

Modular Design of Secure Group Messaging Protocols and the Security of MLS

Joël Alwen AWS Wickr alwenjo@amazon.com

Yevgeniy Dodis New York University dodis@cs.nyu.edu

ABSTRACT

The Messaging Layer Security (MLS) project is an IETF effort aiming to establish an industry-wide standard for *secure group messaging (SGM)*. Its development is supported by several major securemessaging providers (with a combined user base in the billions) and a growing body of academic research.

MLS has evolved over many iterations to become a complex, non-trivial, yet relatively ad-hoc cryptographic protocol. In an effort to tame its complexity and build confidence in its security, past analyses of MLS have restricted themselves to sub-protocols of MLS—most prominently a type of sub-protocol embodying socalled *continuous group key agreement (CGKA)*. However, to date the task of proving or even defining the security of the full MLS protocol has been left open.

In this work, we fill in this missing piece. First, we formally capture the security of SGM protocols by defining a corresponding security game, which is parametrized by a safety predicate that characterizes the exact level of security achieved by a construction. Then, we cast MLS as an SGM protocol, showing how to modularly build it from the following three main components (and some additional standard cryptographic primitives) in a black-box fashion: (a) CGKA, (b) forward-secure group AEAD (FS-GAEAD), which is a new primitive and roughly corresponds to an "epoch" of group messaging, and (c) a so-called PRF-PRNG, which is a two-input hash function that is a pseudorandom function (resp. generator with input) in its first (resp. second) input. Crucially, the security predicate for the SGM security of MLS can be expressed purely as a function of the security predicates of the underlying primitives, which allows to swap out any of the components and immediately obtain a security statement for the resulting SGM construction. Furthermore, we provide instantiations of all component primitives, in particular of CGKA with MLS's TreeKEM sub-protocol

CCS '21, November 15-19, 2021, Virtual Event, Republic of Korea.

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8454-4/21/11...\$15.00

https://doi.org/10.1145/3460120.3484820

Sandro Coretti IOHK sandro.coretti@iohk.io

Yiannis Tselekounis University of Edinburgh y.tselekounis@ed.ac.uk

(which we prove adaptively secure) and of FS-GAEAD with a novel construction (which has already been adopted by MLS).

Along the way we introduce a collection of new techniques, primitives, and results with applications to other SGM protocols and beyond. For example, we extend the Generalized Selective Decryption proof technique (which is central in CGKA literature) and prove adaptive security for another (practical) *more secure* CGKA protocol called RTreeKEM (Alwen et al., CRYPTO '20). The modularity of our approach immediately yields a corollary characterizing the security of an SGM construction using RTreeKEM.

CCS CONCEPTS

• Security and privacy \rightarrow *Cryptography*.

KEYWORDS

Cryptographic protocols; Messaging Layer Security; MLS; TreeKEM; RTreeKEM; Secure Messaging; Forward Secrecy; Backward Secrecy

ACM Reference Format:

Joël Alwen, Sandro Coretti, Yevgeniy Dodis, and Yiannis Tselekounis. 2021. Modular Design of Secure Group Messaging Protocols and the Security of MLS. In Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (CCS '21), November 15–19, 2021, Virtual Event, Republic of Korea. ACM, New York, NY, USA, 21 pages. https://doi.org/10. 1145/3460120.3484820

ACKNOWLEDGMENTS

This research is partially supported by gifts from VMware Labs and Google, NSF grants 1619158, 1319051, 1314568, and the "Quantum Computing for Modern Cryptography" project, funded by the UK Quantum Computing and Simulation Hub from the UKRI EPSRC grant EP/T001062/1.

1 INTRODUCTION

End-to-end encrypted asynchronous *secure group messaging (SGM)* is becoming one of the most widely used cryptographic applications with billions of daily users. To help solidify the foundations of this trend, the IETF is in the final stages of standardizing the *Messaging Layer Security* (MLS) SGM protocol [8]. The effort is being lead by industry (e.g. by Cloudflare, Cisco, Facebook, Google, Mozilla, Twitter, Wickr and Wire) with a lot of help from academia. Thus, MLS holds the potential to be widely deployed with in the coming years.

Due to the expected feature set and the multi-faceted security goals, MLS is quite complex. Furthermore, the description of MLS is

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

quite monolithic. In an effort to come to terms with the complexity (both of the construction and desired security notions) all work formally analyzing MLS has restricted itself in scope. Besides simplifying the PKI [1, 3] the trend has been to focus only on a subset of the complete protocol. In particular, the majority of work has considered so-called *continuous group key agreement (CGKA)* [1, 3, 6]. First introduced by Cohn-Gordon et al. [13] (under the name "Group Ratcheting"), a CGKA protocol is often motivated by the folklore belief that CGKA encapsulates the essence of building SGM (not unlike how KEMs capture the essence of PKE). Indeed, this folklore has motivated a host of new CGKA protocols designed to improve on the CGKA protocol implicit to MLS. Some constructions aim for stronger security properties [3, 5], others for a more flexible communication model [10, 23], while the construction in [1] is optimized for certain common SGM settings.

CGKA abstracts the basic task of asynchronously maintaining secrets in a group with an evolving set of members. To this end, each time a member joins, leaves or *updates* their cryptographic state (in an effort to heal from past compromises and defend against future ones) a new epoch in the ongoing session is initiated. Each epoch is equipped with a symmetric group key that can be derived by all parties that are in the group during that epoch. The CGKA most often identified as underpinning MLS is a protocol called TreeKEM [7, 26]. Although, it has undergone at least 2 major iterations before reaching its current version [8] it remains, at most, a passively secure protocol, i.e., with no authenticity mechanisms included. Moreover, it provides weaker guarantees (for the group key) than is expected of the (encryption keys) in the full MLS protocol. Thus, there are clearly further significant security mechanisms at work in MLS beyond the underlying passively secure CGKA, i.e., the protocol requires additional components.

Indeed, [6] extends the abstraction boundary around TreeKEM including more of the MLS protocol to obtain an *actively* secure CGKA which they dub ITK (for "Insider Secure TreeKEM"). Yet even ITK is only a CGKA and not an SGM protocol, and so here too there clearly remains a gulf between the full MLS protocol and what has actually been formally analyzed in [6].

Two exceptions to the focus on CGKA are: [15], which focuses exclusively on how to optimize the bandwidth required to updated cryptographic state concurrent sessions with overlapping membership sets, and [11], which considers the key derivation paths within MLS. While the paths extend beyond the CGKA (even ITK) into parts of MLS specific to messaging (e.g., key/nonce derivation for message encryption) the work leaves much open when it comes to a holistic understanding of MLS's security. For example, the work makes no statements at all about authenticity. So for example, it leaves out the relatively involved and interdependent mechanisms used by MLS to provide this property. Moreover, the work forgoes any formal modeling of network communication, leaving open the questions around how a network adversary can generally impact the security of a session by, say, arbitrarily manipulating packets in transit.

Our contributions

From a practice-oriented view point, the primary contribution in this work is a rigorous security proof showing that the basic instantiation of MLS is an SGM protocol. More generally, we validate the folklore claim that CGKA embodies the essence of SGM. In doing so, we make progress on several fronts in coming to terms with the complexity involved in describing MLS (and other SGM protocols) as well as defining (and proving) security for such protocols.

Black-Box Construction. In more detail, we give a black-box construction of an SGM protocol from a passively secure CGKA and several other primitives. Instantiated appropriately, the construction recovers the basic SGM protocol in MLS.¹ Besides passive CGKA, the construction also uses a new primitive called a Forward Secure Group AEAD (FS-GAEAD). Intuitively, it combines a forward secure key schedule with an authenticated encryption with associated data (AEAD) scheme, to allow a group of sender's sharing a single secret to send any number of AEAD encrypted messages with no further synchronization or communication such that any receiver can immediately decrypt any ciphertexts they receive regardless of the order in which they are delivered. We further abstract the key schedule of an FS-GAEAD with a new primitive called a Forward Secure KDF (FS-KDF) which captures a generic forward-secure symmetric key schedule supporting the derivation of keys in any order. Due to its improved efficiency profile, our construction of an FS-KDF has since replaced the original FS-KDF in the MLS standard.

History graphs. We introduce the first formal security notion for SGM.² Intuitively, it ensures correctness as well as privacy and authenticity in the form of *post-compromise forward secrecy* (PCFS) [3]. That is, messages in a given epoch are both private and authenticated despite arbitrary state leakage of participants sufficiently in the past and future.

In effort to manage the complexity inherent in such a definition we developed the *history graph* (HG) paradigm for defining security of asynchronous group protocols like MLS and CGKA protocols. The paradigm allows for an intuitive (even visual) understanding of an otherwise complicated notion by representing the (securityrelevant) semantics of an execution as an annotated graph. This allows cleanly separating functionality and communication model from the security details being captured. Those details are formalized as a predicate defined over HG and a particular challenge message (or group key for CGKA). The output of the predicate indicates if we can expect the challenge to be secure given the execution represented by the HG. This modular approach makes our SGM and CGKA definitions easy to adapt to, say, future versions of MLS. It also allows for more immediate comparison of the security enjoyed by different constructions of the same type.

¹That is, the part of MLS that allows parties to manage group membership, update their cryptographic states at will and send/recieve encrypted and authenticated messages. We do not analyze more advanced features such as importing/exporting secrets, meta-data hiding and so called "external commits" where an external party makes changes to the group state.

²We note that, concurrently to our work, [9] introduced a security notion for SGM defined using formal verification tools which we discuss in the Related Work section below.

PKI. Another contribution to the analysis of MLS is a more accurate modeling of the PKI used by MLS. With the notable exceptions of [6, 15]), the PKI in the works of [1, 3] the analysis of (earlier versions of) TreeKEM is based on simplified PKI models that encode strong (but implicit) assumptions. For example, their PKI precludes the adversary from registering ephemeral keys even for recently corrupted parties. Instantiating this faithfully seems to require strong authenticity checks by the PKI that may not rely on any state that can be leaked during a compromise. In contrast, the PKI in our work can be realized under assumptions more compatible with E2E security. Concretely, our PKI can be realized if parties use an out-of-band mechanism to authenticate each other's long-term keys while distributing ephemeral public keys via an untrusted key server. Reflecting the MLS standard, the out-of-band authentication mechanism is not made concrete but common instantiations include peer-to-peer key verification via 2D bracodes or (in the enterprise setting) certification authorities binding long term keys to a users identity.

Security Proofs. Armed with our HG based security notions for SGM and passively secure CGKA we prove security for the black-box SGM protocol there by making progress not just in validating the design of MLS but also the more general folklore motivating the study of CGKA protocols. At a high level the proof follows the outline of proof for the double-ratchet in [2]. However the similarities end there as we require new arguments and hybrids to instantiate this outline in light of the substantial differences between the two protocols and 1-on-1 vs. group settings they operate in.

To recover MLS security we also show that TreeKEM (as implicit to the current version in Draft 11 of MLS) [8] is indeed a passively secure CGKA according to our definition. In particular, in comparison to the TreeKEM analysis in [1], we more accurately model the PKI, capture transcript consistency and secure group management. Meanwhile, in comparison to the analysis in [3] we improve the PKI model and allow for a much more capable network adversary that can, e.g. deliver packets in different orders to different participants; a potentially unrealistic restriction for real world adversaries that controls, say, the network connection of a user (let alone the delivery server used by the session).

