functions for creating and maintaining the history graph.

*valid-kp(id, kp) 1: $ack \leftarrow \text{query}(\text{has-kp}, id, kp)$ to \mathcal{F}_{KS} 2: return ack	*valid-svk(id, svk') 1: if $\text{Ptr}[id] \neq \perp$ then 2: $svk \leftarrow \text{Node}[\text{Ptr}[id]].\text{mem}[id]$ 3: if $svk' \neq \perp \wedge svk = svk'$ then 4: return true 5: $ack \leftarrow \text{query}(\text{has-ssk}, id, svk')$ to \mathcal{F}_{AS} 6: return ack	*update-kp(id, svk) 1: if $\text{flag}_{\text{contHide}} \wedge \text{safe}(\text{Ptr}[id])$ then 2: $ssk \leftarrow \text{query}(\text{get-ssk}, svk)$ to \mathcal{F}_{AS} on behalf of id 3: if $\text{RandCorr}[id] = \text{'good'}$ then 4: $(kp, dk) \leftarrow \text{genKP}(id, svk, ssk)$ 5: else 6: Send (rnd, id) to the adversary and receive r 7: $(kp, dk) \leftarrow \text{genKP}(id, svk, ssk; r)$ 8: else 9: Receive (kp, dk) from \mathcal{S} 10: return (kp, dk)
---	--	--

Figure 16: The helper functions related to keys. The `*update-kp` function highlighted in orange is newly introduced to manage key packages in the functionality.

output-proc(node-id) 1: $id_c \leftarrow \text{Node}[\text{node-id}].\text{orig}$ 2: $kp_c \leftarrow \text{Node}[\text{node-id}].\text{mem}[id_c]$ 3: $(, \text{propSem}) \leftarrow \text{*next-members}(\text{node-id}, id_c, \text{Node}[\text{node-id}].\text{prop}, kp_c)$ 4: return $(\text{Node}[\text{node-id}].\text{orig}, \text{propSem}, \text{Node}[\text{node-id}].\text{mem})$	*output-join(node-id) 1: $gid \leftarrow \text{Node}[\text{node-id}].gid$ 2: $\text{epoch} \leftarrow \text{Node}[\text{node-id}].\text{epoch}$ 3: $\text{mem} \leftarrow \text{Node}[\text{node-id}].\text{mem}$ 4: $id_c \leftarrow \text{Node}[\text{node-id}].\text{orig}$ 5: return $(id_c, gid, \text{epoch}, \text{mem})$
--	---

Figure 17: The helper functions define output of process and join protocols.

<p>*next-members(node-id, id_c, \vec{p}, kp_c)</p> <pre> 1: if Node[node-id] ≠ ⊥ ∧ (id_c, *) ∈ Node[node-id].mem ∧ ∀ p ∈ \vec{p}: Prop[p] ≠ ⊥ ∧ Prop[p].par = node-id then 2: $\vec{p}_{\text{upd}'} \parallel \vec{p}_{\text{rem}'} \parallel \vec{p}_{\text{add}'}$ ← *sort-proposals(\vec{p}) 3: mem ← Node[node-id].mem 4: mem ← (id_c, *); mem ← (id_c, kp_c) 5: L ← { id_c } / set of updated parties 6: foreach p ∈ $\vec{p}_{\text{upd}'}$ do 7: (id_s, 'upd'-kp) ← (Prop[p].orig, Prop[p].act) 8: if ¬((id_s, *) ∈ mem ∧ id_s ∉ L) then return (⊥, ⊥) 9: mem ← (id_s, *); mem ← (id_s, kp) 10: L ← id_s 11: foreach p ∈ $\vec{p}_{\text{rem}'}$ do 12: (id_s, 'rem'-id_t) ← (Prop[p].orig, Prop[p].act) 13: if ¬((id_s, *) ∈ mem ∧ (id_t ∈ mem ∧ id_t ∉ L)) then return (⊥, ⊥) 14: mem ← (id_t, *) 15: foreach p ∈ $\vec{p}_{\text{add}'}$ do 16: (id_s, 'add'-id_t-kp_t) ← (Prop[p].orig, Prop[p].act) 17: if ¬((id_s, *) ∈ mem ∧ (id_t, *) ∉ mem) then return (⊥, ⊥) 18: mem ← (id_t, kp_t) 19: P ← ((Prop[PropID[p]].orig, Prop[PropID[p]].act) : p ∈ $\vec{p}_{\text{upd}'}$ ∥ $\vec{p}_{\text{rem}'}$ ∥ $\vec{p}_{\text{add}'}$) 20: return (mem, P) 21: else 22: return (⊥, ⊥) </pre>	<p>*sort-proposals(\vec{p})</p> <pre> 1: $\vec{p}_{\text{upd}'}, \vec{p}_{\text{rem}'}, \vec{p}_{\text{add}'}$ ← () 2: foreach p ∈ \vec{p} do 3: act_p ← Prop[PropID[p]].act 4: if act_p = 'upd'-* then 5: $\vec{p}_{\text{upd}'}$ ← p 6: elseif act_p = 'rem'-* then 7: $\vec{p}_{\text{rem}'}$ ← p 8: if act_p = 'add'-* then 9: $\vec{p}_{\text{add}'}$ ← p 10: return $\vec{p}_{\text{upd}'}$ ∥ $\vec{p}_{\text{rem}'}$ ∥ $\vec{p}_{\text{add}'}$ </pre>
---	--

Figure 18: The helper functions to determine the group state after applying a commit. *sort-proposals(\vec{p}) orders applying proposals.

<p>*consistent-prop(prop-id, id, act)</p> <pre> 1: assert Prop[prop-id].par = Ptr[id] ∧ Prop[prop-id].orig = id 2: ∧ Prop[prop-id].act = act </pre>	<p>*succeed-com(id, \vec{p}, kp)</p> <pre> 1: return ∀ p ∈ \vec{p}: (prop-id := PropID[p] ≠ ⊥ ∧ Prop[prop-id].stat ≠ 'adv') 2: ∧ *next-members(Ptr[id], id, \vec{p}, kp) ≠ (⊥, ⊥) </pre>
<p>*consistent-com(node-id, id, \vec{p}, mem)</p> <pre> 1: *valid-successor(node-id, id, \vec{p}, mem) 2: assert RandCorr[id] = 'bad' ∧ Node[node-id].orig = id </pre>	<p>*succeed-proc(id, c₀, \widehat{c}, \vec{p})</p> <pre> 1: node-id ← NodeID[c₀] 2: index_{id} ← Node[node-id].indexOf(id) 3: return node-id ≠ ⊥ ∧ Node[node-id] ≠ ⊥ ∧ Node[node-id].par = Ptr[id] 4: ∧ Node[node-id].prop = \vec{p} ∧ Node[node-id].stat ≠ 'adv' 5: ∧ ∀ p ∈ \vec{p}: Prop[PropID[p]].stat ≠ 'adv' 6: ∧ Node[node-id].vcom[index_{id}] = \widehat{c} </pre>
<p>*valid-successor(node-id, id, \vec{p}, mem)</p> <pre> 1: assert Node[node-id] ≠ ⊥ ∧ Node[node-id].mem = mem ∧ Node[node-id].prop ∈ { ⊥, \vec{p} } ∧ Node[node-id].par ∈ { ⊥, Ptr[id] } </pre>	<p>*succeed-wel(id, \widehat{w})</p> <pre> 1: node-id ← Wel[id, \widehat{w}] 2: return node-id ≠ ⊥ 3: ∧ Node[node-id] ≠ ⊥ ∧ Node[node-id].stat ≠ 'adv' 4: ∧ (id, *) ∈ (Node[node-id].mem \ Node[Node[node-id].par].mem) </pre>
<p>*attach(node-id, id, \vec{p})</p> <pre> 1: / Cannot attach to the original honest root node-id = 0 2: assert node-id ≠ 0 3: Node[node-id].par ← Ptr[id] 4: Node[node-id].prop ← \vec{p} </pre>	

Figure 19: The helper functions for consistency and correctness.

auth-invariant**return true iff**

- (a) $\forall \text{node-id}$ with $\text{node-id}_p := \text{Node}[\text{node-id}].\text{par}$, $\text{node-id}_p \neq \perp$ and $\text{id} := \text{Node}[\text{node-id}].\text{orig}$:
if $\text{Node}[\text{node-id}].\text{stat} = \text{'adv'}$ **then**
 sig-inj-allowed($\text{node-id}_p, \text{id}$) \wedge **mac-inj-allowed**(node-id_p) and
- (b) $\forall p$ with $\text{node-id}_p := \text{Prop}[p].\text{par}$ and $\text{id} := \text{Prop}[p].\text{orig}$:
if $\text{Prop}[p].\text{stat} = \text{'adv'}$ **then**
 sig-inj-allowed($\text{node-id}_p, \text{id}$) \wedge **mac-inj-allowed**(node-id_p) and
- (c) $\forall \text{node-id}$ with $\text{Node}[\text{node-id}].\text{par} = \perp$ and $\text{id} := \text{Node}[\text{node-id}].\text{orig}$:
 sig-inj-allowed($\text{node-id}, \text{id}$)

cons-invariant**return true iff**

- (a) $\forall \text{node-id}$ s.t. $\text{Node}[\text{node-id}].\text{par} \neq \perp$:
 $\text{Node}[\text{node-id}].\text{prop} \neq \perp \wedge$
 $\forall p \in \text{Node}[\text{node-id}].\text{prop}$:
 $\text{Prop}[\text{PropID}[p]].\text{par} = \text{Node}[\text{node-id}].\text{par}$ and
- (b) $\forall \text{id}$ s.t. $\text{Ptr}[\text{id}] \neq \perp$: $(\text{id}, *) \in \text{Node}[\text{Ptr}[\text{id}]].\text{mem}$ and
- (c) the history graph contains no cycle

Figure 20: The history graph invariants.

C STATIC METADATA-HIDING CGKA: CONSTRUCTION AND SECURITY PROOF

In this section, we first modify Chained CmpKE of Hashimoto et al. [36] into a protocol, which we call Chained CmpKE^{ctxt}, that further secures the static metadata. We then prove that Chained CmpKE^{ctxt} UC-realizes the ideal functionality $\mathcal{F}_{CGKA}^{ctxt}$. This results in the first CGKA that provably secures the 2nd layer.

C.1 Constructing Chained CmpKE^{ctxt}

We provide the description of Chained CmpKE^{ctxt}. The protocol state and the related helper method are shown in Tabs. 5 to 8. The main protocol is depicted in Figs. 22 and 23, and the associated helper functions are depicted in Figs. 25 to 33.

Chained CmpKE^{ctxt} is almost identical to Chained CmpKE [37]. The main differences are explained below and highlighted in yellow in the figures. For a detailed description of the common parts of the protocols, we refer the readers to [37].

- Chained CmpKE^{ctxt} additionally generates an *encryption secret* `encSecret` and a *welcome secret* `welcomeSecret` from the joiner secret.²⁸ `encSecret` is used to encrypt the contents in the proposal and commit messages, excluding the group identifier `gid`, epoch, and the message type, which are necessary for message delivery (cf. `*enc-prop` and `*enc-commit` functions in Fig. 32). `welcomeSecret` is used to encrypt the

group information and the signature in the welcome message (cf. `*enc-welcome` function in Fig. 33).

- Chained CmpKE^{ctxt} creates a separate welcome message for each new member. In contrast, Chained CmpKE included a member-independent welcome message `w0` that is transmitted to every new group member (see Footnote 9). However, this will allow the server to infer that the recipients with the same `w0` will join a (possibly unknown) group.
- In Chained CmpKE^{ctxt}, parties explicitly sort received proposals. In contrast, Chained CmpKE and MLS's TreeKEM implicitly assume that when parties fetch the proposals, the server sorts the proposals according to predetermined rules²⁹. However, since proposals are encrypted in Chained CmpKE^{ctxt}, the server can no longer sort them. Therefore, we make parties sort fetched proposals via `*dec-and-sort-proposals` shown in Fig. 32 before they commit or process the proposals.
- (Optional for hiding the dynamic metadata) Chained CmpKE^{ctxt} additionally generates an *metadata secret* `metaKey` from the joiner secret. This is used when hiding the *dynamic* metadata (see App. E for more detail).

²⁸An encryption secret and welcome secret are also generated in MLS [14, Table 3] used to secure the static metadata.

²⁹MLS stipulates that proposals are to be applied in the order Update, Remove, Add [14, Sec. 13.2.2].

Table 5: The protocol state. The additional component from [37] are highlighted in yellow.

$G.gid$	The identifier of the group.
$G.epoch$	The current epoch number.
$G.confTransHash$	The confirmed transcript hash.
$G.confTransHash-w.o-'id_c'$	The confirmed transcript hash without the committer identity.
$G.interimTransHash$	The interim transcript hash for the next epoch.
$G.member[*]$	A mapping associating party id with its state.
$G.memberHash$	A hash of the public part of $G.member[*]$.
$G.certSvks[*]$	A mapping associating the set of validated signature verification keys to each party.
$G.pendUpd[*]$	A mapping associating the secret keys for each pending update proposal issued by id.
$G.pendCom[*]$	A mapping associating the new group state for each pending commit issued by id.
$G.id$	The identity of the party.
$G.ssk$	The current signing key.
$G.appSecret$	The current epoch's shared key.
$G.membKey$	The key used to MAC proposal packages.
$G.encSecret$	The key used to encrypt proposal and commit messages.
$G.metaKey$	The key used to hide the dynamic metadata in a higher level protocol.
$G.initSecret$	The next epoch's init secret.

Table 6: The party id's state stored in $G.member[id]$ and helper method.

id	The identity of the party.
ek	The encryption key of a mPKE scheme.
dk	The corresponding decryption key.
svk	The signature verification key of a signature scheme.
sig	The signature for (id, ek, svk) under the signature signing key corresponding to svk.
kp()	Returns (id, ek, svk, sig) (if $G.member[id] \neq \perp$).

Table 7: The helper methods on the protocol state. The additional method from [37] are highlighted in yellow.

$G.clone()$	Returns (independent) copy of G .
$G.memberIDs()$	Returns the list of party ids sorted by dictionary order.
$G.memberIDsvks()$	Returns the list of party ids and its associating svk sorted by dictionary order in the ids.
$G.memberPublicInfo()$	Returns the public part of $G.member[*]$.
$G.groupCont()$	Returns ($G.gid, G.epoch, G.memberHash, G.confTransHash$).
$G.indexOf(id)$	Returns the index of id in the sorted member list returned by $G.memberIDs()$.

Table 8: The protocol state maintained only during the proof. The additional component from [37] are highlighted in yellow.

$G.joinerSecret$	The current epoch's joiner secret.
$G.comSecret$	The current epoch's commit secret.
$G.confKey$	The key used to MAC for commit and welcome messages.
$G.welcomeSecret$	The key used to encrypt group information and signature included in welcome messages.
$G.confTag$	The MAC tag included either in the commit or welcome message.
$G.membTags$	The set of MAC tags included in the proposal messages.




Figure 21: Key generation and verification algorithms.

Input (Create, svk)

```

1: req  $G = \perp \wedge \text{id} = \text{id}_{\text{creator}}$ 
2:  $G.\text{gid} \leftarrow \{0, 1\}^K; G.\text{joinerSecret} \leftarrow \{0, 1\}^K$ 
3:  $G.\text{epoch} \leftarrow 0$ 
4:  $G.\text{member}[*] \leftarrow \perp; G.\text{memberHash} \leftarrow \perp$ 
5:  $G.\text{confTransHash-w.o-}'\text{id}_c' \leftarrow \perp$ 
6:  $G.\text{confTransHash} \leftarrow \perp$ 
7:  $G.\text{certSvks}[*] \leftarrow \emptyset$ 
8:  $G.\text{pendUpd}[*] \leftarrow \perp; G.\text{pendCom}[*] \leftarrow \perp$ 
9:  $G.\text{id} \leftarrow \text{id}$ 
10: try  $\text{ssk} \leftarrow \text{*fetch-ssk-if-nec}(G, \text{svk})$ 
11:  $(\text{kp}, \text{dk}) \leftarrow \text{genKP}(\text{id}, \text{svk}, \text{ssk})$ 
12:  $G \leftarrow \text{*assign-kp}(G, \text{id}, \text{kp})$ 
13:  $G.\text{member}[\text{id}].\text{dk} \leftarrow \text{dk}$ 
14:  $G.\text{ssk} \leftarrow \text{ssk}$ 
15:  $G.\text{memberHash} \leftarrow \text{*derive-member-hash}(G)$ 
16:  $(G, \text{confKey}) \leftarrow \text{*derive-epoch-keys}(G, G.\text{joinerSecret})$ 
17:  $\text{confTag} \leftarrow \text{*gen-conf-tag}(G, \text{confKey})$ 
18:  $G \leftarrow \text{*set-interim-trans-hash}(G, \text{confTag})$ 

```

Input (Propose, 'upd'-svk)

```

1: req  $G \neq \perp$ 
2: try  $\text{ssk} \leftarrow \text{*fetch-ssk-if-nec}(G, \text{svk})$ 
3:  $(\text{kp}, \text{dk}) \leftarrow \text{genKP}(\text{id}, \text{svk}, \text{ssk})$ 
4:  $P \leftarrow (\text{'upd'}, \text{kp})$ 
5:  $p \leftarrow \text{*frame-prop}(G, P)$ 
6:  $G.\text{pendUpd}[p] \leftarrow (\text{ssk}, \text{dk})$ 
7:  $p^{\text{ctxt}} \leftarrow \text{*enc-prop}(G.\text{encSecret}, p)$ 
8: return  $p^{\text{ctxt}}$ 

```

Input (Propose, 'add'-id_t-kp_t)

```

1: req  $G \neq \perp \wedge \text{id}_t \notin G.\text{memberIDs}()$ 
2: req  $\text{kp}_t \neq \perp$ 
3: Send  $(\text{has-kp}, \text{id}_t, \text{kp}_t)$  to  $\mathcal{F}_{\text{KS}}$  and receive  $\text{ack}$ 
4: req  $\text{ack}$  /  $\text{ack} = \text{true}$  implies  $\text{*validate-kp}(\text{kp}_t, \text{id}_t) = \text{true}$ 
5:  $P \leftarrow (\text{'add'}, \text{kp}_t)$ 
6:  $p \leftarrow \text{*frame-prop}(G, P)$ 
7:  $p^{\text{ctxt}} \leftarrow \text{*enc-prop}(G.\text{encSecret}, p)$ 
8: return  $p^{\text{ctxt}}$ 

```

Input (Propose, 'rem'-id_t)

```

1: req  $G \neq \perp \wedge \text{id}_t \in G.\text{memberIDs}()$ 
2:  $P \leftarrow (\text{'rem'}, \text{id}_t)$ 
3:  $p \leftarrow \text{*frame-prop}(G, P)$ 
4:  $p^{\text{ctxt}} \leftarrow \text{*enc-prop}(G.\text{encSecret}, p)$ 
5: return  $p^{\text{ctxt}}$ 

```

Input (Commit, \vec{p}^{ctxt} , svk)

```

1: req  $G \neq \perp$ 
2: try  $\vec{p} \leftarrow \text{*dec-and-sort-proposals}(G.\text{encSecret}, \vec{p}^{\text{ctxt}})$ 
3:  $G' \leftarrow \text{*init-epoch}(G)$ 
4: try  $(G', \text{upd}, \text{rem}, \text{add}) \leftarrow \text{*apply-props}(G, G', \vec{p})$ 
5: req  $(*, \text{'rem'-id}) \notin \text{rem} \wedge (\text{id}, *) \notin \text{upd}$ 
6:  $\text{addedMem} \leftarrow \{ \text{id}_t \mid (*, \text{'add'-id}_t) \in \text{add} \}$  / Recipients of the welcome message
7:  $\text{receivers} \leftarrow G'.\text{memberIDs}() \setminus \text{addedMem}$  / Recipients of the new commit secret
8: try  $(G', \text{comSecret}, \text{kp}, \text{ct}_0, \vec{c} = (\vec{c}_{\text{id}})_{\text{id} \in \text{receivers}}) \leftarrow \text{*rekey}(G', \text{receivers}, \text{id}, \text{svk})$ 
9:  $G' \leftarrow \text{*set-member-hash}(G')$ 
10:  $\text{propIDs} \leftarrow ()$ 
11: foreach  $p \in \vec{p}$  do  $\text{propIDs} \leftarrow H(p)$ 
12:  $C_0 \leftarrow (\text{propIDs}, \text{kp}, \text{ct}_0)$ 
13:  $\text{sig} \leftarrow \text{*sign-commit}(G, C_0)$ 
14:  $G' \leftarrow \text{*set-conf-trans-hash}(G, G', \text{id}, C_0, \text{sig})$ 
15:  $(G', \text{confKey}, \text{joinerSecret}) \leftarrow \text{*derive-keys}(G, G', \text{comSecret})$ 
16:  $\text{confTag} \leftarrow \text{*gen-conf-tag}(G', \text{confKey})$ 
17:  $c_0 \leftarrow \text{*frame-commit}(G, C_0, \text{sig}, \text{confTag})$ 
18:  $G' \leftarrow \text{*set-interim-trans-hash}(G', \text{confTag})$ 
19:  $\vec{c} \leftarrow \emptyset$ 
20: foreach  $\text{id} \in G.\text{memberIDs}()$  do
21:   if  $\text{id} \in \text{receivers}$  then  $\vec{c} \leftarrow (\text{id}, \vec{c}_{\text{id}})$ 
22:   else  $\vec{c} \leftarrow (\text{id}, \perp)$ 
23: if  $\text{add} \neq ()$  then
24:    $(G', \text{w}_0, \vec{w}) \leftarrow \text{*welcome-msg}(G', \text{addedMem}, \text{joinerSecret}, \text{confTag})$ 
25: else
26:    $\text{w}_0 \leftarrow \perp; \vec{w} \leftarrow \emptyset$ 
27:  $G.\text{pendCom}[c_0] \leftarrow (G', \vec{p}, \text{upd}, \text{rem}, \text{add})$ 
28: / Encrypt messages
29:  $(c_0^{\text{ctxt}}, \vec{c}^{\text{ctxt}}) \leftarrow \text{*enc-commit}(G.c_0, \vec{c})$ 
30:  $\vec{w}^{\text{ctxt}} \leftarrow \emptyset$ 
31:  $\text{welcomeSecret} \leftarrow \text{HKDF.Expand}(\text{joinerSecret}, \text{'wel'})$ 
32: foreach  $\widehat{w} \in \vec{w}$  then
33:    $\widehat{w}^{\text{ctxt}} \leftarrow \text{*enc-welcome}(\text{welcomeSecret}, \text{w}_0, \widehat{w})$ 
34:    $\vec{w}^{\text{ctxt}} \leftarrow \widehat{w}^{\text{ctxt}}$ 
35: return  $(c_0^{\text{ctxt}}, \vec{c}^{\text{ctxt}}, \vec{w}^{\text{ctxt}})$ 

```

Figure 22: Static metadata-hiding CGKA protocol Chained CmpKE^{ctxt}: Create, Propose, and Commit. The major changes from [37] are highlighted in yellow.

Input (Process, $c_0^{\text{ctxt}}, \widehat{c}^{\text{ctxt}}, \vec{p}^{\text{ctxt}}$)	Input (Join, $\widehat{w}^{\text{ctxt}}$)
<pre> 1: req $G \neq \perp$ 2: try $(c_0, \widehat{c}) \leftarrow *dec-commit(G.encSecret, c_0^{\text{ctxt}}, \widehat{c}^{\text{ctxt}})$ 3: try $\vec{p} \leftarrow *dec-and-sort-proposals(G.encSecret, \vec{p}^{\text{ctxt}})$ 4: $(id_c, C_0, sig, confTag) \leftarrow *unframe-commit(G, c_0)$ 5: if $id_c = id$ then 6: parse $(G', \vec{p}', upd, rem, add) \leftarrow G.pendCom[c_0]$ 7: req $\vec{p} = \vec{p}'$ 8: return $(id_c, upd rem add, G'.memberIDsvks())$ 9: parse $(propIDs, kp_c, ct_0) \leftarrow C_0$ 10: parse $(id', \widehat{ct}_{id'}) \leftarrow \widehat{c}$ 11: req $G.id = id'$ 12: for $i \in 1, \dots, \vec{p}$ do 13: req $H(\vec{p}[i]) = propIDs[i]$ 14: $G' \leftarrow *init-epoch(G)$ 15: try $(G', upd, rem, add) \leftarrow *apply-props(G, G', \vec{p})$ 16: req $(*, id_c) \notin rem \wedge (id_c, *) \notin upd$ 17: if $(*, 'rem'-id) \in rem$ then 18: $G' \leftarrow \perp$ 19: return $(id_c, upd rem add, \perp)$ 20: else 21: $G' \leftarrow *set-conf-trans-hash(G, G', id_c, C_0, sig)$ 22: $(G', comSecret) \leftarrow *apply-rekey(G', id_c, kp_c, ct_0, \widehat{ct}_{id'})$ 23: $G' \leftarrow *set-member-hash(G')$ 24: $(G', confKey, joinerSecret) \leftarrow *derive-keys(G, G', comSecret)$ 25: req $*vrf-conf-tag(G', confKey, confTag)$ 26: $G' \leftarrow *set-interim-trans-hash(G', confTag)$ 27: return $(id_c, upd rem add, G'.memberIDsvks())$ </pre>	<pre> 1: req $G = \perp$ 2: $(w_0, \widehat{w}) \leftarrow *dec-welcome(\widehat{w}^{\text{ctxt}})$ 3: parse $(ct_0, groupInfo, sig) \leftarrow w_0$ 4: parse $(id', kphash, \widehat{ct}_{id'}) \leftarrow \widehat{w}$ 5: req $id = id'$ 6: try $(G, confTag, id_c) \leftarrow *initialize-group(G, id, groupInfo)$ 7: req $G.confTransHash = H(G.confTransHash-w.o-'id_c', id_c)$ 8: req $G.interimTransHash = H(G.confTransHash, confTag)$ 9: req $SIG.Verify(G.member[id_c].svk, sig, (ct_0, \widehat{ct}, groupInfo))$ 10: try $G \leftarrow *vrf-group-state(G)$ 11: $G.id \leftarrow id$ 12: $svk \leftarrow G.member[id].svk$ 13: Send (get-ssk, svk) to \mathcal{F}_{AS} and receive ssk 14: $G.ssk \leftarrow ssk$ 15: Send (get-dks) to \mathcal{F}_{KS} and receive kbs 16: joinerSecret $\leftarrow \perp$ 17: foreach $(kp, dk) \in kbs$ do 18: if $H(kp) = kphash$ then 19: req $G.member[id].kp() = kp$ 20: $G.member[id].dk \leftarrow dk$ 21: joinerSecret $\leftarrow mDec(dk, ct_0, \widehat{ct})$ 22: req joinerSecret $\neq \perp$ 23: $(G, confKey) \leftarrow *derive-epoch-keys(G, joinerSecret)$ 24: req $*vrf-conf-tag(G, confKey, confTag)$ 25: return $(id_c, G.memberIDsvks())$ </pre>

Figure 23: Static metadata-hiding CGKA protocol Chained CmpKE^{ctxt}: Process and Join. The major changes from [37] are highlighted in yellow.

Input (Key)	Input (Key _{mh} , c_0)	Input (NextKey _{mh} , c_0)
<pre> 1: req $G \neq \perp$ 2: $k \leftarrow G.appSecret$ 3: return k </pre>	<pre> 1: req $G \neq \perp$ 2: $k_{mh} \leftarrow G.metaKey$ 3: return k_{mh} </pre>	<pre> 1: / Returns the pending k_{mh}' 2: / for the next epoch 3: parse $(G', *) \leftarrow G.pendCom[c_0]$ 4: $k_{mh}' \leftarrow G'.metaKey$ 5: return k_{mh}' </pre>

Figure 24: Static metadata-hiding CGKA protocol Chained CmpKE^{ctxt}: Retrieve group secret key. The major changes from [37] are highlighted in yellow. Note that Key_{mh} and NextKey_{mh} are used in the higher level metadata-hiding protocol.

*fetch-ssk-if-nec(G, svk')	*validate-kp(G, kp, id)	*assign-kp(G, kp)
<pre> 1: svk ← G.member[G.id].svk 2: if svk ≠ svk' then 3: Send (get-ssk, svk) to \mathcal{F}_{AS} and 4: receive ssk 5: else 6: ssk ← G.ssk 7: return ssk </pre>	<pre> 1: parse (id', ek, svk, sig) ← kp 2: req id = id' 3: if svk ∉ G.certSvks[id] then 4: Send (verify-cert, id', svk) to \mathcal{F}_{AS} and receive succ 5: req succ 6: G.certSvks[id] ← svk 7: req SIG.Verify(pp_{SIG}, svk, sig, (id, ek, svk)) 8: return G </pre>	<pre> 1: parse (id, ek, svk, sig) ← kp 2: G.member[id].ek ← ek 3: G.member[id].svk ← svk 4: G.member[id].sig ← sig 5: return G </pre>

Figure 25: Helper functions: Key material related.