Beyond MLS and TreeKEM we also prove that RTreeKEM [3] is a passively secure CGKA. For this we extend the *generalized selective decryption (GSD)* technique to account for encrypting keys using the so-called *updatable public-key encryption (UPKE)* [3], used by RTreeKEM. GSD lies at the heart of almost all adaptive security proofs for CGKA protocols. To the best of our knowledge, it remains the only technique for proving adaptive security with meaningful security loss (e.g. quadratic in the group size). In contrast, the proof of the ART CGKA protocol in [13] suffers from an exponential loss in group size. Extending GSD to allow for protocols using UPKE is likely to have applications beyond analyzing RTreeKEM as UPKE is a very natural drop in replacement for standard PKE but with improved forward security [20]. Indeed, recent results in [17] construct post-quantum secure UPKE with an aim towards building a PQ variant of MLS.

Related work. The history graph paradigm was conceived early on in this project. In an effort to keep pace with the rapid development of MLS the paradigm was shared with other researchers before this work was published. Consequently, the paradigm has been used in subsequent work appearing before this one (that cite this one as the source of the paradigm) [5, 6].³ Symbolic representations of an execution, equipped with safety predicates, have been used in [2, 14]. The GSD proof technique, introduced by Panjwani [25] and generalized by Jafargholi et al. [19], was first applied to CGKA protocols in [1].

The most influential precursor to TreeKEM, the asynchronous ratchet tree (ART) protocol, was introduced by Cohn-Gordon et al. [13], focusing on adaptive security (informally sketched) for static groups. It inspired the initial version of TreeKEM [26] which was updated to provide secure group management in [7] and finally to its current form in [8]. Meanwhile, ART uses an older technique called "Treebased DH groups" [27, 28, 30] which is also used by [21] to build key agreement. However, TreeKEM and ART differ significantly from [21], as discussed in [3]. TreeKEM is also related to schemes for (symmetric-key) broadcast encryption [16, 18] and multicast encryption [12, 24, 30]. Cremers et al. [15] note MLS/TreeKEM's disadvantages w.r.t. PCS for multiple groups, and Weider [29] suggests Causal TreeKEM, a variant that requires less ordering of protocol messages. Finally, in concurrent work appearing recently, [9] analyzed an older version of TreeKEM (from MLS Draft 7) using formal verification techniques.

2 PRELIMINARIES

Notation. We use associative arrays which map arbitrary key strings to item strings. For array A we use "." to denote all entries. In particular, we (implicitly) declare variable A to be a (1-dimensional) array and set all of its entries empty string by writing $A[.] \leftarrow \varepsilon$. Similarly, we declare a new 2-dimensional array B with empty entries by $B[.,.] \leftarrow \varepsilon$. We use the shorthand $B[x,.] \leftarrow A$ to denote that $\forall y$ with $A[y] \neq \varepsilon$ we set $B[x, y] \leftarrow A[y]$. For subset $Y' \subseteq$ Y the term $Y' \subseteq A$ returns true iff A contains all elements in Y'; that is $\forall y \in Y' \exists x : A[x] = y$. For vector $\mathbf{x} = (x_1, \dots, x_d)$ with all components $x_i \in X$ we write $A[\mathbf{x}]$ to denote the vector $(A[x_1], \ldots, A[x_d])$. We denote the empty list with [.] and the length of a list L by |L|. The following special keywords are used to simplify the exposition of the security games: req is followed by a condition; if the condition is not satisfied, the oracle/procedure containing the keyword is exited and all actions by it are undone. let is followed by a variable and a condition. After evaluating the expression the variable is assigned with the value that satisfies the condition, if such value exists, and \perp otherwise. **chk** is followed by a condition; if the condition is not satisfied, the oracle/procedure containing the keyword returns false, otherwise the next instruction is executed.

3 FORMAL DEFINITION OF SECURE GROUP MESSAGING

In this section we formally define *secure group messaging* (SGM). Specifically, the formal syntax of SGM schemes is introduced in Section 3.1, and Section 3.2 introduces the security game for SGM schemes.

³While not ideal, we felt this approach was necessary given the strong incentive to make progress on MLS's analysis prior to it being to widely deployed.

3.1 Syntax

This section introduces the formal syntax of an SGM scheme. Parties are identified by unique party IDs ID chosen from an arbitrary fixed set. In the following, *s* and *s'* denote the internal state of the SGM scheme before and after an operation, respectively. An SGM scheme SGM consists of 15 algorithms grouped into the categories above:

- For initialization:
 - $s \leftarrow$ Init(ID): takes as input a party ID ID and generates the initial state.
- For interaction with the PKI:
 - (s', spk) ← Gen-SK(s): generates new signature key pair and outputs the public key;
 - s' ← Rem-SK(s, spk): removes the key pair corresponding to spk from the state;
 - s' ← Get-SK(s, ID', spk'): stores signature public key spk' for party ID';
 - − $(s', kb) \leftarrow$ Gen-KB(s, spk): generates new key bundle (aka initial key material), signed with spk;
 - − (s', ok) ← Get-KB(s, ID', kb'): stores key bundle kb' for party ID' − can reject kb' by outputting ok = false.

Key bundles have the format kb = (wpk, spk, sig), where sig is a signature of wpk under spk, and wpk is so-called *welcome key material*, which can be used to encrypt secret information for joining group members.

- For group creation:
 - s' ← Create(s, spk, wpk): creates group with self, where the party uses key pair corresponding to spk to sign in this group, while control messages for which the recipient is the group creator, will be encrypted under keys in wpk.
- For proposals:
 - − $(s', P) \leftarrow Add(s, ID')$: generates a proposal to add party ID';
 - $(s', P) \leftarrow \text{Remove}(s, \text{ID'})$: generates a proposal to remove party ID';
 - $(s', P) \leftarrow$ Update(s, spk): generates a proposal for self to update the personal key material; the new signing verification key will be spk;
 - $(s', PI) \leftarrow Proc-PM(s, P)$: adds proposal *P* to the state and outputs information PI about *P*.
- For commits:
 - $(s', E, \mathbf{W}, T) \leftarrow \text{Commit}(s, \mathbf{P})$: creates a commit corresponding to a vector \mathbf{P} of proposals and outputs an epoch ID E, an array of welcome messages \mathbf{W} (where $\mathbf{W}[\text{ID}]$ is the welcome message for newly added party ID), and a commit message T (for existing group members);
 - $(s', GI) \leftarrow Proc-CM(s, T)$: used by existing group members to process a commit message T and reach a new epoch; outputs updated group information GI (where GI = \perp if Tis considered invalid);

- $-(s', GI) \leftarrow Proc-WM(s, W)$: used by newly added group members to process welcome message W and join a group; outputs group information GI (where GI = \perp if W is considered invalid).
- For sending and receiving messages:
 - $(s', e) \leftarrow \text{Send}(s, a, m)$: generates a ciphertext *e* encrypting plaintext *m* and authenticating AD *a*;
- $(s', E, S, i, m) \leftarrow \text{Rcv}(s, a, e)$: decrypts ciphertext *e* to plaintext *m* and verifies AD *a*; also outputs a triple (*E*, *S*, *i*) consisting of epoch ID *E*, sender ID S, and message index *i*.

3.2 Security

3.2.1 Bookkeeping. The complexity of the SGM security definition stems from the bookkeeping required to determine which messages are safe to challenge and when the attacker is allowed to inject. The security game keeps track of all the relevant execution data with the help of a so-called *history graph*. The recorded data informs three *safety predicates* (each pertaining to privacy, authenticity, or both) used to determine whether a given execution was legal. These predicates are kept generic and considered parameters of the security definition.

History graphs. A history graph is a directed tree whose nodes correspond to the group state in the various epochs of an execution: The root of the tree is a special node v_{root} that corresponds to the state of not being part of a group. The children of the root node are v_{create}-nodes, which correspond to the creation of groups.⁴ The remaining nodes all correspond to the group state after a particular commit operation. Two nodes v and v' are connected by an edge (v, v') if the commit operation leading to v' was created in v. Concretely, a history graph node consists of the following values: v = (vid, orig, data, pid), where vid is the node's (unique) ID, orig is the party that caused the node's creation, i.e., orig is either the group creator or the committer, data is additional data, and pid is the vector of (IDs of) proposals (see below) that were included in the commit. The history graph is accessed by the security-game oracles via the "HG object" HG, which provides information via several methods (explained later as they are needed). The (ID vid of the) node corresponding to a party ID's current state is stored by the array V-Pt[ID].

Proposals. Similarly to HG and the history graph, the object Props keeps track of proposals. The information recorded about each proposal is a vector p = (pid, vid, op, orig, data), where pid is the proposal's (unique) ID, vid is the (ID of) the HG node corresponding to the epoch in which the proposal was created, $op \in \{add, rem, upd\}$ is the type of proposal, orig is the party that issued the proposal, and data is additional data. Similarly to the HG object, Props will also export several useful methods (also explained later) to the oracles of the security game.

⁴The security game allows multiple groups to be *created* (when multiple parties simultaneously create a group). However, in order to remain in the single-group setting, the attacker is required to pick a single one of these groups as the "canonical group." Parties may only join the canonical group.

PKI bookkeeping. Just as epochs and proposals are kept track of symbolically, so are welcome keys and signature keys. Specifically, each welcome key has a (unique) welcome-key ID wkid, and each signature key has a (unique) signature-key ID skid. The arrays WK-ID[·] and SK-ID[·] map wkid and skid values to the ID of the party that owns the corresponding secret keys; the arrays WK-PK[·] and SK-PK[·] map wkid and skid values to the corresponding public keys. Moreover, the array WK-SK[·] remembers the binding of welcome keys to signature keys; these bindings are created by key bundles: when wpk with wkid is signed under spk with skid in a key bundle, the game sets WK-SK[wkid] \leftarrow skid.

The array CL-KB[\cdot , \cdot] keeps track of the *contact list* for each party. Specifically, CL-KB[ID, ID'] is a *queue* of pairs (wkid, skid) of (IDs of) initial key material that ID would use if it were to add ID' to the group.

Stored and leaked values. In order to stay on top of the information attacker \mathcal{A} accumulates via state compromise, the security game maintains the following arrays/sets. Stored values: V-St[ID] contains pairs (vid, flag) of (i) (IDs of) the epochs for which ID currently stores information and (ii) a flag recording whether said information can be used to infer information about subsequent states. This array is organized as a queue with maximum capacity r, where r is a parameter of the definition and of SGM protocols: it stands for the maximum number of "open" epochs a party keeps at any point in time.⁵ P-St[ID] contains the (IDs of) the proposals currently stored by ID. WK-St $[\cdot]$ and SK-St $[\cdot]$ contain the (IDs of) the welcome and signing keys, respectively, currently stored by ID. Leaked values: V-Lk is a set that contains triples (vid, ID). Each such tuple means that the attacker learned the state of ID in epoch vid, that for that state flag = true, which implies that leaked information can be used to infer information about subsequent epochs. WK-Lk and SK-Lk contain the (IDs of) the welcome and signing keys, respectively, for which the attacker has learned the secret keys. AM-Lk records application messages for which the attacker has learned the key material.

Deletions, and lack thereof. The SGM definition also requires that parties who fail to delete old values can not use them to their advantage after they are removed from the group. Correspondingly, the array Del keeps track of which parties *are* deleting values as they are supposed to (those with Del[ID] = true) and which ones are not. Formally, a party with Del[ID] = false will simply move such values to a special "trash tape," instead of deleting them; the contents of the trash tape will be revealed to \mathcal{A} upon state leakage. Consequently, there are several "trash arrays" that keep track of what is stored on the trash tape of each party: V-Tr (for undeleted information about epochs), AM-Tr (for undeleted information about application messages), WK-Tr (for undeleted information about welcome keys), and SK-Tr (for undeleted information about signature keys).

Challenges. The array Chall[vid] stores, for each vid, pairs (S, i), indicating that the *i*th message from S in epoch vid was a challenge message.

Epoch IDs. The receiving algorithm Rcv must correctly output a "sequence number" for each message received. A natural way to identify application messages is by the triple of epoch ID, sender, and index (as above). However, the SGM scheme cannot be expected to output epoch identifiers vid used by the security game. Instead, the scheme gets to label the epochs itself whenever a commit is created by outputting an *epoch ID E*. Algorithm Rcv will use the same *E* to refer to messages sent in the corresponding epoch. The security game stores the *E* used by the SGM scheme in array Ep-ID[vid].

Bad randomness. The game keeps track of which update proposals and commits were created with randomness known to the attacker with the help of the Boolean array $BR[\cdot]$. This information must be recorded because such proposals/commits will not contribute to PCS.

In the oracles below that correspond to randomized SGM algorithms, the attacker gets to possibly supply the random coins r used by the affected party. Whenever \mathcal{A} does not wish to specify said coins, it sets $r = \bot$, in which case uniformly random coins (unknown to \mathcal{A}) are used.

// General	// App. Messages
$b \leftarrow \{0, 1\}$	$AM-Tr[\cdot] \leftarrow \emptyset$
idCtr++	$AM\text{-}Lk \gets \emptyset$
$\forall ID : s[ID] \leftarrow Init(ID)$	<pre>// Miscellaneous</pre>
<pre>// Communication</pre>	$Chall[\cdot] \leftarrow \emptyset$
$CM[\cdot, \cdot] \leftarrow \epsilon$	$Del[\cdot] \leftarrow true$
$WM[\cdot, \cdot] \leftarrow \epsilon$	$BR[\cdot] \leftarrow false$
$PM[\cdot, \cdot] \leftarrow \epsilon$	// PKI
$AM[\cdot, \cdot, \cdot, \cdot] \leftarrow \epsilon$	$CL-KB[\cdot, \cdot] \leftarrow \epsilon$
// History Graph	$WK-SK[\cdot] \leftarrow \epsilon$
vid _{root} ← HG.init	WK-ID[\cdot] $\leftarrow \epsilon$
$V-Pt[\cdot] \leftarrow vid_{root}$	$WK-PK[\cdot] \leftarrow \epsilon$
$V\text{-}St[\cdot] \leftarrow [(vid_{root}, false)]$	$WK\text{-}St[\cdot] \leftarrow \emptyset$
$V-Tr[\cdot] \leftarrow \emptyset$	WK - $Tr \leftarrow \emptyset$
$V-Lk \leftarrow \emptyset$	$WK-Lk \leftarrow \emptyset$
$P-St[\cdot] \leftarrow \emptyset$	$SK-ID[\cdot] \leftarrow \epsilon$
$Ep-ID[\cdot] \leftarrow \epsilon$	$SK-PK[\cdot] \leftarrow \epsilon$
	$SK-St[\cdot] \leftarrow \emptyset$
	$SK-Tr \leftarrow \emptyset$
	$SK-Lk \leftarrow \emptyset$

Figure 1: Initialization of the security game for secure groupmessaging schemes.

3.2.2 Initialization. At the onset of the execution, the SGM security game initializes (cf. Figure 1) all of the bookkeeping variables listed above. Additionally, it randomly chooses the challenge bit *b*, initializes a counter idCtr that serves to provide IDs for epochs (vid), for proposals (pid), for welcome keys (wkid), as well as for signature keys (skid), and sets up communication arrays (explained where they are used) dedicated to control and application messages.

The initialization also initializes the state of all possible parties by running the Init algorithm and storing the result in the state array $s[\cdot]$.⁶

Finally, the first node of the history graph is created via a call to the HG.init method. This causes HG to create the root node $v_{\text{root}} = (.\text{vid} \leftarrow \text{idCtr}++, .\text{orig} \leftarrow \bot, .\text{data} \leftarrow \bot, .\text{pid} \leftarrow \bot)$ and return vid_{root} = $v_{\text{root}}.$ vid.

⁵The intuition behind the flag is that only the information stored about *newest* epoch should allow the attacker (upon state leakage) to compute key material for subsequent epochs (which is unavoidable). Information corresponding to older epochs, which are only kept open to receive delayed application messages, should not lend itself to compromise the security of subsequent epochs.

⁶Of course, this is really done on an on-demand basis.

3.2.3 Oracles. All adversary oracles described below proceed according to the same pattern: (a) verifying the validity of the oracle call, (b) retrieving values needed for (c), (c) running the corresponding SGM algorithm, and (d) updating the bookkeeping. Validity checks (a) are described informally in the text below; a formal description is provided in Section D of the Appendix. Note that, most of the time, (b) and (c) are straight-forward and are not mentioned in the descriptions. To improve readability, lines (c) are highlighted.

<pre>// ID generates new sig. key</pre>	<pre>// ID stores sig. key of ID'</pre>
gen-new-SigK(ID)	get-SK(ID, ID', skid)
$(s[ID], spk) \leftarrow Gen-SK(s[ID])$	req SK-ID[skid] = ID'
skid \leftarrow idCtr++	$spk \leftarrow SK-PK[skid]$
$SK-ID[skid] \leftarrow ID$	$s[ID] \leftarrow Get-SK(s[ID], ID', spk)$
$SK-PK[skid] \leftarrow spk$	
SK-St[ID] +← skid	SK-St[ID] +← skid
return (skid, spk)	<pre>// ID generates new key bundle</pre>
// ID removes sig. key	gen-new-KB(ID, skid)
rem-SigK(ID, skid)	req SK-ID[skid] = ID
req SK-ID[skid] = ID	req skid ∈ SK-St[ID]
req skid ∈ SK-St[ID]	$spk \leftarrow SK-PK[skid]$
req skid ∉ HG.SKsUsed(ID)	$(s[ID], kb) \leftarrow Gen-KB(s[ID], spk)$
$spk \leftarrow SK-PK[skid]$	
$s[ID] \leftarrow \text{Rem-SK}(s[ID], \text{spk})$	(wpk, ., .) ← kb
if ¬Del[ID]	wkid ← idCtr++
SK-Tr[ID] +← skid	WK-ID[wkid] \leftarrow ID
$SK-St[ID] \rightarrow skid$	WK-PK[wkid] \leftarrow wpk
•	WK-St[ID] +← wkid
	WK-SK[wkid] +← skid

Figure 2: PKI-related oracles of the security game for secure group-messaging schemes. The compatibility functions are described in the accompanying text; a formal description is provided in Section D.

return (wkid, kb)

3.2.4 *PKI Oracles.* The PKI oracles offer the following functionality to attacker \mathcal{A} :

- *New signature keys:* Have a party ID create a new signature key pair. This essentially boils down to ID running algorithm Gen-SK, which outputs a signature public key spk. The SGM game then generates an skid for spk and updates arrays SK-ID, SK-PK, and SK-St correspondingly.
- *Remove signature keys:* Have a party ID delete a signature public key spk, identified by the corresponding skid. In order for a call to this oracle to be legal, skid must correspond to an spk (a) currently stored by ID and (b) not used by ID in either the current epoch or any pending proposals or epochs.⁷ The oracle simply runs Rem-SK on spk and subsequently updates arrays SK-ID, SK-PK, and SK-St to reflect the removal of spk. Additionally, if ID does not delete old values, i.e., if Del[ID] = false, skid is added to SK-Tr[ID].
- Store new signature key: This oracle lets the attacker instruct a party ID to store signature public key, identified by its skid, of a party ID'. The oracle ensures that SK-ID[skid] = ID',

i.e., that skid was really created by ID'. This models the fact that the bindings between signature keys and identities are incorruptible.

• *Generate new key bundle:* The attacker can make a party ID create a new key bundle and store it on the PKI server. To that end, \mathcal{A} specifies the skid of the signature public key spk that is supposed to be used inside the key bundle. The call is only valid if the key pair corresponding to skid is currently stored by ID.

The oracle runs Gen-KB on spk. The first component of the newly generated key bundle is a new welcome key wpk. Thus, the security game generates a new wkid associated with wpk and updates arrays WK-ID, WK-PK, and WK-St accordingly. The oracle also binds wkid to skid.

- *Store new key bundle:* This oracle lets the attacker instruct a party ID to store a key bundle kb = (wpk, spk, sig) belonging to another party ID'. Two conditions must be satisfied for the oracle call to be valid:
- The signature public key spk must belong to ID', i.e., spk = SK-PK[skid] for some skid with SK-ID[skid] = ID'. This models the fact that the long-term secrets binding signature public keys to identities are incorruptible.
- (2) The public key spk must be stored by ID.

If the oracle call is valid, algorithm Get-KB is run on ID' and kb. If the secret key corresponding to spk has not been leaked, i.e., skid \notin SK-Lk, then Get-KB must only accept the key bundle if wpk exists and is bound to spk, i.e., WK-SK[wkid] = skid. If Get-KB outputs ok, skid \notin SK-Lk, but WK-SK[wkid] \neq skid, then the adversary has attempted to create a forgery against the signing key with ID skid, and if it is successful, it wins the game. If wpk is adversarially generated (i.e., wkid = \perp), the game generates a new wkid for wpk, updates WK-PK accordingly, and immediately marks wkid as leaked. Finally, ID's (symbolic) contact list is updated, by adding the pair (wkid, skid) to the end of queue CL-KB[ID, ID'].

3.2.5 Main oracles. The main oracles of the SGM game are split into four figures: oracles related to (i) group creation and proposals (Figure 3), (ii) commits (Figure 4), (iii) sending, receiving and corruptions (Figure 5). The validity of all oracle calls is checked by the corresponding compatibility functions (cf. Section D in the Appendix.).

Group creation. The attacker can instruct a party ID to create a new group by calling the group-creation oracle (Figure 3); such a call must also specify the skid of the signature key under which ID initially signs their messages, as well as the wkid of the welcome key material. A call to the group creation oracle is valid if (i) ID is not in a group yet, (ii) skid, wkid, are currently stored by ID, and (iii) wkid is signed under skid.

The bookkeeping is updated as follows: A call is made to the HG.create(ID, skid) method. This causes HG to create a node v = (.vid \leftarrow idCtr++, .orig \leftarrow ID, .data \leftarrow skid, .**pid** $\leftarrow \perp$) as a child of v_{root} and return vid = v.vid. Then, the bad-randomness array is filled (depending on whether \mathcal{A} specified coins r or not), ID's pointer is set to vid, and V-St[ID] is set to a queue containing

 $^{^7 \}rm Pending$ epochs are child epochs of V-Pt[ID]; they are explained in more detail later on.

only (vid, true), where the second component records that from the information ID currently stores about vid, one can compute information about (future) child states of vid.