*init-epoch(G)	*apply-props(G, G', \vec{p})
<pre> 1: G' ← G.clone() 2: G'.epoch ← G.epoch + 1 3: G'.pendUpd[*], G'.pendCom[*] ← ⊥ 4: return G' </pre>	<pre> 1: upd, rem, add ← () 2: foreach p ∈ \vec{p} do 3: try (ids, P) ← *unframe-prop(G, p) 4: parse (type, val) ← P 5: if type = 'upd' then 6: req id_s ∈ G.memberIDs() 7: req (id_s, *) ∉ upd ∧ rem = () ∧ add = () 8: try G' ← *validate-kp(G', val, id_s) 9: G' ← *assign-kp(G', val) 10: if id_s = G.id then 11: parse (ssk, dk) ← G.pendUpd[p] 12: G'.ssk ← ssk 13: G'.member[G.id].dk ← dk 14: svk ← G'.member[id_s].svk 15: upd ← (id_s, 'upd'-svk) 16: elseif type = 'rem' then 17: parse id_t ← val 18: req id_t ≠ id_s ∧ id_t ∈ G.memberIDs() 19: req (id_t, *) ∉ upd ∧ add = () 20: G'.member[id_t] ← ⊥ 21: rem ← (id_s, 'rem'-id_t) 22: elseif type = 'add' then 23: (id_t, *, *, *) ← val 24: req id_t ∉ G.memberIDs() 25: try G' ← *validate-kp(G', val, id_t) 26: G' ← *assign-kp(G', val) 27: add ← (id_s, 'add'-val) 28: else 29: return ⊥ 30: return (G', upd, rem, add) </pre>
<pre> *rekey(G', receivers, id, svk) 1: try ssk ← *fetch-ssk-if-nec(G', svk) 2: (kp, dk) ← genKP(id, svk, ssk) 3: G' ← *assign-kp(G', kp) 4: G'.ssk ← ssk 5: G'.member[id].dk ← dk 6: comSecret ← $\{0, 1\}^k$ 7: ek ← (G.member[id'].ek)_{id' ∈ receivers} 8: (ct₀, ct) = (ct_{id'})_{id' ∈ receivers} ← mEnc(pp_{mPKE}, ek, comSecret) 9: return (G', comSecret, kp, ct₀, ct) </pre>	<pre> *apply-rekey(G', id_c, kp_c, ct₀, ct) 1: dk ← G'.member[G'.id].dk 2: comSecret ← mDec(dk, ct₀, ct) 3: try G' ← *validate-kp(G', kp_c, id_c) 4: G' ← *assign-kp(G', kp_c) 5: return (G', comSecret) </pre>

Figure 26: Helper functions: Commit and Process related.

```

*welcome-msg( $G'$ , addedMem, joinerSecret, confTag)
1:  $\vec{ek} \leftarrow (G'.member[id_t].ek)_{id_t \in \text{addedMem}}$  do
2:  $(ct_0, \vec{ct}) = (\widehat{ct}_{id_t})_{id_t \in \text{addedMem}} \leftarrow \text{mEnc}(\text{pp}_{\text{mPKE}}, \vec{ek}, \text{joinerSecret})$ 
3:  $\text{groupInfo} \leftarrow (G'.gid, G'.epoch,$ 
4:    $G'.memberPublicInfo(), G'.memberHash,$ 
5:    $G'.confTransHash-w.o-'id_c', G'.confTransHash,$ 
6:    $G'.interimTransHash, \text{confTag}, G'.id)$ 
7:  $\text{sig} \leftarrow \text{SIG.Sign}(\text{pp}_{\text{SIG}}, G'.\text{ssk}, (ct_0, \text{groupInfo}))$ 
8:  $w_0 \leftarrow (ct_0, \text{groupInfo}, \text{sig})$ 
9:  $\vec{w} \leftarrow \emptyset$ 
10: foreach  $id_t \in \text{addedMem}$  do
11:    $\text{kphash}_t \leftarrow H(G'.member[id_t].\text{kp}())$ 
12:    $\vec{w} \leftarrow \widehat{w}_{id_t} = (id_t, \text{kphash}_t, \widehat{ct}_{id_t})$ 
13: return  $(G', w_0, \vec{w})$ 

```

```

*initialize-group( $G, id, \text{groupInfo}$ )
1: parse  $(gid, epoch, member, memberHash,$ 
    $\text{confTransHash-w.o-'id_c'}, \text{confTransHash},$ 
    $\text{interimTransHash}, \text{confTag}, id_c) \leftarrow \text{groupInfo}$ 
2:  $(G.gid, G.epoch, G.member, G.memberHash,$ 
    $G.\text{confTransHash-w.o-'id_c'}, G.\text{confTransHash},$ 
    $G.\text{interimTransHash}) \leftarrow (gid, epoch, member,$ 
    $memberHash, \text{confTransHash-w.o-'id_c'},$ 
    $\text{confTransHash}, \text{interimTransHash})$ 
3:  $G.\text{certSvks}[*] \leftarrow \emptyset$ 
4:  $G.\text{pendUpd}[*] \leftarrow \perp; G.\text{pendCom}[*] \leftarrow \perp$ 
5:  $G.id \leftarrow id$ 
6: return  $(G, \text{confTag}, id_c)$ 

```

```

*vrf-group-state( $G$ )
1: req  $G.\text{memberHash} = \text{*derive-member-hash}(G)$ 
2:  $\text{mem} \leftarrow G.\text{memberIDs}()$ 
3: foreach  $id \in \text{mem}$  do
4:    $\text{kp} \leftarrow G.\text{member}[id].\text{kp}()$ 
5:   try  $G \leftarrow \text{*validate-kp}(G, \text{kp}, id)$ 
6: return  $G$ 

```

Figure 27: Helper functions: Join related.

```

*gen-conf-tag( $G, \text{confKey}$ )
1: return  $\text{MAC.TagGen}(\text{confKey}, G.\text{confTransHash})$ 

*vrf-conf-tag( $G, \text{confKey}, \text{confTag}$ )
1: return  $\text{MAC.TagVerify}(\text{confKey}, \text{confTag}, G.\text{confTransHash})$ 

```

Figure 28: Helper function: Confirmation tag.

```

*set-member-hash( $G$ )
1:  $G.\text{memberHash} \leftarrow \text{*derive-member-hash}(G)$ 
2: return  $G$ 

*derive-member-hash( $G$ )
1:  $\text{KP} \leftarrow (); \text{mem} \leftarrow G.\text{memberIDs}()$  / mem is sorted by dictionary order
2: foreach  $id \in \text{mem}$  do
3:    $\text{KP} \leftarrow G.\text{member}[id].\text{kp}()$ 
4: return  $H(\text{KP})$ 

*set-conf-trans-hash( $G, G', id_c, C_0, \text{sig}$ )
1:  $\text{comCont} \leftarrow (G.gid, G.epoch, 'commit', C_0, \text{sig})$ 
2:  $G'.\text{confTransHash-w.o-'id_c'} \leftarrow H(G.\text{interimTransHash}, \text{comCont})$ 
3:  $G'.\text{confTransHash} \leftarrow H(G'.\text{confTransHash-w.o-'id_c'}, id_c)$ 
4: return  $G'$ 

*set-interim-trans-hash( $G', \text{confTag}$ )
1:  $G'.\text{interimTransHash} \leftarrow H(G'.\text{confTransHash}, \text{confTag})$ 
2: return  $G'$ 

```

Figure 29: Helper function: Member hash and transcript hash.

```

*derive-keys( $G, G', \text{comSecret}$ )
1:  $s \leftarrow \text{HKDF.Extract}(G.\text{initSecret}, \text{comSecret})$ 
2:  $\text{joinerSecret} \leftarrow \text{HKDF.Expand}(s, 'joi')$ 
3:  $(G', \text{confKey}) \leftarrow \text{*derive-epoch-keys}(G', \text{joinerSecret})$ 
4: return  $(G', \text{confKey}, \text{joinerSecret})$ 

*derive-epoch-keys( $G', \text{joinerSecret}$ )
1:  $\text{confKey} \leftarrow \text{HKDF.Expand}(\text{joinerSecret}, G'.\text{groupCont}() \parallel \text{'conf'})$ 
2:  $G'.\text{appSecret} \leftarrow \text{HKDF.Expand}(\text{joinerSecret}, G'.\text{groupCont}() \parallel \text{'app'})$ 
3:  $G'.\text{membKey} \leftarrow \text{HKDF.Expand}(\text{joinerSecret}, G'.\text{groupCont}() \parallel \text{'memb'})$ 
4:  $G'.\text{encSecret} \leftarrow \text{HKDF.Expand}(\text{joinerSecret}, G'.\text{groupCont}() \parallel \text{'enc'})$ 
5:  $G'.\text{metaKey} \leftarrow \text{HKDF.Expand}(\text{joinerSecret}, G'.\text{groupCont}() \parallel \text{'meta'})$ 
6:  $G'.\text{initSecret} \leftarrow \text{HKDF.Expand}(\text{joinerSecret}, G'.\text{groupCont}() \parallel \text{'init'})$ 
7: return  $(G', \text{confKey})$ 

```

Figure 30: Helper function: Key scheduling. The major changes from [37] are highlighted in yellow.

```

*frame-prop( $G, P$ )
1: propCont  $\leftarrow (G.groupCont(), G.id, 'proposal', P)$ 
2: sig  $\leftarrow SIG.Sign(pp_{SIG}, G.ssk, propCont)$ 
3: membTag  $\leftarrow MAC.TagGen(G.membKey, (propCont, sig))$ 
4:  $\bar{p} \leftarrow (G.id, P, sig, membTag)$ 
5: return ( $G.gid, G.epoch, 'proposal', \bar{p}$ )

*unframe-prop( $G, p$ )
1: parse ( $gid, epoch, contType, \bar{p}$ )  $\leftarrow p$ 
2: parse ( $ids, P, sig, membTag$ )  $\leftarrow \bar{p}$ 
3: req contType = 'proposal'  $\wedge$  gid =  $G.gid \wedge$  epoch =  $G.epoch$ 
4: propCont  $\leftarrow (G.groupCont(), ids, 'proposal', P)$ 
5: req  $G.member[ids] \neq \perp$ 
    $\wedge SIG.Verify(pp_{SIG}, G.member[ids].svk, sig, propCont)$ 
    $\wedge MAC.TagVerify(G.membKey, membTag, (propCont, sig))$ 
6: return ( $ids, P$ )

*sign-commit( $G, C_0$ )
1: comCont  $\leftarrow (G.groupCont(), G.id, 'commit', C_0)$ 
2: sig  $\leftarrow SIG.Sign(pp_{SIG}, G.ssk, comCont)$ 
3: return sig

*frame-commit( $G, C_0, sig, confTag$ )
1:  $\bar{c}_0 \leftarrow (G.id, C_0, sig, confTag)$ 
2: return ( $G.gid, G.epoch, 'commit', \bar{c}_0$ )

*unframe-commit( $G, c_0$ )
1: parse ( $G.gid, G.epoch, 'commit', CT_{\bar{c}_0}$ )  $\leftarrow c_0$ 
2: parse ( $id_c, C_0, sig, confTag$ )  $\leftarrow \bar{c}_0$ 
3: req contType = 'commit'  $\wedge$  gid =  $G.gid \wedge$  epoch =  $G.epoch$ 
4: comCont  $\leftarrow (G.groupCont(), id_c, 'commit', C_0)$ 
5: svk_c  $\leftarrow G.member[id_c].svk$ 
6: req  $G.member[id_c] \neq \perp \wedge SIG.Verify(pp_{SIG}, svk_c, sig, comCont)$ 
7: return ( $id_c, C_0, sig, confTag$ )

```

Figure 31: Helper function: Frame and unframe packets.

```

*enc-prop(encSecret,  $p$ )
1: parse ( $gid, epoch, 'proposal', \bar{p}$ )  $\leftarrow p$ 
2:  $CT_{\bar{p}} \leftarrow SKE.Enc(encSecret, \bar{p})$ 
3: return  $p^{ctxt} := (gid, epoch, 'proposal', CT_{\bar{p}})$ 

*dec-prop(encSecret,  $p^{ctxt}$ )
1: parse ( $gid, epoch, 'proposal', CT_{\bar{p}}$ )  $\leftarrow p^{ctxt}$ 
2:  $\bar{p} \leftarrow SKE.Dec(encSecret, CT_{\bar{p}})$ 
3: return  $p := (gid, epoch, 'proposal', \bar{p})$ 

*enc-commit(encSecret,  $c_0, \bar{c}$ )
1: parse ( $gid, epoch, 'commit', \bar{c}_0$ )  $\leftarrow c_0$ 
2:  $CT_{\bar{c}_0} \leftarrow SKE.Enc(encSecret, \bar{c}_0)$ 
3:  $\hat{c}^{ctxt} \leftarrow \emptyset$ 
4: foreach  $\hat{c} \in \bar{c}$  do
5:    $\hat{c}^{ctxt} \leftarrow SKE.Enc(encSecret, \hat{c})$ 
6:  $c_0^{ctxt} \leftarrow (G.gid, G.epoch, 'commit', CT_{\bar{c}_0})$ 
7: return ( $c_0^{ctxt}, \hat{c}^{ctxt}$ )

*dec-commit(encSecret,  $c_0^{ctxt}, \hat{c}^{ctxt}$ )
1: parse ( $gid, epoch, 'commit', CT_{\bar{c}_0}$ )  $\leftarrow c_0^{ctxt}$ 
2:  $\bar{c}_0 \leftarrow SKE.Dec(encSecret, CT_{\bar{c}_0})$ 
3:  $\hat{c} \leftarrow SKE.Dec(encSecret, \hat{c}^{ctxt})$ 
4:  $c_0 \leftarrow (gid, epoch, 'commit', \bar{c}_0)$ 
5: return ( $c_0, \hat{c}$ )

*dec-and-sort-proposals(encSecret,  $\vec{p}^{ctxt}$ )
1:  $\vec{p}, \vec{p}_{rem}, \vec{p}_{upd}, \vec{p}_{add} \leftarrow ()$ 
2: foreach  $p^{ctxt} \in \vec{p}^{ctxt}$  do
3:   try  $p \leftarrow *dec-prop(encSecret, p^{ctxt})$ 
4:   try type  $\leftarrow *extract-proposal-type(p)$ 
5:   if type = 'upd' then
6:      $\vec{p}_{upd} \leftarrow p$ 
7:   elseif type = 'rem' then
8:      $\vec{p}_{rem} \leftarrow p$ 
9:   elseif type = 'add' then
10:     $\vec{p}_{add} \leftarrow p$ 
11:   else return  $\perp$ 
12: / Return sorted proposal list
13: return  $\vec{p}_{upd} \parallel \vec{p}_{rem} \parallel \vec{p}_{add}$ 

*extract-proposal-type( $p$ )
1: parse ( $gid, epoch, contType, \bar{p}$ )  $\leftarrow p$ 
2: parse ( $ids, P, sig, membTag$ )  $\leftarrow \bar{p}$ 
3: parse (type, val)  $\leftarrow P$ 
4: return type

```

Figure 32: Helper functions unique to Chained CmPKE^{ctxt}: Encrypt and decrypt proposals and commits. The major changes from [37] are highlighted in yellow. *sort-proposals decrypts proposals and sort them following the order of MLS specification [14, Sec. 13.2.2].

<pre> *enc-welcome(welcomeSecret, w₀, ŵ) 1: parse (id_ℓ, kphash_ℓ, ct̂) ← w₀ 2: parse (ct₀, groupInfo, sig) ← ŵ 3: welcomeSecret_{id_ℓ} ← HKDF.Expand(welcomeSecret, id_ℓ) 4: CT ← SKE.Enc(welcomeSecret_{id_ℓ}, (groupInfo, sig)) 5: return ŵ^{ctxt} := (id_ℓ, (kphash_ℓ, ct₀, ct̂), CT) </pre>	<pre> *dec-welcome(ŵ^{ctxt}) 1: parse (id_ℓ, (kphash, ct₀, ct̂), CT) ← ŵ^{ctxt} 2: Send (get-dks) to F_{KS} and receive kbs 3: joinerSecret, kp_{id}, dk_{id} ← ⊥ 4: foreach (kp, dk) ∈ kbs do 5: if H(kp) = kphash then 6: (kp_{id}, dk_{id}) ← (kp, dk) 7: joinerSecret ← mDec(dk, ct₀, ct̂) 8: break 9: req joinerSecret ≠ ⊥ 10: welcomeSecret ← HKDF.Expand(joinerSecret, 'wel') 11: welcomeSecret_{id_ℓ} ← HKDF.Expand(welcomeSecret, id_ℓ) 12: (groupInfo, sig) ← SKE.Dec(welcomeSecret_{id_ℓ}, CT) 13: w₀ ← (ct₀, groupInfo, sig) 14: ŵ ← (id_ℓ, kphash_ℓ, ct̂) 15: return (w₀, ŵ) </pre>
--	--

Figure 33: Helper functions unique to Chained CmpKE^{ctxt}: Encrypt and decrypt welcome messages. The major changes from [37] are highlighted in yellow.

C.2 Safety Predicates and Leakage Functions

As explained in App. B.2, to formally prove that Chained CmpKE^{ctxt} UC-realizes the ideal functionality $\mathcal{F}_{CGKA}^{\text{ctxt}}$, we must first specify the safety predicates **safe**, **mac-inj-allowed**, and **sig-inj-allowed**, and the leakage functions ***leak-create**, ***leak-prop**, ***leak-com**, ***leak-wel**, and ***leak-proc**.

These are formally defined in Figs. 34 and 35. The definition of the safety predicates **safe**, **mac-inj-allowed** and **sig-inj-allowed** is identical to those considered by Chained CmpKE [37, Fig. 28], modulo the syntactical difference in the definition of node pointers (see Sec. 3.2).

The leakage functions ***leak-create**, ***leak-prop**, ***leak-com**, ***leak-wel**, and ***leak-proc** are new to the definition of $\mathcal{F}_{CGKA}^{\text{ctxt}}$ and dictate the amount of static metadata leaked to the server. Since the concrete leakage information depends on the concrete choice of CGKA, our choice only captures the leakage of Chained CmpKE^{ctxt}. Namely, other protocols such as the ciphertext variant of TreeKEM used in MLSCiphertext may require a slightly more complex leakage function. We provide some discussion in Sec. 7. Below, we explain the specific choice of our leakage functions in more detail.

***leak-create**(id, svk): This defines the information leaked when a new group is created. In the previous (non-metadata-hiding) Chained CmpKE, the simulator was given the identity of the party id creating the group and its corresponding signature verification key svk to simulate the group identifier gid . In other words, gid leaked information on (id, svk) . In Chained CmpKE^{ctxt}, the simulator is asked to simulate gid without giving any other information. This models the fact that when a new group is created, the server cannot guess who created it. As mentioned in the introduction, this could potentially open the door to a DoS attack on the server. One way to solve this would be to use standard anonymous credentials [26]. We leave such direction of research as an important future work.

***leak-prop**(id, act): This defines the information that is leaked from proposals p . In Chained CmpKE, the simulator was given the identity of the party id creating the proposal, the commit node $Ptr[id]$ that locates id in the history graph, and the action act of the proposal to simulate the proposal p . Here, when there is no fork in the main group, then $Ptr[id]$ essentially pinpoints the specific epoch of the group. In case there is a fork, then epoch alone is not sufficient to identify where party id is located, in which case we need the exact proposal node $Ptr[id]$. In other words, proposal p leaked the proposing party id and the action act , where note that p must always include $Ptr[id]$ (which is $(gid, epoch)$ essentially) for the server to deliver p to the appropriate member in a specific epoch.

In Chained CmpKE^{ctxt}, we define ***leak-prop**(id, act) to only output $Ptr[id]$ and the length of $|id|$ and $|act|$. This models the fact that proposals p only leak the specific group by which p is supposed to be processed. Moreover, in case the size of each action act is different, this leaks the action type. However, considering that the addition or removal of a group member is noticeable from the server (since it changes the size of commit messages), hiding the type of action may not add as much security as one expects.

***leak-com**(id, \vec{p}, svk): This defines the information that is leaked from commits $(c_0, \vec{c} = (\widehat{c}_{i,d'})_{i,d'})$. In Chained CmpKE, the simulator was given the identity of the party id creating the commit; the commit node $Ptr[id]$ that locates id in the history graph; the list of proposals \vec{p} that is going to be committed; the new signing key svk of id ; and the list of current members mem in the group to simulate the commit (c_0, \vec{c}) . In other words, a commit leaks the committer identity id , its signing key svk , and the current list of members mem .

In Chained CmpKE^{ctxt}, we define ***leak-com**(id, \vec{p}, svk) to only output $(Ptr[id], |id|, \vec{p}, |svk|, |mem|)$. Note that similarly to proposals, $Ptr[id]$ cannot be hidden. Moreover, \vec{p} is guaranteed to hide the static metadata from above, we can provide this to the simulator. Finally, the size of the current member $|mem|$ would necessarily leak from commit messages when using Chained CmpKE since \vec{c} scales linearly with the group size. While we could apply some padding to hide the group size from the commit, the server can infer the group size when selective downloading is performed; if N distinct \widehat{c} were downloaded out of $M \geq N$ commits (that includes the padded garbage commits), then the server can guess that the group size was N .

***leak-wel**($node-id_{cur}, node-id_{next}, id_t$): This defines the information that is leaked from welcome messages \widehat{w} . In Chained CmpKE, the simulator was given the key package kp_t of the added member id_t ; the commit node of the next epoch $node-id_{next}$; and all the information stored on the commit node $Node[node-id_{next}]$ to simulate the welcome message \widehat{w} . In other words, welcome messages leak the group identifier gid , the next epoch, the identity id of the party who created the welcome message, the size of the member of the next epoch, and so on.

In Chained CmpKE^{ctxt}, we define ***leak-wel**($node-id_{cur}, node-id_{next}, id_t$) to only output $(kp_t, |gid|, |epoch|, |id|, |mem|)$. Here, we include kp_t in clear since Chained CmpKE^{ctxt} (and MLSCiphertext) includes the hash of the key package of id_t in the welcome message in the clear. Namely, in case id_t creates many key packages, the simulator must know which kp_t is used to simulate the welcome message.

We note that the reason why our protocol and MLSCiphertext includes a hash $kphash_t$ of kp_t in the welcome message is for efficiency. The added party, which may have created many key packages for different groups, can check which key package to use by simply finding the key package that equals the hash value $kphash_t$. Otherwise, the party must go through all the decryption keys associated with each key package it has and try to see whether the welcome message can be decrypted correctly. We also note that id_t is always leaked from the welcome message since otherwise, the server will not know the destination of the welcome message, i.e., it cannot deliver it.

***leak-proc**(id): This defines the information that is leaked from a party-dependent commit (c_0, \vec{c}) and a list of proposals \vec{p} . In Chained CmpKE, the simulator was given the identity of the party id processing the commit and the commit node $Ptr[id]$ that locates id in the history graph to decide whether (c_0, \vec{c}, \vec{p}) should be processed correctly by a party id .

In Chained CmpKE^{ctxt}, we define ***leak-proc**(id) to only output $(Ptr[id], index_{id})$. As explained above, $Ptr[id]$ cannot be hidden

since this information is required for id to retrieve the uploaded commit and proposals from the server. Moreover, since selective downloading is performed, \widehat{c} must leak the position of the party id in the group (assuming that id can process (c_0, \widehat{c}) correctly).

Knowledge of party's secrets.

know(node-id, id) \iff

(a) $id \in \text{Node}[\text{node-id}].\text{exp} \vee$
(b) $\text{*secrets-injected}(\text{node-id}, id) \vee$
(c) $(\text{Node}[\text{node-id}].\text{par} \neq \perp \wedge \text{know}(\text{Node}[\text{node-id}].\text{par}, id))$
 $\wedge \neg \text{*secrets-replaced}(\text{node-id}, id) \vee$
(d) $\exists \text{node-id}' : (\text{Node}[\text{node-id}'].\text{par} = \text{node-id} \wedge \text{know}(\text{node-id}', id))$
 $\wedge \neg \text{*secrets-replaced}(\text{node-id}', id)$

***secrets-injected**(node-id, id) \iff

(a) $(\text{Node}[\text{node-id}].\text{orig} = id \wedge \text{Node}[\text{node-id}].\text{stat} \neq \text{'good'}) \vee$
(b) $\exists p \in \text{Node}[\text{node-id}].\text{prop}$ with $\text{prop-id} := \text{PropID}[p] :$
 $(\text{Prop}[\text{prop-id}].\text{act} = \text{'upd'-*} \wedge \text{Prop}[\text{prop-id}].\text{orig} = id$
 $\wedge \text{Prop}[\text{prop-id}].\text{stat} \neq \text{'good'}) \vee$
(c) $\exists p \in \text{Node}[\text{node-id}].\text{prop}$ with $\text{prop-id} := \text{PropID}[p] :$
 $(\text{Prop}[\text{prop-id}].\text{act} = \text{'add'-id-*} \wedge \text{svk} \in \text{ExposedSvk})$

***secrets-replaced**(node-id, id) \iff
 $\text{Node}[\text{node-id}].\text{orig} = id \vee$
 $\exists p \in \text{Node}[\text{node-id}].\text{prop}$ with $\text{prop-id} := \text{PropID}[p] :$
 $\text{Prop}[\text{prop-id}].\text{act} \in \{ \text{'add'-id-*}, \text{'rem'-id} \} \vee$
 $\exists p \in \text{Node}[\text{node-id}].\text{prop}$ with $\text{prop-id} := \text{PropID}[p] :$
 $(\text{Prop}[\text{prop-id}].\text{act} = \text{'upd'-*} \wedge \text{Prop}[\text{prop-id}].\text{orig} = id)$

Knowledge of epoch secrets.

know(node-id, 'epoch') $\iff \text{Node}[\text{node-id}].\text{exp} \neq \emptyset \vee \text{*can-traverse}(\text{node-id})$

***can-traverse**(node-id) \iff

(a) $\exists p \in \text{Node}[\text{node-id}].\text{prop} :$
 $(\text{Prop}[\text{PropID}[p]].\text{act} = \text{'add'-id-*} \wedge \text{svk} \in \text{ExposedSvk})$
(b) $\text{*reused-welcome-key-leaks}(\text{node-id}) \vee$
(c) $\text{Node}[\text{node-id}].\text{stat} = \text{'bad'}$
 $\wedge \exists p \in \text{Node}[\text{node-id}].\text{prop} : \text{Prop}[\text{PropID}[p]].\text{act} = \text{'add'-*} \vee$
(d) $(\text{Node}[\text{node-id}].\text{par} = \perp \vee \text{know}(\text{Node}[\text{node-id}].\text{par}, \text{'epoch'})) \wedge$
 $\exists (id, *) \in \text{Node}[\text{node-id}].\text{mem} : \text{know}(\text{node-id}, id)$

***reused-welcome-key-leaks**(node-id) \iff
 $\exists id, p \in \text{Node}[\text{node-id}].\text{prop} : \text{Prop}[p].\text{act} = \text{'add'-id-*} \wedge$
 $\exists \text{node-id}_d : \text{node-id}_d \text{ is a descendant of node-id} \wedge id \in \text{Node}[\text{node-id}_d].\text{exp} \wedge$
 $\text{no node node-id}_h \text{ exists on node-id-node-id}_d \text{ path}$
 $\text{s.t. } \text{*secrets-replaced}(\text{node-id}_h, id) = \text{true}$

Safe and can-inject.

safe(node-id) $\iff \text{know}(\text{node-id}, \text{'epoch'})$

sig-inj-allowed(node-id, id) $\iff \text{Node}[\text{node-id}].\text{mem}[id] \in \text{ExposedSvk}$

mac-inj-allowed(node-id) $\iff \text{know}(\text{node-id}, \text{'epoch'})$

Figure 34: The safety predicate for Chained CmpKE^{ctxt}. This is identical to [37, Fig. 28] modulo the syntactical difference in the definition of node pointers.

<hr/> <p>*leak-create(id, svk)</p> <p>1: if $\text{flag}_{\text{contHide}}$ then</p> <p>2: return \perp</p> <p>3: else</p> <p>4: return (id, svk)</p> <hr/> <p>*leak-prop(id, act)</p> <p>1: if $\text{flag}_{\text{contHide}} \wedge \text{safe}(\text{Ptr}[id])$ then</p> <p>2: return (Ptr[id], id , act)</p> <p>3: else</p> <p>4: return (Ptr[id], id, act)</p>	<hr/> <p>*leak-com(id, \vec{p}, svk)</p> <p>1: $\text{mem} \leftarrow \text{Node}[\text{Ptr}[id]].\text{mem}$</p> <p>2: if $\text{flag}_{\text{contHide}} \wedge \text{safe}(\text{Ptr}[id])$ then</p> <p>3: return (Ptr[id], id , \vec{p}, svk , mem)</p> <p>4: else</p> <p>5: return (Ptr[id], id, \vec{p}, svk, mem)</p> <hr/> <p>*leak-proc(id)</p> <p>1: $\text{index}_{id} \leftarrow \text{Node}[\text{Ptr}[id]].\text{indexOf}(id)$</p> <p>2: if $\text{flag}_{\text{contHide}} \wedge \text{safe}(\text{Ptr}[id])$ then</p> <p>3: return (Ptr[id], index_{id})</p> <p>4: else</p> <p>5: return (Ptr[id], id)</p>	<hr/> <p>*leak-wel(node-id_{cur}, node-id_{next}, id_t)</p> <p>1: $\text{gid} \leftarrow \text{Node}[\text{node-id}_{\text{next}}].\text{gid}$</p> <p>2: $\text{epoch} \leftarrow \text{Node}[\text{node-id}_{\text{next}}].\text{epoch}$</p> <p>3: $\text{id}_c \leftarrow \text{Node}[\text{node-id}_{\text{next}}].\text{orig}$</p> <p>4: $\text{mem} \leftarrow \text{Node}[\text{node-id}_{\text{next}}].\text{mem}$</p> <p>5: if $\text{flag}_{\text{contHide}} \wedge \text{safe}(\text{node-id}_{\text{next}})$ then</p> <p>6: return (kp_t, gid , epoch , id_c , mem)</p> <p>7: else</p> <p>8: return (kp_t, node-id_{next}, Node[node-id_{next}])</p>
---	--	---

Figure 35: Leakage functions for Chained CmpKE^{ctxt}.

C.3 Security of Chained CmpKE^{ctxt}

The following theorem establishes that Chained CmpKE^{ctxt} UC-realizes the ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$.

THEOREM C.1. *Assuming that mPKE is IND-CCA secure, SIG is sEUF-CMA secure, and SKE is IND-CCA secure and has key-committing property, our CGKA protocol Chained CmpKE^{ctxt} adaptively and securely realizes the ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ in the $(\mathcal{F}_{\text{AS}}, \mathcal{F}_{\text{KS}}, \mathcal{G}_{\text{RO}})$ -hybrid model.*

Here, $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ is defined with respect to the safety predicates and leakage functions in Figs. 34 and 35. Moreover, calls to the hash function H , HKDF, and MAC are replaced by calls to the global random oracle \mathcal{G}_{RO} .

C.3.1 Proof Overview. The proof consists of a sequence of hybrids where we gradually modify the protocol Chained CmpKE^{ctxt} into the ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. As explained in Sec. 3.2, with some work, we can reuse a great part of the proof by Hashimoto et al. [36].

The key observation is that when the leakage functions ***leak-create**, ***leak-prop**, ***leak-com**, ***leak-wel** and ***leak-proc** are defined to leak all the static metadata (i.e., $\text{flag}_{\text{contHide}} = \text{false}$), then $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ is almost identical to the ideal functionality $\mathcal{F}_{\text{CGKA}}$ that does not hide the static metadata. The only difference is that the proposals and commits in $\mathcal{F}_{\text{CGKA}}$ leak the static metadata, while $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ encrypts them. To this end, using the fact that the simulator \mathcal{S} knows the encryption key (which follows from $\text{flag}_{\text{contHide}} = \text{false}$), the simulator \mathcal{S} for $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ will internally run the simulator $\tilde{\mathcal{S}}$ for $\mathcal{F}_{\text{CGKA}}$ defined by Hashimoto et al. [36].

In a bit more detail, our proof goes through the same Hybrids 1 to 6 as in [36, Theorem E.1]. At Hybrid 6, the application keys for the nodes with $\text{safe} = \text{true}$ will become indistinguishable from random — this is the place where [36, Theorem E.1] ends. We then continue to complete the proof by further adding a new Hybrid 7, where $\text{flag}_{\text{contHide}}$ is switched to true . This hybrid is part of our security proof that takes care of the static metadata. In each Hybrid i for $i \in [6]$, we define the simulator \mathcal{S}_i that internally runs $\tilde{\mathcal{S}}_i$ defined in [36, Theorem E.1]. Informally, whenever $\tilde{\mathcal{S}}_i$ takes as input a (non-static metadata-hiding) proposal, commit, or a welcome message, \mathcal{S}_i first decrypts the (static metadata-hiding) proposal, commit, and welcome message before feeding it into $\tilde{\mathcal{S}}_i$. On the other hand, if $\tilde{\mathcal{S}}_i$ outputs a (non-static metadata-hiding) content, then \mathcal{S}_i encrypts them before outputting them to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ or the environment \mathcal{Z} . Here, since $\text{flag}_{\text{contHide}} = \text{false}$, we can assume \mathcal{S}_i knows all the encryption/decryption key to perform the above procedure.

While the high-level idea of piggybacking on the prior proof of Hashimoto et al. [36] sounds straightforward, several technical issues make the above non-trivial. This non-triviality is caused by the difference of the *semantics* of the nodes in the history graph in our new ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ and the prior ideal functionality $\mathcal{F}_{\text{CGKA}}$. For instance, when the environment \mathcal{Z} invokes id on input $(\text{Commit}, \vec{p}, \text{svk})$, the simulator in [36] (roughly) outputs $(\text{ack}, \text{rt}, \text{e}_0, \vec{e})$, where $\text{rt} \in \mathbb{N}$. Our simulator must reinterpret this and output an appropriate node-id such that the history graph defined with respect to e_0 remains consistent with the history graph defined with respect to node-id. In particular, the most non-trivial part of reusing the proof by Hashimoto et al. [36] is to check that

the proof moves from Hybrid 2 to 3 in [36, Theorem E.1], which concerns the consistency of the history graph, translates to our setting.

We also note that while the proof of Hashimoto et al. [36] required an mPKE that is IND-CCA secure *with adaptive corruptions*, we only require a standard mPKE. As explained in Sec. 3.2, this is due to the added restriction on the adversary (see Fig. 14 and App. B.2), which is necessary for any natural static metadata-hiding CGKA in order not to trivially win the security game. Specifically, the exponential loss appearing in the proof of moving from Hybrid 5 to Hybrid 6 in [36] disappears by considering the restricted adversary defined in $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. Moreover, while Hashimoto et al. [36] used a *committing* mPKE to construct Chained CmpKE, they mentioned in Footnote 5 that they are not required for the proof to go through. In this work, we thus rely on the more simpler mPKE.

Finally, we prove the indistinguishability of the real and ideal protocols assuming that the adversary *cannot* inject bad randomness. That is, RandCorr is always set to ‘good’ in the ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. To be more precise, our proof up till Hybrid 6 allows the adversary to inject bad randomness. We only restrict the adversary when moving to Hybrid 7 — this is where we explicitly use the group secret to perform symmetric key encryption to hide the static metadata. This restriction is also done in recent work by Alwen et al. [9], where they disallow the adversary to manipulate the randomness used for the symmetric key encryption. We leave security analysis of such types of attacks as future work.

C.3.2 Proof of Thm. C.1. We now provide the full proof of Thm. C.1.

PROOF. We consider the following sequence of hybrids. While the environment \mathcal{Z} interacts with Chained CmpKE^{ctxt} in Hybrid 1, it interacts with the ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ in Hybrid 7. As explained above, Hybrids 1 to 6 are defined identically to [36], where the only difference is in the definition of the simulator \mathcal{S} in each hybrid. Below, we first define all the hybrids and then explain how the simulators are defined.

Hybrid 1. This is the real world, where we make a syntactic change.

We consider a simulator \mathcal{S}_1 that interacts with a dummy functionality $\mathcal{F}_{\text{dummy}}$. $\mathcal{F}_{\text{dummy}}$ sits between the environment \mathcal{Z} and \mathcal{S}_1 , and relays any message from \mathcal{Z} to \mathcal{S}_1 without modifying them. \mathcal{S}_1 internally simulates the real-world parties and adversary \mathcal{A} by using the messages sent from $\mathcal{F}_{\text{dummy}}$. From \mathcal{A} ’s point of view, \mathcal{S}_1 is the environment \mathcal{Z} .

Hybrid 2. This hybrid concerns the authentication service and key service. We replace $(\mathcal{F}_{\text{AS}}, \mathcal{F}_{\text{KS}})$ with these ideal version $(\mathcal{F}_{\text{AS}}^{\text{IW}}, \mathcal{F}_{\text{KS}}^{\text{IW}})$. Since these functions are not accessible by \mathcal{Z} , this modification is undetectable by \mathcal{Z} . Thus, the view of \mathcal{Z} in Hybrid 1 and Hybrid 2 are identical.