// ID creates group; signs with skid

```
create-group(ID, skid, wkid, r)
      req *compat-create(ID, skid, wkid)
      s[ID] \leftarrow Create(s[ID], SK-PK[skid], WK-PK[wkid]; r)
      vid \leftarrow HG.create(ID, skid)
      BR[vid] \leftarrow r \neq \bot
      V-Pt[ID] \leftarrow vid
      V-St[ID] \leftarrow [(vid, true)]
     return vid
// ID proposes to add ID'
prop-add-user(ID, ID')
      req *compat-prop(add, ID, ID', \perp)
      (s[ID], P) \leftarrow Add(s[ID], ID')
      (wkid', skid') \leftarrow CL-KB[ID, ID'].deq
      pid ← Props.new(add, ID, (ID', wkid', skid'))
      \mathsf{PM}[\mathsf{pid}] \leftarrow P
     return (pid, P)
// ID proposes to remove ID'
prop\text{-}rem\text{-}user(\mathsf{ID},\mathsf{ID'})
      req *compat-prop(rem, ID, ID', ⊥)
      (s[ID], P) \leftarrow \text{Remove}(s[ID], ID')
      pid \leftarrow Props.new(rem, ID, ID')
      \mathsf{PM}[\mathsf{pid}] \leftarrow P
     return (pid, P)
// ID proposes update
prop-up-user(ID, skid, r)
      \textbf{req *compat-prop}(\mathsf{upd},\mathsf{ID},\bot,\mathsf{skid})
      (s[ID], P) \leftarrow Update(s[ID], SK-PK[skid]; r)
      pid ← Props.new(upd, ID, skid)
      \mathsf{PM}[\mathsf{pid}] \leftarrow P
     return (pid, P)
// Proposal pid delivered to ID
dlv-PM(ID, pid)
      req *compat-dlv-PM(ID, pid)
      (s[ID], PI) \leftarrow Proc-PM(s[ID], PM[pid])
      if ¬Props.checkPI(pid, PI)
         win
      P-St[ID] + \leftarrow pid
// Proposal P' injected to ID
inj-PM(ID, P')
      req *compat-inj-PM(ID, P')
      (s[ID], PI) \leftarrow Proc-PM(s[ID], P')
      \mathsf{vid} \leftarrow \mathsf{V}\text{-}\mathsf{Pt}[\mathsf{ID}]
      ID_0 \leftarrow PI.orig
      \overrightarrow{req} \neg (``auth-compr(vid))
           \wedge *SK-compr(vid, ID<sub>0</sub>))
```

 $\mathbf{if} \; \mathsf{PI} \neq \bot$ win

Figure 3: Group-creation and proposal-related oracles of the security game for secure group-messaging schemes. The compatibility functions are described in the accompanying text; a formal description is provided in Section D of the Appendix.

Creating proposals. The oracles prop-{add, rem, up}-user (Figure 3) allow the attacker to instruct a party ID to issue add/remove/update proposals. Calls to these proposal oracles are valid if ID is a group member and:

• (add proposals) the target ID' is not a group member already and ID has initial key material for ID' stored;

- (*remove proposals*) the target ID' *is* a group member;
- (update proposals) the skid with which ID is supposed to sign is currently stored by ID.

Note that the last bullet means that if a party wishes to change its active signature key, they must have generated one and registered it with the PKI before issuing the corresponding proposal.

Bookkeeping is updated by calling Props.new(op, ID, data), which records proposal data $p = (.pid \leftarrow idCtr+, .vid \leftarrow V-Pt[ID], .op \leftarrow$ op, .orig \leftarrow ID, .data \leftarrow data), where op \in {add, rem, upd} is the proposal type and where data stores (add proposals) data = (ID', wkid', skid'), where (wkid', skid') is the head of the contact list CL-KB[ID, ID']; (remove proposals) data = ID'; (update proposals) data = skid. Furthermore, the proposal message P output by corresponding proposal algorithm (Add, Remove, or Update) is stored in the communication array PM for proposal messages.

Delivering and injecting proposals. There are two oracles for getting a proposal to a party ID (Figure 3): The first one, dlv-PM, is for honest proposal delivery. The attacker specifies a pid-which must belong to the epoch ID is currently in-and the corresponding proposal is fed to the SGM algorithm Proc-PM. Proc-PM is required to output proposal information PI, which is checked by function Props.checkPI. Function checkPI ensures that PI correctly identifies the *type*, the *originator*, as well as the *data* of the proposal.⁸ The last action performed by the oracle is to update P-St[ID] to include pid, indicating that ID now stores the new proposal.

The second oracle, **inj-PM**, is used by \mathcal{A} to inject proposals to a party ID. More precisely, \mathcal{A} is allowed to submit proposals P' that (a) either belong to an epoch different from vid := V-Pt[ID] or (b) are completely made up (i.e., there exists no pid with PM[pid] = P'). A call to this second oracle is only allowed if the adversary is not currently able to forge messages for epoch vid. This is the case if either the key material used to authenticate in vid is compromised (*auth-compr(vid)) or if the signing key used by the supposed originator ID_{Ω} of the proposal (as determined by the output of Proc-PM) is compromised (*SK-compr(vid, ID_O)). The functions *auth-compr and *SK-compr are two of the safety predicates mentioned above.

Creating commits. The commit-related oracles (Figure 4) deal with commits and delivery of commit and welcome messages.

In order to have a party ID create a commit based on a set of proposals (with IDs) pid, the attacker calls the commit oracle. The specified **pid** must be a subset of all proposals currently stored by ID, and all proposals must belong to ID's current epoch. In addition, a (rather permissive) sensibility check is run on the vector of proposal specified by **pid**: the proposals are processed in the given order (changing the group roster accordingly), and for each proposal it is checked whether the originator would be in the group and, additionally, whether (add proposal) the target would be in the group; (remove proposal) the target would not be in the group.

Subsequently, the SGM algorithm Commit is run on the given set of proposals, which results in an epoch ID E, a vector of welcome messages W, and a commit message T being output. Bookkeeping calls HG.commit(ID, pid), which creates a new HG node

⁸Of course, checkPI checks that PI contains the actual welcome and signature keys (and not wkid and skid).

 $v = (.vid \leftarrow idCtr++, .orig \leftarrow ID, .data \leftarrow \bot, .pid \leftarrow pid)$, as a child of ID's current epoch V-Pt[ID] and returns vid = v.vid. Then, the bad-randomness array is filled (depending on whether \mathcal{A} specified coins *r* or not), and *E* is stored in array Ep-ID. Finally, the commit messages (for current group members) are stored in array CM, and the welcome messages (for new group members) in WM.

// ID commits proposals pid	// Control msg. T' injected to ID
commit(ID, pid, r)	ini-CM(ID, T')
reg *compat-commit(ID, pid)	req *compat-inj-CM(ID, T')
$(s[ID], E, W, T) \leftarrow Commit(s[ID], PM[nid]);$	r) $(s[ID] GI) \leftarrow Proc-CM(s[ID] T')$
	$vid \leftarrow V_Pt[ID]$
vid \leftarrow HG commit(ID nid)	$ID_{0} \leftarrow Glorig$
$BR[vid] \leftarrow r \neq 1$	reg ¬(*auth-compr(vid)
Ep-ID[vid] $\leftarrow E$	\wedge * SK-compr (vid, ID _O))
for $ D' \in HG$ roster(V-Pt[ID])	if Gl ≠ 1
$ $ CM[ID', vid] $\leftarrow T$	win
for $ID' \in Props.addedIDs(nid)$	// Welcome msg. delivered to ID
$ WM[ID', vid] \leftarrow W[ID']$	dlw-WM(ID vid)
return vid	req *compat-dlv-WM(ID, vid)
// Control msg. delivered to ID	$W \leftarrow WM[ID, vid]$
dlv-CM(ID, vid)	$(s[ID], GI) \leftarrow Proc-WM(s[ID], W)$
req *compat-dlv-CM(ID, vid)	if ¬HG,checkGl(vid,Gl)
$T \leftarrow CM[ID, vid]$	win
$(s[ID], GI) \leftarrow Proc-CM(s[ID], T)$	$V-Pt[ID] \leftarrow vid$
if ¬HG.checkGl(vid, GI)	$V-St[ID] \leftarrow [vid]$
win	wkid \leftarrow HG.addedWK(ID, vid)
if HG.isRemoved(ID, vid)	if ¬Del[ID]
$V-Pt[ID] \leftarrow vid_{root}$	WK-Tr[ID] +← wkid
else	WK-St[ID] $\rightarrow \leftarrow$ wkid
$ V-Pt[ID] \leftarrow vid$	// Welcome msg. W^\prime injected to ID
$i[ID] \leftarrow 0$	inj-WM(ID, W')
$P-St[ID] \leftarrow \emptyset$	<pre>req *compat-inj-WM(ID, W')</pre>
if Del[ID]	$(s[ID], GI) \leftarrow Proc-WM(s[ID], W')$
$V-St[ID].last.flag \leftarrow false$	$E \leftarrow Gl.eplD$
else	$req \exists vid : Ep-ID[vid] = E$
V-Ir[ID] + V-St[ID].tirst	$ID_O \leftarrow GI.orig$
V-St[ID].enq((V-Pt[ID], true))	req ¬(*auth-compr(vid)
	$\land *SK-compr(vid, ID_O))$
	if Gl ≠ ⊥
	win

Figure 4: Commit-related oracles of the security game for secure group-messaging schemes. The compatibility functions are described in the accompanying text; a formal description is provided in Section D of the Appendix.

Delivering and injecting commit messages. As with proposals, there are two oracles that allow attacker \mathcal{A} to get a commit message (CM) to a current group member-for honestly generated CMs and for adversarially generated CMs (Figure 4). The former oracle, dlv-**CM**, can be used by \mathcal{A} to deliver to a party ID any CM T that corresponds to a child epoch vid of ID's current epoch V-Pt[ID], provided all proposals that lead to vid have been delivered to ID. Oracle dlv-CM runs SGM algorithm Proc-CM on T, which results in group information GI being output. This information is checked by function HG.checkGI. Function checkGI checks that GI correctly reports the new group roster, the originator, as well as the parties added and removed. Next, if ID has been removed as part of the commit (checked by function HG.isRemoved), its pointer V-Pt[ID] is set to vidroot; otherwise, V-Pt[ID] is set to the new epoch vid. Furthermore, the index counter *i*[ID] is reset to 0, and P-St[ID], the set of proposals stored, is set to the empty set. If ID deletes old values, i.e., if Del[ID] = true, the game changes the flag in

(·, flag) = V-St[ID].last to false since it is no longer the newest state. If ID does not delete old values, this change is not made, and the first element of V-St[ID] is put into V-Tr[ID] (because it is about to be removed from the queue). Finally, the new epoch vid is added to the end of queue V-St[ID]. Recall that this queue has a capacity of *r*, which means that the first element of the queue is removed. This captures that information about the oldest epoch in ID's state is now deleted by ID.

The second oracle, **inj-CM**, is used by \mathcal{A} to inject CMs to a party ID. More precisely, \mathcal{A} is allowed to submit any CM T' that (a) either belongs to an epoch different from any child epoch of vid := V-Pt[ID] or (b) is completely made up. A call to this second oracle is only allowed if the adversary is not currently able to forge messages for epoch vid. Similarly to proposals, whether this is the case is determined via the (generic) safety predicates ***auth-compr** and ***SK-compr**.

Delivering and injecting welcome messages. The oracles **dlv-WM** and **inj-WM** can be used by \mathcal{A} to deliver and inject welcome messages, respectively. The work analogously to oracles **dlv-CM** and **inj-CM** above.

Sending messages and challenges. Oracle **send** allows \mathcal{A} to instruct any current group member S to send a message *m* and associated data (AD) *a*. The oracle runs algorithm Send on *m* and *a*, which creates a ciphertext *e*. The oracle increments S's message counter, and stores the triple (*a*, *m*, *e*) in AM.

The challenge oracle **chall** works quite similarly, except that \mathcal{A} specifies two equal-length messages m_0 and m_1 , and m_b is passed to Send (where *b* is the bit chosen during initialization). Furthermore, the pair (S, *i*[S]) is recorded in array Chall.

Delivering and injecting application messages. The attacker can get application messages to group members in two ways, by using either **dlv-AM** or **inj-AM**—for honestly and adversarially generated AMs, respectively. Oracle **dlv-AM** takes as input a tuple (vid, S, *i*, R) identifying the epoch, the sender, the index, and the recipient of the AM to be delivered. Then, the corresponding AD and ciphertext are input to Rcv, which subsequently outputs (E', S', i', m'). These values are checked, and if R misidentifies any of them, the attacker immediately wins the security game. Subsequently, AM[vid, S, *i*, R] is set to received, indicating that R should no longer keep around any key material that can be used to decipher *e*. If R does not delete old values, then a corresponding entry is place in array AM-Tr[R]. Note that the oracle does not return any values since the attacker knows which values are output by Rcv.

Oracle **inj-AM** allows \mathcal{A} to inject any AD/ciphertext pair (a', e') to a group member R as long as they do not correspond to an honestly generated AM. The game requires that if Rcv accepts e' and outputs (E', S', i', m'), it must be the case that (1) the key material for the i'^{th} message sent by S' in epoch E' has not been compromised, and (2) no message has been output for these identifiers before (i.e., a successful replay attack is considered a break of authenticity). Whether (1) is the case is determined by the (generic) safety predicates ***AM-sec** and ***SK-compr**.

Corruption oracle. The corruption oracle **corr** allows the attacker to leak the state of any party ID. It does bookkeeping for the following:

```
// S sends message m with AD a
                                                                    // Disable deletions for \ensuremath{\mathsf{ID}}
send(S, a, m)
                                                                    no-del(ID)
       req *compat-send(S)
                                                                            disable deletions for ID
       (s[S], e) \leftarrow Send(s[S], a, m)
                                                                           Del[ID] \leftarrow false
                                                                    // Safety for privacy
       i[S]++
       \begin{array}{l} \textbf{for } \mathsf{R} \in \mathsf{HG.roster}(\mathsf{V}\text{-}\mathsf{Pt}[\mathsf{S}]) \setminus \{\mathsf{S}\} \\ \mid & \mathsf{AM}[\mathsf{V}\text{-}\mathsf{Pt}[\mathsf{S}],\mathsf{S},i[\mathsf{S}],\mathsf{R}] \leftarrow \end{array}
                                                                    *priv-safe
                                                                           return \forall vid, S, i : ((S, i) \in
                (a, m, e)
                                                                              Chall[vid] \implies *AM-sec(vid, S, i))
       return e
                                                                    // S sends challenge with AD a
// Deliver i^{\mathrm{th}} msg of S in vid to R
                                                                    chall(S, a, m_0, m_1)
                                                                            req *compat-chall(S)
dlv-AM(vid, S, i, R)
                                                                            (s[S], e) \leftarrow \text{Send}(s[S], a, m_b)
       req *compat-dlv-AM(vid, S, i, R)
                                                                            i[S]++
       (a, m, e) \leftarrow \mathsf{AM}[\mathsf{S}, \mathsf{vid}, i, \mathsf{R}]
                                                                            for R \in HG.roster(V-Pt[S]) \setminus \{S\}
       (s[R], E', S', i', m') \leftarrow \operatorname{Rev}(s[R], a, e)
                                                                                  AM[V-Pt[S], S, i[S], R] ↔
                                                                                    (a, m, e)
       if (E', S', i', m') \neq
                                                                            Chall[V-Pt[S]] + \leftarrow (S, i[S])
         (Ep-ID[vid], S, i, m)
                                                                           return e
              win
                                                                    // a', e' get injected to R
       if ¬Del[R]
                                                                    inj-AM(a', e', R)
             AM-Tr[R] \leftarrow (vid, S, i)
                                                                           req *compat-inj-AM(a', e', R)
       AM[vid, S, i, R] \leftarrow received
                                                                           (s[R], E', S', i', m') \leftarrow \operatorname{Rev}(s[R], a', e')
// Attacker leaks state of ID
corr(ID)
                                                                           if m' \neq \bot
       for (vid, flag) \in V-St[ID] \cup V-Tr[ID]
              if flag
↓ V-Lk +← (vid, ID)
                                                                                  let vid' : Ep-ID[vid'] = E'
                                                                                  if \neg (\neg^* AM \operatorname{-sec}(\operatorname{vid}', S', i'))
              AM-Rcvd ←
                 \{(S, i) | AM[vid, S, i, ID] =
                                                                                    \land *SK-compr(vid', S')
                 received}
                                                                                      \vee AM[vid', S', i', R] =
              AM-Ik+←
                                                                                     received
               (vid, AM-Rcvd, AM-Tr[ID])
                                                                                         win
       W \leftarrow WK-St[ID] \cup WK-Tr[ID]
                                                                                  if ¬Del[vid]
       WK-Lk + \leftarrow \{wkid \in W \mid
                                                                                         AM-Tr[R] \leftarrow (vid', S', i')
         WK-ID[wkid] = ID}
                                                                                  AM[vid', S', i', R] \leftarrow
       S \leftarrow SK-St[ID] \cup SK-Tr[ID]
                                                                                     received
       SK-Lk + \leftarrow \{ skid \in S \mid
                                                                            return (E', S', i', m')
         SK-ID[skid] = ID}
       HG.corrHanging(ID)
       return s[ID]
```

Figure 5: Send/receive, corruption, no-deletion, and privacysafety oracles of the security game for secure group-messaging schemes. The compatibility functions are described in the accompanying text; a formal description is provided in Section D of the Appendix.

- Epochs: All epochs ID currently stores information for are kept track of by V-St[ID]; recall that for (vid, flag) ∈ V-St[ID] the variable flag records whether or not the information stored about vid allows to infer key material of child epochs. The security game copies the pairs in V-St[ID] with flag = true into the set V-Lk. Additionally, for all (vid, ·) ∈ V-St[ID], records the current set of messages already received by ID in epoch vid. This helps keep track of which application messages are affected by this instance of state leakage. Finally, the values in V-Tr[ID] are also recorded.
- *Welcome keys:* By leaking ID's state, *A* learns all welcome secret keys currently stored by ID: either because ID is supposed to be storing them (recorded by WK-St[ID]) or because they are in ID's trash (i.e., in WK-Tr[ID]). The corresponding wkids are added to the set WK-Lk.
- Signature keys: Handled analogously to welcome keys.
- *Application-message trash:* The attacker also learns all values in the array AM-Tr[ID].

Pending proposals and commits. If at the time ID is corrupted, there are outstanding (i.e., uncommitted) update proposals (with ID) pid by ID or hanging (i.e, created but not processed yet) commits (with ID) vid by ID, the attacker learns the corresponding secrets. To that end HG.corrHanging(ID) sets BR[pid] ← true resp. BR[vid] ← true for all of them. This will reflect the fact that these update proposals / commits cannot be used for PCS.

Disabling deletions. When the attacker calls oracle **no-del** for a party ID, that party stops deleting old values and stores them on a special trash tape instead (which is leaked along with the rest of ID's state to the attacker upon state compromise).

3.2.6 *Privacy-related safety.* At the end of the execution of the SGM security game, the procedure ***priv-safe** ensures that the attacker has only challenged messages that are considered secure by the (generic) safety predicate ***AM-sec**. If the condition is not satisfied, the attacker loses the game.

3.2.7 Advantage. Let $\Pi = \{*AM\text{-sec}, *auth\text{-compr}, *SK\text{-compr}\}\)$ be the set of generic safety predicates used in the SGM definition. The attacker \mathcal{A} is parameterized by it's running time, t, and the number of challenge queries, q, and referred to as (t, q)-attacker. The advantage of \mathcal{A} against an SGM scheme Γ w.r.t. to predicates Π is denoted by $\operatorname{Adv}_{SGM\Pi}^{\Gamma}(\mathcal{A})$.

Definition 3.1. An SGM scheme Γ is (t, q, ε) -secure w.r.t. predicates Π , if for all (t, q)-attackers,

 $\operatorname{Adv}_{\operatorname{SGM},\Pi}^{\Gamma}(\mathcal{A}) \leq \varepsilon$.

4 MODULARIZING MLS AND PROVING ITS SECURITY

This section provides a summary of the MLS protocol, our modularization of it, as well as the security proofs for the full protocol based only on the security of the *generic* components—and for the components themselves (based on specific common cryptographic assumptions).

Our modularization splits MLS into the following three component primitives: *continuous group key-agrement (CGKA), forwardsecure group AEAD (FS-GAEAD)*, and *PRF-PRNGs*. FS-GAEAD corresponds to a single epoch of secure group messaging, the PRF-PRNG can be thought of as an entropy pool from which key material for FS-GAEAD is extracted, and the entropy pool itself is continually refreshed by values from CGKA.

4.1 Continuous Group Key Agreement

CGKA schemes have a structure that is very close to that of an SGM protocol: they proceed in epochs using the propose-and-commit paradigm and provide algorithms for group creation/joining, interaction with the PKI, issuing proposals, and creating/processing commits. However, instead of being used to send/receive messages, in each epoch, CGKA scheme outputs a high-entropy *update secret*. The full CGKA syntax can be found in Section A of the Appendix.

The part of MLS responsible for CGKA is called TreeKEM; an improved version called RTreeKEM is also known and considered in this work, despite not currently being part of the MLS standard.

Security definition. CGKA schemes are expected to guarantee the secrecy of the generated group keys. Furthermore, just like SGM schemes, they must provide *post-compromise forward secrecy (PCFS)*. Once more, both update proposals and commits must contribute to post-compromise security (PCS). CGKA schemes must also be able to deal with bad randomness as well as parties who do not delete old values.

An important difference between CGKA and SGM schemes is that the former are designed in a fully authenticated setting. That is, the attacker in the CGKA security game is *not* permitted to *inject* control messages. This turns out to be sufficient as the "authentication layer" can be added by the higher-level protocol using the CGKA scheme; it also greatly simplifies the security analysis.

Similarly to SGM, the CGKA security game allows adversary \mathcal{A} to control the execution of and attack a single group. In particular, \mathcal{A} controls who creates the group, who is added and removed, who updates, who commits, etc. The attacker is also allowed to leak the state of any party (whether currently part of the group or not) at any time. The privacy of keys is captured by considering a challenge that outputs either the actual key output by the CGKA scheme or a truly random one—which one is determined by an internal random bit *b*, which must be guessed by \mathcal{A} at the end.

The CGKA security game follows the history-graph paradigm to keep track of all the relevant execution data. The recorded data informs a *safety predicate* ***CGKA-priv** used at the end of the game to exclude trivial wins by \mathcal{A} . More details can be found in Section A of the Appendix.