Hybrid 3. This hybrid concerns the correctness and consistency guarantees. We replace $\mathcal{F}_{\text{dummy}}$ with a variant of $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$, denoted as $\mathcal{F}_{\text{CGKA},3}^{\text{ctxt}}$, where **safe** (resp. **sig-inj-allowed** and **mac-inj-allowed**) always returns false (resp. true) and $\text{flag}_{\text{contHide}}$ is set to false. In other words, all application secrets are set by the simulator, injections are always allowed, and the confidentiality of messages is not considered. The simulator \mathcal{S}_3 is identical to \mathcal{S}_2 .

Hybrid 4. This hybrid concerns the security of the signature scheme. We modify $\mathcal{F}_{\text{CGKA},3}^{\text{ctxt}}$ to use the original **sig-inj-allowed** predicate, denoted as $\mathcal{F}_{\text{CGKA},4}^{\text{ctxt}}$. $\mathcal{F}_{\text{CGKA},4}^{\text{ctxt}}$ halts if a proposal, a commit, or a welcome message is injected when the sender’s signing key is not exposed. The simulator \mathcal{S}_4 is identical to \mathcal{S}_3 .

Hybrid 5. This hybrid concerns the security of the MAC scheme. We modify $\mathcal{F}_{\text{CGKA},4}^{\text{ctxt}}$ to use the original **mac-inj-allowed** predicate, denoted as $\mathcal{F}_{\text{CGKA},5}^{\text{ctxt}}$. $\mathcal{F}_{\text{CGKA},5}^{\text{ctxt}}$ halts if a proposal or a commit are injected when the corresponding MAC key is not exposed. The simulator \mathcal{S}_5 is identical to \mathcal{S}_4 .

Hybrid 6. This hybrid concerns the confidentiality of the application secrets. We modify $\mathcal{F}_{\text{CGKA},5}^{\text{ctxt}}$ where it uses the original **safe** predicate, denoted as $\mathcal{F}_{\text{CGKA},6}^{\text{ctxt}}$. The simulator \mathcal{S}_6 is identical to \mathcal{S}_5 except that it sets only those application secrets for which **safe** is false. (Note that this functionality roughly corresponds to the previous ideal functionality $\mathcal{F}_{\text{CGKA}}$ where the security of the static metadata is not considered.)

Hybrid 7. This hybrid concerns the confidentiality of proposal, commit and welcome messages. We modify $\mathcal{F}_{\text{CGKA},6}^{\text{ctxt}}$ so that `flagcontHide` is set to true, denoted as $\mathcal{F}_{\text{CGKA},7}^{\text{ctxt}}$. The simulator \mathcal{S}_7 is identical to \mathcal{S}_6 except that it simulates the protocol executions with the leakage information defined by ***leak-create**, ***leak-prop**, ***leak-com**, ***leak-wel**, and ***leak-proc**. The functionality $\mathcal{F}_{\text{CGKA},7}^{\text{ctxt}}$ corresponds to the ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$.

We show the indistinguishability of Hybrids 2 to 7 in Lems. C.2 and C.7. This completes the proof of the main theorem. \square

C.3.3 From Hybrids 2 to 6: Proof of Lem. C.2. We first prove the indistinguishability of hybrid 2 to Hybrid 6.

LEMMA C.2. *Hybrid 2 and Hybrid 6 are indistinguishable assuming the IND-CCA security of mPKE, sEUF-CMA security of SIG, and the key-committing property of SKE.*

Moreover, we assume the adversary can inject bad randomness throughout these hybrids, i.e., RandCorr can be set to ‘bad’.

PROOF. The proof consists of five parts. We first explain how the simulator simulates the protocols with the simulator defined in [36]; second, we explain how to prove the correspondence of the history graphs maintained by $\mathcal{F}_{\text{CGKA},i}$ and $\mathcal{F}_{\text{CGKA},i}^{\text{ctxt}}$, which is the goal of this proof; third, we explain the detailed description of the simulator; fourth, we prove the correspondence between the two history graphs; finally, we prove the supporting propositions to finish the proof of this lemma.

Part 1: Preparation. As explained in the overview, \mathcal{S}_i internally executes the simulator $\tilde{\mathcal{S}}_i$ presented in the proof by Hashimoto et al. [36]. To be more precise, \mathcal{S}_i outsources the simulation of Chained CmpKE^{ctxt} to $\tilde{\mathcal{S}}_i$ that simulates Chained CmpKE by appropriately modifying the inputs and outputs of $\tilde{\mathcal{S}}_i$.

To formalize the description of \mathcal{S}_i , we make one modification to $\tilde{\mathcal{S}}_i$. Without loss of generality, while $\tilde{\mathcal{S}}_i$ is internally simulating Chained CmpKE, we assume it executes and stores the following two secrets whenever the parties are invoked on a Commit or a Join:

- $G.\text{encSecret} \leftarrow \text{HKDF.Expand}(\text{joinerSecret}, G.\text{groupCont}()\| \text{'enc'})$, and
- $\text{welcomeSecret} \leftarrow \text{HKDF.Expand}(\text{joinerSecret}, \text{'wel'})$.

Here, the inputs to HKDF are states of the parties simulated by $\tilde{\mathcal{S}}_i$. $\tilde{\mathcal{S}}_i$ then outputs these stored secrets to \mathcal{S}_i . \mathcal{S}_i uses these secrets to either encrypt proposal and commit messages output by $\tilde{\mathcal{S}}_i$ or decrypt proposal, commit, and welcome messages sent by $\mathcal{F}_{\text{CGKA},i}^{\text{ctxt}}$. Note that these secret keys are never used in the original Chained CmpKE — these are encryption keys used only to secure the static metadata included in proposal, commit, and welcome messages of Chained CmpKE^{ctxt}. It can be checked that the proof in [36] still holds even if we consider this slightly modified simulator $\tilde{\mathcal{S}}_i$ as above since `encSecret` and `welcomeSecret` leak no information on `joinerSecret` in the global random oracle model. We also modify $\tilde{\mathcal{S}}_i$ so that it outputs commit query (`Commit, ...`) and welcome query (`Welcome, ...`) separately like \mathcal{S}_i . That is, $\tilde{\mathcal{S}}_i$ first receives (`Commit, id, \tilde{p} , svk`) and simulates the commit \tilde{c}_0 ; then it receives (`Welcome, kp_t , Ptr[id], \tilde{c}_0`) and simulates the welcome message, where kp_t is the receiver’s key package, `Ptr[id]` is the committer’s current node, and \tilde{c}_0 is the new epoch associated with the welcome message. The above modification in the proof in [36] is without loss of generality since the information $\tilde{\mathcal{S}}_i$ receives to simulate the welcome message corresponds exactly to the information required to invoke the simulated party id to output a welcome message.

Part 2: Goal of Proof. We are now ready to explain the description of \mathcal{S}_i that internally executes $\tilde{\mathcal{S}}_i$. Looking ahead, the main goal of the proof is to relate the history graph created within $\tilde{\mathcal{S}}_i$ to those maintained by $\mathcal{F}_{\text{CGKA},i}^{\text{ctxt}}$. Then, since [36] proved that the history graph made within $\tilde{\mathcal{S}}_i$ and those maintained by the ideal functionality $\mathcal{F}_{\text{CGKA}}$ in Hybrid i (i.e., $\mathcal{F}_{\text{CGKA},i}$) are identical, we will be able to indirectly prove that the history graph made within \mathcal{S}_i and those maintained by the ideal functionality $\mathcal{F}_{\text{CGKA},i}^{\text{ctxt}}$ are identical as well. Indistinguishability between each adjacent hybrid then follows straightforwardly from the prior proof and the correspondence of the history graphs maintained by $\mathcal{F}_{\text{CGKA},i}$ and $\mathcal{F}_{\text{CGKA},i}^{\text{ctxt}}$.

Below, we put tilde on top, e.g., $\tilde{\text{Ptr}}[*]$, $\tilde{\text{Node}}[*]$, $\tilde{\text{Wel}}[*]$, to denote the components created by the history graph within $\tilde{\mathcal{S}}_i$ (which is also known to \mathcal{S}_i). As explained above, the history graphs created within $\tilde{\mathcal{S}}_i$ and maintained by $\mathcal{F}_{\text{CGKA},i}^{\text{ctxt}}$ are identical so we may use them interchangeably. We put tilde on proposals \tilde{p} and commits \tilde{c}_0 to denote that they are *non-encrypted* variants used by $\tilde{\mathcal{S}}_i$.

Using these terminologies, we can restate our goal as to create a one-to-one correspondence between the nodes $\tilde{\text{Node}}[\tilde{c}]$, $\tilde{\text{Prop}}[\tilde{p}]$, and $\tilde{\text{Wel}}[\tilde{w}]$ maintained by the history graph in $\mathcal{F}_{\text{CGKA},i}$ and the nodes $\text{Node}[\text{node-id}]$, $\text{Prop}[\text{prop-id}]$, and $\text{Wel}[\text{node-id}]$ maintained by the history graph in $\mathcal{F}_{\text{CGKA},i}^{\text{ctxt}}$. We will make the meaning of “one-to-one correspondence” more clear later, but we informally mean that the graph topology and the syntactical information stored in each node are consistent with each other. We assume \mathcal{S}_i maintains lists L_p and L_c that (roughly) assign unique counters to proposal and commit nodes created within \mathcal{S}_i , respectively. This allows us to relate which node in $\mathcal{F}_{\text{CGKA},i}^{\text{ctxt}}$ relates to $\mathcal{F}_{\text{CGKA},i}$.

Part 3: Description of Simulator \tilde{S}_i . For a detailed motivation and description of \tilde{S}_i , we refer the readers to the original proof by Hashimoto et al. [36]. We only use their result proving that the history graphs created within \tilde{S}_i and maintained by $\mathcal{F}_{\text{CGKA},i}^{\text{ctxt}}$ are identical.

Below, we explain how \tilde{S}_i answers each queries made by the ideal functionality $\mathcal{F}_{\text{CGKA},i}^{\text{ctxt}}$ for $i \in [3 : 6]$, where we omit the case $i = 2$ since it is already defined (see the definition of Hybrids 1 and 2). Here, keep in mind that for $i \in [3 : 6]$ we have $\text{flag}_{\text{contHide}} = \text{false}$, i.e., the leakage functions leak all the static metadata. Since the description of \tilde{S}_i in [36] is identical for each $i \in [3 : 6]$, we omit the subscript i below for simplicity. In other words, we only need to check that the history graph maintained by $\mathcal{F}_{\text{CGKA},3}^{\text{ctxt}}$ has a one-to-one correspondence with $\mathcal{F}_{\text{CGKA},3}$. As a result, the indistinguishability of Hybrids 2 to 6 then inherits from [36].

(1) Create query from $\text{id}_{\text{creator}}$. This concerns the case when \mathcal{Z} queries (Create, svk) to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. If $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ outputs (Create, $\text{id}_{\text{creator}}$, svk) to \mathcal{S} , \mathcal{S} invokes \tilde{S} on the same input. From \tilde{S} 's point of view, \mathcal{S} acts identically to the ideal functionality $\mathcal{F}_{\text{CGKA}}$. \tilde{S} simply runs the simulated party $\text{id}_{\text{creator}}$ on input (Create, svk).³⁰ \mathcal{S} then samples a random group identifier $\text{gid} \leftarrow \{0, 1\}^K$ and outputs it to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. This creates the node $\text{Node}[0]$ within the history graph maintained by $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. Finally, since \tilde{S} creates a root node $\text{Node}[\text{root}_0]$, \mathcal{S} records $L_c[\text{root}_0] \leftarrow 0$.

(2) Propose query from id . This concerns the case when \mathcal{Z} queries (Propose, act) for some act $\in \{\text{'upd' -svk, 'add' -id}_t\text{-kp}_t, \text{'rem' -id}_t\}$ to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. If $\text{Ptr}[\text{id}] \neq \perp$, then $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ outputs (Propose, node-id := $\text{Ptr}[\text{id}], \text{id}, \text{act}$) to \mathcal{S} , and \mathcal{S} invokes \tilde{S} on the same input. From \tilde{S} 's point of view, \mathcal{S} acts identically to the ideal functionality $\mathcal{F}_{\text{CGKA}}$. \tilde{S} then runs the simulated party id on input (Propose, act). If $\text{RandCorr}[\text{id}] = \text{'bad'}$ and act = 'upd'-svk, it asks \mathcal{S} to provide the randomness to run party id . \mathcal{S} then queries this request to its own adversary \mathcal{A} to receive the adversarially chosen randomness. Here, recall that randomness is only used by \tilde{S} to generate an update key package kp. If party id , internally simulated by \tilde{S} , returns \perp , then \tilde{S} returns ($\text{ack} := \text{false}, \perp$) to \mathcal{S} . \mathcal{S} then outputs ($\text{ack} := \text{false}, \perp, \perp$) to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$.³¹

Otherwise, if id returns \tilde{p} , then \tilde{S} returns ($\text{ack} := \text{true}, \tilde{p}$). \mathcal{S} then encrypts \tilde{p} as $p^{\text{ctxt}} \leftarrow \text{*enc-prop}(G.\text{encSecret}, \tilde{p})$ using $G.\text{encSecret}$ of id (provided by \tilde{S} , see Part 1). It runs the encryption on the adversarially controlled randomness if $\text{RandCorr}[\text{id}] = \text{'bad'}$. There are two cases to consider.

Case 1. If $\widetilde{\text{Prop}}[\tilde{p}] = \perp$, i.e., it did not exist in the history graph maintained by \tilde{S} , then \mathcal{S} outputs ($\text{ack} := \text{true}, \perp, p^{\text{ctxt}}$).

Case 2. If $\widetilde{\text{Prop}}[\tilde{p}] \neq \perp$, i.e., it exists in the history graph maintained by \tilde{S} , then \mathcal{S} searches for any $p^{\text{ctxt}'}$ such that $\tilde{p} \leftarrow \text{*dec-prop}(G.\text{encSecret}, p^{\text{ctxt}'})$ and $\text{prop-id} = \text{PropID}[p^{\text{ctxt}'}] \neq \perp$. It then chooses any such $p^{\text{ctxt}'}$ and outputs ($\text{ack} := \text{true}, \text{prop-id}, p^{\text{ctxt}}$).

³⁰We note that in the previous definition of $\mathcal{F}_{\text{CGKA}}$, the group identifies gid was omitted. However, this can be amended to $\mathcal{F}_{\text{CGKA}}$ without loss of generality.

³¹Recall that due to our modification in the definition of the key service, \tilde{S} does not need to output svk_t (see App. B.1.2).

By Proposition C.5, Item 1, which we prove later, such $p^{\text{ctxt}'}$ exists and prop-id is unique regardless of the choice of $p^{\text{ctxt}'}$.

If act = 'upd', \tilde{S} outputs the updated key package kp^{ctxt} . \mathcal{S} passes it to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$.

(3) Commit query from id . This concerns the case when \mathcal{Z} queries (Commit, \tilde{p}^{ctxt} , svk) to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. If $\text{Ptr}[\text{id}] \neq \perp$, then $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ outputs (Commit, node-id := $\text{Ptr}[\text{id}], \text{id}, \tilde{p}^{\text{ctxt}}$, svk) to \mathcal{S} . \mathcal{S} then decrypts the proposals $\tilde{p} := \text{*dec-and-sort-proposals}(G.\text{encSecret}, \tilde{p}^{\text{ctxt}})$ using $G.\text{encSecret}$ of id . If decryption fails, \mathcal{S} outputs ($\text{ack} := \text{false}, \perp, \perp, \perp$). Otherwise, \mathcal{S} invokes \tilde{S} on input (Commit, \tilde{p}, svk). \tilde{S} then runs the simulated party id on input (Commit, \tilde{p}, svk), where it asks \mathcal{S} to provide the randomness to run party id if $\text{RandCorr}[\text{id}] = \text{'bad'}$. \mathcal{S} then queries this request to its own adversary \mathcal{A} to receive the adversarially chosen randomness. If party id returns \perp , then \tilde{S} outputs ($\text{ack} := \text{false}, \perp, \perp, \perp$) to \mathcal{S} , and \mathcal{S} outputs ($\text{ack} := \text{false}, \perp, \perp, \perp$) to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$.

Otherwise, if party id returns $(\tilde{c}_0, \tilde{c}, \tilde{w}_0, \tilde{w})$, \tilde{S} decides what to output depending on the following checks (which are identical to the checks in [36]). We also explain what \mathcal{S} does when given the output of \tilde{S} .

Case 1. If $\widetilde{\text{Node}}[\tilde{c}_0] = \perp$, $w_0 \neq \perp$, and if there exists some $rt' \in \mathbb{N}$ and \tilde{w}'_0 such that $\widetilde{\text{Wel}}[\tilde{w}'_0] = \text{root}_{rt'}$ and \tilde{w}'_0 includes the same confTag as \tilde{w}_0 , then \tilde{S} chooses any such (\tilde{w}'_0, rt') (guaranteed to exist uniquely from [36, Proposition E.6]) and returns ($\text{ack} := \text{true}, rt := rt' \neq \perp, \tilde{c}_0, \tilde{c}$) to \mathcal{S} . \mathcal{S} then encrypts the (non-encrypted) commit message as $(c_0^{\text{ctxt}}, \tilde{c}^{\text{ctxt}}) \leftarrow \text{*enc-commit}(G.\text{encSecret}, \tilde{c}_0, \tilde{c})$ using $G.\text{encSecret}$ of id (provided by \tilde{S} , see Part 1). By Proposition C.6, Item 3 (which we prove later), there exists a unique node-id such that $\text{Node}[\text{node-id}] = \widetilde{\text{Node}}[\text{root}_{rt}]$. That is, the syntactical information stored on the node, e.g., gid, epoch, prop, orig, and so on, are identical. \mathcal{S} then outputs ($\text{ack} := \text{true}, \text{node-id}, c_0^{\text{ctxt}}, \tilde{c}^{\text{ctxt}}$) to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. Afterwards, when \mathcal{S} is invoked on $(\text{Welcome}, \text{kp}_t, \text{node-id} := \text{Ptr}[\text{id}], \tilde{c}_0)$ by $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$, it invokes \tilde{S} on the same input. \tilde{S} then returns ($\text{ack} := \text{true}, \tilde{w}_0, \tilde{w}$) to \mathcal{S} , where \tilde{w} is the welcome message for added member $\text{id}_t \in \text{mem}$ with kp_t . Finally, \mathcal{S} encrypts (\tilde{w}_0, \tilde{w}) as \tilde{w}^{ctxt} using $G.\text{welcomeSecret}$ of id and returns ($\text{ack} := \text{true}, \tilde{w}^{\text{ctxt}}$) to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$.

Case 2. Otherwise, if $\widetilde{\text{Node}}[\tilde{c}_0] = \perp$ and either $\tilde{w}_0 = \perp$ or there does not exist \tilde{w}'_0 such that $\widetilde{\text{Wel}}[\tilde{w}'_0] = \text{root}_{rt'}$ for some $rt' \in \mathbb{N}$ and \tilde{w}'_0 includes the same confTag as \tilde{w}_0 , then \tilde{S} returns ($\text{ack} := \text{true}, rt := \perp, \tilde{c}_0, \tilde{c}$) to \mathcal{S} . \mathcal{S} then encrypts the commit as above and returns ($\text{ack} := \text{true}, \perp, c_0^{\text{ctxt}}, \tilde{c}^{\text{ctxt}}$) to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. The welcome message is handled similarly to Case 1.

Case 3. Finally, if $\widetilde{\text{Node}}[\tilde{c}_0] \neq \perp$, then \tilde{S} returns ($\text{ack} := \text{true}, rt := \perp, \tilde{c}_0, \tilde{c}$) to \mathcal{S} . \mathcal{S} then searches for any $c_0^{\text{ctxt}'}$ such that $\tilde{c}_0 \leftarrow \text{*dec-commit}_0(G.\text{encSecret}, c_0^{\text{ctxt}'})$ and $\text{node-id} = \text{NodeID}[c_0^{\text{ctxt}'}] \neq \perp$ (see Proposition C.6 for the definition of *dec-commit_0).

³²Recall that due to our modification in the definition of $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$, \tilde{S} needs to output kp if act = 'upd'. However, this can be amended to \tilde{S} without loss of generality.

It then chooses any such $c_0^{\text{ctxt}'}$ and outputs $(ack := \text{true}, \text{node-id}, c^{\text{ctxt}}, \tilde{c}^{\text{ctxt}})$. By Proposition C.6, Item 1 (which we prove later), such $c_0^{\text{ctxt}'}$ exists and node-id is unique regardless of the choice of $c_0^{\text{ctxt}'}$.

If act = 'upd', \tilde{S} outputs the committer's updated key package kp, and \mathcal{S} passes it to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ accordingly.

Finally, when $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ queries (Propose, Ptr[id], p^{ctxt}) to \mathcal{S} during the *fill-prop check, it must first decrypt p^{ctxt} to obtain the (non-encrypted) proposal \tilde{p} . We use [36, Proposition E.8] that establishes that any member at the same node $\tilde{\text{Ptr}}[\text{id}]$ stores the same secret key $G.\text{encSecret}$. Namely, \mathcal{S} retrieves the unique $G.\text{encSecret}$ assigned to the node $\tilde{\text{Ptr}}[\text{id}]$ to decrypt p^{ctxt} . If $\text{Prop}[\tilde{p}] = \perp$, then \mathcal{S} invokes \tilde{S} on input (Propose, \tilde{p}) and receives (orig, act)³³. It then outputs $(\perp, \text{orig}, \text{act})$ to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. Otherwise, if $\text{Prop}[\tilde{p}] \neq \perp$, then \mathcal{S} searches for any $p^{\text{ctxt}'}$ such that $\tilde{p} \leftarrow *dec\text{-prop}(G.\text{encSecret}, p^{\text{ctxt}'})$ and prop-id = $\text{PropID}[p^{\text{ctxt}'}] \neq \perp$. It then choses any such $p^{\text{ctxt}'}$ and outputs $(ack := \text{true}, \text{prop-id}, p^{\text{ctxt}'})$. By Proposition C.5, Item 1, such $p^{\text{ctxt}'}$ exists and prop-id is unique regardless of the choice of $p^{\text{ctxt}'}$. It retrieves $(\text{orig}, \text{act}) \leftarrow (\text{Prop}[\tilde{p}], \text{orig}, \text{Prop}[\tilde{p}].\text{act})$ and returns (prop-id, orig, act) to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$.

(4) Process query from id. This concerns the case when \mathcal{Z} queries (Process, $c_0^{\text{ctxt}}, \tilde{c}^{\text{ctxt}}, \tilde{p}^{\text{ctxt}}$) to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. If $\text{Ptr}[\text{id}] \neq \perp$, then $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ outputs (Process, node-id, id, $c_0^{\text{ctxt}}, \tilde{c}^{\text{ctxt}}, \tilde{p}^{\text{ctxt}}$) to \mathcal{S} . \mathcal{S} then decrypts the input messages as $(\tilde{c}_0, \tilde{c}) := *dec\text{-commit}(G.\text{encSecret}, c_0^{\text{ctxt}}, \tilde{c}^{\text{ctxt}})$ and $\tilde{p} := *dec\text{-and-sort-proposals}(G.\text{encSecret}, \tilde{p}^{\text{ctxt}})$ as in the real protocol. If decryption fails, \mathcal{S} outputs $(ack := \text{false}, \perp, \perp, \perp)$. Otherwise, \mathcal{S} invokes \tilde{S} on input (Process, $\tilde{c}_0, \tilde{c}, \tilde{p}$). \tilde{S} then (deterministically) runs the simulated party id on input (Process, id, $\tilde{c}_0, \tilde{c}, \tilde{p}$). If party id returns \perp , then \tilde{S} outputs $(ack := \text{false}, \perp, \perp, \perp)$ to \mathcal{S} . \mathcal{S} then outputs $(ack := \text{false}, \perp, \perp, \perp)$ to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$.

Otherwise, if party id returns $(id_c, \text{upd}||\text{rem}||\text{add})$, \tilde{S} decides what to output depending on the following checks (which are the checks identical to [36]). We also explain what \mathcal{S} does when given the output of \tilde{S} . Note that *fill-prop queries from $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ to \mathcal{S} are answered exactly as in commit queries described above.

Case 1. If $\text{Node}[\tilde{c}_0] = \perp$ and if there exists \tilde{w}_0 that includes the same confTag as \tilde{c}_0 such that $\text{Wel}[\tilde{w}_0] = \text{root}_{rt'}$ for some $rt' \in \mathbb{N}$, then \tilde{S} chooses any such (\tilde{w}_0, rt') (guaranteed to exist uniquely from [36, Proposition E.6]) and returns $(ack := \text{true}, rt := rt', \perp, \perp)$ to \mathcal{S} . By Proposition C.6, Item 3, there exists a unique node-id such that $\text{Node}[\text{node-id}] = \text{Node}[\text{root}_{rt}]$. \mathcal{S} then outputs $(ack := \text{true}, \text{node-id}, \perp, \perp)$ to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$.

Case 2. If $\text{Node}[\tilde{c}_0] = \perp$ and no such \tilde{w}_0 exists, then \tilde{S} retrieves the associating long-term public key svk_c of id_c (which is guaranteed to exist when process succeeds in the real protocol) and returns $(ack := \text{true}, \perp, \text{orig}' := \text{id}_c, \text{svk}' := \text{svk}_c)$ to \mathcal{S} . \mathcal{S} then outputs $(ack := \text{true}, \perp, \text{orig}', \text{svk}')$ to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$.

Case 3. Finally, if $\text{Node}[\tilde{c}_0] \neq \perp$, then \tilde{S} simply returns $(ack := \text{true}, \perp, \perp, \perp)$ to \mathcal{S} . \mathcal{S} then searches for any $c_0^{\text{ctxt}'}$ such that

$\tilde{c}_0 \leftarrow *dec\text{-commit}_0(G.\text{encSecret}, c_0^{\text{ctxt}'})$ and $\text{node-id} = \text{NodeID}[c_0^{\text{ctxt}'}] \neq \perp$ (see Proposition C.6 for the definition of $*dec\text{-commit}_0$). It then choses any such $c_0^{\text{ctxt}'}$ and outputs $(ack := \text{true}, \perp, \perp, \perp)$. By Proposition C.6, Item 1, such $c_0^{\text{ctxt}'}$ exists and node-id is unique regardless of the choice of $c_0^{\text{ctxt}'}$.

(5) Join query from id. This concerns the case when \mathcal{Z} queries (Join, \tilde{w}^{ctxt}) to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. If $\text{Ptr}[\text{id}] = \perp$, then $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ outputs (Join, id, \tilde{w}^{ctxt}) to \mathcal{S} . \mathcal{S} then decrypts the input messages as $(\tilde{w}_0, \tilde{w}) := *dec\text{-welcome}(\tilde{w}^{\text{ctxt}})$ as in the real protocol. If decryption fails, \mathcal{S} outputs $(ack := \text{false}, \perp, \perp, \perp, \perp, \perp)$. Otherwise, \mathcal{S} invokes \tilde{S} on input (Join, id, \tilde{w}_0, \tilde{w}). \tilde{S} then (deterministically) runs the simulated party id on input (Join, \tilde{w}_0, \tilde{w}). If party id returns \perp , then \tilde{S} outputs $(ack := \text{false}, \perp, \perp, \perp)$ to \mathcal{S} . \mathcal{S} then outputs $(ack := \text{false}, \perp, \perp, \perp, \perp, \perp)$ to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$.

Otherwise, if party id returns (id_c, mem) , \tilde{S} decides what to output depending on the following checks (which are the checks identical to [36]). We also explain what \mathcal{S} does when given the output of \tilde{S} .

Case 1. If $\text{Wel}[\tilde{w}_0] \neq \perp$, \tilde{S} returns $(ack := \text{true}, \perp, \perp, \perp)$ to \mathcal{S} . \mathcal{S} then returns $(ack := \text{true}, \perp, \perp, \perp, \perp, \perp)$ to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$.

Case 2. Otherwise, \tilde{S} checks if there exists a non-root \tilde{c}_0 such that $\text{Node}[\tilde{c}_0] \neq \perp$ and \tilde{c}_0 includes the same confTag as the one included in \tilde{w}_0 . If such \tilde{c}_0 exists (guaranteed to be unique by [36]), then \tilde{S} returns $(ack := \text{true}, \tilde{c}'_0 := \tilde{c}_0, \perp, \perp)$ to \mathcal{S} . \mathcal{S} then retrieves $G.\text{encSecret}$ stored on the node $\text{Node}[\tilde{c}_0]$ and searches for any $c_0^{\text{ctxt}'}$ such that $\tilde{c}_0 \leftarrow *dec\text{-commit}_0(G.\text{encSecret}, c_0^{\text{ctxt}'})$ and $\text{node-id} = \text{NodeID}[c_0^{\text{ctxt}'}] \neq \perp$ (see Proposition C.6 for the definition of $*dec\text{-commit}_0$). It then choses any such $c_0^{\text{ctxt}'}$ and outputs $(ack := \text{true}, \text{node-id}, \perp, \perp, \perp, \perp)$ to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$.

Case 3. Otherwise, if no such \tilde{c}_0 exists, then \tilde{S} further checks if there exists \tilde{w}'_0 such that $\text{Wel}[\tilde{w}'_0] \neq \perp$ that includes the same confTag as the one included in \tilde{w}_0 . If so, \tilde{S} chooses any such \tilde{w}'_0 and returns $(ack := \text{true}, \tilde{c}'_0 := \text{Wel}[\tilde{w}'_0], \perp, \perp)$. Here, [36, Proposition E.6] establishes that the value of $\text{Wel}[\tilde{w}'_0]$ is guaranteed to be the same root value, i.e., root_{rt} for some $rt \in \mathbb{N}$, for all such \tilde{w}'_0 . Then, by Proposition C.6, Item 3, there exists a unique node-id such that $\text{Node}[\text{node-id}] = \text{Node}[\text{root}_{rt}]$. \mathcal{S} then outputs $(ack := \text{true}, \text{node-id}, \perp, \perp, \perp, \perp)$ to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$.

Case 4. Finally, if no such \tilde{c}_0 or \tilde{w}'_0 exist, then \tilde{S} returns $(ack := \text{true}, \perp, \text{orig}' := \text{id}_c, \text{mem}' := \text{mem})$ to \mathcal{S} . This corresponds to the case $\text{Wel}[\tilde{w}_0]$ is initialized by root_{rt} for a new $rt \in \mathbb{N}$, whose value is known to \tilde{S} . \mathcal{S} then outputs $(ack := \text{true}, \perp, \text{gid}', \text{epoch}', \text{orig}', \text{mem}')$ to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$, where gid' and epoch' are included in \tilde{w}'_0 in the clear.

(6) Key query from id. This concerns the case when \mathcal{Z} queries Key to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. If $\text{Ptr}[\text{id}] \neq \perp$ and $\text{Node}[\text{Ptr}[\text{id}]].\text{key} = \perp$, then $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ outputs (Key, id) to \mathcal{S} . (Recall that **safe** is always set to false in this hybrid.) \mathcal{S} invokes \tilde{S} on input (Key, id) and receives the group secret key k. \mathcal{S} returns k to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$.

Part 4: Correspondence Between the Two History Graphs. Before proving the supporting Propositions C.5 and C.6, we prove that the

³³Without loss of generality, we assume \tilde{S} returns $\text{act} \in \{\text{'upd'-'kp'}, \text{'add'-'id'-'kp'}$, 'rem'-'id' }

history graph created within \widetilde{S} (and hence the ideal functionality $\mathcal{F}_{\text{CGKA}}$) has a “one-to-one” correspondence between the history graph maintained by $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. To formalize this, we first define what it means for the two history graphs to be *isomorphic*.

Definition C.3. Let HG and $\widetilde{\text{HG}}$ be the history graphs maintained by $\mathcal{F}_{\text{CGKA}}$ and $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$, respectively. We say the two history graphs are *isomorphic* if the following holds:

- Let $S_p := \{\text{prop-id}\}_{\text{prop-id}}$ and $\widetilde{S}_p := \{p\}_p$ be the sets of all proposal nodes in HG and $\widetilde{\text{HG}}$, respectively. Then, there is a bijection f_p between these two sets such that $\text{Prop}[\text{prop-id}]$ and $\widetilde{\text{Prop}}[f_p(\text{prop-id})]$ store the same values (modulo the syntactical difference that $\cdot\text{par}$ stores node-id or c_0/root_{rt}).
- Let $S_c := \{\text{node-id}\}_{\text{node-id}}$ and $\widetilde{S}_c := \{c_0\}_{c_0} \cup \{\text{root}_{rt}\}_{rt}$ be the sets of all commit nodes in HG and $\widetilde{\text{HG}}$, respectively. Then, there is a bijection f_c between these two sets such that $\text{Node}[\text{node-id}]$ and $\widetilde{\text{Node}}[f_c(\text{node-id})]$ store the same values (modulo the syntactical difference that $\cdot\text{par}$ stores node-id or c_0/root_{rt} and $\cdot\text{prop}$ stores the encrypted or non-encrypted proposal).

The following lemma is the key lemma to prove Lem. C.2. Since isomorphic graphs agree on the same syntactical information, this also implies that the nodes for which the predicate **safe** is true are identical. Notably, all the proofs used to move from Hybrids 2 to 6 in Hashimoto et al. [36] with respect to $\mathcal{F}_{\text{CGKA}}$ carryovers to those of $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$.