Instantiation: (R)TreeKEM. The (R)TreeKEM CGKA protocols are based on so-called (binary) ratchet trees (RTs). In an RT, group members are arranged at the leaves, and all nodes have an associated public-key encryption (PKE) key pair, except for the root. The tree invariant is that each user knows all secret keys on their direct path, i.e., on the path from their leaf node to the root. In order to perform a commit operation and produce a new update secret I, a party first generates fresh key pairs on every node of their direct path. Then, for every node v' on its *co-path*—the sequence of siblings of nodes on the direct path-it encrypts specific information under the public key of v' that allows each party in the subtree of v' to learn all new secret keys from v's parent up to the root. In TreeKEM, standard (CPA-secure) PKE is used. As shown in previous work [3], this leads to subpar PCFS. RTreeKEM improves on this by using so-called updatable PKE (UPKE), in which public and secret keys are "rolled forward" every time a message is encrypted and decrypted, respectively.

Security results. In order to establish *adaptive* security of the RTreeKEM protocol, this work applies the so-called Generalized-Selective-Decryption (GSD) paradigm to UPKE. Specifically, we prove GSD security of IND-CPA-secure UPKE in the random oracle model, by first defining a GSD game that models UPKE based executions, in the presence of bad randomness and group splitting attacks, and then appropriately adapting the framework of [1]. The security of RTreeKEM itself is established by reduction to the GSD security of the underlying UPKE scheme. The safety predicate ***CGKA-priv** considers the commit secret of an epoch vid secure if (informally) the following conditions are satisfied: (1) There is no "corrupted"

ancestor epoch in the history graph, where an epoch can be corrupted if either a party gets added with leaked initial keys or if the state of a party in the epoch is leaked; (2) no information about vid is known to the attacker as the result of *splitting* the group (cf. full version [4]); (3) good randomness was used by the commiter creating vid. More details can be found in the full version [4].

TreeKEM is secure w.r.t. a weaker security predicate than RTreeKEM. Namely, it is required that either there is no post-challenge compromise (the notion of PCS in [3]), or if there is, compromise happens after the party has already updated it's state (the FSU notion of [3]). Then, security proof for TreeKEM is similar to that of RTreeKEM and proceeds in two steps: first, we consider a reduction from GSD for standard public-key encryption to the IND-CPA security of the underlying scheme, and then we reduce the security of TreeKEM to that of GSD.

4.2 Forward-Secure Group AEAD

FS-GAEAD schemes provide the convenient abstraction of an "epoch of group messaging." That is, they capture the sending and receiving of application messages within a single epoch of a full SGM scheme. In an execution of FS-GAEAD, all participating group members are initialized with the same random group key (i.e., that key is assumed to be generated and distributed among the group members by the higher-level protocol).

An FS-GAEAD scheme protects the authenticity and privacy of messages sent. Furthermore, it provides forward secrecy, i.e., the security of messages received will not be affected by state compromise. Note that FS-GAEAD is not required to provide any form of post-compromise security, which allows to design completely deterministic schemes (apart from the initial key).

Security definition. The security of FS-GAEAD is captured via a corresponding game, in which *n* parties share the same initial key. The attacker can have parties send and receive messages arbitrarily, ask for challenges, and leak the state of any party at any time. The game keeps track of the entire execution, but crucially of which messages have been received by which parties. When the state of some party ID is leaked, the set of messages received by ID is stored: these are the messages that must remain secure even given ID's leaked state. This information is used by a safety predicate ***FS-sec** to avoid trivial wins by \mathcal{A} , be it w.r.t. privacy or authenticity. More details can be found in Section B of the Appendix and the full version [4].

Instantiation and security results. We show how to build FS-GAEAD from a *forward-secure key-derivation function (FS-KDF)*. An FS-KDF keeps state *s*—initially set to a uniformly random string—and, upon request, derives keys corresponding to labels lab. After each such request, the FS-KDF also updates its own state. For each initial state, there a unique key corresponding to each label, irrespective of the order in which the labels were queried. The FS-KDF is forward-secret because even if its state is leaked, all keys output up to that point remain secure. An FS-KDF itself can be obtained via a tree construction based on a normal PRG.

Given an FS-KDF, the construction of FS-GAEAD is as follows: messages and AD are encrypted/authenticated with a normal AEAD

scheme, where the key for the i^{th} message by party ID is derived using the label (ID, i).

The safety predicate ***FS-sec** achieved by our construction considers a message secure if no party's state is leaked before it receives the message. For more information see Section B of the Appendix and the full version [4].

4.3 PRF-PRNGs

A *PRF-PRNG* resembles both a pseudo-random function (PRF) and a pseudorandom number generator with input (PRNG)—hence the name. On the one hand, as a PRNG would, a PRF-PRNG (1) repeatedly accepts inputs *I* and context information *C* and uses them to refresh its state σ and (2) occasionally uses the state, provided it has sufficient entropy, to derive a pseudo-random pair of output *R* and new state; for the purposes of secure messaging, it suffices to combine properties (1) and (2) into a single procedure. On the other hand, a PRF-PRNG can be used as a PRF in the sense that if the state has high entropy, the answers to various pairs (*I*, *C*) on *the same state* are indistinguishable from random and independent values.

Security definition. The intuitive security properties for PRF-PRNGs mentioned above must also hold in the presence of state compromise. In particular, a PRF-PRNG must satisfy PCFS (cf. Section 3.2) and deal with splitting (cf. Section C of the Appendix and the full version [4]). Therefore, the security game for PRF-PRNGs follows the same history-graph approach as the definitions of SGM and CKGA. However, since the game only consists of the state of the PRF-PRNG and there are no parties, it suffices to keep track of a much smaller amount information. Formal definitions for PRF-PRNGs are provided in Section C of the Appendix.

Instantiation and security results. The MLS protocol uses HKDF [22] as PRF-PRNG. This work models HKDF as a random oracle and shows that it achieves achieves security w.r.t. safety predicate ***PP-secure**, which captures that in order for a value *R* to be considered secure in a particular epoch *e*, (1) *e* must have an ancestor *e'* (possibly itself) that was reached via a random input *I* not known to \mathcal{A} , and (2) there must have been no corruptions on the path from *e'* to *e*'s parent. More details can be found in the full version [4].

4.4 Plugging things together

The protocol is based on the following primitives: A CGKA scheme K = (K-Gen-IK, K-Create, K-Add, K-Remove, K-Update, K-Commit, K-Proc-Com, K-Join), an FS-GAEAD scheme <math>F = (F-Init, F-Send, F-Rcv), a CPA-secure public key encryption scheme PKE = (E-KeyGen, E-Enc, E-Dec), an existentially unforgeable signature scheme S = (S-KeyGen, S-Sign, S-Ver), a message authentication scheme M = (M-Tag, M-Ver), a PRF-PRNG PP, a collision resistant hash function H. The initialization and PKI algorithms of our construction are depicted on Figure 6. The helper functions are formally presented in the full version. Below we describe the main SGM algorithms depicted on Figure 7, in which we use different colors to highlight the use of the underlying primitives: CGKA, FS-GAEAD, PKE, Signatures, MAC, PRF-PRNG, Hash.

Group creation. The group creation operation, Create, receives (possibly bad) randomness *r*, a signature verification key, spk, which

is the key that will be used by the group members to verify messages sent by the group creator, as well as the welcome key material wpk. It adds the group creator's id to the roster (s.G \leftarrow [ME]), executes the CGKA group creation operation, absorbs the output *I* into the PRF-PRNG PP. PP outputs the new PRF-RPNG state s. σ , the FS-GAEAD key k_e , a MAC key s. k_m (used to authenticate control messages), and the current epoch id, s.C-epid. Next, the FS-GEAD init operation F-Init, is executed with inputs k_e , the group size, which is 1, and the id ME of the group creator.

Proposals. To add a party ID_a , Add recovers the key bundle kb' for that ID from the contact list, and runs the CGKA add-proposal algorithm with keys from kb', getting a CGKA proposal \bar{P} , which it then uses to construct and authenticate (with MAC and signature) the SGM proposal message P'. Remove and Update are similar (where Update takes the new signing key of the updater as input). When processing any proposal, algorithm Proc-PM simply attempts to authenticate the proposal and stores it locally. Proc-PM returns proposal information: the operation op, the origin of the proposal, orig, and the data data, as computed by *get-propInfo(P') (cf. Section E of the Appendix).

Committing. To commit to a set **P** of proposals, Commit calls the CGKA commit operation to obtain CGKA welcome messages W_{pub} and **WPrv** as well as CGKA control message \overline{T} and update secret *I*. \overline{T} is used to construct and authenticate (with MAC and signature) the SGM control message T' (which includes a hash of the proposals). Then, Commit creates and authenticates (with MAC and signature) the SGM welcome messages for joining parties by adding an encryption of the PRF-PRNG state s. σ to the CGKA welcome messages. The MAC key used for authenticating control and welcome messages is derived by absorbing *I* into the PRF-PRNG (with context information that depends on T'). Note that s. σ is not updated yet, however.

Process commit/welcome messages. Algorithm Proc-CM first verifies the authentication information (MAC and signatures) of a commit message $T' = ("com", epid, ID, h, \overline{T})$. However, in order to obtain the MAC key, (1) the CGKA control message \overline{T} has to be processed by Proc-Com, which (2) recovers the update secret *I*, which in turn (3) must be absorbed into the PRF-PRNG (this time updating its state). The last step also produces the shared key for the new FS-GAEAD session.

Algorithm Proc-WM, used by joining parties to processes welcome messages W' proceeds similarly, except that the state of the PRF-PRNG must first be obtained by decrypting the corresponding ciphertext in W'.

Send message. To send associated data a and plaintext m, Send passes a and m to the current epoch's FS-GAEAD send operation, and, additionally, also signs the resulting ciphertext, a, and the identifier s.C-epid of the current epoch. Receiving works analogously, with the caveat that the appropriate FS-GAEAD session must be used.

4.4.1 *Security. Main idea:* The authenticity property of our construction relies on the EU-CMA security of signatures, the unforgeability of MACs, and the authenticity of the FS-GAEAD scheme. In particular, if the sender's signing key is secure, then an injection w.r.t. that key fails with overwhelming probability. If this is not the case, an honest commit operation provides post-compromise authenticity, by producing a secure CGKA update secret, *I*, which feeds the PRF-PRNG with good randomness, which in turn outputs secure MAC and FS-GAEAD keys, used for the authentication of future control (via MAC security) and application (via FS-GAEAD security) messages (this requires the adversary to remain passive for one epoch). PCFS with respect to privacy is similar and relies on the PCFS security CGKA and the FS of FS-GAEAD.

Simplified properties. Proving SGM security is facilitated by considering three simplified properties, namely correctness, authenticity, and privacy. In the full version of the paper [4] we prove that these properties together imply full SGM security. Besides modularity, simplified properties facilitate the transition from *selective* to *fully* adaptive security, as the individual "simplified" games, defined in the full version [4], consider selective adversaries that commit to the challenge (e.g., the challenge or the last healing, epoch, the message sender and index used for the challenge) at the beginning of the game. In the reduction from the simplified properties to full SGM security, the adversarial strategy is being guessed and the success probability is bounded by values that relate to the running time of the adversary. After proving the simplified properties theorem, one can prove authenticity and privacy, individually, against selective adversaries that commit to their strategy before the security game begins.