LEMMA C.4. *The history graph HG maintained by $\mathcal{F}_{\text{CGKA}}$ and $\widetilde{\text{HG}}$ maintained by $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ are isomorphic.*

PROOF OF LEM. C.4. We prove by induction. When the history graphs are both empty, then they are isomorphic with the two bijections f_p and f_c defined to be the zero function. Now, assume the history graphs HG and $\widetilde{\text{HG}}$ are currently isomorphic, i.e., there exists bijections f_p and f_c between the propose and commit nodes. We check that the isomorphism is maintained even after a **Propose**, a **Commit**, a **Process**, or a **Join** query is performed. We omit queries to **Create** and **key** as they trivially maintain the isomorphism.

Consistency of isomorphism after Propose. Let us consider *Case 1* in (2), i.e., $\widetilde{\text{Prop}}[\widetilde{p}] = \perp$. By Proposition C.5, Item 2, we have $\text{PropID}[p^{\text{ctxt}}] = \perp$. Thus, both $\mathcal{F}_{\text{CGKA}}$ and $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ creates a new proposal node. By assumption on the induction and how $\cdot\text{create-prop}$ is defined, we have isomorphism even after **Propose**. In particular, letting propCtr be the current stored by $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$, the new bijection f'_p extends the domain and range of f_p by adding $f'_p(\text{propCtr}) = \widetilde{p}$.

Next, let us consider *Case 2* in (2), i.e., $\widetilde{\text{Prop}}[\widetilde{p}] \neq \perp$. In this case $\mathcal{F}_{\text{CGKA}}$ and $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ performs a consistency check $\cdot\text{consistent-prop}$. Since no new proposal node is created, we only need to show that the output of this consistency check is identical for both ideal functionalities. There are two cases: $\text{PropID}[p^{\text{ctxt}}] = \perp$ or $\text{PropID}[p^{\text{ctxt}}] = \text{prop-id}'$ for some $\text{prop-id}' \neq \perp$. In the former case, by Proposition C.5, Item 1, \mathcal{S} finds $p^{\text{ctxt}'} \neq p^{\text{ctxt}}$ and outputs $\text{prop-id} = \text{PropID}[p^{\text{ctxt}'}]$. $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ performs the consistency check using this prop-id . In the latter case, $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ uses the existing $\text{prop-id}' = \text{PropID}[p^{\text{ctxt}}]$ and ignores prop-id output by \mathcal{S} .

By Proposition C.5, Item 1, the values of prop-id and $\text{prop-id}'$ are identical. We thus only need to focus on $\text{prop-id}'$.

Recall $\text{prop-id}' = \text{PropID}[p^{\text{ctxt}}]$ is created either during **Propose** or $\cdot\text{fill-prop}$. Due to the key-committing property of SKE, it must be created by a node $\widetilde{\text{Ptr}}[\text{id}']$ that has the same $G.\text{encSecret}$ as $\widetilde{\text{Ptr}}[\text{id}]$. Using [36, Proposition E.8], we must have $\widetilde{\text{Ptr}}[\text{id}'] = \widetilde{\text{Ptr}}[\text{id}]$. In other words, $\text{prop-id}'$ was created by the node $\text{Ptr}[\text{id}]$. Thus, by the assumption on the induction, the consistency check by $\mathcal{F}_{\text{CGKA}}$ using \widetilde{p} and those by $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ using $\text{prop-id}'$ are identical.

Consistency of isomorphism after Commit. Let us consider *Case 1* in (3). This is the case where \mathcal{S} assigns a detached root to c_0^{ctxt} . By Proposition C.6, Item 3, there exists a unique node-id such that $\text{Node}[\text{node-id}] = \widetilde{\text{Node}}[\text{root}_{rt}]$. Thus, both $\mathcal{F}_{\text{CGKA}}$ and $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ perform the same consistency check and the attach procedure. At the end of **Commit**, the new bijection f'_c is identical to f_c except that $f'_c(\text{node-id}) = \widetilde{c}_0$ rather than $f_c(\text{node-id}) = \text{root}_{rt}$.

Next, let us consider *Case 2* in (3). This is the case where \mathcal{S} assigns c_0^{ctxt} to a new node. By Proposition C.5, Item 2, we have $\text{PropID}[p^{\text{ctxt}}] = \perp$. Thus, both $\mathcal{F}_{\text{CGKA}}$ and $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ create a new commit node. By assumption on the induction and how $\cdot\text{create-child}$ is defined, we have isomorphism even after **Commit**. In particular, letting nodeCtr be the current stored by $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$, the new bijection f'_c extends the domain and range of f_c by adding $f'_c(\text{nodeCtr}) = \widetilde{c}_0$.

Finally, let us consider *Case 3* in (3). In this case $\mathcal{F}_{\text{CGKA}}$ and $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ performs a consistency check $\cdot\text{consistent-com}$. Since no new commit node is created, we only need to show that the output of this consistency check is identical for both ideal functionalities. Similarly to the argument made in **Propose**, there are two cases to consider: $\text{NodeID}[c_0^{\text{ctxt}}] = \perp$ or $\text{NodeID}[c_0^{\text{ctxt}}] = \text{node-id}'$ for some $\text{node-id}' \neq \perp$. In the former case, by Proposition C.6, Item 1, \mathcal{S} finds $c_0^{\text{ctxt}'} \neq c_0^{\text{ctxt}}$ and outputs $\text{node-id} = \text{NodeID}[c_0^{\text{ctxt}'}]$. $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ performs the consistency check using this node-id . In the latter case, $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ uses the existing $\text{node-id}' = \text{NodeID}[c_0^{\text{ctxt}}]$ and ignores node-id output by \mathcal{S} . By Proposition C.6, Item 1, the values of node-id and $\text{node-id}'$ are identical. We thus only need to focus on $\text{node-id}'$.

Recall $\text{node-id}' = \text{NodeID}[c_0^{\text{ctxt}}]$ is created either during **Commit** or **Process**. Following an almost exact argument made for **PropID** above, we can show $\text{node-id}'$ was created by the node $\text{Ptr}[\text{id}]$. Thus, by the assumption on the induction, the consistency check by $\mathcal{F}_{\text{CGKA}}$ using \widetilde{c} and those by $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ using $\text{node-id}'$ are identical.

Consistency of isomorphism after Process. All three cases considered in **Process** are covered by the cases considered in **Commit**. Namely, \mathcal{S} chooses to either attach a detached root to c_0^{ctxt} , create a new commit node, or attach the commit to an already existing node.

Isomorphism after Join. First, we show that if $\widetilde{\text{Wel}}[\widetilde{w}_0] = \perp$, then $\text{Wel}[\text{id}, \widetilde{w}] = \perp$. Assume $\text{Wel}[\text{id}, \widetilde{w}] \neq \perp$. Since $\text{Wel}[\text{id}, \widetilde{w}]$ is only created during a **Commit**, this implies that there exists $\widetilde{w}'_0 \neq \widetilde{w}_0$ such that $\widetilde{\text{Wel}}[\widetilde{w}'_0] \neq \perp$. This further implies that \widetilde{w} decrypts to two different values \widetilde{w}'_0 and \widetilde{w}_0 , which contradicts the committing property of the SKE. Therefore, we have $\text{Wel}[\text{id}, \widetilde{w}] = \perp$ when $\widetilde{\text{Wel}}[\widetilde{w}_0] = \perp$.

With this in mind, *Cases 1* and *2* in (5) do not alter the propose or commit nodes of the history graphs maintained by both $\mathcal{F}_{\text{CGKA}}$ and

$\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. These cases correspond to attaching a welcome message to an existing node on the history graph. Moreover, in *Case 3* in (5), as we already established that $\text{Node}[\text{node-id}] = \widetilde{\text{Node}}[\text{root}_r] \neq \perp$. Therefor, this also does not alter the history graph.

The final *Case 4* in (5) corresponds to the case where a new detached root is created within $\mathcal{F}_{\text{CGKA}}$. By Proposition C.6, Item 2, $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ also creates a new detached root. Since *create-root is called on the same input, it is clear that the newly created history graphs remain isomorphic. Namely, the new bijection f'_c extends the domain and range of f_c by adding $f'_c(\text{nodeCtr}) = \text{rt}$, where rt and nodeCtr are the counters used to add a new commit node in $\mathcal{F}_{\text{CGKA}}$ and $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$, respectively. \square

Part 5: Proving the Supporting Propositions. To finish the proof of Lem. C.2, we finally prove Propositions C.5 and C.6. Below, we rely on [36, Proposition E.8] that establishes that any member at the same node $\widetilde{\text{Ptr}}[\text{id}]$ stores the same secret key $G.\text{encSecret}$.

PROPOSITION C.5. *We have the following:*

- (1) If $\widetilde{\text{Prop}}[\widetilde{p}] \neq \perp$, then there exists some p^{ctxt} such that $\widetilde{p} \leftarrow \text{*dec-prop}(G.\text{encSecret}, p^{\text{ctxt}})$ and $\text{PropID}[p^{\text{ctxt}}] \neq \perp$, where $G.\text{encSecret}$ is the secret key maintained in the node $\widetilde{\text{Prop}}[\widetilde{p}]$. Moreover, for all such p^{ctxt} , the value of $\text{PropID}[p^{\text{ctxt}}]$ is identical.
- (2) If $\widetilde{\text{Prop}}[\widetilde{p}] = \perp$, then does not exist any p^{ctxt} and encSecret such that $\widetilde{p} \leftarrow \text{*dec-prop}(\text{encSecret}, p^{\text{ctxt}})$ and $\text{PropID}[p^{\text{ctxt}}] \neq \perp$.

PROOF. Let us first consider Item 1. $\widetilde{\text{Prop}}[\widetilde{p}]$ is only created during a *Propose* or a **fill-prop*. If it was created during a *Propose*, then by definition such p^{ctxt} exists. If it was created during a **fill-prop*, then \mathcal{S} was given p^{ctxt} such that $\widetilde{p} \leftarrow \text{*dec-prop}(G.\text{encSecret}, p^{\text{ctxt}})$. At the end of **fill-prop*, $\text{PropID}[p^{\text{ctxt}}]$ is created. Moreover, due to the key-committing property of the SKE, the proposal node assigned to such proposals p^{ctxt} are identical.

Next, let us consider Item 2. Due to the key-committing property of the SKE, if such a p^{ctxt} existed, then we must have $\widetilde{\text{Prop}}[\widetilde{p}] \neq \perp$. Thus, Item 2 follows by taking the contrapositive. \square

PROPOSITION C.6. *Let us define the function $\text{*dec-commi}t_0$ such that given c_0^{ctxt} and encSecret , it parses $(\text{gid}, \text{epoch}, \text{'commit'}, \text{CT}_{\widetilde{c}_0}) \leftarrow c_0^{\text{ctxt}}$, runs $\widetilde{c}_0 \leftarrow \text{SKE.Dec}(\text{encSecret}, \text{CT}_{\widetilde{c}_0})$, and outputs $\widetilde{c}_0 := (\text{gid}, \text{epoch}, \text{'commit'}, \widetilde{c}_0)$.*

Then, we have the following:

- (1) If $\widetilde{\text{Node}}[\widetilde{c}_0] \neq \perp$, then there exists some c_0^{ctxt} such that $\widetilde{c}_0 \leftarrow \text{*dec-commi}t_0(G.\text{encSecret}, c_0^{\text{ctxt}})$ and $\text{NodeID}[c_0^{\text{ctxt}}] \neq \perp$, where $G.\text{encSecret}$ is the secret key maintained in the node $\widetilde{\text{Node}}[\widetilde{c}_0]$. Moreover, for all such c_0^{ctxt} , the value of $\text{NodeID}[c_0^{\text{ctxt}}]$ is identical.
- (2) If $\widetilde{\text{Node}}[\widetilde{c}_0] = \perp$, then there does not exist any c_0^{ctxt} and encSecret such that $\widetilde{c}_0 \leftarrow \text{*dec-commi}t_0(\text{encSecret}, c_0^{\text{ctxt}})$ and $\text{NodeID}[c_0^{\text{ctxt}}] \neq \perp$.
- (3) If $\widetilde{\text{Node}}[\widetilde{c}_0] = \perp$ and there exists \widetilde{w}_0 that includes the same confTag as \widetilde{c}_0 such that $\widetilde{\text{Wel}}[\widetilde{w}_0] = \text{root}_r$ for some $r \in \mathbb{N}$, then there exists a unique node-id such that $\text{Node}[\text{node-id}] = \widetilde{\text{Node}}[\text{root}_r]$. That is, the syntactical information stored on the node, e.g., gid , epoch , prop , orig , and so on, are identical.

PROOF. Let us first consider Item 1. $\widetilde{\text{Node}}[\widetilde{c}]$ is only created during a *Commit* or a *Process*. More concretely, $\widetilde{\text{Node}}[\widetilde{c}_0]$ is only created if $\widetilde{\mathcal{S}}$ assigns a detached root to \widetilde{c}_0 or a new commit node during a *Commit* or a *Process*. In either cases, a corresponding c_0^{ctxt} is output by \mathcal{S} and $\text{NodeID}[c_0^{\text{ctxt}}]$ is generated within $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. Moreover, due to the key-committing property of the SKE, the commit node assigned to such commits c_0^{ctxt} are identical.

Next, let us consider Item 2. Due to the key-committing property of the SKE, if such a c^{ctxt} existed, then we must have $\widetilde{\text{Node}}[\widetilde{c}_0] \neq \perp$. Thus, Item 2 follows by taking the contrapositive.

Finally, let us consider Item 3. Such \widetilde{w}_0 is generated when $\widetilde{\mathcal{S}}$ creates a new root detached root during a *Join* (corresponding to *Case 4* in (5)). Since \mathcal{S} creates a new commit node with the same syntactical information that $\widetilde{\mathcal{S}}$ used, Item 3 follows. \square

This concludes the proof of Lem. C.2. \square

C.3.4 From Hybrid 6 to 7: Proof of Lem. C.7. In this section, we prove the indistinguishability of remaining hybrid 6 and hybrid 7. As explained in the overview, we assume the adversary *cannot* inject bad randomness throughout these hybrids, i.e., RandCorr is always ‘good’.

Also, we slightly modify the description of \mathcal{S}_6 . We assume \mathcal{S}_6 executes the code of $\widetilde{\mathcal{S}}_6$ by itself instead of invoking $\widetilde{\mathcal{S}}_6$.

To prove Lem. C.7, we first formally define the behavior of simulator \mathcal{S}_7 in Hybrid 7; then we analyze the simulation provided by \mathcal{S}_7 provides an indistinguishable view to \mathcal{Z} as in Hybrid 6.

Part 1: Description of the Simulator \mathcal{S}_7 . We first explain the description of \mathcal{S}_7 . For simulation, \mathcal{S}_7 stores group encryption key encSecret for epoch node-id in the list $L_{\text{encSecret}}$.

\mathcal{S}_7 simulates group states G (which includes e.g., membership list and group secret keys) of node-id depending on the status of the parent epoch of node-id denoted by node-id_p .

Case (1). If $\text{safe}(\text{node-id}_p) = \text{false}$ when node-id is created, \mathcal{S}_7 knows the group state of node-id_p . Thus, \mathcal{S}_7 generates the group state of node-id from it following the protocol as in \mathcal{S}_6 .

Case (2). If $\text{safe}(\text{node-id}_p) = \text{true}$ when node-id is created, \mathcal{S}_7 defers the generation of group state of node-id to the following timing. Note that due to the restricted environment, these are the only subcase of Case (2), and only one of the following cases occurs for each node-id.

Case (2-1) safe(node-id) = true when proposal/commit messages are generated/processed at node-id: \mathcal{S}_7 sets $G.\text{gid} \leftarrow \text{Node}[\text{node-id}].\text{gid}$ and $G.\text{epoch} \leftarrow \text{Node}[\text{node-id}].\text{epoch} + 1$. Also, it chooses a random encSecret and stores $L_{\text{encSecret}}[\text{node-id}] \leftarrow \text{encSecret}$. In this case, it does not generate the other group state.

Case (2-2) Party id at node-id is corrupted (i.e., safe(node-id) becomes false) before proposal/commit messages are generated/processed at node-id: When \mathcal{A} corrupts id at node-id, \mathcal{S}_7 receives the following information from the functionality.

- id’s current mPKE secret key dk (stored in CurrDK array)
- id’s current signature signing key ssk (stored in SSK array)
- The epoch pointer node-id = $\text{Ptr}[\text{id}]$ (i.e., id’s current node in the history graph)

- The information stored in $\text{Node}[\text{node-id}]$
- Using this information, \mathcal{S}_7 simulates the group states. It first initializes the group state G as follows:
- $G.\text{gid} \leftarrow \text{Node}[\text{node-id}].\text{gid}$
 - $G.\text{epoch} \leftarrow \text{Node}[\text{node-id}].\text{epoch}$
 - $G.\text{member} \leftarrow \text{Node}[\text{node-id}].\text{mem}$
 - $G.\text{joinerSecret} \leftarrow_{\$} \{0, 1\}^K$
 - $G.\text{confTransHash-w.o-}'id_c' \leftarrow_{\$} \{0, 1\}^K$
 - $G.\text{confTransHash} \leftarrow H(G.\text{confTransHash-w.o-}'id_c', id_c)$, where $id_c := \text{Node}[\text{node-id}].\text{orig}$
 - $G.\text{certSvks}[*], G.\text{pendUpd}[*], G.\text{pendCom}[*] \leftarrow \emptyset$
 - $G.\text{id} \leftarrow id$
 - $G.\text{ssk} \leftarrow \text{ssk}$
 - $G.\text{member}[id].\text{dk} \leftarrow dk$

Then, \mathcal{S}_7 executes the following functions in order to generate group secrets and hash values.

- (1) $G.\text{memberHash} \leftarrow *derive\text{-member}\text{-hash}(G)$
- (2) $(G, \text{confKey}) \leftarrow *derive\text{-epoch}\text{-keys}(G, G.\text{joinerSecret})$
- (3) $\text{confTag} \leftarrow *gen\text{-conf}\text{-tag}(G, \text{confKey})$
- (4) $G \leftarrow *set\text{-interim}\text{-trans}\text{-hash}(G, \text{confTag})$

Since the previous epoch is secure and due to the restricted environment, \mathcal{A} cannot corrupt the past secure epochs (cf. lines 9-10 in Fig. 14). Thus, \mathcal{A} cannot distinguish the simulated group state from a real one since \mathcal{A} cannot check whether the simulated group states are consistent for the previous epoch. Finally, \mathcal{S}_7 stores $L_{\text{encSecret}}[\text{node-id}] \leftarrow G.\text{encSecret}$.

\mathcal{S}_7 answers each query made by the ideal functionality $\mathcal{F}_{\text{CGKA},7}^{\text{ctxt}}$ as in \mathcal{S}_6 except for the differences shown below. We only describe \mathcal{S}_7 when **safe** is true for the epoch \mathcal{Z} queries. This is because if **safe** is false, \mathcal{S}_7 knows the corresponding group state (see above) and obtains the same information \mathcal{S}_6 receives from $\mathcal{F}_{\text{CGKA},7}^{\text{ctxt}}$; thus \mathcal{S}_7 can simulate the protocol messages as \mathcal{S}_6 does.

Simulation on input (Create). \mathcal{S}_7 chooses gid at random and returns it to the functionality. Note that \mathcal{S}_7 does not generate other group information in this timing (see above for when and how group state is generated). $h(1 - \delta)$ -correct **Simulation on input**

(Propose, node-id, |id|, |act|). Let $\text{encSecret} := L_{\text{encSecret}}[\text{node-id}]$. \mathcal{S}_7 computes $\text{CT}_{\bar{p}} \leftarrow \text{SKE}.\text{Enc}(\text{encSecret}, 0^{\ell_{\bar{p}}})$ and sets $p^{\text{ctxt}} := (\text{gid}, \text{epoch}, \text{'proposal'}, \text{CT}_{\bar{p}})$, where gid and epoch is the group identity and the current epoch number of node-id .³⁴ It returns $(ack := \text{true}, \text{prop-id} := \perp, p^{\text{ctxt}})$ to the functionality. Note that $\ell_{\bar{p}}$, the length of \bar{p} , can be computed from the received information and public parameters (which determine the length of the hash value, mPKE ciphertext, signature, etc.).

Simulation on input (Commit, node-id, |id|, \bar{p}^{ctxt} , |svk|, |mem|). For each $p^{\text{ctxt}} \in \bar{p}^{\text{ctxt}}$, \mathcal{S}_7 checks whether it satisfies the following conditions.

- $\text{PropID}[p^{\text{ctxt}}] \neq \text{node-id}$, i.e., the received proposal was issued at a different epoch. We call the event $E_{\text{rej-1}}$.
- $\text{PropID}[p^{\text{ctxt}}] = \perp$, i.e., the received proposal was generated by the adversary. We call the event $E_{\text{rej-2}}$.

³⁴The simulator always knows gid and epoch for each node-id .

If one of the above condition holds for some $p^{\text{ctxt}} \in \bar{p}^{\text{ctxt}}$, \mathcal{S}_7 outputs $ack := \text{false}$.

Otherwise³⁵, \mathcal{S}_7 generates $\text{CT}_{\bar{c}_0} \leftarrow \text{SKE}.\text{Enc}(\text{encSecret}, 0^{\ell_{\bar{c}_0}})$ and $\bar{c}_i^{\text{ctxt}} \leftarrow \text{SKE}.\text{Enc}(\text{encSecret}, 0^{\ell_{\bar{c}_i}})$ for $\text{encSecret} := L_{\text{encSecret}}[\text{node-id}]$ and each $i \in [|\text{mem}|]$. Then, it sets $c_0^{\text{ctxt}} := (\text{gid}, \text{epoch}, \text{'commit'}, \text{CT}_{\bar{c}_0})$ and $\bar{c}^{\text{ctxt}} := \{\bar{c}_i^{\text{ctxt}}\}_{i \in [|\text{mem}|]}$, and returns $(ack := \text{true}, \text{node-id} := \perp, c_0^{\text{ctxt}}, \bar{c}^{\text{ctxt}})$ to the functionality, where gid and epoch is the group information of the epoch node-id . \mathcal{S}_7 stores $L_{\bar{c}^{\text{ctxt}}}[\text{node-id}] \leftarrow \bar{c}^{\text{ctxt}}$. Note that both the length of \bar{c}_0 and \bar{c} ($\ell_{\bar{c}_0}$ and $\ell_{\bar{c}}$) can be computed from the received information and public parameters.

Simulation on input (Welcome, kp_t , |gid|, |epoch|, |id|, |mem|). \mathcal{S}_7 receives this message if **safe** is true for the next epoch, where new members will join. Let ek_t be the mPKE encryption key in kp_t . \mathcal{S}_7 computes the following:

- $\text{ct} \leftarrow \text{mEnc}(ek_t, 0)$
- $\text{welcomeSecret}_{id_t} \leftarrow_{\$} \{0, 1\}^K$
- $\text{CT} \leftarrow \text{SKE}.\text{Enc}(\text{welcomeSecret}_{id_t}, 0^{\ell_{(\text{groupInfo}, \text{sig})}})$.

Note that $\ell_{(\text{groupInfo}, \text{sig})}$ can be computed from the received leakage information and public system parameters. \mathcal{S}_7 sets $\widehat{w}^{\text{ctxt}} := (id_t, (H(kp_t), \text{ct}), \text{CT})$ and returns $(ack := \text{true}, \widehat{w}^{\text{ctxt}})$ to the functionality. \mathcal{S}_7 also stores $L_{\widehat{w}^{\text{ctxt}}} \leftarrow (id_t, H(kp_t), \text{ct}, \text{CT})$.

Simulation on input (Process, node-id, index, c_0^{ctxt} , \bar{c}^{ctxt} , \bar{p}^{ctxt}). \mathcal{S}_7 first checks the following conditions and outputs $ack := \text{false}$ if one of the conditions holds.

- $\text{PropID}[p^{\text{ctxt}}] = \text{node-id}'$, $\text{NodeID}[c_0^{\text{ctxt}}] = \text{node-id}'$ or $\bar{c}^{\text{ctxt}} \in \text{Node}[\text{node-id}'].\text{vcom}$ such that $\text{node-id}' \neq \text{node-id}$, i.e., the received proposal/commit message was generated at a different epoch. We call the event $E_{\text{rej-1}}$.
- $\text{PropID}[p^{\text{ctxt}}] = \perp$, $\text{NodeID}[c_0^{\text{ctxt}}] = \perp$ or $\bar{c}^{\text{ctxt}} \notin \text{Node}[\text{node-id}'].\text{vcom}$ for all $\text{node-id}'$, i.e., the received proposal/commit was generated by the adversary. We call the event $E_{\text{rej-2}}$.

Otherwise³⁶, \mathcal{S}_7 returns $(ack := \text{true}, \perp, \perp, \perp)$.

Simulation on input (Join, id, $\widehat{w}^{\text{ctxt}}$). Let $(id_t, \text{kphash}_t, \text{ct}_t, \text{CT}_t) := \widehat{w}^{\text{ctxt}}$. (In the following, we assume $id_t = \text{id}$ and id succeeds to fetch kp_t such that $\text{kphash}_t = h(kp_t)$; otherwise both \mathcal{S}_6 and \mathcal{S}_7 return $ack := \text{false}$.) If $\widehat{w}^{\text{ctxt}}$ is generated by an honest committer (i.e., $*\text{succeed-wel}$ returns true), \mathcal{S}_7 returns $(ack := \text{true}, \perp, \perp, \perp, \perp, \perp)$. Else, \mathcal{S}_7 processes the welcome message as in \mathcal{S}_6 (i.e., following protocol description).

Part 2: Indistinguishability of Two Hybrids. We then prove the following lemma.

LEMMA C.7. *Hybrid 6 and Hybrid 7 are indistinguishable assuming the INDCCA security of mPKE and the IND-CCA security of SKE.*

PROOF. To show Lem. C.7, we consider the following sub-hybrids between Hybrid 6 and Hybrid 7.

Hybrid 6-0 := Hybrid 6. This is identical to Hybrid 6. We use the functionality $\mathcal{F}_{\text{CGKA},6}^{\text{ctxt}}$, and the simulator $\mathcal{S}_{6-0} := \mathcal{S}_6$.

³⁵This implies \bar{p}^{ctxt} only contains proposals such that $\text{PropID}[p^{\text{ctxt}}] = \text{node-id}$ and thus $*\text{succeed-com}$ returns true.

³⁶This implies $(c_0^{\text{ctxt}}, \bar{c}^{\text{ctxt}}, \bar{p}^{\text{ctxt}})$ is honestly issued at node-id and thus $*\text{succeed-proc}$ returns true.

Hybrid 6-1 In this hybrid, the simulator S_{6-1} is defined exactly as S_{6-0} except that it always outputs $ack := \text{false}$ to the functionality when the event $E_{\text{rej-1}}$ occurs.

Hybrid 6-2 In this hybrid, the simulator S_{6-2} is defined exactly as S_{6-1} except that it replaces encSecret and $\text{welcomeSecret}_{id_t}$ with a random value if **safe** is true.

Hybrid 6-3. In this hybrid, the simulator S_{6-3} is defined exactly as S_{6-2} except that it always outputs $ack := \text{false}$ when the event $E_{\text{rej-2}}$ occurs.

Hybrid 6-4. In this hybrid, the simulator S_{6-4} is defined exactly as S_{6-3} except that it replaces a mPKE ciphertext in welcome messages with a ciphertext of 0 if **safe** is true for the next epoch where new members join with the welcome messages.

Hybrid 6-5. In this hybrid, the simulator S_{6-5} is defined exactly as S_{6-4} except that it replaces SKE ciphertexts in welcome messages with encryption of zero-string if **safe** is true for the epoch where new members join with the welcome messages.

Hybrid 6-6. In this hybrid, the simulator S_{6-6} is defined exactly as S_{6-5} except that it replaces SKE ciphertexts in proposal and commit messages with encryption of zero-string if **safe** is true for the epoch where these messages are issued. Note that S_{6-6} is identical to S_7 .

Hybrid 6-7 = Hybrid 7. We replace the functionality $\mathcal{F}_{\text{CGKA},6}^{\text{ctxt}}$ with the functionality $\mathcal{F}_{\text{CGKA},7}^{\text{ctxt}}$. The simulator S_7 is defined by the above description. Since S_7 has simulated the protocol with the information given from $\mathcal{F}_{\text{CGKA},7}^{\text{ctxt}}$, Hybrid 6-6 and Hybrid 6-7 are identical.

From Lems. C.8 to C.11, C.13 and C.15 provided below, Hybrid 6-0 and Hybrid 6-6 are indistinguishable. Therefore, we conclude that Hybrid 6 and Hybrid 7 are indistinguishable. \square

C.3.5 From Hybrid 6-0 to 6-1: Proof of Lem. C.8.

LEMMA C.8. *Hybrid 6-0 and Hybrid 6-1 are indistinguishable assuming the key-committing property of SKE.*

PROOF. The difference between S_{6-0} and S_{6-1} is S_{6-1} always outputs $ack := \text{false}$ when the event $E_{\text{rej-1}}$ occurs. In other words, if S_{6-0} outputs $ack := \text{true}$ when $E_{\text{rej-1}}$ occurs, \mathcal{Z} can distinguish the two hybrids. We show that, if \mathcal{Z} can distinguish the two hybrids, then there exists an adversary \mathcal{B} that breaks the key-committing property of SKE.

\mathcal{B} simulates for \mathcal{Z} the interaction with $\mathcal{F}_{\text{AS}}^{\text{IW}}$ and $\mathcal{F}_{\text{KS}}^{\text{IW}}$, and the protocol execution of each party as in S_{6-0} . Assume $E_{\text{rej-1}}$ occurs for some SKE ciphertext CT and \mathcal{Z} distinguish the two hybrids when CT is processed. By the condition of $E_{\text{rej-1}}$, there exists the encryption key k (encSecret or $\text{welcomeSecret}_{id_t}$) used to generate CT and the other encryption key $k' \neq k$ that correctly decrypts CT. This means CT can be correctly decrypted with both k and k' , which implies \mathcal{B} can break the key-committing property of SKE. This contradicts the assumption that SKE has key-committing property. Therefore, both S_{6-0} and S_{6-1} always outputs $ack := \text{false}$ when $E_{\text{rej-1}}$ occurs. \square

C.3.6 From Hybrid 6-1 to 6-2: Proof of Lem. C.9.

LEMMA C.9. *Hybrid 6-1 and Hybrid 6-2 are indistinguishable assuming mPKE is Chained CmpKE conforming GSD secure.*

PROOF. The proof is identical to the proof in [37, Lemma E.29]. To prove the indistinguishability of the two hybrids, we gradually replace each encSecret and $\text{welcomeSecret}_{id_t}$ with a random value if **safe** is true when a propose/commit/welcome message is created. Similar to the proof of randomness of appSecret provided in [37, Lemma E.29], we can show that, if \mathcal{Z} can distinguish whether encSecret or $\text{welcomeSecret}_{id_t}$ are real or random, it can be used to break the Chained CmpKE conforming GSD security of mPKE. Note that the value of **safe** is fixed when a propose/commit/welcome message is created at the epoch (cf. `*mark-content-hidden-epoch`), and the adversary is restricted from colluding to change this value (cf. line 11 in Fig. 14). Thus, we can construct a reduction \mathcal{B} as in [37, Lemma E.29]. Therefore, if mPKE is Chained CmpKE conforming GSD secure, Hybrid 6-1 and Hybrid 6-2 are indistinguishable. \square

C.3.7 From Hybrid 6-2 to 6-3: Proof of Lem. C.10.

LEMMA C.10. *Hybrid 6-2 and Hybrid 6-3 are indistinguishable assuming SKE is the IND-CCA secure.*

PROOF. The difference between S_{6-2} and S_{6-3} is S_{6-3} always outputs $ack := \text{false}$ when the event $E_{\text{rej-2}}$ occurs³⁷. In other words, \mathcal{Z} distinguishes the two hybrids if the adversary produces malicious ciphertexts that can be decrypted correctly. Such a malicious ciphertext may contain (1) a malicious plain message (i.e., the message generated by the adversary) or (2) an honest plain message (i.e., the message generated by an honest party at the same epoch). For the former case, we already proved that the adversary cannot produce acceptable messages without knowing MAC key (cf. Hybrid 5). Thus, the previous simulator also outputs $ack := \text{false}$ in the former case. Therefore, we care about the latter case. We show that, if \mathcal{Z} can distinguish the two hybrids when the adversary produces a malicious ciphertext that contains an honestly generated plain message, then we can construct an adversary \mathcal{B} that breaks the IND-CCA security of SKE.