Safety predicates. Using safety predicates, we provide a generic theorem (cf. Theorem 4.1), that considers *any* CGKA, FS-GAEAD, and PRF-PRNG scheme. Those schemes come along with their security predicates, Π_{CGKA} , Π_{FS} and Π_{PP} , respectively, and we prove that as long as the attacker's actions are not violating those predicates, then the resulting SGM construction is secure w.r.t. the SGM predicate $\Pi_{SGM} = \Pi_{SGM}(\Pi_{CGKA}, \Pi_{FS}, \Pi_{PP})$. Our SGM safety predicate Π_{SGM} (explained below) is depicted in Figure 8 and operates over history graph information, generated by the SGM security game. Here, we consider $\Pi_{CGKA} = {^*CGKA-priv}, \Pi_{FS} = {^*FS-sec}, \Pi_{PP} = {^*PP-secure}$ (cf. Figure 8).

Proof idea. The adversary breaks authenticity if it manages to make a non-trivial injection, which implies that either of the following holds: (1) injects a (proposal, welcome, commit or key bundle) message in epoch vid that is signed with a non-compromised signing key of the party ID (determined by ***SK-compr**(vid, ID)), (2) injects a (proposal, welcome, commit) message when ***auth-compr**(vid) holds, which implies that *PP-secure(vid, *Proj-PP(SGM-Data)), i.e., PP is secure, and (3) injects an application message when the FS-GAEAD state is not trivially compromised (determined by *FS-sec). Clearly, security against (1) reduces to the EU-CMA security of S. For (2) the output of the PRF-PRNG, PP, in epoch vid is secure, therefore that MAC key (output by PP) is secure and we can rely on the security of PP and the unforgeability of M. This requires the following hybrids: (A) In the first hybrid, if the attacker finds a collision against H, the execution aborts (reduces to the collision resistance property of H). This hybrid is required for the protection of the PP state in the presence of group splitting attacks in which the adversary can split the group, corrupt in one branch to recover the PP state, and challenge on another branch in which the PP state

is related to the corrupted one. Here, collision resistance ensures that different control messages have different hash values, therefore lead to independent PP states. (B) In the next hybrid, in the last healing epoch before the challenge, substitute the CGKA update secret (which feeds PP) with a uniformly random value (required for the reduction to the PRF-PRNG security). Note that, by the definitions of *PP-secure, *Proj-PP, a good healing epoch before the challenge epoch, exists, and this epoch satisfies *CGKA-priv. (C) In the next hybrid we use the CPA security of PKE. In particular, in the commit operation that creates the target epoch, encrypt the zero message instead of encrypting the PP state, as part of the welcome message (requires a reduction to the CPA security of PKE). (D) Next, substitute the output of the PRF-PRNG PP in a commit message for the target epoch, with a uniformly random value (reduces to PRF-PRNG security). Finally, since the output of PP is substituted by a uniformly random value, so does that MAC key, thus we have a reduction to the unfogeability of the MAC scheme M for case (2). Case (3) is similar, however in the last step we have a reduction to the FS-GAEAD authenticity property of F. For privacy we consider the same sequence of hybrids, however the final reduction is against the FS-GAEAD privacy property of F. For that reduction we use the fact that *FS-sec is satisfied.

We ultimately prove the following theorem:

THEOREM 4.1. Let SGM = SGM(K, F, PKE, S, M, PP, H) be the SGM scheme presented above and let $\Pi_{SGM} = \Pi_{SGM}(\Pi_{CGKA}, \Pi_{FS}, \Pi_{PP})$ be predicates such that: (1) K is a CGKA scheme with respect to predicate Π_{CGKA} , (2) F is an FS-GAEAD scheme with respect to predicate Π_{FS} , (3) PKE is a CPA secure public-key encryption scheme, (4) S is an existentially unforgeable signature scheme, (5) M is a message authentication code, (6) PP is a PRF-PRNG with respect to predicate Π_{PP} , (7) H is a collision resistant hash. Then, SGM is an SGM scheme with respect to predicate Π_{SGM} .



Figure 6: The SGM Construction : Initialization and PKI Algorithms.

// Create a group Create(spk, wpk; r) | s.G \leftarrow [ME] wsk \leftarrow WK-wk[wpk] $(\cdot, isk) \leftarrow wsk$ $(s.\gamma, I) \leftarrow K-Create(ME, isk; r)$ $(s.σ, k_e, s.k_m, s.C-epid) \leftarrow PP(0, I, 0)$ s.Ep-SPK[s.C-epid, ME] ← spk $s.v[s.C-epid] \leftarrow F-Init(k_e, 1, ME)$ // Add proposal $\begin{array}{c} \mathsf{Add}(\mathsf{ID}_a) \\ | & \mathsf{kb}' \leftarrow \mathsf{CL}\text{-}\mathsf{KB}[\mathsf{ID}_a] \end{array}$ $\begin{array}{l} ((epk, ipk), spk) \leftarrow kb' \\ (s.\gamma, \bar{P}) \leftarrow K-Add(s.\gamma, ID_a, ipk) \\ P' \leftarrow ("add", s.C-epid, ME, (ID_a, kb'), \bar{P}) \end{array}$ $t \leftarrow M\text{-Tag}(s.k_m, P')$ sig \leftarrow S-Sign(s.C-ssk, (P', t)) return (P', t, sig) // Remove proposal $\begin{array}{c} \text{Remove}(\text{ID}_{r}) \\ | (s.\gamma, \overline{P}) \leftarrow \text{K-Remove}(s.\gamma, \text{ID}_{r}) \end{array}$ $P' \leftarrow ("rem", s.C-epid, ME, ID_r, \bar{P})$ $t \leftarrow M\text{-Tag}(s.k_m, P')$ sig \leftarrow S-Sign(s.C-ssk, (P', t)) return (P', t, sig)// Update proposal Update(spk; r) | req SK-sk[spk] $\neq \varepsilon$ $(s.\gamma, \bar{P}) \leftarrow K-Update(s.\gamma; r)$ $P' \leftarrow ("upd", s.C-epid, ME, spk, \bar{P})$ $t \leftarrow M\text{-Tag}(s.k_m, P')$ $sig \leftarrow S-Sign(s.C-ssk, (P', t))$ return (P', t, sig) ceive A Messag $\begin{array}{c} \mathsf{Rcv}(a,e) \\ \mid & (e', \operatorname{sig}) \leftarrow e \end{array}$ $(epid, a, \bar{e}) \leftarrow e'$ $(s.v[epid], ID_S, i, m) \leftarrow F-Rcv(s.v[epid], (epid, a), \bar{e})$ $\operatorname{req} \mathsf{ID}_{\mathcal{S}} \neq \bot$ $spk \leftarrow s.Ep-SPK[epid, ID_S]$ req S-Ver(spk, e', sig) return (epid, ID_S, *i*, *m*) // Commit Commit(P; r) $req \ P \subseteq s. \mathsf{Props}$ // Get CGKA Proposals $\bar{\mathbf{Q}}.\mathbf{q} \to \bar{\mathbf{Q}}$ // - Prepare Commit Message -(s. γ , W_{pub} , WPrv, \bar{T} , I) \leftarrow K-Commit(s. γ , \bar{P} ; r) $\textbf{T'} \leftarrow (\texttt{"com"}, \textbf{s.C-epid}, \textsf{ME}, \textbf{H(P)}, \bar{\textbf{T}})$ v ← H(T') // New MAC Key & Epoch ID $\frac{(.,.,k_m,\text{epid}) \leftarrow \text{PP(s.}\sigma,I,\text{v})}{t \leftarrow \text{M-Tag}(k_m,\text{T}')}$ $\mathsf{sig} \gets \mathsf{S}\text{-}\mathsf{Sign}(\mathsf{s}.\mathsf{C}\text{-}\mathsf{ssk},(\mathsf{T}',t))$ // Commit Message $\mathtt{T} \leftarrow (\mathtt{T}', \mathit{t}, \mathrm{sig})$ // - Prepare Welcome Messages (ID, wpk) ← *added(P) *new-spks(s.Ep-SPK[s.C-epid, ·], P) for $i \in ID$ $| e \leftarrow E-Enc(wpk[i].epk, s.\sigma)$ w′ ← w ← ("wel", ME, ID[*i*], v, W_{pub},... ..., WPrv[i], e, wpk[i], spk) $t \leftarrow M\text{-Tag}(k_m, (W', epid))$ sig \leftarrow S-Sign(spk[ME], (W', t)) // Welcome Messag $W[i] \leftarrow (W', t, sig)$ return (epid, W, T)

// Process a proposal Proc-PM(P) $(\mathbf{P}', t, sig) \leftarrow \mathbf{P}$ $(., \mathsf{epid}, \mathsf{ID}, ., .) \leftarrow \mathsf{P'}$ req epid = s.C-epid $spk \leftarrow s.Ep-SPK[epid, ID]$ req M-Ver(s.k_m, P', t) req S-Ver(spk, (P', t), sig) s.Props + \leftarrow P return *get-propInfo(P') // Send A Message Send(a, m)| $E \leftarrow s.C-epid$ $\frac{(\mathbf{s}.\boldsymbol{v}[E],\bar{\boldsymbol{e}}) \leftarrow \mathsf{F}}{\boldsymbol{e}' \leftarrow (E,a,\bar{\boldsymbol{e}})}$ -Send(s.v[E], (E, a), m) $sig \leftarrow S-Sign(s.C-ssk, e')$ return (e', sig) // Process Commit Message Proc-CM(T) $(T', t, sig) \leftarrow T$ $(., \mathsf{epid}, \overset{\sim}{\mathsf{ID}}, \mathbf{h}, \bar{\mathsf{T}}) \gets \mathsf{T'}$ // Matching Epochs? req epid = s.C-epid $spk \leftarrow s.Ep-SPK[epid, ID]$ req S-Ver(spk, (T', t), sig) $P \leftarrow s.Props[h]$ // Call CGKA $(s.\gamma, Gl, I) \leftarrow K\text{-Proc-Com}(s.\gamma, \overline{T})$ $(s.\sigma, k_e, s.k_m, s.C-epid) \leftarrow PP(s.\sigma, I, H(T'))$ req M-Ver(s.k_m, T', t) $s.G \leftarrow GI.G$ pos ← *roster-pos(ME, s.G)
// Start FS-GAEAD $s.v[s.C-epid] \leftarrow F-I$ k_{e} , s.G, pos s.Ep-SPK[s.C-epid, .] \leftarrow *new-spks(s.Ep-SPK[epid, .], P) $\begin{array}{l} \mathsf{spk} \leftarrow \mathsf{s}.\mathsf{Ep}\text{-}\mathsf{SPK}[\mathsf{s}.\mathsf{C}\text{-}\mathsf{epid},\mathsf{ME}] \\ \mathsf{s}.\mathsf{C}\text{-}\mathsf{ssk} \leftarrow \mathsf{SK}\text{-}\mathsf{sk}[\mathsf{spk}] \end{array}$ return Gl // Join a Group Proc-WM(W) $(W', t, sig) \leftarrow W$ $(., ID_s, ., v, W_{pub}, W_{priv}, e, wpk, spk) \leftarrow$ w' $spk \leftarrow spk[ID_S]$ req $spk \in CL-S[ID_S] \land S-Ver(spk, W', sig)$ (esk, isk) \leftarrow WK-wk[wpk] s. $\sigma \leftarrow$ E-Dec(esk, e) $(s.\gamma, GI, I) \leftarrow K$ -Join(ME, ID_s, W_{pub}, W_{priv} , isk) $\begin{array}{l} (\textbf{s.}\sigma, k_{e}, \textbf{s.}k_{m}, \textbf{s.}C\text{-epid}) \leftarrow \mathsf{PP}(\textbf{s.}\sigma, I, \textbf{v}) \\ \textbf{req} \ \mathsf{M}\text{-}\mathsf{Ver}(\textbf{s.}k_{m}, (\mathsf{W}', \textbf{s.}C\text{-epid}), t) \end{array}$ $s.G \leftarrow GI.G$ $pos \leftarrow *roster-pos(ME, s.G)$ $s.v[s.C-epid] \leftarrow F-Init(k_e, | s.G$ $s.Ep-SPK[s.C-epid, .] \leftarrow spk$ $spk \leftarrow s.Ep-SPK[s.C-epid, ME]$