\mathcal{B} simulates for \mathcal{Z} the interaction with $\mathcal{F}_{\text{AS}}^{\text{IW}}$ and $\mathcal{F}_{\text{KS}}^{\text{IW}}$, and the protocol execution of each party as in S_{6-2} except for encrypting messages. At the beginning of the game, \mathcal{B} chooses $i \in [Q]$ at random (where Q is the total number of honestly generated epochs) and hopes that \mathcal{Z} distinguishes the hybrids when the event occurs in the i -th honest epoch. That is, we assume that, when a party processes some ciphertext CT with the i -th honest encryption key, S_{6-2} outputs $ack := \text{true}$, but S_{6-3} outputs $ack := \text{false}$. (\mathcal{B} succeeds to guess such an epoch with probability $1/Q$.) \mathcal{B} embed the challenge SKE key of the IND-CCA game to the i -th encSecret . Note that if **safe** is true for the corresponding epoch, the encryption key is chosen at random due to the modification we made in Hybrid 6-2; thus the challenge key can be embedded. In addition, the adversary cannot corrupt the challenge key after the key is used for encryption since the adversary's corruption is restricted (cf. `*mark-content-hidden-epoch` function). When \mathcal{B} encrypts messages with the i -th encSecret (i.e., the challenge key), it uses the encryption oracle \mathcal{LR} by setting M_0 as the actual message and M_1 as a random message with the same length. \mathcal{B} uses its decryption

³⁷The event $E_{\text{rej-2}}$ only occurs in epochs where **safe** is true.

oracle when it wants to decrypt a ciphertext³⁸. If the decryption result is identical to the message M_b , \mathcal{B} outputs the bit b to the challenger of the IND-CCA game. Else, the decryption result is different from the messages sent to the encryption oracle, it outputs $ack := \text{false}$. Note that for other keys, \mathcal{B} simulates protocol as in the previous hybrid.

We can verify that \mathcal{B} wins the IND-CCA game if the adversary can produce a ciphertext that contains an honestly generated message. Note that the messages M_{1-b} are information-theoretically hidden from the adversary (and \mathcal{Z}). Hence, if \mathcal{Z} can distinguish the two hybrids, \mathcal{B} can break the IND-CCA security of SKE. This contradicts the assumption that SKE is the IND-CCA secure. Therefore, \mathcal{Z} cannot distinguish the two hybrids. In other words, both \mathcal{S}_{6-2} and \mathcal{S}_{6-3} outputs $ack = \text{false}$ when $E_{\text{rej-2}}$ occurs. \square

C.3.8 From Hybrid 6-3 to 6-4: Proof of Lem. C.11.

LEMMA C.11. *Hybrid 6-3 and Hybrid 6-4 are indistinguishable assuming mPKE is IND-CCA secure.*

PROOF. We assume \mathcal{Z} creates at most W welcome messages by `Commit` query. To show Lem. C.11, we consider the following sub-hybrids between Hybrid 6-3 and Hybrid 6-4.

Hybrid 6-3-0 := Hybrid 6-3. This is identical to Hybrid 6-3. The simulator $\mathcal{S}_{6-3-0} = \mathcal{S}_{6-3}$ generates protocol messages following the protocol procedures.

Hybrid 6-3- i . i runs through $[W]$. The simulator \mathcal{S}_{6-3-i} is defined exactly as $\mathcal{S}_{6-3-(i-1)}$ except that, when it simulates the i -th welcome message, if the `safe` predicates is true for the next epoch the welcome message is associated with, it encrypts zero-strings instead of the joiner secret with the new member's mPKE encryption key. Note that we count welcome messages in the order the simulator creates. We show in Lem. C.12 that Hybrid 6-3- $(i-1)$ and Hybrid 6-3- i are indistinguishable.

Hybrid 6-4. This is identical to Hybrid 6-3- W . In this hybrid, mPKE ciphertexts in welcome messages issued for secure epochs (i.e., `safe` is true) are replaced with the encryption of zero-string.

The indistinguishability between Hybrid 6-3-0 and Hybrid 6-3- Q is derived by applying Lem. C.12 for all $i \in [W]$. Therefore, we conclude that Hybrid 6-3 and Hybrid 6-3 are indistinguishable. \square

C.3.9 Proof of Lem. C.12: From Hybrid 6-3- $(i-1)$ to 6-3- i .

LEMMA C.12. *Hybrid 6-3- $(i-1)$ and Hybrid 6-3- i are indistinguishable assuming mPKE is IND-CCA secure³⁹.*

PROOF. The difference between $\mathcal{S}_{6-3-(i-1)}$ and \mathcal{S}_{6-3-i} is, when generating the i -th welcome message, if the next epoch is secure (i.e., `safe` is true for the next epoch), \mathcal{S}_{6-3-i} replaces the mPKE ciphertext in the i -th welcome message with a ciphertext of 0 encrypted with the new member's mPKE encryption key (the simulator receives the new member's key package from the functionality). We show that, if \mathcal{Z} can distinguish the two hybrids, then there exists an adversary

\mathcal{B} that breaks IND-CCA security of mPKE. We first explain the description of \mathcal{B} and then evaluate \mathcal{B} 's advantage.

\mathcal{B} simulates for \mathcal{Z} the interaction with $\mathcal{F}_{\text{AS}}^{\text{IW}}$ and $\mathcal{F}_{\text{KS}}^{\text{IW}}$, and the protocol execution as in $\mathcal{S}_{6-3-(i-1)}$ except for generating the i -th welcome message. Let ek^* be the challenge mPKE encryption key provided by the IND-CCA game. Observe that honest key packages are generated on `register-kp` queries to $\mathcal{F}_{\text{KS}}^{\text{IW}}$. We assume \mathcal{Z} issues at most Q `register-kp` queries. At the beginning of the game, \mathcal{B} chooses an index $J \in [Q]$ at random, embeds the challenge key in the J -th `register-kp` query, and hopes that the J -th key package will be used to add the i -th welcome message. (\mathcal{B} succeeds to guess with probability $1/Q$.) For other `register-kp` queries, \mathcal{B} generates key packages following the description of $\mathcal{F}_{\text{KS}}^{\text{IW}}$. If \mathcal{B} decrypts a ciphertext with the challenge decryption key, \mathcal{B} uses the decryption oracle provided by the IND-CCA game.

Assume the i -th welcome message is created with the J -th key package and `safe` is true for the new epoch where the i -th welcome message is created. Let `joinerSecret` be the corresponding joiner secret. \mathcal{B} outputs $M_0 := \text{joinerSecret}$ and $M_1 := 0$ to IND-CCA game challenger and receives the challenge ciphertext (ct_0^*, \widehat{ct}^*) . \mathcal{B} uses it as the mPKE ciphertext in the i -th welcome message. Note that, since `safe` is true, the adversary has not corrupted the challenge key, and it is restricted from colluding to compute the challenge key (cf. `*mark-content-hidden-epoch` function). Thus, \mathcal{B} never corrupts the challenge key.

We finally analyze \mathcal{B} 's advantage. If the challenger returns the ciphertext of the joiner secret, \mathcal{Z} 's view is identical to Hybrid 6-3- $(i-1)$; else, i.e., the challenger returns the ciphertext of 0, \mathcal{Z} 's view is identical to Hybrid 6-3- i . Hence, if \mathcal{Z} distinguishes Hybrid 6-3- $(i-1)$ and Hybrid Hybrid 6-3- i with non-negligible probability, \mathcal{B} wins the IND-CCA game with non-negligible probability by using \mathcal{Z} 's output. This contradicts the assumption that mPKE is IND-CCA secure. Therefore, Hybrid 6-3- $(i-1)$ and Hybrid Hybrid 6-3- i are indistinguishable. \square

C.3.10 From Hybrid 6-4 to 6-5: Proof of Lem. C.13.

LEMMA C.13. *Hybrid 6-4 and Hybrid 6-5 are indistinguishable assuming SKE is the IND-CCA secure.*

PROOF. We assume \mathcal{Z} creates at most W welcome messages by `Commit` query. To show Lem. C.13, we consider the following sub-hybrids between Hybrid 6-4 and Hybrid 6-5.

Hybrid 6-4-0 := Hybrid 6-4. This is identical to Hybrid 6-4. The simulator $\mathcal{S}_{6-4-0} = \mathcal{S}_{6-4}$ generates protocol messages following the protocol procedures.

Hybrid 6-4- i . i runs through $[W]$. The simulator \mathcal{S}_{6-4-i} is defined exactly as $\mathcal{S}_{6-4-(i-1)}$ except that, when it simulates the i -th welcome message, if `safe` is true for the next epoch the welcome message is associated with, it generates a ciphertext as $\text{CT} \leftarrow \text{SKE.Enc}(\text{welcomeSecret}_{\text{id}_r}, 0^{\ell(\text{groupInfo}, \text{sig})})$. Note that we count welcome messages in the order the simulator creates. We show in Lem. C.14 that Hybrid 6-4- $(i-1)$ and Hybrid 6-4- i are indistinguishable.

Hybrid 6-3. This is identical to Hybrid 6-2- W . In this hybrid, SKE ciphertext in welcome messages issued for secure epochs (i.e., `safe` is true) is replaced with a random string.

³⁸By definition of the simulator, \mathcal{B} only sends ciphertexts that are not produced by the encryption oracle.

³⁹We use the IND-CCA game with $N = 1$ (cf. Fig. 8).

The indistinguishability between Hybrid 6-4-0 and Hybrid 6-4- Q is derived by applying Lem. C.14 for all $i \in [W]$. Therefore, we conclude that Hybrid 6-4 and Hybrid 6-3 are indistinguishable. \square

C.3.11 Proof of Lem. C.14: From Hybrid 6-4-($i - 1$) to 6-4- i .

LEMMA C.14. *Hybrid 6-4-($i - 1$) and Hybrid 6-4- i are indistinguishable assuming SKE is the IND-CCA secure.*

PROOF. The difference between $\mathcal{S}_{6-4-(i-1)}$ and \mathcal{S}_{6-4-i} is, when generating the i -th welcome message, if the next epoch is secure (i.e., **safe** is true for the next epoch), \mathcal{S}_{6-4-i} replaces the SKE ciphertext in the i -th welcome message with an encryption of zero-string. We show that, if \mathcal{Z} can distinguish the two hybrids, then there exists an adversary \mathcal{B} that breaks the IND-CCA security of SKE. We first explain the description of \mathcal{B} and then evaluate \mathcal{B} 's advantage.

\mathcal{B} simulates for \mathcal{Z} the interaction with \mathcal{F}_{AS}^{IW} and \mathcal{F}_{KS}^{IW} , and the protocol execution as in $\mathcal{S}_{6-4-(i-1)}$ except for generation of SKE ciphertext in the i -th welcome message if **safe** is true for the new epoch corresponding to the i -th welcome message. \mathcal{B} embed the challenge SKE key of the IND-CCA game to $\text{welcomeSecret}_{id_i}$, which is used to encrypt the i -th welcome message. If **safe** is true for the corresponding epoch, welcomeSecret is chosen at random due to the modification we made in Hybrid 6-1; thus the challenge key can be embedded. In addition, the adversary cannot corrupt the challenge key after messages are encrypted with the key since the adversary's corruption is restricted (cf. `*mark-content-hidden-epoch` function).

When \mathcal{B} encrypts the contents M in the i -th welcome message, it queries $\text{CT} := \mathcal{LR}(M_0 := (\text{groupInfo}, \text{sig}), M_1 := 0^{\ell(\text{groupInfo}, \text{sig})})$ and uses the challenge ciphertext CT as the ciphertext in the i -th welcome message. When \mathcal{B} decrypts ciphertexts different from challenge ciphertexts with the challenge key, it uses its decryption oracle.

We finally analyze \mathcal{B} 's advantage. If the oracle \mathcal{LR} returns a ciphertext of the real contents M_0 , \mathcal{Z} 's view is identical to Hybrid 6-4-($i - 1$); else, i.e., the oracle \mathcal{LR} returns a ciphertext of zero-string, \mathcal{Z} 's view is identical to Hybrid 6-4- i . Hence, if \mathcal{Z} distinguishes Hybrid 6-4-($i - 1$) and Hybrid Hybrid 6-4- i with non-negligible probability, \mathcal{B} breaks the IND-CCA security of SKE with non-negligible probability by using \mathcal{Z} 's output. This contradicts the assumption that SKE is IND-CCA secure. Therefore, Hybrid 6-4-($i - 1$) and Hybrid Hybrid 6-4- i are indistinguishable. \square

C.3.12 From Hybrid 6-5 to 6-6: Proof of Lem. C.15.

LEMMA C.15. *Hybrid 6-5 and Hybrid 6-6 are indistinguishable assuming SKE is IND-CCA secure.*

PROOF. We assume \mathcal{Z} creates at most Q epochs (i.e., commit nodes). To show Lem. C.15, we consider the following sub-hybrids between Hybrid 6-5 and Hybrid 6-6.

Hybrid 6-5-0 := Hybrid 6-5. This is identical to Hybrid 6-5. We use the functionality $\mathcal{F}_{CGKA,6}^{\text{ctxt}}$ and the simulator $\mathcal{S}_{6-5-0} = \mathcal{S}_{6-5-0}$.

Hybrid 6-5- i . i runs through $[Q]$. The simulator \mathcal{S}_{6-5-i} is defined exactly as $\mathcal{S}_{6-5-(i-1)}$ except that, when a party issues proposal or commit messages at the i -th epoch, if **safe** is true for the epoch, it encrypts a zero-string instead of the real

contents in the non-encrypted proposal/commit message. Note that we count epochs in the order in which Propose or Commit is called. We show in Lem. C.16 that Hybrid 6-5-($i - 1$) and Hybrid 6-5- i are indistinguishable.

Hybrid 6-6. This is identical to Hybrid 6-5- Q . In this hybrid, ciphertexts in proposal and commit messages issued at secure epochs (epochs such that **safe** is true) are replaced with random bit-strings.

The indistinguishability between Hybrid 6-5-0 and Hybrid 6-5- Q is derived by applying Lem. C.16 for all $i \in [Q]$. Therefore, we conclude that Hybrid 6-5 and Hybrid 6-6 are indistinguishable. \square

C.3.13 Proof of Lem. C.16: From Hybrid 6-5-($i - 1$) to 6-5- i .

LEMMA C.16. *Hybrid 6-5-($i - 1$) and Hybrid 6-5- i are indistinguishable assuming SKE is IND-CCA secure.*

PROOF. The difference between $\mathcal{S}_{6-5-(i-1)}$ and \mathcal{S}_{6-5-i} is, when generating proposal or commit messages at the i -th epoch such that **safe** is true, \mathcal{S}_{6-5-i} encrypts a zero-string instead of the real contents. We show that, if \mathcal{Z} can distinguish the two hybrids, then there exists an adversary \mathcal{B} that breaks the IND-CCA security of SKE. We first explain the description of \mathcal{B} . Then we evaluate \mathcal{B} 's advantage.

\mathcal{B} simulates for \mathcal{Z} the interaction with \mathcal{F}_{AS}^{IW} and \mathcal{F}_{KS}^{IW} , and the protocol execution of each party as in $\mathcal{S}_{6-5-(i-1)}$ except for encryption of proposal and commit messages generated at the i -th epoch if **safe** is true. \mathcal{B} embed the challenge SKE key of the IND-CCA game to encSecret of the i -th epoch (if **safe** is true for the i -th epoch, encSecret is chosen at random due to the modification we made in Hybrid 6-2; thus the challenge key can be embedded). In addition, the adversary cannot corrupt the challenge key after messages are encrypted with the key since the adversary's corruption is restricted (cf. `*mark-content-hidden-epoch` function). When \mathcal{B} encrypts contents M in a proposal or commit message at the i -th epoch, it queries $\text{CT} := \mathcal{LR}(M_0 := M, M_1 := 0^{|M|})$. Note that due to the modification we made in Hybrid 6-3, \mathcal{B} can reject any messages which do not generate with the i -th encryption key or generated by the adversary. Thus it does not need to use the decryption oracle.

We finally analyze \mathcal{B} 's advantage. If the encryption oracle \mathcal{LR} returns ciphertexts of M , \mathcal{Z} 's view is identical to Hybrid 6-5-($i - 1$); else, i.e., the \mathcal{LR} oracle returns ciphertexts of zero-string, \mathcal{Z} 's view is identical to Hybrid 6-5- i . In addition, the i -th encryption secret is hidden from the adversary. Hence, if \mathcal{Z} distinguishes Hybrid 6-5-($i - 1$) and Hybrid Hybrid 6-5- i with non-negligible probability, \mathcal{B} wins the IND-CCA game with non-negligible probability by using \mathcal{Z} 's output. This contradicts the assumption that SKE is IND-CCA secure. Therefore, Hybrid 6-5-($i - 1$) and Hybrid Hybrid 6-5- i are indistinguishable. \square

D METADATA-HIDING CGKA: DEFINITION

In this section, we propose a UC security model capturing the security of the 1st, 2nd, and 3rd layers (i.e., group secret keys, static and *dynamic* metadata) by defining a new ideal functionality $\mathcal{F}_{CGKA}^{\text{mh}}$.

D.1 An Overview of \mathcal{F}_{CGKA}^{mh}

The ideal functionality \mathcal{F}_{CGKA}^{mh} is formally defined in Figs. 36 to 38, 40 and 41, along with the helper functions in Fig. 39 to aid the readability. As it can be checked, \mathcal{F}_{CGKA}^{mh} shares a large portion of its code with $\mathcal{F}_{CGKA}^{ctxt}$. This is because \mathcal{F}_{CGKA}^{mh} by definition also models an ideal functionality of a CGKA securing the 1st & 2nd layers. For better readability, we outsourced the description of \mathcal{F}_{CGKA}^{mh} that is non-essential to the dynamic metadata to $\mathcal{F}_{CGKA}^{ctxt}$. This should not be misunderstood as \mathcal{F}_{CGKA}^{mh} making oracle calls to $\mathcal{F}_{CGKA}^{ctxt}$ — the former simply reuses the codes of $\mathcal{F}_{CGKA}^{ctxt}$.

While \mathcal{F}_{CGKA}^{mh} allows the adversary to corrupt the parties (and the server) as in $\mathcal{F}_{CGKA}^{ctxt}$, we do not model adversarial controlled randomness (i.e., RandCorr is always set to ‘good’). As the first work to formally capture the dynamic metadata layer, we opted for simplicity and better readability. We leave it as future work to allow such types of attacks. We now explain the ideal functionality \mathcal{F}_{CGKA}^{mh} in more detail.

States. Similarly to $\mathcal{F}_{CGKA}^{ctxt}$, \mathcal{F}_{CGKA}^{mh} maintains a history graph. As a custom, we assume one honest group is created by the designated party $id_{creator}$. Such group is assigned to the main root node- $id = 0$. Since \mathcal{F}_{CGKA}^{mh} does not allow the adversary to manipulate the party’s randomness, RandCorr[id] cannot be switched to ‘bad’.

Other than the history graph, \mathcal{F}_{CGKA}^{mh} also maintains three databases: PropDB[*,*], ComDB[*,*], and WelDB[*]. The proposal database PropDB[gid, epoch] stores proposals issued at (gid, epoch). The welcome database WelDB[id] stores a welcome message sent to id. Since \mathcal{F}_{CGKA}^{mh} processes a join query only when Ptr[id] = \perp (i.e., a party id is not assigned to any group), WelDB[id] stores only one welcome message for each id, and overwrites the old one if a new welcome message is published to id. Finally, the commit database ComDB[gid, epoch] stores the status of the group gid at epoch. Depending on the status of the group, it takes one of the following five values:

- ComDB[gid, epoch] = \perp : It indicates that the epoch has not been initialized yet.
- ComDB[gid, epoch] = (\top , node-id): It indicates that the epoch has been initialized by an honest party located at the commit node node-id (of the history graph), and no commit has been issued at epoch.
- ComDB[gid, epoch] = ((c_0, \vec{c}) , node-id): It indicates that the epoch has been initialized by an honest party located at the commit node node-id, and a commit (c_0, \vec{c}) has been issued at epoch.
- ComDB[gid, epoch] = (\top , ‘adv’): It indicates that the epoch has been initialized by an adversary, and no commit has been issued at epoch.
- ComDB[gid, epoch] = ((c_0, \vec{c}) , ‘adv’): It indicates that the epoch has been initialized by an adversary, and a commit (c_0, \vec{c}) has been issued at epoch.

Looking ahead, (assuming the server Sv is honest) ComDB[gid, epoch] is initialized with \perp . If some party at epoch $- 1$ creates a commit, ComDB[gid, epoch] is initialized to (\top , *), where * depends on whether the party was honest or corrupt. Once the proposals in at epoch, i.e., PropDB[gid, epoch], are committed, then ComDB[gid, epoch]

stores that commit and this freezes the epoch, and initializes a new ComDB[gid, epoch + 1] = \perp . This models the fact that a commit is generated only once per epoch when the server is honest. We refer the readers to see Sec. 4.2 for a pictorial example.

Interfaces. \mathcal{F}_{CGKA}^{mh} offers similar interfaces to the parties as $\mathcal{F}_{CGKA}^{ctxt}$. A party can create a group, create proposal, commit, or welcome messages, process these messages, and obtain group secret keys. In addition to these functionalities, it also offers the functionalities of publishing and fetching messages, which models the message delivery between a (possibly malicious) party and the server. To formally capture this, \mathcal{F}_{CGKA}^{mh} makes the existence of the server explicit — recall that the server was implicit in $\mathcal{F}_{CGKA}^{ctxt}$ and abstracted away as a malicious network that injects arbitrary message. Specifically \mathcal{F}_{CGKA}^{mh} models an *honest-but-curious* server. When the server is corrupted by the adversary, we end up with the same functionality as $\mathcal{F}_{CGKA}^{ctxt}$.

In the following, we explain each of the functionalities in more detail.

D.2 Functions Used by Legitimate Parties

Create and register group (See Fig. 36). The designated party $id_{creator}$ first creates the group as in $\mathcal{F}_{CGKA}^{ctxt}$ (cf. *Group creation* in App. B.2). $id_{creator}$ then registers the group to the server via the subroutine RegisterGroup. $\mathcal{F}_{CGKA}^{ctxt}$ first checks that the party $id_{creator}$ is located at epoch = 0 (i.e., Ptr[id_{creator}] = 0). Then, \mathcal{F}_{CGKA}^{mh} informs the adversary that a new group with (gid, 0) is being registered.

- If the server is honest (i.e., ServerStat = ‘good’), the ideal server checks that the group with gid has not been registered yet. If so, the server initializes the group as ComDB[gid, epoch] ← (\top , 0)⁴⁰ and returns *accept* = true to the party. Else, the registration is rejected and the party receives *accept* = false. The ideal server outputs (*accept*, gid, epoch = 0) (to the environment \mathcal{Z}). Since parties are assumed to access the server via a client-anonymous authenticated channel, the server never outputs the accessing party’s identity.
- If the server is malicious, the adversary specifies the protocol result *accept'*, and it is outputted to id. This means the malicious server can decide the protocol result arbitrarily.

Create and publish proposals (See Fig. 36). When a party id create a proposal, \mathcal{Z} invokes \mathcal{F}_{CGKA}^{mh} on input (Propose, act). \mathcal{F}_{CGKA}^{mh} first creates a proposal message p as in $\mathcal{F}_{CGKA}^{ctxt}$ (cf. *Creating proposals* in App. B.2). Then, id publishes p to the server via the subroutine PublishProposal.

- If the server is honest (i.e., ServerStat = ‘good’), \mathcal{F}_{CGKA}^{mh} gives (gid, epoch, p) to the adversary, where (gid, epoch) is the destination group identity and epoch. The party accesses the server via the client-anonymous authenticated channel, and therefore, the adversary does not receive the party’s identity. The adversary reports whether the protocol succeeds or not by setting *accept'* to true or false. \mathcal{F}_{CGKA}^{mh} then determines the party and server’s output as follows.

⁴⁰The RegisterGroup subroutine only occurs when the main group (which is assigned to node-id = 0) is registered.

- If $\text{ComDB}[\text{gid}, \text{epoch}] = (\top, \text{node-id})$ and $\text{node-id} = \text{Ptr}[\text{id}]$, the ideal party always outputs $\text{accept} := \text{true}$. This models correctness: if a party holds the same group state used to initialize node-id , the server must accept the published proposal.
- If $\text{ComDB}[\text{gid}, \text{epoch}] = (\top, \text{'adv'})$ and $\text{accept}' = \text{true}$, the ideal party outputs $\text{accept} := \text{true}$. This models the fact that when the destination $(\text{gid}, \text{epoch})$ is initialized by the adversary, we let the adversary decide if the ideal server should accept the published proposal or not. Looking at the example from Sec. 4.2, if $\text{ComDB}[\text{gid}, \text{epoch}]$ was initialized with a fake group statement gvk , then the adversary can deliberately make the publish proposal fail.
- Otherwise, the ideal party outputs $\text{accept} := \text{false}$. This case occurs if $\text{accept}' = \text{false}$ (i.e., the adversary decides to reject), $\text{ComDB}[\text{gid}, \text{epoch}] \neq (\top, *)$ (i.e., the destination has not been initialized or a commit has been issued), or $\text{ComDB}[\text{gid}, \text{epoch}] = (\top, \text{node-id})$ but $\text{node-id} \neq \text{Ptr}[\text{id}]$ (i.e., the party is located in a different commit node).

The ideal server outputs the destination $(\text{gid}, \text{epoch})$ and the published proposal p to \mathcal{Z} . Since parties access the server via a client-anonymous authenticated channel, the server never outputs the accessing party's identity.

- If the server is malicious, $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ gives $\text{Ptr}[\text{id}]$ to the adversary, where $\text{Ptr}[\text{id}] = \text{node-id}$ is the current commit node on the history graph, which id is located at. The reason why we don't give $(\text{gid}, \text{epoch})$ as above is that when the server is malicious, it can arbitrary fork the group. In such a case, $(\text{gid}, \text{epoch})$ is not enough to identify the location of the party. Note that in the real protocol, the malicious server gets to learn which fork a party is in since the party will try to access the server using the group state. Hence, $\text{Ptr}[\text{id}]$ correctly models what the malicious server learns in the real world. Finally, the adversary specifies the protocol result accept' , and it is outputted to id . This means the malicious server can decide the protocol result arbitrarily.

Create and publish commits (See Fig. 36). When a party id wants to issue a commit message, \mathcal{Z} invokes $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ on input $(\text{Commit}, \text{svk})$. The description composes of two parts.

Fetch proposals. id must first fetch the list of proposals from the server via the subroutine FetchProposals .

- If the server is honest (i.e., $\text{ServerStat} = \text{'good'}$), $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ gives $(\text{gid}, \text{epoch})$ to the adversary, where $(\text{gid}, \text{epoch})$ is the destination. Since the party accesses to the server via a client-anonymous authenticated channel, and therefore the server does not receive the party's identity. The adversary reports whether the protocol succeeds or not by setting accept' to true or false, and specifies the proposals \vec{p}' . $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ then determines the party and server's output as follows.
 - If $\text{ComDB}[\text{gid}, \text{epoch}] = (\top, \text{node-id})$ and $\text{node-id} = \text{Ptr}[\text{id}]$, the ideal party always outputs $\text{accept} := \text{true}$ and $\vec{p} := \text{PropDB}[\text{gid}, \text{epoch}]$. This models correctness: if a party holds the same group state used to initialize node-id , the server must accept and returns the proposals stored on the database.

- If $\text{ComDB}[\text{gid}, \text{epoch}] = (\top, \text{'adv'})$ and $\text{accept}' = \text{true}$, the ideal party outputs $\text{accept} := \text{true}$ and $\vec{p} := \text{PropDB}[\text{gid}, \text{epoch}]$. This models the fact that when the destination $(\text{gid}, \text{epoch})$ is initialized by the adversary, we let the adversary decide if the ideal server should return the stored proposals or not.
- Otherwise, the ideal party outputs $\text{accept} := \text{false}$. This case occurs if $\text{accept}' = \text{false}$ (i.e., the adversary decides to reject) or $\text{ComDB}[\text{gid}, \text{epoch}] = \perp$ (i.e., the destination has not been initialized).

In all three cases, the ideal server outputs the destination $(\text{gid}, \text{epoch})$ to \mathcal{Z} . Since parties access the server via a client-anonymous authenticated channel, the server never outputs the accessing party's identity.

- If the server is malicious, $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ gives $\text{Ptr}[\text{id}]$ to the adversary. The adversary specifies the protocol result accept' and the proposals \vec{p}' , and they are outputted to id . This means the malicious server can return an arbitrary message to the parties. This is consistent with $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ which allows an adversary to inject malicious proposals.

If id successfully fetches the proposals (i.e., id receives $\text{accept} = \text{true}$), then id creates a corresponding commit (c_0, \vec{c}) and welcome messages $\vec{w} = \{\vec{w}\}$ as in $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ (cf. *Committing to proposals* in App. B.2).

Publish commit and welcome messages. id finally publishes the commit and welcome messages to the server. We first describe how it publishes a commit. Before publishing the commit message, the party permutes \vec{c} to \vec{c}_{perm} using the function *permute-commit . This effectively makes the party's identity and index of \vec{c} unlikable. Note that the *permute-commit function uses a random permutation (if **safe** is true for the id 's current epoch) or a permutation the adversary chooses (if **safe** is false). This means if **safe** is true the permuted index will look random from the adversary.

- If the server is honest (i.e., $\text{ServerStat} = \text{'good'}$), $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ gives $(\text{gid}, \text{epoch}, c_0, \vec{c}_{\text{perm}})$ to the adversary. The party accesses the server via a client-anonymous authenticated channel, and therefore, the adversary does not receive the party's identity. The adversary reports whether the protocol succeeds or not by setting accept' to true or false. $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ then determines the party and server's output as follows.
 - If $\text{ComDB}[\text{gid}, \text{epoch}] = (\top, \text{node-id})$ and $\text{node-id} = \text{Ptr}[\text{id}]$, the ideal party always outputs $\text{accept} := \text{true}$. This models correctness: if a party holds the same group state used to initialize node-id , the server must accept the published commit message.
 - If $\text{ComDB}[\text{gid}, \text{epoch}] = (\top, \text{'adv'})$ and $\text{accept}' = \text{true}$, the ideal party outputs $\text{accept} := \text{true}$. This models the fact that if the destination $(\text{gid}, \text{epoch})$ is initialized by the adversary, we let the adversary decide if the ideal server accepts the published commit.
 - Otherwise, the ideal party outputs $\text{accept} := \text{false}$. This case occurs if $\text{accept}' = \text{false}$ (i.e., the adversary decides to reject), $\text{ComDB}[\text{gid}, \text{epoch}] \neq (\top, *)$ (i.e., the destination has not been initialized or a commit message has been issued), or $\text{ComDB}[\text{gid}, \text{epoch}] = (\top, \text{node-id})$ but

node-id \neq Ptr[id] (i.e., the party is located in a different commit node).

In all three cases, the ideal server outputs the destination (gid, epoch) and the published commit $(c_0, \tilde{c}_{\text{perm}})$ to \mathcal{Z} . Since parties access the server via a client-anonymous authenticated channel, the server never outputs the accessing party's identity.

- If the server is malicious, $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ gives (Ptr[id], $c_0, \tilde{c}_{\text{perm}}$) to the adversary. The adversary specifies the protocol result accept' , and it is outputted to id. This means the malicious server can decide the protocol result arbitrarily.

Finally, id publishes the welcome messages $\widehat{w} \in \widehat{w}$ via the subroutine `PublishWelcome`. To hide the relationship among welcome messages, id publish \widehat{w} separately. $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ informs the adversary of publishing a welcome message by giving (id_t, \widehat{w}) , where id_t is the intended recipient of \widehat{w} . Note that the party accesses the server via a client-anonymous authenticated channel, and therefore the adversary does not receive the party's identity. The adversary reports whether the protocol succeeds or not by setting accept' to true or false. $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ then determines the party and server's output as follows.

- If the server is honest (i.e., `ServerStat` = 'good'), $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ stores $\text{WelDB}[id_t] \leftarrow \widehat{w}$ and outputs $\text{accept} := \text{true}$ to id. The ideal server outputs (id_t, \widehat{w}) . This models correctness: the honest server always accepts the welcome message and does not know who sent it.
- If the server is malicious, $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ outputs accept' specified by the adversary. This means the malicious server can decide the protocol result arbitrarily.

Fetch and process commits (See Fig. 36) To process, party id needs to fetch the commit and proposals from the server via the subroutine `FetchCommit`. Before fetching a commit and proposals, the party permutes its index to $\text{index}_{\tilde{c}}$ by the function `*permuted-commit-index` to make the party's identity and index of \tilde{c} unlikable. Here, the function uses a random permutation (if `safe` is true for the id's current epoch) or a permutation the adversary chooses (if `safe` is false). This models the fact that if `safe` is true, then the permuted index seems random to the adversary.

- If the server is honest (i.e., `ServerStat` = 'good'), $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ gives (gid, epoch, $\text{index}_{\tilde{c}}$) to the adversary, where $\text{index}_{\tilde{c}}$ is the permuted id's index. Note that the party accesses to the server via a client-anonymous authenticated channel, and therefore, the adversary does not receive the party's identity. The adversary reports whether the protocol succeeds or not by setting accept' to true or false, and specifies the commit message $(c'_0, \tilde{c}', \tilde{p}')$. $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ then determines the party and server's output as follows.
 - If $\text{ComDB}[\text{gid}, \text{epoch}] = ((c'_0, \tilde{c}'), \text{node-id})$ and $\text{node-id} = \text{Ptr}[\text{id}]$, the ideal party always outputs $\text{accept} := \text{true}$ and the commit message $(c_0, \tilde{c}, \tilde{p}) := (c'_0, \tilde{c}'[\text{index}_{\tilde{c}}], \text{PropDB}[\text{gid}, \text{epoch}])$. This models correctness: if a party holds the same group state used to initialize node-id, the server must accept and returns the stored commit (which was received during protocol `PublishCommit`)

- If $\text{ComDB}[\text{gid}, \text{epoch}] = ((c'_0, \tilde{c}'), \text{'adv'})$ and $\text{accept}' = \text{true}$, the ideal party outputs $\text{accept} := \text{true}$ and the commit $(c_0, \tilde{c}, \tilde{p}) := (c'_0, \tilde{c}'[\text{index}_{\tilde{c}}], \text{PropDB}[\text{gid}, \text{epoch}])$. This models the fact that if the destination (gid, epoch) was initialized by the adversary, the adversary gets to decide if the ideal server accepts.
- Otherwise, the ideal party outputs $\text{accept} := \text{false}$. This case occurs if $\text{accept}' = \text{false}$ (i.e., the adversary decides to reject), $\text{ComDB}[\text{gid}, \text{epoch}] \neq ((c'_0, \tilde{c}'), *)$ (i.e., the destination has not been initialized or a commit message has not been issued), or $\text{ComDB}[\text{gid}, \text{epoch}] = (\top, \text{node-id})$ but $\text{node-id} \neq \text{Ptr}[\text{id}]$ (i.e., the party is located in a different commit node).

In all three cases, the ideal server outputs (gid, epoch, $\text{index}_{\tilde{c}}$) to \mathcal{Z} .

- If the server is malicious, $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ gives (Ptr[id], $\text{index}_{\tilde{c}}$) to the adversary. The adversary specifies the protocol result accept' and the commit $(c'_0, \tilde{c}', \tilde{p}')$, and it is outputted to id. This means the malicious server can return an arbitrary message to the parties. This is consistent with $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ which allows an adversary to inject malicious commits and proposals.

Finally, if id successfully fetches a commit and a list of proposals (i.e., id receives $\text{accept} = \text{true}$), id then processes them as in $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ (cf. *Processing commits* in App. B.2).

Join a group (See Fig. 36). To join a group, party id fetches a welcome message from the server via the subroutine `FetchWelcome`. $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ first informs the adversary that id fetches a welcome message. In this case, the server and the adversary learn who is accessing. The adversary reports whether the protocol succeeds or not by setting accept' to true or false, and specifies the welcome message \widehat{w}' . $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ then determines the party and server's output as follows.

- If the server is honest (i.e., `ServerStat` = 'good'), $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ returns to id the welcome message $\widehat{w} := \text{WelDB}[\text{id}]$ stored in its database or $\text{accept} = \text{false}$ if $\text{WelDB}[\text{id}] = \perp$. This models correctness: the honest server must return the message received during the subroutine `PublishWelcome`. The ideal server outputs the accessing party's identity id, which models the fact that any concrete protocol allows the server to know who was fetching the welcome message.
- If the server is malicious, the specified welcome message \widehat{w}' is outputted to id. This models that the malicious server can send an arbitrary message to parties.

Finally, if id successfully fetches a welcome message (i.e., id receives $\text{accept} = \text{true}$), $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ then processes the welcome message as in $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ (cf. *Joining a group* in App. B.2).

Group keys (See Fig. 36). Parties can fetch the current group secret via the `Key` query. The returned group secret k is random if the protocol guarantees its confidentiality (identified by the `safe` predicate). Otherwise, k is set by the adversary. Unlike $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$, $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ only provides an interface to retrieve one group secret k . That is, it does not have the interface to retrieve the group metadata secret k_{mh} and the next key function. This is because, those were special keys only used to secure the dynamic metadata (see App. E), and

in particular, \mathcal{F}_{CGKA}^{mh} only requires one group secret k that will be used to exchange the actual message.

Corruption. Similar to $\mathcal{F}_{CGKA}^{ctxt}$, the adversary can obtain party id's internal states via the (exposed, id) query. To prevent the so-called commitment problem, \mathcal{F}_{CGKA}^{mh} fixes the value of the safety predicate **safe** when a new group or epoch is initialized. This is controlled by the function `*mark-next-db-initialized-epoch` in `Create` and `Commit`. This is because the group secret generated at a specific commit node is explicitly used in the real protocol and we must restrict the adversary to not corrupting these nodes in order not to trivially win the security game. In addition, the adversary can corrupt the server via the `CorruptServer` query, and take over the role of the server. For simplicity, we assume once the server becomes malicious, it remains malicious and will never become honest. Finally, as mentioned in the overview, we do not consider adversary-controlled randomness in the current model.

D.3 Functions Used by the Adversary

\mathcal{F}_{CGKA}^{mh} offers the publish and fetch message interfaces to the adversary so that it can impersonate an honest party to a server. Similar to the case an honest party accesses the server, \mathcal{F}_{CGKA}^{mh} defines the ideal function when the adversary (malicious party) accesses the server. Since the party is malicious, \mathcal{F}_{CGKA}^{mh} never requires correctness. In contrast, \mathcal{F}_{CGKA}^{mh} defines the security requirements: It defines the conditions under which the adversary can access the server. In case the conditions do not hold, the functionality models the fact that the adversary cannot access the server. Note that when both party and server are malicious, \mathcal{F}_{CGKA}^{mh} does nothing since the adversary can perform all the protocols by itself.

Publish or fetch proposal or commit messages by the adversary. To publish a proposal or commit, the adversary sends `PublishProposalAdv` or `PublishCommitAdv` to \mathcal{F}_{CGKA}^{mh} . $\mathcal{F}_{CGKA}^{ctxt}$ determines the server's output as follows.

- If $\text{ComDB}[\text{gid}, \text{epoch}] = (\top, \text{node-id})$ and $\text{accept}' = \text{true}$, \mathcal{F}_{CGKA}^{mh} checks whether authenticity is guaranteed for the honestly generated epoch node-id by the safety predicate **adv-access-allowed**(node-id). If the predicate returns false, then \mathcal{F}_{CGKA}^{mh} halts. This models the fact that the adversary can publish messages for an honestly initialized epoch only when adversarial access is allowed. If access is allowed, the server stores the published message from the adversary in its database.
- If $\text{ComDB}[\text{gid}, \text{epoch}] = (\top, \text{'adv'})$ and $\text{accept}' = \text{true}$, the ideal server outputs $\text{accept} := \text{true}$. This models the fact that if the destination (gid, epoch) was initialized by the adversary, the adversary gets to decide whether the ideal server accepts.
- Otherwise, the ideal server outputs $\text{accept} := \text{false}$. This case occurs if $\text{accept}' = \text{false}$ (i.e., the adversary decides to reject) or $\text{ComDB}[\text{gid}, \text{epoch}] \neq (\top, *)$ (i.e., the destination has not been initialized or a commit message has been issued).

To fetch proposals, the adversary sends `FetchProposalsAdv` to \mathcal{F}_{CGKA}^{mh} . \mathcal{F}_{CGKA}^{mh} determines the server's output as follows.

- If $\text{ComDB}[\text{gid}, \text{epoch}] = (*, \text{node-id})$ and $\text{accept}' = \text{true}$, \mathcal{F}_{CGKA}^{mh} checks whether authenticity is guaranteed for the honestly generated epoch node-id by the safety predicate **adv-access-allowed**(node-id). If the predicate returns false, then \mathcal{F}_{CGKA}^{mh} halts. This models the fact that the adversary can fetch proposals from an honestly initialized epoch only when adversarial access is allowed. If the access is allowed, the adversary obtains $\vec{p} := \text{PropDB}[\text{gid}, \text{epoch}]$.
- If $\text{ComDB}[\text{gid}, \text{epoch}] = (*, \text{'adv'})$ and $\text{accept}' = \text{true}$, the ideal server outputs $\text{accept} := \text{true}$. This models the fact that if the destination (gid, epoch) was initialized by the adversary, the adversary can decide to accept.
- Otherwise, the ideal party outputs $\text{accept} := \text{false}$. This case occurs if $\text{accept}' = \text{false}$ (i.e., the adversary decides to reject) or $\text{ComDB}[\text{gid}, \text{epoch}] = \perp$ (i.e., the destination has not been initialized).

Finally, to publish a proposal or a commit the adversary sends (`FetchCommitAdv`, gid, epoch, index) to \mathcal{F}_{CGKA}^{mh} . \mathcal{F}_{CGKA}^{mh} determines the server's output as follows.

- If $\text{ComDB}[\text{gid}, \text{epoch}] = ((c'_0, \vec{c}'), \text{node-id})$ and $\text{accept}' = \text{true}$, \mathcal{F}_{CGKA}^{mh} checks whether authenticity is guaranteed for the honestly generated epoch node-id by the safety predicate **adv-access-allowed**(node-id). If the predicate returns false, then \mathcal{F}_{CGKA}^{mh} halts. This models the fact that the adversary can fetch messages for an honestly initialized epoch only when adversarial access is allowed. If access is allowed, the adversary obtains the stored $(c_0, \vec{c}, \vec{p}) := (c'_0, \vec{c}'[\text{index}_{\vec{c}}], \text{PropDB}[\text{gid}, \text{epoch}])$.
- If $\text{ComDB}[\text{gid}, \text{epoch}] = ((c'_0, \vec{c}'), \text{'adv'})$ and $\text{accept}' = \text{true}$, the adversary obtains the stored $(c_0, \vec{c}, \vec{p}) := (c'_0, \vec{c}'[\text{index}_{\vec{c}}], \text{PropDB}[\text{gid}, \text{epoch}])$. This models the fact that if the destination (gid, epoch) was initialized by the adversary, the adversary can decide to accept.
- Otherwise, the ideal party outputs $\text{accept} := \text{false}$. This case occurs if $\text{accept}' = \text{false}$ (i.e., the adversary decides to reject) or $\text{ComDB}[\text{gid}, \text{epoch}] \neq ((c'_0, \vec{c}'), *)$ (i.e., the destination has not been initialized or a commit message has not been issued).

Publish or fetch welcome messages by the adversary. To publish or fetch messages, the adversary sends `PublishWelcome` or `FetchWelcome` to \mathcal{F}_{CGKA}^{mh} . In this case, \mathcal{F}_{CGKA}^{mh} determines the server's output as in the case an honest party accesses the server. The protocol succession means the adversary succeeds to publish or fetch a welcome message.

Initialization

```

1 : / Line 1 to 7 below are identical to selective downloading  $\mathcal{F}_{CGKA}^{\text{ctxt}}$ 
2 :  $\text{Ptr}[*], \text{Prop}[*], \text{Node}[*], \text{Wel}[*] \leftarrow \perp$ 
3 :  $\text{propCtr}, \text{nodeCtr} \leftarrow 1$ 
4 :  $\text{flag}_{\text{selDL}} = \text{true}, \text{DesignatedCom}[*] \leftarrow \perp$ 
5 :  $\text{flag}_{\text{contHide}} = \text{true}$ 
6 :  $\text{PropID}[*], \text{NodeID}[*] \leftarrow \perp$ 
7 :  $\text{RandCorr}[*] \leftarrow \text{'good'}$ 
8 : / Initialize below for dynamic metadata-hiding
9 :  $\text{flag}_{\text{dbinit}} = \text{true}$ 
10 :  $\text{PropDB}[*,*], \text{ComDB}[*,*], \text{WelDB} \leftarrow \perp$ 
11 :  $\text{ServerStat} \leftarrow \text{'good'}$ 

```

Input from the party $\text{id}_{\text{creator}}$ **Input (Create, svk)**

```

1 :  $\text{req Ptr}[\text{id}_{\text{creator}}] = \perp$ 
2 : / Run identical code as Input (Create, svk) of  $\mathcal{F}_{CGKA}^{\text{ctxt}}$ .
3 : / excluding the final return (gid, 0, mem) line
4 :  $\langle\langle \text{Insert (Create, svk) of } \mathcal{F}_{CGKA}^{\text{ctxt}} \rangle\rangle$ 
5 :  $\text{req RegisterGroup}(\text{gid}, 0)$ 
6 :  $\text{*mark-next-db-initialized-epoch}(\text{Ptr}[\text{id}_{\text{creator}}])$ 

```

Subroutine RegisterGroup(gid, epoch)

```

1 :  $\text{channelType} \leftarrow (\text{'anon'}, \perp, \text{Sv})$ 
2 :  $\text{req Ptr}[\text{id}] = 0$ 
3 : Send (channelType, RegisterGroup, gid, 0) to  $\mathcal{S}$  and
   receive  $\text{accept}'$ 
4 : / If Sv is good, then accept if  $\text{ComDB}[\text{gid}, *]$  is empty
5 : if  $\text{ServerStat} = \text{'good'}$  then
6 :   if  $\text{ComDB}[\text{gid}, *] = \perp$  then
7 :      $\text{accept} \leftarrow \text{true}$ 
8 :     / Main group is assigned to node-id = 0.    $\text{ComDB}[\text{gid}, 0] \leftarrow (\top, 0)$ 
9 :   else
10 :     $\text{accept} \leftarrow \text{false}$ 
11 :   Send (channelType, RegisterGroup, (accept, gid, 0)) to Sv
12 : / If Sv is corrupt, then let  $\mathcal{S}$  decide if Sv accepts
13 : else
14 :    $\text{accept} \leftarrow \text{accept}'$ 
15 : return accept

```

Input Key

```

1 : / Run identical code as Input (Key) of  $\mathcal{F}_{CGKA}^{\text{ctxt}}$ .
2 : / including the final return  $\text{Node}[\text{Ptr}[\text{id}]].\text{key}$  line
3 :  $\langle\langle \text{Inset (Key) of } \mathcal{F}_{CGKA}^{\text{ctxt}} \rangle\rangle$ 

```

Inputs from a party id**Input (Propose, act), $\text{act} \in \{\text{'upd'-'svk'}, \text{'add'-'id}_\ell, \text{'rem'-'id}_\ell\}$**

```

1 :  $\text{req Ptr}[\text{id}] \neq \perp$ 
2 :  $\text{gid} \leftarrow \text{Node}[\text{Ptr}[\text{id}]].\text{gid}; \text{epoch} \leftarrow \text{Node}[\text{Ptr}[\text{id}]].\text{epoch}$ 
3 : / Run identical code as Input (Propose, act) of  $\mathcal{F}_{CGKA}^{\text{ctxt}}$ .
4 : / excluding the final return  $p$  line
5 :  $\langle\langle \text{Insert (Propose, act) of } \mathcal{F}_{CGKA}^{\text{ctxt}} \rangle\rangle$ 
6 :  $\text{req PublishProposal}(\text{gid}, \text{epoch}, p)$ 
7 : return  $p$ 

```

Input (Commit, svk)

```

1 :  $\text{req Ptr}[\text{id}] \neq \perp$ 
2 :  $\text{gid} \leftarrow \text{Node}[\text{Ptr}[\text{id}]].\text{gid}; \text{epoch} \leftarrow \text{Node}[\text{Ptr}[\text{id}]].\text{epoch}$ 
3 :  $(\text{accept}, \vec{p}) \leftarrow \text{FetchProposals}(\text{gid}, \text{epoch})$ 
4 :  $\text{req accept}$ 
5 : / Run identical code as Input (Commit,  $\vec{p}$ , svk) of  $\mathcal{F}_{CGKA}^{\text{ctxt}}$ .
6 : / excluding the final return  $(c_0, \vec{c}, \vec{w} = \{\vec{w}\})$  line
7 :  $\langle\langle \text{Insert (Commit, } \vec{p}, \text{svk) of } \mathcal{F}_{CGKA}^{\text{ctxt}} \rangle\rangle$ 
8 : / Permute the order of party sensitive commitment  $\vec{c}$ 
9 :  $\vec{c}_{\text{perm}} \leftarrow \text{*permute-commit}(\text{Ptr}[\text{id}], \vec{c})$ 
10 :  $\text{try PublishCommit}(\text{gid}, \text{epoch}, c_0, \vec{c}_{\text{perm}})$ 
11 : foreach  $\vec{w} \in \vec{w}$  do
12 :    $\text{parse}(\text{id}_\ell, *) \leftarrow \vec{w}$ 
13 :    $\text{try PublishWelcome}(\text{id}_\ell, \vec{w})$ 
14 :  $\text{*mark-next-db-initialized-epoch}(\text{NodeID}[c_0])$ 
15 : return  $(c_0, \vec{c}_{\text{perm}}, \vec{w})$ 

```

Input Process

```

1 :  $\text{req Ptr}[\text{id}] \neq \perp$ 
2 :  $\text{gid} \leftarrow \text{Node}[\text{Ptr}[\text{id}]].\text{gid}; \text{epoch} \leftarrow \text{Node}[\text{Ptr}[\text{id}]].\text{epoch}$ 
3 : / Permuted index of party id required for selective downloading.
4 :  $\text{index}_{\vec{c}} \leftarrow \text{*permuted-commit-index}(\text{Ptr}[\text{id}], \text{id})$ 
5 :  $(\text{accept}, c_0, \vec{c}, \vec{p}) \leftarrow \text{FetchCommit}(\text{gid}, \text{epoch}, \text{index}_{\vec{c}})$ 
6 :  $\text{req accept}$ 
7 : / Run identical code as Input (Process,  $c_0, \vec{c}, \vec{p}$ ) of  $\mathcal{F}_{CGKA}^{\text{ctxt}}$ .
8 : / including the final return  $\text{*output-proc}(\text{node-id})$  line
9 :  $\langle\langle \text{Insert (Process, } c_0, \vec{c}, \vec{p}) \text{ of } \mathcal{F}_{CGKA}^{\text{ctxt}} \rangle\rangle$ 

```

Input Join

```

1 :  $\text{req Ptr}[\text{id}] = \perp$ 
2 :  $(\text{accept}, \vec{w}) \leftarrow \text{FetchWelcome}(\text{id})$ 
3 :  $\text{req accept}$ 
4 : / Run identical code as Input (Join,  $\vec{w}$ ) of  $\mathcal{F}_{CGKA}^{\text{ctxt}}$ .
5 : / including the final return  $\text{*output-join}(\text{node-id})$  line
6 :  $\langle\langle \text{Insert (Join, } \vec{w}) \text{ of } \mathcal{F}_{CGKA}^{\text{ctxt}} \rangle\rangle$ 

```

Figure 36: The ideal MH-CGKA functionality $\mathcal{F}_{CGKA}^{\text{mh}}$: Create, Propose, Commit, Process, and Join interface for honest parties. For better readability, we outsource the lines identical to $\mathcal{F}_{CGKA}^{\text{ctxt}}$ to Fig. 13. Note that the Create function is invoked only once by the previously designated party $\text{id}_{\text{creator}}$. (Only one $\text{id}_{\text{creator}}$ exists.)

Subroutine PublishProposal($gid, epoch, p$)

```

1 : channelType  $\leftarrow$  ('anon',  $\perp$ , Sv) / Connect to the server via anonymous channel.
2 : if ServerStat = 'good' then
3 :   Send (channelType, PublishProposal, gid, epoch, p) to S and
       receive accept'
4 :   / If the party assigned to gid-epoch accesses, the server must accept.
5 :   if ComDB[gid, epoch] = ( $\top$ , node-id)  $\wedge$  node-id = Ptr[id] then
6 :     PropDB[gid, epoch]  $\# \leftarrow$  p
7 :     accept  $\leftarrow$  true
8 :     / If gid-epoch was initialized by the adversary, S decide to accept or reject.
9 :   elseif ComDB[gid, epoch] = ( $\top$ , 'adv')  $\wedge$  accept' then
10 :    PropDB[gid, epoch]  $\# \leftarrow$  p
11 :    accept  $\leftarrow$  true
12 :    / Otherwise, the server must reject.
13 :   else
14 :     accept  $\leftarrow$  false
15 :   Send (channelType, PublishProposal, (accept, gid, epoch, p)) to Sv
16 :   / If Sv is corrupt, then S decides if Sv accepts.
17 :   else
18 :     Send (channelType, PublishProposal, Ptr[id], p) to S and
       receive accept'
19 :     accept  $\leftarrow$  accept'
20 :   return accept

```

Subroutine FetchProposals($gid, epoch$)

```

1 : channelType  $\leftarrow$  ('anon',  $\perp$ , Sv) / Connect to the server via anonymous channel.
2 : if ServerStat = 'good' then
3 :   Send (channelType, FetchProposals, gid, epoch) to S and
       receive accept'
4 :   / If the party assigned to gid-epoch accesses, the server must accept.
5 :   if ComDB[gid, epoch] = (*, node-id)  $\wedge$  node-id = Ptr[id] then
6 :     (accept,  $\vec{p}$ )  $\leftarrow$  (true, PropDB[gid, epoch])
7 :     / If gid-epoch was initialized by the adversary, S decide to accept or reject.
8 :   elseif ComDB[gid, epoch] = (*, 'adv')  $\wedge$  accept' then
9 :     (accept,  $\vec{p}$ )  $\leftarrow$  (true, PropDB[gid, epoch])
10 :    / Otherwise, the server must reject.
11 :   else
12 :     (accept,  $\vec{p}$ )  $\leftarrow$  (false,  $\perp$ )
13 :   Send (channelType, FetchProposals, (accept, gid, epoch)) to Sv
14 :   / If Sv is corrupt, then S decides what Sv returns to id.
15 :   else
16 :     Send (channelType, FetchProposals, Ptr[id]) to S and
       receive (accept',  $\vec{p}'$ )
17 :     (accept,  $\vec{p}$ )  $\leftarrow$  (accept',  $\vec{p}'$ )
18 :   return (accept,  $\vec{p}$ )

```

Subroutine PublishCommit($gid, epoch, c_0, \vec{c}_{perm}$)

```

1 : channelType  $\leftarrow$  ('anon',  $\perp$ , Sv) / Connect to the server via anonymous channel.
2 : if ServerStat = 'good' then
3 :   Send (channelType, PublishCommit, gid, epoch, c_0,  $\vec{c}_{perm}$ ) to S and
       receive accept'
4 :   / If the party uses the state assigned to gid-epoch, the server must accept.
5 :   if ComDB[gid, epoch] = ( $\top$ , node-id)  $\wedge$  node-id = Ptr[id] then
6 :     ComDB[gid, epoch]  $\leftarrow$  ((c_0,  $\vec{c}_{perm}$ ), node-id)
7 :     / If an honest party initialize the next epoch, the corresponding node is marked.
8 :     ComDB[gid, epoch + 1]  $\leftarrow$  ( $\top$ , NodeID[c_0])
9 :     accept  $\leftarrow$  true
10 :    / If gid-epoch was initialized by the adversary, S decide to accept or reject.
11 :   elseif ComDB[gid, epoch] = ( $\top$ , 'adv')  $\wedge$  accept' then
12 :     ComDB[gid, epoch]  $\leftarrow$  ((c_0,  $\vec{c}_{perm}$ ), 'adv')
13 :     / If an honest party initialize the next epoch, the corresponding node is marked.
14 :     ComDB[gid, epoch + 1]  $\leftarrow$  ( $\top$ , NodeID[c_0])
15 :     accept  $\leftarrow$  true
16 :     / Otherwise, the server must reject.
17 :   else
18 :     accept  $\leftarrow$  false
19 :   Send (channelType, PublishCommit, (accept, gid, epoch, c_0,  $\vec{c}_{perm}$ )) to Sv
20 :   / If Sv is corrupt, then S decides if Sv accepts.
21 :   else
22 :     Send (channelType, PublishCommit, gid, epoch, c_0,  $\vec{c}_{perm}$ ) to S and
       receive accept'
23 :     accept  $\leftarrow$  accept'
24 :   return accept

```

Subroutine FetchCommit($gid, epoch, index_{\vec{c}}$)

```

1 : channelType  $\leftarrow$  ('anon',  $\perp$ , Sv) / Connect to the server via anonymous channel.
2 : if ServerStat = 'good' then
3 :   Send (channelType, FetchCommit, gid, epoch, index_{\vec{c}}) to S and
       receive accept'
4 :   / If the party assigned to gid-epoch accesses, the server must accept.
5 :   if ComDB[gid, epoch] = ((c'_0,  $\vec{c}'$ ), node-id)  $\wedge$  node-id = Ptr[id] then
6 :     (accept, c_0,  $\vec{c}$ ,  $\vec{p}$ )  $\leftarrow$  (true, c'_0,  $\vec{c}'$ [index_{\vec{c}}], PropDB[gid, epoch])
7 :     / If gid-epoch was initialized by the adversary, S decide to accept or reject.
8 :   elseif ComDB[gid, epoch] = ((c'_0,  $\vec{c}'$ ), 'adv')  $\wedge$  accept' then
9 :     (accept, c_0,  $\vec{c}$ ,  $\vec{p}$ )  $\leftarrow$  (true, c'_0,  $\vec{c}'$ [index_{\vec{c}}], PropDB[gid, epoch])
10 :    / Otherwise, the server must reject.
11 :   else
12 :     (accept, c_0,  $\vec{c}$ ,  $\vec{p}$ )  $\leftarrow$  (false,  $\perp$ ,  $\perp$ ,  $\perp$ )
13 :   Send (channelType, FetchCommit, (accept, gid, epoch, index_{\vec{c}})) to Sv
14 :   else
15 :     Send (channelType, FetchCommit, Ptr[id], index_{\vec{c}}) to S and
       receive (accept', c'_0,  $\vec{c}'$ ,  $\vec{p}'$ )
16 :     (accept, c_0,  $\vec{c}$ ,  $\vec{p}$ )  $\leftarrow$  (accept', c'_0,  $\vec{c}'$ ,  $\vec{p}'$ )
17 :   return (accept, c_0,  $\vec{c}$ ,  $\vec{p}$ )

```

Figure 37: The ideal metadata-hiding CGKA functionality $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$: subroutines for publish and fetch proposal and commit messages. They are used in Propose, Commit and Process interface shown in Fig. 36. If $\vec{c} = \perp$, we define $\vec{c}[\text{index}] = \perp$ for any index index.

Input from a party id or adversary \mathcal{S} **Input PublishWelcome(id_t, \widehat{w})**

```

1 : / Connect to the server via anonymous channel.
2 : channelType  $\leftarrow$  ('anon',  $\perp$ , Sv)
3 : Send (channelType, PublishWelcome,  $id_t, \widehat{w}$ ) to  $\mathcal{S}$  and receive accept'
4 : / If Sv is honest, then store  $\widehat{w}$  in WelDB
5 : if ServerStat = 'good' then
6 :   accept  $\leftarrow$  true
7 :   WelDB[ $id_t$ ]  $\leftarrow$   $\widehat{w}$ 
8 :   Send (channelType, PublishWelcome, (accept,  $id_t, \widehat{w}$ )) to Sv
9 : / If Sv is corrupt, then  $\mathcal{S}$  decides if Sv accepts.
10 : else
11 :   accept  $\leftarrow$  accept'
12 : / If invoked by  $\mathcal{S}$ , then no output is required
13 : if ServerStat = 'good'  $\vee$  invoked by party id then
14 :   return accept

```

Input FetchWelcome(id)

```

1 : / Connect to the server via authenticated channel.
2 : channelType  $\leftarrow$  ('auth', id, Sv)
3 : Send (channelType, FetchWelcome) to  $\mathcal{S}$  and receive (accept',  $\widehat{w}'$ )
4 : if ServerStat = 'good' then
5 :   if WelDB[id]  $\neq$   $\perp$  then
6 :     (accept,  $\widehat{w}$ )  $\leftarrow$  (true, WelDB[id])
7 :   else
8 :     (accept,  $\widehat{w}$ )  $\leftarrow$  (false,  $\perp$ )
9 :   Send (channelType, FetchWelcome, (accept, id)) to Sv
10 : / If Sv is corrupt, then let  $\mathcal{S}$  decide Sv's output
11 : else
12 :   (accept,  $\widehat{w}$ )  $\leftarrow$  (accept',  $\widehat{w}'$ )
13 : / If invoked by  $\mathcal{S}$ , then no output is required
14 : if ServerStat = 'good'  $\vee$  invoked by party id
15 :   return (accept,  $\widehat{w}$ )

```

Figure 38: The ideal metadata-hiding CGKA functionality \mathcal{F}_{CGKA}^{mh} : Subroutines for publish and fetch welcome messages. They are used in Commit and Join interface shown in Fig. 36.

***permuted-commit-index(node-id, id)**

```

1 : assert Node[node-id]  $\neq$   $\perp$ 
2 : / If flagselDL is false, there is no need to permute index.
3 : if  $\neg$ flagselDL then
4 :   return  $\perp$ 
5 : mem  $\leftarrow$  Node[node-id].mem
6 : assert id  $\in$  mem
7 : index  $\leftarrow$  Node[node-id].index(id)
8 : if Node[node-id].perm =  $\perp$  then
9 :   *set-index-permutation(node-id)
10 :  $\phi$   $\leftarrow$  Node[node-id].perm
11 : return  $\phi$ (index)

```

***permute-commit(node-id, \vec{c})**

```

1 : assert Node[node-id]  $\neq$   $\perp$ 
2 : / If flagselDL is false, no party-dependent commitment can exist.
3 : if  $\neg$ flagselDL then
4 :   assert  $\vec{c} = \perp$ 
5 :   return  $\perp$ 
6 : if Node[node-id].perm =  $\perp$  then
7 :   *set-index-permutation(node-id)
8 :  $\phi$   $\leftarrow$  Node[node-id].perm
9 :  $\vec{c}_{perm} \leftarrow$  ()
10 : for index = 1, ...,  $|\vec{c}|$  do
11 :    $\vec{c}_{perm} \# \leftarrow \vec{c}[\phi(\text{index})]$ 
12 : return  $\vec{c}_{perm}$ 

```

***set-index-permutation(node-id)**

```

1 : mem  $\leftarrow$  Node[node-id].mem
2 : / If the input node is not corrupted,
   / sample a random permutation over  $S_{|mem|}$ .
3 : if random-index(node-id) then
4 :    $\phi \leftarrow S_{|mem|}$ 
5 :   Node[node-id].perm  $\leftarrow$   $\phi$ 
6 : else
7 :   Send (Permutation, node-id) to  $\mathcal{S}$  and
   receive  $\phi$ 
8 : assert  $\phi \in S_{|mem|}$ 
9 : Node[node-id].perm  $\leftarrow$   $\phi$ 

```

Figure 39: Helper function: Permute indices. $S_{|mem|}$ is the set of permutations on the range $[|mem|]$.

Input from the adversary \mathcal{S}

Input (PublishProposalAdv, gid, epoch, p)

```

1: channelType  $\leftarrow$  ('anon',  $\perp$ , Sv)
2: Send (channelType, PublishProposalAdv, gid, epoch,  $p$ ) to  $\mathcal{S}$  and
   receive accept'
3: / The honest server accepts if  $\mathcal{S}$  decides to accept and ComDB[gid, epoch] = ( $\top$ , *).
4: if ServerStat = 'good' then
5:   if ComDB[gid, epoch] = ( $\top$ , node-id)  $\wedge$  accept' then
6:     / If gid-epoch was initialized by an honest party,  $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$  checks authenticity is guaranteed.
7:     assert adv-access-allowed(node-id)
8:     PropDB[gid, epoch]  $\# \leftarrow p$ 
9:     accept  $\leftarrow$  true
10:  else if ComDB[gid, epoch] = ( $\top$ , 'adv')  $\wedge$  accept' then
11:    PropDB[gid, epoch]  $\# \leftarrow p$ 
12:    accept  $\leftarrow$  true
13:  else
14:    accept  $\leftarrow$  false
15:  Send (channelType, PublishProposal, (accept, gid, epoch,  $p$ )) to Sv
16:  return accept
```

Input (FetchProposalsAdv, gid, epoch)

```

1: channelType  $\leftarrow$  ('anon',  $\perp$ , Sv)
2: Send (channelType, FetchProposalsAdv, gid, epoch) to  $\mathcal{S}$  and
   receive accept'
3: / The honest server accepts if  $\mathcal{S}$  decides to accept and ComDB[gid, epoch]  $\neq \perp$ .
4: if ServerStat = 'good' then
5:   else if ComDB[gid, epoch] = (*, node-id)  $\wedge$  accept' then
6:     / If gid-epoch was initialized by an honest party,  $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$  checks authenticity is guaranteed.
7:     assert adv-access-allowed(node-id)
8:     (accept,  $\vec{p}$ )  $\leftarrow$  (true, PropDB[gid, epoch])
9:   else if ComDB[gid, epoch] = (*, 'adv')  $\wedge$  accept' then
10:    (accept,  $\vec{p}$ )  $\leftarrow$  (true, PropDB[gid, epoch])
11:   else
12:    (accept,  $\vec{p}$ )  $\leftarrow$  (false,  $\perp$ )
13:   Send (channelType, FetchProposals, (accept, gid, epoch) to Sv
14:   return (accept,  $\vec{p}$ )
```

Input (PublishCommitAdv, gid, epoch, c_0 , \vec{c})

```

1: channelType  $\leftarrow$  ('anon',  $\perp$ , Sv)
2: Send (channelType, PublishCommitAdv, gid, epoch,  $c_0$ ,  $\vec{c}$ ) to  $\mathcal{S}$  and
   receive accept'
3: / The honest server accepts if  $\mathcal{S}$  decides to accept and ComDB[gid, epoch] = ( $\top$ , *).
4: if ServerStat = 'good' then
5:   ComDB[gid, epoch]
6:   else if ComDB[gid, epoch] = ( $\top$ , node-id)  $\wedge$  accept' then
7:     / If gid-epoch was initialized by an honest party,  $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$  checks authenticity is guaranteed.
8:     assert adv-access-allowed(node-id)
9:     ComDB[gid, epoch]  $\leftarrow$  (( $c_0$ ,  $\vec{c}$ ), node-id)
10:    / If injection succeeds, the next epoch is marked as adversarial initialized.
11:    ComDB[gid, epoch + 1]  $\leftarrow$  ( $\top$ , 'adv')
12:   else if ComDB[gid, epoch] = ( $\top$ , 'adv')  $\wedge$  accept' then
13:     ComDB[gid, epoch]  $\leftarrow$  (( $c_0$ ,  $\vec{c}$ ), 'adv')
14:    / If injection succeeds, the next epoch is marked as adversarial initialized.
15:    ComDB[gid, epoch + 1]  $\leftarrow$  ( $\top$ , 'adv')
16:   else
17:     accept  $\leftarrow$  false
18:   Send (channelType, PublishCommit, (accept, gid, epoch,  $c_0$ ,  $\vec{c}$ )) to Sv
19:   return accept
```

Input (FetchCommitAdv, gid, epoch, index)

```

1: channelType  $\leftarrow$  ('anon',  $\perp$ , Sv)
2: Send (channelType, FetchCommitAdv, gid, epoch, index) to  $\mathcal{S}$  and
   receive accept'
3: / The honest server accepts if  $\mathcal{S}$  decides to accept and ComDB[gid, epoch] = (( $c'_0$ ,  $\vec{c}'$ ), *).
4: if ServerStat = 'good' then
5:   if ComDB[gid, epoch] = (( $c'_0$ ,  $\vec{c}'$ ), node-id)  $\wedge$  accept' then
6:     / If gid-epoch was initialized by an honest party,  $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$  checks authenticity is guaranteed.
7:     assert adv-access-allowed(node-id)
8:     (accept,  $c_0$ ,  $\vec{c}$ ,  $\vec{p}$ )  $\leftarrow$  (true,  $c'_0$ ,  $\vec{c}'$ [index], PropDB[gid, epoch])
9:   if ComDB[gid, epoch] = (( $c'_0$ ,  $\vec{c}'$ ), 'adv')  $\wedge$  accept' then
10:    (accept,  $c_0$ ,  $\vec{c}$ ,  $\vec{p}$ )  $\leftarrow$  (true,  $c'_0$ ,  $\vec{c}'$ [index], PropDB[gid, epoch])
11:   else
12:    (accept,  $c_0$ ,  $\vec{c}$ ,  $\vec{p}$ )  $\leftarrow$  (false,  $\perp$ ,  $\perp$ ,  $\perp$ )
13:   Send (channelType, PublishProposal, (accept, gid, epoch, index) to Sv
14:   return (accept,  $c_0$ ,  $\vec{c}$ ,  $\vec{p}$ )
```

Figure 40: The ideal metadata-hiding CGKA functionality $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$: Interface for the adversary \mathcal{S} . If $\vec{c} = \perp$, $\vec{c}[\text{index}]$ is defined to be \perp for any index.

Input (Expose, id)

```

1: if Ptr[id] ≠ ⊥ then
2:   Node[Ptr[id]].exp +← id
3:   *update-stat-after-exp(id) / Pending secrets are marked as exposed.
4:   svk ← Node[Ptr[id]].mem[id]
5:   Send (exposed, id, svk) to  $\mathcal{F}_{AS}$ 
6:   Send (Ptr[id], Node[Ptr[id]]) to  $\mathcal{S}$  / All information stored in Node[Ptr[id]] is sent to  $\mathcal{S}$ .
7:   Send (exposed, id) to  $\mathcal{F}_{KS}$ 
8:   restrict  $\forall$ node-id :
9:     if Node[node-id].chall = true then safe(node-id) = true
10:    if Node[node-id].conthide = true then safe(node-id) = true
11:    if Node[node-id].dbinit = true then safe(node-id) = true

```

Input CorruptServer

```

1: / Once corrupted, the server remains corrupted.
2: ServerStat ← 'adv'
*mark-next-db-initialized-epoch(node-id)


---


1: if safe(node-id) then
2:   Node[node-id].dbinit ← true
3: else
4:   Node[node-id].dbinit ← false

```

Figure 41: The metadata-hiding CGKA functionality \mathcal{F}_{CGKA}^{mh} : Corruptions from the adversary \mathcal{S} . The difference between those of $\mathcal{F}_{CGKA}^{ctxt}$ is highlighted in yellow. \mathcal{S} can call `CorruptServer` only once.

E METADATA-HIDING CGKA: CONSTRUCTION AND SECURITY PROOF

E.1 Constructing the Wrapper Protocol W^{mh}

In this section, we propose a metadata CGKA W^{mh} and prove that it UC-realizes $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ in the $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ -hybrid model. We call W^{mh} as a *wrapper* protocol since it works as a wrapper around any static metadata-hiding CGKA that UC-realizes $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$, and turns it into full metadata-hiding CGKA. In particular, the sole functionality of W^{mh} is to take care of how the proposals, commits, and welcome messages are uploaded and downloaded from the server in a dynamic metadata-hiding manner. The construction is very simple and can be implemented only from a standard signature scheme.

$G.\text{gid}$	The identifier of the group.
$G.\text{epoch}$	The current epoch number.
$G.\text{mem}$	A list of (id, svk)-pair. The list is sorted in lexicographic order by ids.
$(G.\text{gsk}, G.\text{gvk})$	The group signature key used to authenticate group membership to the server.
$G.\text{permKey}$	The PRP key used to permute member index in membership list.
$G.\text{indexOf}(id)$	Returns the index of id in the list $G.\text{mem}$.

Table 9: The party’s protocol state and helper method.

$\text{PropDB}[*,*]$	It stores a list of proposal message p issued at $(\text{gid}, \text{epoch})$.
$\text{ComDB}[*,*]$	It stores a group signature key gvk used to authenticate membership on $(\text{gid}, \text{epoch})$ and possibly a commit message (c_0, \vec{c}) to move to $(\text{gid}, \text{epoch})$.
$\text{WelDB}[*]$	It stores a welcome message \widehat{w} for id .

Table 10: The server’s protocol state.

E.1.1 Protocol States. Each party holds a group state G . It consists of the variables listed in Tab. 9. The $G.\text{mem}$ list stores the group member’s identity and its signature key. The list is sorted in lexicographic order by the party’s identity. Parties can fetch the index of their identities in $G.\text{mem}$ via the method $G.\text{indexOf}(*)$.

The server keeps three databases: $\text{PropDB}[*,*]$, $\text{ComDB}[*,*]$, and $\text{WelDB}[*]$. $\text{PropDB}[*,*]$ and $\text{ComDB}[*,*]$ both use $(\text{gid}, \text{epoch})$ as indices to store proposals and commits, respectively. We refer the readers to Sec. 4.2 for a pictorial example, where we merge the PropDB and ComDB into one database in the figures.

- $\text{PropDB}[\text{gid}, \text{epoch}] = \vec{p}$: The proposal database stores the list of proposals \vec{p} created at $(\text{gid}, \text{epoch})$. These proposals, once committed, are used to move any party at epoch to the next epoch $' = \text{epoch} + 1$, where the group state is updated accordingly to the proposals.
- $\text{ComDB}[\text{gid}, \text{epoch}] = (\text{gvk}, c_0, \vec{c})$ or $(\text{gvk}, \perp, \perp)$: The commit database stores a group signing key (also called a group statement) gvk . This will be used by a group member at epoch to anonymously prove that they are indeed a group member. $\text{ComDB}[\text{gid}, \text{epoch}]$ is initialized with $(\text{gvk}, \perp, \perp)$ and is later updated to $(\text{gvk}, c_0, \vec{c})$ when some party creates a commit (c_0, \vec{c}) with the proposals stored in $\text{PropDB}[\text{gid}, \text{epoch}]$. Once $\text{ComDB}[\text{gid}, \text{epoch}]$ has a commit stored, then no other commits can be made at this epoch.
- $\text{WelDB}[\text{id}] = \widehat{w}$: It stores a welcome message for id . We assume $\text{WelDB}[\text{id}]$ stores only one welcome message for id . This is due to the fact that previous CGKA UC-security models assume a party can join at most one group (i.e., $\text{Ptr}[\text{id}]$ identifies the unique group that id is a member of). In our protocol, when the server receives a new welcome message (id, \widehat{w}) , it overwrites it as $\text{WelDB}[\text{id}] := \widehat{w}$. We emphasize that this restriction is required only to prove UC-security,

and functionality-wise, WelDB[id] can store as many welcome messages as it needs.

E.1.2 Protocol Algorithms. The main interface and the associated helper functions are depicted in Fig. 42. The protocol W^{mh} is defined in the $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ -hybrid model and internally calls $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ to create a group, generate and process protocol messages. To register groups and publish or fetch protocol messages, it uses the subroutines depicted in Figs. 43 to 49.

Group Creation. A group can be created by the group creator $\text{id}_{\text{creator}}$ by invoking $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ on input (Create, svk). The group creator first invokes $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ with the same input. Then, it runs the helper function *init-states to initialize the group state of the wrapper protocol. Note that *init-states internally queries (Key_{mh}) to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ to use the group secret k_{mh} . Finally, it executes the group registration protocol RegisterGroup shown in Fig. 43. The group creator sends the new group's group identity gid, epoch counter epoch, and group signature key gvk via a client-anonymous authenticated channel. Upon receiving the group creation message, the server checks that a group with the same identity does not exist and the epoch counter is equal to 0. If the check passes, the server stores (gvk, \perp , \perp) in ComDB[gid, 0]. Finally, the server notifies the party of the success or failure of the group creation. The group creator checks the protocol result. If group creation fails, it unwinds all state changes and outputs \perp .

Proposal. A party located at (gid, epoch) can generate a proposal message p by invoking $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ on input (Propose, act). Then, it executes the publish proposal protocol PublishProposal shown in Fig. 44, where party accesses the server via a client-anonymous authenticated channel. The party and the server perform a challenge-response type *membership authentication protocol*, that allows a party to anonymously prove that it is a valid group member for (gid, epoch). In more detail, upon receiving the PublishProposal proposal message, the server chooses a random challenge message $ch \leftarrow \{0, 1\}^k$ and sends it to the party. The party then signs the challenge with its group signing key $G.\text{ssk}$ and sends the signature σ along with the destination (gid, epoch) and the proposal p . The server checks that ComDB[gid, epoch] = (gvk, \perp , \perp) and the signature σ is valid with respect to gvk. If the check passes, the server stores the proposal p in PropDB[gid, epoch]. Note that the server rejects the proposal if ComDB[gid, epoch] contains a commit since this indicates that a new epoch has been created. We later show in App. E.2 that the above membership authentication protocol can be made non-interactive if we allow the server to perform an additional simple check on the database.

Commit. A party located at (gid, epoch) can perform a commit by invoking $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ on input (Commit, svk). The party first executes the subroutine FetchProposals shown in Fig. 45 to download the list of proposals created in (gid, epoch). FetchProposals consists of a challenge-response type membership authentication protocol almost identical to PublishProposal explained above. Once the server accepts, it is convinced that the calling anonymous party is indeed a valid group member at epoch so it sends the list of proposals \vec{p} stored in the database PropDB[gid, epoch].

Once the party succeeds in fetching the proposals \vec{p} , it generates a commit and welcome messages (c_0, \vec{c}, \vec{w}) by invoking $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$

on input (Commit, svk, \vec{p}). The list \vec{c} is then randomly permuted to \vec{c}_{perm} by *permute-commit . This procedure allows for a shuffle of the order of the member for each new epoch and makes the selective downloading performed during FetchCommit unlinkable between different epochs. The party then publishes $(c_0, \vec{c}_{\text{perm}}, \vec{w})$ by performing two different uploads. It first executes PublishCommit in Fig. 46 to publish the commit $(c_0, \vec{c}_{\text{perm}})$. The party accesses the server via a client-anonymous authenticated channel and performs the authentication protocol with the server similar to PublishProposal. After generating the response signature σ , the party further generates a group signature key (gvk', gsk') for the next epoch. This is generated from the group secret k_{mh}' at the next epoch, which can be obtained by querying (NextKey_{mh}, c_0) to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. The party finally sends $(\sigma, \text{gid}, \text{epoch}, c_0, \vec{c}_{\text{perm}}, \text{gvk}')$ to the server. If the signature σ is a valid signature with respect to the group signing key stored in ComDB[gid, epoch] = (gvk, \perp , \perp), then it updates the database by ComDB[gid, epoch] \leftarrow (gvk, c_0 , \vec{c}_{perm}). In case a commit was already stored, the server rejects the commit. Moreover, the server initializes a new entry in the database as ComDB[gid, epoch+1] \leftarrow (gvk', \perp , \perp). This creates a new epoch to which the parties can upload new proposals. The server returns whether it succeeded or not to the party.

If the party succeeds to publish the commit and welcome messages exit (i.e., $\vec{w} \neq \emptyset$), it then further executes PublishWelcome shown in Fig. 48 for each $\widehat{w} \in \vec{w}$. The party accesses the server via a client-anonymized authenticated channel and uploads $(\text{id}_t, \widehat{w})$, where id_t is extracted from \widehat{w} . The server stores \widehat{w} in WelDB[id_t]. Here, there is no membership authentication protocol.

Process. A party located at (gid, epoch) can try to process the current commit and proposals by invoking $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ on input (Process). The party first executes the FetchCommit subroutine shown in Fig. 47 to download a commit and their associating proposals. The FetchCommit protocol is similar to the FetchProposals protocol. The party and server engage in the membership authentication protocol explained above. Notably, the party sends an index that specifies the index of the party-dependent commit the party wants to download. Since the party sends an epoch-dependent permuted index using the function $\text{*permuted-commit-index}$, the index from different epochs remain unlinkable. If the server accepts the party, it returns the list of proposals \vec{p} stored on the database PropDB[gid, epoch], the party independent commit c_0 , and the party dependent commit $\vec{c} := \vec{c}_{\text{perm}}[\text{index}]$ stored on the database ComDB[gid, epoch].

Upon fetching the content (c_0, \vec{c}, \vec{p}) , the party processes them by invoking $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ on input (c_0, \vec{c}, \vec{p}) . Finally, it also updates the internal state by calling *update-states function.

Join. A party can join a group by invoking $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ on input (Join). The party first executes the FetchWelcome subroutine shown in Fig. 49 to download a welcome message. The party accesses the server via a standard authenticated channel (i.e., the party discloses its identity id). This is necessary for the server to identify which welcome message to provide. The server either returns the welcome message \widehat{w} designated to id or notifies that there is no welcome message. Once the party receives the welcome message $\widehat{w} \neq \perp$, it processes \widehat{w} by invoking $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ on input (Join, \widehat{w}). Finally, it setups the initial state by calling *init-states function.

E.2 Making Membership Authentication Protocol Non-Interactive

We discuss two simple ways to make the membership authentication protocol ran to publish and fetch the contents from the server non-interactive.

Recall that when the group member tries to access the server, they always engage in a membership authentication protocol to prove anonymously to the server that they are indeed a valid group member. Our W^{mh} protocol realized this membership authentication protocol in an *interactive* challenge-response type manner using a digital signature scheme. The server sends a challenge and asks the group member to sign the challenge. Assuming that the same challenge is never reused, this allowed the server to be convinced that the communicating party indeed has a group signing key.

We can remove this interaction between the server in some scenarios by asking the server to perform an extra check on the database. Namely, during the membership authentication protocol run during the protocols `PublishProposal` and `PublishCommit`, we allow the parties to sign on to the content they wish to upload to the server, rather than being provided a challenge from the server. For instance, in the protocol `PublishProposal`, the party id signs the message $(\text{gid}, \text{epoch}, p)$. The server checks if the signature is valid with respect to the group signing key at $(\text{gid}, \text{epoch})$, and additionally checks if such p is included in the proposal database `PropDB[gid, epoch]`. If all check passes, it updates the database with p . Notice that the server did not need to check if p was in the database in the previous interactive protocol. This non-interactive variant securely realizes the desired functionality since due to the EUF-CMA security of the signature scheme, a non-group member cannot upload any proposals that haven't been signed. The only possible attack would be to resend the observed signature-proposal pair to the server. However, since the server is modified to never accept the same proposal, this attack fails. The same idea can be used to make the protocol `PublishCommit` non-interactive.

Unfortunately, it is not clear how to make the protocols `FetchProposals` and `FetchCommit` non-interactive using the above idea. This is because the party does not have any contents to sign when it is performing a fetch/download. A potential idea is to weaken our UC-security model to allow non-group members to fetch/download the contents from the server, while still disallowing them to publish/upload any contents on the server. In particular, the protocols `FetchProposals` and `FetchCommit` will simply consist of a party querying the server $(\text{gid}, \text{epoch})$, and the server responding with the proposals and commits on the database without checking membership. This seems like a reasonable compromise considering that the contents uploaded on the server can be provided to an adversary without compromising the security of the group — this follows since the contents are commits and proposals generated from $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$, which by definition secures the group secret key and static metadata. One possible issue is that without any authentication on the fetch/download, it will allow an adversary to learn if a group with the group identifier gid exists. Our interactive protocol does not leak such information since it will output `accept = false` when authentication fails.

In summary, by reasonably weakening the ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ in a meaningful way, we can make the wrapper protocol W^{mh} fully non-interactive. It remains an interesting future work to investigate whether this weakening has any practical impact on the protocol.

E.3 Security of the Wrapper Protocol W^{mh}

The following theorem proves that the wrapper protocol W^{mh} UC-realizes the ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ in the $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ -hybrid model.

THEOREM E.1. *Assuming that SIG' is EUF-CMA secure, PRF is a secure pseudorandom function, and PRP is a secure pseudorandom permutation, the protocol W^{mh} UC-realizes the ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ in the $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ -hybrid model, where the safety predicates and leakage functions for $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ are defined in Figs. 34, 35 and 50.*

PROOF. We consider the following sequence of hybrids. While the environment \mathcal{Z} interacts with W^{mh} in Hybrid 1, it interacts with the ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ in Hybrid 6. Below, we first define all the hybrids and then explain how the simulators are defined.

Hybrid 1. This is the real-world execution of the protocol, where we make a syntactic change. We consider a simulator \mathcal{S}_1 that interacts with a dummy functionality $\mathcal{F}_{\text{dummy}}$. $\mathcal{F}_{\text{dummy}}$ sits between the environment \mathcal{Z} and \mathcal{S}_1 , and simply routes all messages without any modification. \mathcal{S}_1 internally simulates the real-world parties and adversary \mathcal{A} by routing all the messages sent from $\mathcal{F}_{\text{dummy}}$; from \mathcal{A} 's point of view, \mathcal{S}_1 is the environment \mathcal{Z} .

Hybrid 2. In this hybrid, we replace the dummy functionality $\mathcal{F}_{\text{dummy}}$ by the ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ except that we replace the functions used within `Create`, `Propose`, `Commit`, `Process`, and `Join` by those defined in Figs. 51 to 54. We call this modified ideal functionality $\mathcal{F}_{\text{CGKA},2}^{\text{mh}}$. In words, $\mathcal{F}_{\text{CGKA},2}^{\text{mh}}$ includes all the descriptions of the ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ for static metadata-hiding CGKA and outsources any other checks performed by the wrapper protocol W^{mh} to the simulator \mathcal{S}_2 . Namely, these correspond to the party-server interaction. In this hybrid, all consistency and security regarding the static metadata are guaranteed. The description of \mathcal{S}_2 is provided in Lem. E.2.

Hybrid 3. In this hybrid, we add all the missing consistency checks regarding the wrapper protocol W^{mh} of the ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ into $\mathcal{F}_{\text{CGKA},2}^{\text{mh}}$. We call this ideal functionality $\mathcal{F}_{\text{CGKA},3}^{\text{mh}}$. More precisely, $\mathcal{F}_{\text{CGKA},3}^{\text{mh}}$ is identical to $\mathcal{F}_{\text{CGKA}}^{\text{mh}}$ except that **random-index** (resp. **adv-access-allowed**) always returns false (resp. true). It only takes care of the correctness guarantees and does not guarantee any security properties regarding the dynamic metadata. Simulator \mathcal{S}_3 is defined identically to \mathcal{S}_2 .

Hybrid 4. In this hybrid, we change how the randomness used to derive the permutation key and group signature key are generated. The simulator \mathcal{S}_4 is identical to \mathcal{S}_3 except that, rather than generating `permKey` $\leftarrow \text{PRF}(k_{\text{mh}}, \text{'perm'})$ and `authKey` $\leftarrow \text{PRF}(k_{\text{mh}}, \text{'auth'})$ (which occurs during `Create` or `Commit`), if **safe** is true for that epoch, then it samples

a random permKey and authKey. In this hybrid, we use $\mathcal{F}_{\text{CGKA},4}^{\text{mh}} := \mathcal{F}_{\text{CGKA},3}$.

Hybrid 5. In this hybrid, we add the missing security guarantee on the randomness of the party's index. Namely, we modify $\mathcal{F}_{\text{CGKA},4}^{\text{mh}}$ to use the original **random-index** predicate, denoted as $\mathcal{F}_{\text{CGKA},5}^{\text{mh}}$. $\mathcal{F}_{\text{CGKA},5}^{\text{mh}}$ permutes the indices with a random permutation if predicate **random-index** is true. Simulator \mathcal{S}_5 is identical to \mathcal{S}_4 .

Hybrid 6. In this hybrid, we add the missing security guarantee when an adversary tries to access the honest server. Namely, we modify $\mathcal{F}_{\text{CGKA},5}^{\text{mh}}$ to use the original predicate **adv-access-allowed**, denoted as $\mathcal{F}_{\text{CGKA},6}^{\text{mh}}$. $\mathcal{F}_{\text{CGKA},6}^{\text{mh}}$ halts if an adversary succeeds to fetch or publish a proposal or commit message without knowing the corresponding group secret key. Simulator \mathcal{S}_6 is identical to \mathcal{S}_5 . At this point, $\mathcal{F}_{\text{CGKA},6}^{\text{mh}}$ is identical to the ideal functionality $\mathcal{F}_{\text{CGKA}}$.

We show indistinguishability of Hybrids 1 to 6 in Lems. E.2 to E.4, E.6 and E.8. This completes the proof of the main theorem. \square

E.3.1 From Hybrid 1 to 2: Lem. E.2.

LEMMA E.2. *Hybrid 1 and Hybrid 2 are perfectly indistinguishable.*

PROOF. We first provide the description of \mathcal{S}_2 . Whenever \mathcal{S}_2 is invoked on an input from $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ called within $\mathcal{F}_{\text{CGKA},2}^{\text{mh}}$,⁴¹ it simply relays the same input to the adversary \mathcal{A} . Since W^{mh} is constructed in the $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ -hybrid model, these inputs are exactly what \mathcal{A} was provided in Hybrid 1. \mathcal{S}_2 returns to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ whatever provided by \mathcal{A} . It remains to describe how \mathcal{S}_2 answers RegisterGroup, and publish and fetch queries sent from $\mathcal{F}_{\text{CGKA},2}^{\text{mh}}$. To this end, we first make a detour and explain how k_{mh} is set in Hybrids 1 and 2.

In Hybrid 1, k_{mh} for a commit node node-id is created when \mathcal{S}_1 invokes some party id with $\text{Ptr}[\text{id}] = \text{node-id}$ on a Create or a Commit. When id queries Create, id performs a query Key_{mh} to $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ within *init-states ; when id queries Commit, id performs a query $\text{NextKey}_{\text{mh}}$ within PublishCommit. If $\text{safe}(\text{node-id}) = \text{true}$, then $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ samples a random key k_{mh} . Otherwise, $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ asks \mathcal{A} to provide the key. These are the only two places where a new k_{mh} is set – Key_{mh} and $\text{NextKey}_{\text{mh}}$ are also queried during Propose, Process, and Join but k_{mh} will already be defined.

In Hybrid 2, however, \mathcal{S}_2 can no longer directly invoke a party id when the static metadata is hidden. For example, when \mathcal{Z} invokes party $\text{id}_{\text{creator}}$ on a Create while ***leak-create** is activated (i.e., the predicate **safe** is true), then \mathcal{S}_2 is only provided with $(|\text{id}_{\text{creator}}|, |\text{svk}|)$ from $\mathcal{F}_{\text{CGKA},2}^{\text{mh}}$. Without knowing $\text{id}_{\text{creator}}$, \mathcal{S}_2 cannot invoke $\text{id}_{\text{creator}}$ to execute *init-states . This in particular means that it cannot set k_{mh} like \mathcal{S}_1 did above by invoking Key_{mh} or $\text{NextKey}_{\text{mh}}$ of $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$.

With this in mind, we now finish the description of \mathcal{S}_2 . \mathcal{S}_2 maintains a list $L_{k_{\text{mh}}}$ to store the key k_{mh} generated at node node-id, i.e., $L_{k_{\text{mh}}}[\text{node-id}] = k_{\text{mh}}$. Notice that when \mathcal{S}_2 needs to simulate a RegisterGroup, or a fetch or publish query, it is given either (gid, epoch) if the server Sv is honest and $\text{Ptr}[\text{id}]$ if Sv is

malicious. Since there is no fork in the group when Sv is honest, (gid, epoch) uniquely defines node-id. Notably, even if \mathcal{S}_2 does not know who the calling party id is, it knows which group and epoch (or $\text{Ptr}[\text{id}]$ in case of a fork) id is included in, and thus, knows which $k_{\text{mh}} = L_{k_{\text{mh}}}[\text{node-id}]$ to use if it exists.

As explained above, when \mathcal{Z} invokes some party to perform Propose, Process, or Join, k_{mh} is already defined. Thus, \mathcal{S}_2 simply uses $k_{\text{mh}} = L_{k_{\text{mh}}}[\text{node-id}]$ to perform the same simulation as \mathcal{S}_1 , where recall id is not used during a publish or fetch protocol. When the environment \mathcal{Z} invokes $\text{id}_{\text{creator}}$ on (Create, svk), this is when a new k_{mh} is generated, i.e., $L_{k_{\text{mh}}}[\text{Ptr}[\text{id}_{\text{creator}}]]$ is still undefined. \mathcal{S}_2 checks if $\text{safe}(0) = \text{false}$ – which it can do since it can simulate an identical semantics of the history graph maintained within $\mathcal{F}_{\text{CGKA},2}^{\text{mh}}$ – and asks \mathcal{A} for k_{mh} . This is identical to what \mathcal{S}_1 did. Otherwise, \mathcal{S}_2 samples a random k_{mh} on its own. In either case, \mathcal{S}_2 stores $L_{k_{\text{mh}}}[0] \leftarrow k_{\text{mh}}$ and uses k_{mh} to perform the RegisterGroup protocol. The only difference between the previous hybrid is that \mathcal{S}_2 samples k_{mh} when $\text{safe}(0) = \text{true}$, rather than letting $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ sample it. However, since $\text{*mark-next-db-initialized-epoch}(0)$ is called and due to the restriction on \mathcal{A} , \mathcal{A} cannot compromise k_{mh} after Create was invoked. Therefore, from the view of \mathcal{Z} , k_{mh} is distributed identically in both hybrids. Thus, the simulation of Create is perfectly indistinguishable from Hybrid 1. The case when the environment \mathcal{Z} invokes id on (Commit, svk) is proven analogously to (Create, svk). This completes the proof. \square

E.3.2 From Hybrid 2 to 3: Lem. E.3.

LEMMA E.3. *Hybrid 2 and Hybrid 3 are indistinguishable assuming the correctness of SIG' .*

PROOF. The only difference between Hybrids 2 and 3 is that $\mathcal{F}_{\text{CGKA},3}^{\text{mh}}$ in Hybrid 3 mandates correctness of the RegisterGroup, and publish and fetch protocols when the server Sv is honest. For instance, in Hybrid 2, during the PublishProposal protocol, \mathcal{S}_2 simulated the interaction between the party id and the honest Sv using the key k_{mh} stored in the list $L_{k_{\text{mh}}}$. \mathcal{S}_2 then returned *accept'* output by Sv to $\mathcal{F}_{\text{CGKA},2}^{\text{mh}}$, and in particular, this notified the environment \mathcal{Z} that \mathcal{S} output *accept'*. In particular, $\mathcal{F}_{\text{CGKA},2}^{\text{mh}}$ did not model any notion of correctness of the PublishProposal protocol.

On the other hand, in Hybrid 3, if $\text{ComDB}[\text{gid}, \text{epoch}] = (\top, \text{node-id})$ and $\text{node-id} = \text{Ptr}[\text{id}]$ (i.e., the group statement gvk is created honestly and id holds the corresponding signing key gsk), then $\mathcal{F}_{\text{CGKA},3}^{\text{mh}}$ always sends *accept* = true to \mathcal{Z} . This models the fact that if the new epoch (node-id) is initialized by an honest party, then any member assigned on the same commit node (i.e., $\text{node-id} = \text{Ptr}[\text{id}]$) can upload a proposal. It is clear that this holds assuming that SIG' is correct, i.e., a properly generated signature is always accepted. $\mathcal{F}_{\text{CGKA},3}^{\text{mh}}$ always sends *accept* = false to \mathcal{Z} if (1) $\text{ComDB}[\text{gid}, \text{epoch}] \neq (\top, \text{node-id})$ or (2) $\text{ComDB}[\text{gid}, \text{epoch}] = (\top, \text{node-id})$ and $\text{node-id} \neq \text{Ptr}[\text{id}]$. For case (1), we can verify that the real server performs the same check as $\text{ComDB}[\text{gid}, \text{epoch}]$ is equal to (gvk, \perp , \perp) or not, and thus \mathcal{S}_2 returned *accept'* = false. For case (2), if \mathcal{S}_2 returned *accept'* = true in case (2), we can construct an adversary that breaks the EUF-CMA security of SIG' by using such \mathcal{S}_2 . Thus \mathcal{S}_2 returned *accept'* = false assuming

⁴¹Note that this is not strictly true since $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ is not defined within $\mathcal{F}_{\text{CGKA},2}^{\text{mh}}$. $\mathcal{F}_{\text{CGKA},2}^{\text{mh}}$ merely has most of the codes included in $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$. We use this informal wording for simplicity.

the EUF-CMA security of SIG' . Therefore, the view to \mathcal{Z} remain identical.

All other RegisterGroup, and publish and fetch protocols in Hybrid 3 can be checked to behave identically to those of Hybrid 2 conditioned on SIG' being correct and EUF-CMA secure. \square

E.3.3 From Hybrid 3 to 4: Lem. E.4.

LEMMA E.4. *Hybrid 3 and Hybrid 4 are indistinguishable assuming PRF is a secure pseudorandom function.*

PROOF. We assume \mathcal{Z} creates at most Q epochs (i.e., commit nodes node-id in the history graph). To show Lem. E.4, we consider the following sub-hybrids between Hybrid 3 and Hybrid 4.

Hybrid 3-0 := Hybrid 3. This is identical to Hybrid 3. We use the functionality $\mathcal{F}_{CGKA,3}$, and the simulator $\mathcal{S}_{3-0} := \mathcal{S}_3$. When simulating protocol, \mathcal{S}_{3-0} uses the group key k_{mh} stored in the list $L_{k_{mh}}$ as PRF key.

Hybrid 3- i . i runs through $[Q]$. The simulator \mathcal{S}_{3-i} is defined exactly as $\mathcal{S}_{3-(i-1)}$ except that when $\mathcal{S}_{3-(i-1)}$ generates a random k_{mh} for the i -th epoch where **safe** is true, it chooses a random authKey and permKey instead of deriving them from PRF and k_{mh} . Note that we count epochs (i.e., node-id) in the order in which Create or Commit is invoked. We show in Lem. E.5 that Hybrid 3- $(i-1)$ and Hybrid 3- i are indistinguishable.

Hybrid 3- Q := Hybrid 4. In this hybrid, all authKey and permKey generated at epochs such that **safe** is true are chosen at random.

The indistinguishability between Hybrid 3 and Hybrid 4 is established by applying the following Lem. E.5 for all $i \in [Q]$.

LEMMA E.5. *Hybrid 3- $(i-1)$ and Hybrid 3- i are indistinguishable assuming PRF is a secure pseudo-random function.*

PROOF. The difference between Hybrid 3- $(i-1)$ and Hybrid 3- i is whether authKey and permKey at the i -th epoch (i.e., node-id) are chosen uniformly at random if **safe** is true. (If **safe** is false for the i -th epoch, two hybrids proceed identically.) Due to the modification we made in Hybrid 2, the group key k_{mh} is chosen uniformly at random. Moreover, due to the restriction on the adversary, k_{mh} generated when **safe** was true cannot be corrupted by the adversary. Thus, by the pseudorandomness of the PRF, the output of the PRF is indistinguishable from a random string for the input labels 'perm' and 'auth'. This implies Hybrid 3- $(i-1)$ and Hybrid 3- i are indistinguishable. More formally, if \mathcal{Z} can distinguish the two hybrids, then there exists an adversary \mathcal{B} that breaks the pseudorandomness of the PRF. We first explain the description of \mathcal{B} ; then we evaluate \mathcal{B} 's advantage.

\mathcal{B} simulates for \mathcal{Z} the protocol executions between a party and the server as in $\mathcal{S}_{3-(i-1)}$, except for the evaluation of the PRF. To generate authKey and permKey, \mathcal{B} queries authKey := \mathcal{F} ('auth') and permKey := \mathcal{F} ('perm') to its oracle \mathcal{F} .

We evaluate the success probability of \mathcal{B} . If the oracle \mathcal{F} evaluates a pseudo-random function, then \mathcal{Z} 's view is identical to Hybrid 3- $(i-1)$. If \mathcal{F} evaluates a truly random function, then \mathcal{Z} 's view is identical to Hybrid 3- i . Hence, if \mathcal{Z} distinguishes Hybrid 3- $(i-1)$ and Hybrid 3- i with non-negligible probability, \mathcal{B} breaks the pseudorandomness of the PRF with non-negligible probability. This

contradicts the assumption that PRF is secure. Therefore, Hybrid 3- $(i-1)$ and Hybrid 3- i are indistinguishable. \square

\square

E.3.4 From Hybrid 4 to 5: Lem. E.6. Hybrid 5 concerns the randomness of the access index of group members. If **random-index** is true, then the ideal functionality $\mathcal{F}_{CGKA,5}^{mh}$ randomizes the commit vector \vec{c} and the index of each party in the group on behalf of the simulator. We prove in Lem. E.6 that assuming that PRP is secure, Hybrid 4 and Hybrid 5 are indistinguishable.

LEMMA E.6. *Hybrid 4 and Hybrid 5 are indistinguishable assuming PRP is a secure pseudorandom permutation.*

PROOF. We assume \mathcal{Z} creates at most Q epochs (i.e., commit nodes node-id in the history graph). To show Lem. E.6, we consider the following sub-hybrids between Hybrid 4 and Hybrid 5.

Hybrid 4-0 := Hybrid 4. This is identical to Hybrid 4. We use the functionality $\mathcal{F}_{CGKA,4}$, and the simulator $\mathcal{S}_{4-0} := \mathcal{S}_4$. In this hybrid, $\mathcal{F}_{CGKA,4}$ permutes the indices index and commit vectors \vec{c} with a permutation chosen by \mathcal{S}_{4-0} , and \mathcal{S}_{4-0} simulates the PublishCommit and FetchCommit protocols using the indices and commit messages provided from $\mathcal{F}_{CGKA,4-i}$.

Hybrid 4- i . i runs through $[Q]$. We use the functionality $\mathcal{F}_{CGKA,4-i}$ which is defined exactly as $\mathcal{F}_{CGKA,4-(i-1)}$ except that the indices index and commit vectors \vec{c} are permuted with a truly random permutation if **random-index** is true for the i -th epoch. The simulator \mathcal{S}_{4-i} is defined exactly as $\mathcal{S}_{4-(i-1)}$. Note that we count epochs (i.e., node-id) in the order in which Create or Commit is invoked. We show in Lem. E.7 that Hybrid 4- $(i-1)$ and Hybrid 4- i are indistinguishable.

Hybrid 4- Q := Hybrid 5. We use the functionality $\mathcal{F}_{CGKA,4-Q}$ which permutes the indices index and commit vectors \vec{c} with a random permutation for all epochs where **random-index** is true. $\mathcal{F}_{CGKA,4-Q}$ is identical to $\mathcal{F}_{CGKA,5}$.

Indistinguishability between Hybrid 4 and Hybrid 5 is established by applying the following Lem. E.7 for all $i \in [Q]$.

LEMMA E.7. *Hybrid 4- $(i-1)$ and Hybrid 4- i are indistinguishable assuming PRP is a secure pseudorandom permutation.*

PROOF. The difference between Hybrid 4- $(i-1)$ and Hybrid 4- i is whether index and \vec{c} issued at the i -th epoch where **random-index** is true are permuted by a truly random permutation. (If **random-index** is false for the i -th epoch, two hybrids proceed identically.) Due to the modification we made in Hybrid 4, the i -th permutation key permKey is chosen uniformly at random if **random-index** (=safe) is true. Thus, by the pseudorandomness of the PRP, the output of the PRP is indistinguishable from that of a truly random permutation. This implies Hybrid 4- $(i-1)$ and Hybrid 4- i are indistinguishable. More formally, if \mathcal{Z} can distinguish the two hybrids, then we can construct an adversary \mathcal{B} that breaks the pseudorandomness of the PRP. We first explain the description of \mathcal{B} ; then we evaluate \mathcal{B} 's advantage.

\mathcal{B} simulates for \mathcal{Z} the protocol executions between a party and the server as in $\mathcal{S}_{4-(i-1)}$, except for the evaluation of PRP at the i -th epoch. Rather than sampling a permutation key permKey for

the i -th epoch to evaluate the permutation, \mathcal{B} queries its challenge oracle to compute $\text{index}' \leftarrow \mathcal{P}(\text{index})$.

We evaluate the success probability of \mathcal{B} . If the oracle \mathcal{P} evaluates a pseudorandom permutation, then \mathcal{Z} 's view is identical to Hybrid 4- $(i-1)$. Otherwise, if \mathcal{P} evaluates a truly random permutation, then \mathcal{Z} 's view is identical to Hybrid 4- i . Hence, if \mathcal{Z} distinguishes Hybrid 4- $(i-1)$ and Hybrid 4- i with non-negligible probability, \mathcal{B} breaks the pseudorandomness of the PRP with non-negligible probability. This contradicts the assumption that PRP is secure. Therefore, Hybrid 4- $(i-1)$ and Hybrid 4- i are indistinguishable. \square

\square

E.3.5 From Hybrid 5 to 6: Lem. E.8. Hybrid 6 concerns the group membership check performed by the honest server. When the server is honest (i.e., $\text{ServerStat} = \text{'good'}$) and **adv-access-allowed** is false (i.e., group secrets are not compromised), the functionality $\mathcal{F}_{\text{CGKA},6}^{\text{mh}}$ halts if an adversary succeeds to publish or fetch a proposal or commit messages. That is an adversary that does not know the group secret should not be able to publish or fetch contents from the honest server. We prove in Lem. E.8 that if \mathcal{Z} can distinguish the two hybrids, then we can construct an adversary that breaks the EUF-CMA security of SIG' . In other words, assuming that SIG' is EUF-CMA secure, Hybrid 5 and Hybrid 6 are indistinguishable. Concretely, we show the following.

LEMMA E.8. *Hybrid 5 and Hybrid 6 are indistinguishable assuming SIG' is EUF-CMA secure.*

PROOF. We show that, if \mathcal{Z} can distinguish the two hybrids, then there exists an adversary \mathcal{B} that breaks the EUF-CMA security of SIG' . We first explain the description of \mathcal{B} and how \mathcal{B} extracts a valid signature forgery using \mathcal{Z} ; we then show the validity of the forged signature and finally evaluate \mathcal{B} 's advantage.

\mathcal{B} simulates for \mathcal{Z} the protocol executions between a party and the server as in \mathcal{S}_6 , except for the signature generation during the publish and fetch protocols. At the beginning of the game, \mathcal{B} chooses an index $i \in [Q]$ at random, where Q is the total number of epochs (i.e., commit nodes node-id) that \mathcal{Z} creates. \mathcal{B} aborts if **adv-access-allowed** is true for the i -th epoch, and otherwise embeds the challenge signing key svk^* in the i -th group signing key gvk . Here, if **adv-access-allowed** is false (i.e., **safe** is true) for the i -th epoch, the group signing key is generated with a uniformly random authKey due to the modification we made in Hybrid 4. Therefore, \mathcal{B} perfectly simulates svk^* as in Hybrid 5. Moreover, note that once **safe** is true, the restricted adversary cannot compromise the i -th group signing key. Whenever an honest party at the i -th epoch publishes or fetches proposal/commit messages, \mathcal{B} uses the signing oracle to sign the challenge message sent from the server. For the other epochs, \mathcal{B} generates gvk and signs as in the previous hybrid. If the adversary succeeds in making the server accept at execution of a publish and fetch proposal/commit protocols in the i -th epoch, \mathcal{B} retrieves the pair of challenge message and response signature (ch, σ) from the successful transcript, and submits (ch, σ) as the forgery.

Let us analyze the success probability of \mathcal{B} . First, by noticing the only way that \mathcal{Z} can distinguish the two hybrids is by triggering the **assert** condition on **adv-access-allowed**, there must exist at

least one epoch for which the adversary succeeds in the protocol $\text{PublishProposalAdv}$, where **adv-access-allowed** is false. Since the choice $i \leftarrow_{\$} [Q]$ of \mathcal{B} is information-theoretically hidden from \mathcal{Z} and the adversary, the guess made by \mathcal{B} is correct with probability at least $1/Q$. Conditioning on the guess being correct, (ch, σ) is a valid message and signature pair for gvk . Moreover, since the server is honest and the challenge message space is exponentially large, the server never picks the same challenge message ch . Therefore, (ch, σ) is a pair of message and signature that \mathcal{B} did not query to the signing oracle and constitutes a valid forgery. In summary, if \mathcal{Z} distinguishes the two hybrids with non-negligible probability ϵ , \mathcal{B} wins the game with probability at least ϵ/Q , which is also non-negligible. This contradicts the assumption that SIG' is EUF-CMA secure. Thus, Hybrid 5 and Hybrid 6 are indistinguishable. \square

Input (Create, svk)

```

1: req  $G = \perp$ 
2: try query (Create, svk) to  $\mathcal{F}_{CGKA}^{\text{ctxt}}$ 
3:  $G \leftarrow *init\text{-states}(gid, 0, \{ (id_{creator}, svk) \})$ 
4: req RegisterGroup(gid, 0)

```

Input (Propose, act), $act \in \{ \text{'upd'-'svk', 'add'-'id}_t\text{-kp}_t, \text{'rem'-'id}_t \}$

```

1: req  $G \neq \perp$ 
2: try  $p \leftarrow$  query (Propose, act) to  $\mathcal{F}_{CGKA}^{\text{ctxt}}$ 
3: req PublishProposal( $G.gid, G.epoch, p$ )
4: return  $p$ 

```

Input (Commit, svk)

```

1: req  $G \neq \perp$ 
2:  $(accept, \vec{p}) \leftarrow$  FetchProposals( $G.gid, G.epoch$ )
3: req accept
4: try  $(c_0, \vec{c}, \vec{w}) \leftarrow$  query (Commit, svk,  $\vec{p}$ ) to  $\mathcal{F}_{CGKA}^{\text{ctxt}}$ 
5:  $\vec{c}_{perm} \leftarrow *permute\text{-commit}(G, \vec{c})$ 
6: try PublishCommit( $G.gid, G.epoch, c_0, \vec{c}_{perm}$ )
7: foreach  $\vec{w} \in \vec{w}$  do
8:   parse  $(id_t, *) \leftarrow \vec{w}$ 
9:   try PublishWelcome( $id_t, \vec{w}$ )
10: return  $(c_0, \vec{c}_{perm}, \vec{w})$ 

```

Input Process

```

1: req  $G \neq \perp$ 
2:  $index_{\vec{c}} \leftarrow *permuted\text{-commit-index}(G, id)$ 
3: try  $(c_0, \vec{c}, \vec{p}) \leftarrow$  FetchCommit( $G.gid, G.epoch, index_{\vec{c}}$ )
4: try  $(id_c, propSem, mem) \leftarrow$  query (Process,  $c_0, \vec{c}, \vec{p}$ ) to  $\mathcal{F}_{CGKA}^{\text{ctxt}}$ 
5:  $G' \leftarrow *update\text{-states}(G, mem)$ 
6: return  $(id_c, propSem, mem)$ 

```

Input (Key)

```

1: req  $G \neq \perp$ 
2: try  $k \leftarrow$  query (Key) to  $\mathcal{F}_{CGKA}^{\text{ctxt}}$ 
3: return  $k$ 

```

Input Join

```

1: req  $G = \perp$ 
2: try  $\widehat{w} \leftarrow$  FetchWelcome(id)
3: try  $(id_c, gid, epoch, mem) \leftarrow$  query (Join,  $\widehat{w}$ ) to  $\mathcal{F}_{CGKA}^{\text{ctxt}}$ 
4:  $G \leftarrow *init\text{-states}(gid, epoch, mem)$ 
5: return  $(id_c, gid, epoch, mem)$ 

```

*init-states(gid, epoch, mem)

```

1:  $G.gid \leftarrow gid; G.epoch \leftarrow epoch; G.mem \leftarrow mem$ 
2:  $k_{mh} \leftarrow$  query Key $_{mh}$  to  $\mathcal{F}_{CGKA}^{\text{ctxt}}$ 
3:  $G.permKey \leftarrow$  PRF( $k_{mh}, \text{'perm'}$ )
4:  $authKey \leftarrow$  PRF( $k_{mh}, \text{'auth'}$ )
5:  $(G.gvk, G.gsk) \leftarrow$  SIG'.KeyGen( $pp_{SIG'}, authKey$ )
6: return  $G$ 

```

*update-states(G, mem)

```

1:  $G'.gid \leftarrow G.gid; G'.epoch \leftarrow G.epoch + 1$ 
2:  $G'.mem \leftarrow mem$ 
3:  $k_{mh} \leftarrow$  query Key $_{mh}$  to  $\mathcal{F}_{CGKA}^{\text{ctxt}}$ 
4:  $G'.permKey \leftarrow$  PRF( $k_{mh}, \text{'perm'}$ )
5:  $authKey \leftarrow$  PRF( $k_{mh}, \text{'auth'}$ )
6:  $(G'.gvk, G'.gsk) \leftarrow$  SIG'.KeyGen( $pp_{SIG'}, authKey$ )
7: return  $G'$ 

```

*permute-commit(G, \vec{c})

```

1:  $\vec{c}_{perm} \leftarrow ()$ 
2: for index = 1, ...,  $|\vec{c}|$  do
3:    $\vec{c}_{perm} \# \leftarrow \vec{c}[\text{PRP}(G.permKey, index)]$ 
4: return  $\vec{c}_{perm}$ 

```

*permuted-commit-index(G, id)

```

1: index  $\leftarrow G.indexOf(id)$ 
2: return PRP( $G.permKey, index$ )

```

Figure 42: Metadata-hiding CGKA protocol W^{mh} in the $\mathcal{F}_{CGKA}^{\text{ctxt}}$ -hybrid model: Create, Propose, Commit, Process, Join, Key, and some helper functions.

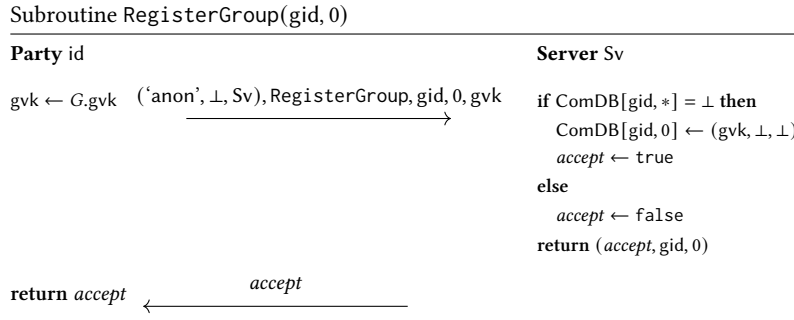


Figure 43: Subroutines for metadata-hiding CGKA protocol W^{mh} in the $\mathcal{F}_{CGKA}^{ctxt}$ -hybrid model: Register a new group and initialize group states for the current epoch. Party id and the server Sv are connected via a client-anonymized authenticated channel.

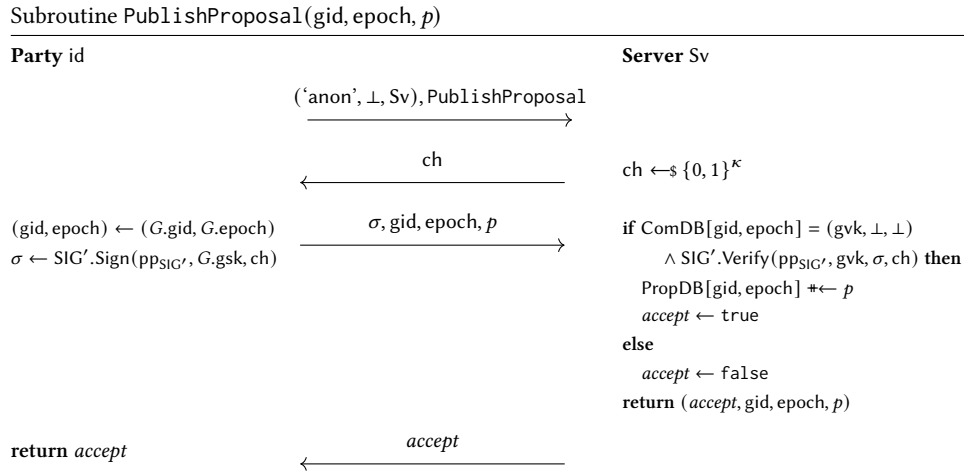


Figure 44: Subroutines for metadata-hiding CGKA protocol W^{mh} in the $\mathcal{F}_{CGKA}^{ctxt}$ -hybrid model: Publish proposal messages. Party id and the server Sv are connected via a client-anonymized authenticated channel.

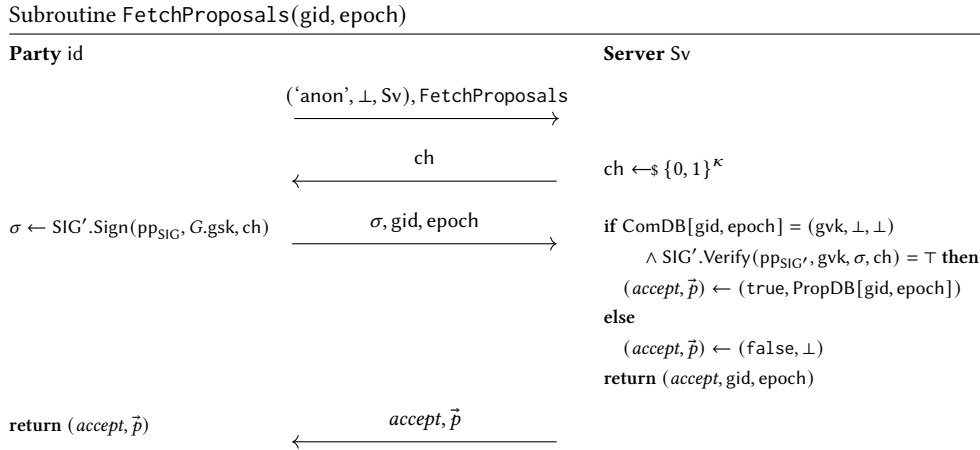


Figure 45: Subroutines for metadata-hiding CGKA protocol W^{mh} in the $\mathcal{F}_{CGKA}^{ctxt}$ -hybrid model: Fetch proposal messages. Party id and the server Sv are connected via a client-anonymized authenticated channel.

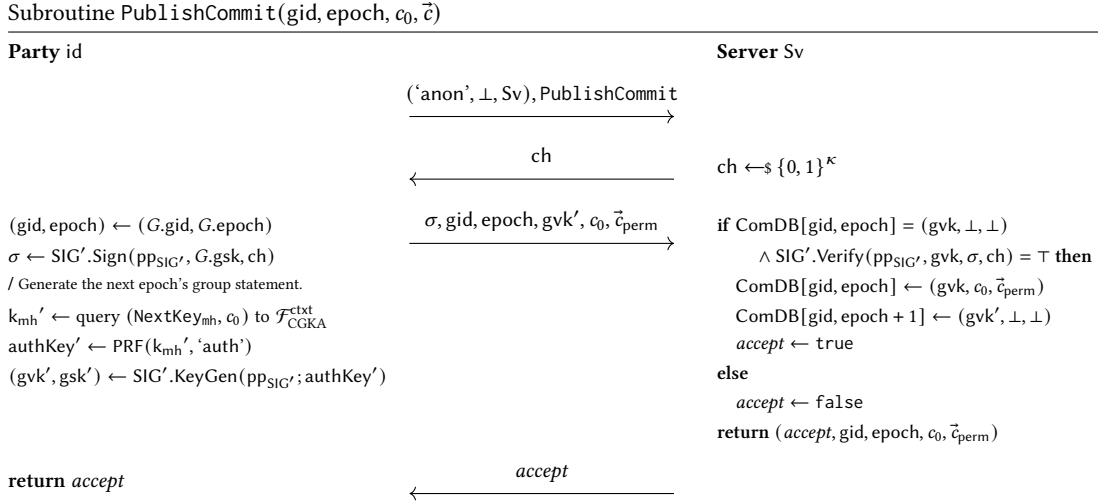


Figure 46: Subroutines for metadata-hiding CGKA protocol W^{mh} in the $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ -hybrid model: Publish commit messages. Party id and the server Sv are connected via a client-anonymized authenticated channel.

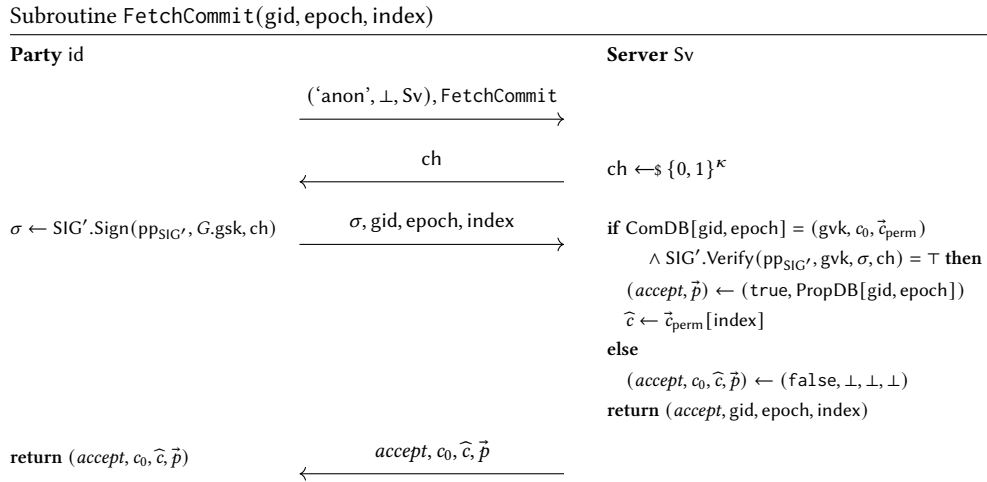


Figure 47: Subroutines for metadata-hiding CGKA protocol W^{mh} in the $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ -hybrid model: Fetch commit messages. Party id and the server Sv are connected via a client-anonymized authenticated channel.

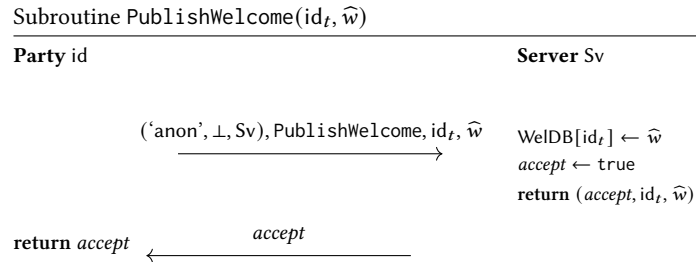


Figure 48: Subroutines for metadata-hiding CGKA protocol W^{mh} in the $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ -hybrid model: Publish and fetch welcome messages. Party id and the server Sv are connected via a client-anonymized authenticated channel.

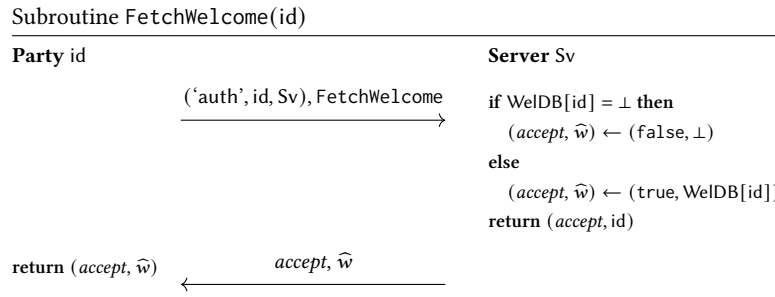


Figure 49: Subroutines for metadata-hiding CGKA protocol W^{mh} in the $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ -hybrid model: Fetch welcome messages. Party id and the server Sv are connected via an authenticated channel. Note that the channel does not need to be anonymized.

Random index and group authentication.

$\text{random-index}(c) \iff \text{safe}(c)$
 $\text{adv-access-allowed}(c) \iff \neg \text{safe}(c)$

Figure 50: Additional safety predicates unique to the wrapper protocol W^{mh} . The other safety predicates used to implicitly define the underlying ideal functionality $\mathcal{F}_{\text{CGKA}}^{\text{ctxt}}$ are identical to those provided in App. C.2.

Subroutine RegisterGroup(gid, epoch)

```

1 : channelType ← ('anon', ⊥, Sv) / Connect to the server via anonymous channel
2 : req epoch = 0
3 : Send (channelType, RegisterGroup, gid, 0) to S and receive accept
4 : if ServerStat = 'good' then
5 :   Send (accept, gid, 0) to Sv
6 : return accept

```

Subroutine PublishProposal(gid, epoch, p)

```

1 : channelType ← ('anon', ⊥, Sv) / Connect to the server via anonymous channel
2 : if ServerStat = 'good' then
3 :   Send (channelType, PublishProposal, gid, epoch, p) to S and
4 :     receive accept
5 :   Send (accept, gid, epoch, p) to Sv
6 : else
7 :   Send (channelType, PublishProposal, Ptr[id], p) to S and
8 :     receive accept
9 : return accept

```

Subroutine FetchProposals(gid, epoch)

```

1 : channelType ← ('anon', ⊥, Sv) / Connect to the server via anonymous channel
2 : if ServerStat = 'good' then
3 :   Send (channelType, FetchProposals, gid, epoch) to S and
4 :     receive (accept, p̄)
5 :   Send (accept, gid, epoch) to Sv
6 : else
7 :   Send (channelType, FetchProposals, Ptr[id]) to S and
8 :     receive (accept, p̄)
9 : return (accept, p̄)
10 :

```

Subroutine PublishCommit(gid, epoch, c₀, c̄_{perm})

```

1 : channelType ← ('anon', ⊥, Sv) / Connect to the server via anonymous channel
2 : if ServerStat = 'good' then
3 :   Send (channelType, PublishCommit, gid, epoch, c0, c̄perm) to S and
4 :     receive accept
5 :   Send (accept, gid, epoch, c0, c̄perm) to Sv
6 : else
7 :   Send (channelType, PublishCommit, Ptr[id], c0, c̄perm) to S and
8 :     receive accept
9 : return accept

```

Subroutine FetchCommit(gid, epoch, index_{c̄})

```

1 : channelType ← ('anon', ⊥, Sv) / Connect to the server via anonymous channel
2 : if ServerStat = 'good' then
3 :   Send (channelType, FetchCommit, gid, epoch, indexc̄) to S and
4 :     receive (accept, c0, c̄, p̄)
5 :   Send (accept, gid, epoch, indexc̄) to Sv
6 : else
7 :   Send (channelType, FetchCommit, Ptr[id], indexc̄) to S and
8 :     receive (accept, c0, c̄, p̄)
9 : return (accept, c0, c̄, p̄)

```

Figure 51: Subroutines for publish and fetch proposal and commit messages used in Hybrid 2.

Input from a party id or adversary S

Input PublishWelcome(id_t, w̄)

```

1 : / Connect to the server via anonymous channel
2 : channelType ← ('anon', ⊥, Sv)
3 : Send (channelType, PublishWelcome, idt, w̄) to S and
4 :   receive accept
5 : if ServerStat = 'good' then
6 :   Send (accept, idt, w̄) to Sv
7 : return accept

```

Input FetchWelcome(id)

```

1 : / Connect to the server via authenticated channel
2 : channelType ← ('auth', id, Sv)
3 : Send (channelType, FetchWelcome) to S and receive (accept, w̄)
4 : if ServerStat = 'good' then
5 :   Send (accept, id) to Sv
6 : return (accept, w̄)

```

Figure 52: Subroutines for publish and fetch welcome messages used in Hybrid 2.

*permuted-commit-index(node-id, id)	*permute-commit(node-id, \vec{c})
<pre> 1: assert Node[node-id] $\neq \perp$ 2: / If flag_{selDL} is false, there is no need to permute index. 3: if \negflag_{selDL} then 4: return \perp 5: mem \leftarrow Node[node-id].mem 6: assert id \in mem 7: Send (*permuted-commit-index, Ptr[id], id) to \mathcal{S} and receive index$_{\vec{c}}$ 8: return index$_{\vec{c}}$ </pre>	<pre> 1: assert Node[node-id] $\neq \perp$ 2: / If flag_{selDL} is false, no party-dependent commitment can exist. 3: if \negflag_{selDL} then 4: assert $\vec{c} = \perp$ 5: return \perp 6: Send (*permute-commit, Ptr[id], \vec{c}) to \mathcal{S} and receive \vec{c}_{perm} 7: return \vec{c}_{perm} </pre>

Figure 53: Helper functions: permute functions used in Hybrid 2.

Input from the adversary \mathcal{S} Input (PublishProposalAdv, gid, epoch, p)

```

1: channelType  $\leftarrow$  ('anon',  $\perp$ , Sv)
2: Send (channelType, PublishProposalAdv, gid, epoch,  $p$ ) to  $\mathcal{S}$  and
   receive accept
3: if ServerStat = 'good' then
4:   Send (channelType, PublishProposal, (accept, gid, epoch,  $p$ )) to Sv
5: return accept

```

Input (FetchProposalsAdv, gid, epoch)

```

1: channelType  $\leftarrow$  ('anon',  $\perp$ , Sv)
2: Send (channelType, FetchProposalsAdv, gid, epoch) to  $\mathcal{S}$  and
   receive (accept,  $\vec{p}$ )
3: if ServerStat = 'good' then
4:   return (accept,  $\vec{p}$ )

```

Input (PublishCommitAdv, gid, epoch, c_0 , \vec{c})

```

1: channelType  $\leftarrow$  ('anon',  $\perp$ , Sv)
2: Send (channelType, PublishCommitAdv, gid, epoch,  $c_0$ ,  $\vec{c}$ ) to  $\mathcal{S}$  and
   receive accept
3: if ServerStat = 'good' then
4:   return accept

```

Input (FetchCommitAdv, gid, epoch, index)

```

1: channelType  $\leftarrow$  ('anon',  $\perp$ , Sv)
2: Send (channelType, FetchCommitAdv, gid, epoch, index) to  $\mathcal{S}$  and
   receive (accept,  $c_0$ ,  $\vec{c}$ ,  $\vec{p}$ )
3: if ServerStat = 'good' then
4:   return (accept,  $c_0$ ,  $\vec{c}$ ,  $\vec{p}$ )

```

Figure 54: Publish functions used by the adversary in Hybrid 2.

CONTENTS

Abstract	1	A.3	Secret Key Encryption	15
1 Introduction	1	A.4	Digital Signatures	16
1.1 Goal of This Work	2	A.5	Message Authentication Codes	16
1.2 Our Contributions	3	A.6	HKDF	16
2 Background	4	A.7	Pseudorandom Function	16
2.1 Syntax of CGKA	4	A.8	Pseudorandom Permutation	16
2.2 Default UC Security Model of CGKA	4	B	Static Metadata-Hiding CGKA: Definition	17
3 Hiding Static Metadata in CGKA	5	B.1	Background	17
3.1 UC Security Model for Static Metadata	5	B.2	UC Security Model and $\mathcal{F}_{CGKA}^{ctxt}$	19
3.2 Proof of Static Metadata-Hiding CGKA	6	C	Static Metadata-Hiding CGKA: Construction and Security Proof	29
4 Constructing Metadata-Hiding CGKA	7	C.1	Constructing Chained CmpKE ^{ctxt}	29
4.1 Goal of the Wrapper Protocol W^{mh}	7	C.2	Safety Predicates and Leakage Functions	37
4.2 High Level Description of W^{mh}	7	C.3	Security of Chained CmpKE ^{ctxt}	40
5 Formal Model for Metadata-Hiding CGKA	9	D	Metadata-Hiding CGKA: Definition	49
5.1 Modeling an Honest but Curious Sever	9	D.1	An Overview of \mathcal{F}_{CGKA}^{mh}	50
5.2 UC Security Model for Dynamic Metadata	9	D.2	Functions Used by Legitimate Parties	50
5.3 Proof of Dynamic Metadata-Hiding CGKA	10	D.3	Functions Used by the Adversary	53
6 Instantiation and Efficiency	10	E	Metadata-Hiding CGKA: Construction and Security Proof	59
6.1 Instantiation	10	E.1	Constructing the Wrapper Protocol W^{mh}	59
6.2 Efficiency	11	E.2	Making Membership Authentication Protocol Non-Interactive	61
7 Limitation of Efficient Metadata-Hiding CGKA	11	E.3	Security of the Wrapper Protocol W^{mh}	61
7.1 Chained CmpKE and TreeKEM	11	Contents		71
7.2 Server-Aided Variants of TreeKEM	12			
References	13			
A Definition of General Cryptographic Primitives	14			
A.1 Notation	14			
A.2 Decomposable Multi-Recipient Public Key Encryption	15			