 $s.C\text{-ssk} \leftarrow SK\text{-sk[spk]}$ $s.Props \leftarrow \epsilon$

```
// Determines if authenticity of epoch vid compromised
*auth-compr(vid)
     return ¬*PP-secure(vid, *Proj-PP(SGM-Data))
// Determines if ID's signature key is compromised in epoch vid
*SK-compr(vid, ID)
 chk HG.getSKIDs(vid, ID) \cap SK-Lk = Ø
// Privacy Predicate
*AM-sec (vid, S, i)
      if *PP-secure(vid, *Proj-PP(SGM-Data))
       return *FS-sec((S, i), *Proj-FS(SGM-Data, vid))
     return false
// Project History Graph to PRF-PRNG Game
*Proj-PP (SGM-Data = (\mathcal{V}, \mathcal{P}, V-Lk, AM-Lk, BR, WK-Lk))
      for (vid, orig, data, pid) \in \mathcal{V}
      \mathcal{V}' + \leftarrow \text{vid}
     for vid \in \mathcal{V}'
       | BI[vid] \leftarrow \neg^*CGKA-priv(vid', *Proj-CGKA(SGM-Data))
     return (V', V-Lk, BI)
// Project History Graph to CGKA Game
*Proj-CGKA (SGM-Data = (V, P, V-Lk, AM-Lk, BR, WK-Lk))
      \begin{array}{l} \textbf{for} \ (\text{vid}, \text{orig}, \text{data}, \textbf{\textit{pid}}) \in \mathcal{V} \\ \mid \mathcal{V}' + \leftarrow \ (\text{vid}, \text{orig}, \textbf{pid}) \end{array} 
     for p = (pid, vid, orig, op, data) \in \mathcal{V}
           if op = add
                 let data = (ID', wkid', skid')
data' ← (ID', wkid')
                 p' \leftarrow (pid, vid, op, orig, data')
           if op = rem
                p' \leftarrow p
           if op = upd
           \begin{array}{c|c} p' \leftarrow \mu \\ p' \leftarrow \mu \\ \mathcal{P}' + \leftarrow p' \end{array}
     return (\mathcal{V}', \mathcal{P}', V-Lk, BR, WK-Lk)
// Project History Graph to FS-GAEAD Game
*Proj-FS (SGM-Data = (\mathcal{V}, \mathcal{P}, V-Lk, AM-Lk, BR, WK-Lk), vid)
      for (vid, ID, AM-Rcvd[ID], AM-Tr[ID]) \in AM-Lk
       AM-Lk' + \leftarrow (ID, AM-Rcvd[ID], AM-Tr[ID])
     return AM-Lk'
Figure 8: Safety oracles and safety predicate of the security
game for SGM schemes.
```

Figure 7: The SGM Construction: main algorithms.

REFERENCES

- J. Alwen, M. Capretto, M. Cueto, C. Kamath, K. Klein, G. Pascual-Perez, K. Pietrzak, and M. Walter. Keep the dirt: Tainted treekem, an efficient and provably secure continuous group key agreement protocol. *IACR Cryptol. ePrint Arch.*, 2019:1489, 2019.
- [2] J. Alwen, S. Coretti, and Y. Dodis. The double ratchet: Security notions, proofs, and modularization for the Signal protocol. In Y. Ishai and V. Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 129–158. Springer, Heidelberg, May 2019.
- [3] J. Alwen, S. Coretti, Y. Dodis, and Y. Tselekounis. Security analysis and improvements for the IETF MLS standard for group messaging. In D. Micciancio and T. Ristenpart, editors, *CRYPTO 2020, Part I*, volume 12170 of *LNCS*, pages 248–277. Springer, Heidelberg, Aug. 2020.
- [4] J. Alwen, S. Coretti, Y. Dodis, and Y. Tselekounis. Modular design of secure group messaging protocols and the security of mls. Cryptology ePrint Archive, Report 2021/1083, 2021. https://ia.cr/2021/1083.
- [5] J. Alwen, S. Coretti, D. Jost, and M. Mularczyk. Continuous group key agreement with active security. In R. Pass and K. Pietrzak, editors, *TCC 2020*, 2020.
- [6] J. Alwen, D. Jost, and M. Mularczyk. On the insider security of MLS. IACR Cryptol. ePrint Arch., 2020:1327, 2020.
- [7] R. Barnes. Subject: [MLS] Remove without double-join (in TreeKEM), 2018. https://mailarchive.ietf.org/arch/msg/mls/Zzw2tqZC1FCbVZA9LKERsMIQXik.
- [8] R. Barnes, B. Beurdouche, J. Millican, E. Omara, K. Cohn-Gordon, and R. Robert. The Messaging Layer Security (MLS) Protocol. Internet-Draft draft-ietf-mlsprotocol-11, Internet Engineering Task Force, Dec. 2020. Work in Progress.
- [9] K. Bhargavan, B. Beurdouche, and P. Naldurg. Formal Models and Verified Protocols for Group Messaging: Attacks and Proofs for IETF MLS. Research report, Inria Paris, Dec. 2019.
- [10] A. Bienstock, Y. Dodis, and P. Rösler. On the price of concurrency in group ratcheting protocols. In R. Pass and K. Pietrzak, editors, *TCC*, 2020.
- [11] C. Brzuska, E. Cornelissen, and K. Kohbrok. Cryptographic security of the mls rfc, draft 11. Cryptology ePrint Archive, Report 2021/137, 2021.
- [12] R. Canetti, J. A. Garay, G. Itkis, D. Micciancio, M. Naor, and B. Pinkas. Multicast security: A taxonomy and some efficient constructions. In *IEEE INFOCOM'99*, pages 708–716, New York, NY, USA, Mar. 21–25, 1999.
- [13] K. Cohn-Gordon, C. Cremers, L. Garratt, J. Millican, and K. Milner. On ends-toends encryption: Asynchronous group messaging with strong security guarantees. In D. Lie, M. Mannan, M. Backes, and X. Wang, editors, ACM CCS 2018, pages 1802–1819. ACM Press, Oct. 2018.
- [14] K. Cohn-Gordon, C. J. F. Cremers, B. Dowling, L. Garratt, and D. Stebila. A formal security analysis of the signal messaging protocol. In 2017 IEEE European Symposium on Security and Privacy, EuroS&P 2017, pages 451–466, 2017.

- [15] C. Cremers, B. Hale, and K. Kohbrok. Revisiting post-compromise security guarantees in group messaging. *IACR Cryptol. ePrint Arch.*, 2019:477, 2019.
- [16] Y. Dodis and N. Fazio. Public key broadcast encryption for stateless receivers. In J. Feigenbaum, editor, *Digital Rights Management*, pages 61–80, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg.
- [17] E. Eaton, D. Jao, and C. Komlo. Towards post-quantum updatable public-key encryption via supersingular isogenies. Cryptology ePrint Archive, Report 2020/1593, 2020. https://eprint.iacr.org/2020/1593.
- [18] A. Fiat and M. Naor. Broadcast encryption. In D. R. Stinson, editor, CRYPTO'93, volume 773 of LNCS, pages 480–491. Springer, Heidelberg, Aug. 1994.
- [19] Z. Jafargholi, C. Kamath, K. Klein, I. Komargodski, K. Pietrzak, and D. Wichs. Be adaptive, avoid overcommitting. In J. Katz and H. Shacham, editors, *CRYPTO 2017*, *Part I*, volume 10401 of *LNCS*, pages 133–163. Springer, Heidelberg, Aug. 2017.
- [20] D. Jost, U. Maurer, and M. Mularczyk. Efficient ratcheting: Almost-optimal guarantees for secure messaging. In Y. Ishai and V. Rijmen, editors, EUROCRYPT 2019, Part I, volume 11476 of LNCS, pages 159–188. Springer, Heidelberg, May 2019.
- [21] Y. Kim, A. Perrig, and G. Tsudik. Group key agreement efficient in communication. IEEE Trans. Computers, 53(7):905–921, 2004.
- [22] H. Krawczyk. Cryptographic extraction and key derivation: The HKDF scheme. In T. Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 631–648. Springer, Heidelberg, Aug. 2010.
- [23] Matthew A. Weidner. Group Messaging for Secure Asynchronous Collaboration. Master's thesis, University of Cambridge, June 2019.
- [24] S. Mittra. Iolus: A framework for scalable secure multicasting. In Proceedings of ACM SIGCOMM, pages 277–288, Cannes, France, Sept. 14–18, 1997.
- [25] S. Panjwani. Tackling adaptive corruptions in multicast encryption protocols. In S. P. Vadhan, editor, TCC 2007, volume 4392 of LNCS, pages 21–40. Springer, Heidelberg, Feb. 2007.
- [26] E. Rescorla. Subject: [MLS] TreeKEM: An alternative to ART. MLS Mailing List, 2018. https://mailarchive.ietf.org/arch/msg/mls/WRdXVr8iUwibaQu0tH6sDnqU1no.
- [27] D. G. Steer, L. Strawczynski, W. Diffie, and M. J. Wiener. A secure audio teleconference system. In S. Goldwasser, editor, CRYPTO '88, 1988.
- [28] D. Wallner, E. Hardner, and R. Agee. Key management for multicast: Issues and architectures. IETF RFC2676, 1999. https://tools.ietf.org/html/rfc2627.
- [29] M. Weidner. Group messaging for secure asynchronous collaboration. MPhil Dissertation, 2019. https://mattweidner.com/acs-dissertation.pdf.
- [30] C. K. Wong, M. Gouda, and S. S. Lam. Secure group communications using key graphs. *IEEE/ACM Transactions on Networking*, 8(1):16–30, Feb. 2000.

A CONTINUOUS GROUP KEY AGREEMENT

A.1 Syntax

This section introduces the formal syntax of a CGKA scheme. Parties are identified by unique party IDs ID chosen from an arbitrary fixed set. In the following, γ and γ' denote the internal state of the CGKA scheme before and after an operation, respectively. A CGKA scheme CGKA consists of eight algorithms:

- For PKI:
 - − (ipk, isk) ← Gen-IK: generates and outputs a new initial key pair.
- For group creation:
 - γ' ← Create(ID, isk): initializes the state and creates group with self, using as init secret key isk.
- For proposals:
 - (γ' , *P*) ← Add(γ , ID', ipk'): generates a proposal to add party ID' using initial public key ipk';
 - (γ', P) ← Remove(γ, ID'): generates a proposal to remove party ID';
 - − $(\gamma', P) \leftarrow Update(\gamma)$: generates a proposal for self to update.
- For commits:
 - $(\gamma', I, W_{\text{pub}}, \mathbf{W}_{\text{priv}}, T) \leftarrow \text{Commit}(\gamma, \mathbf{P})$: creates a commit corresponding to a vector **P** of proposals and outputs the resulting *commit secret I*, a *public* welcome message W_{pub} as well as a vector \mathbf{W}_{priv} of *private* welcome messages (for newly added parties), and a control message *T* (for existing group members);
 - (γ' , Gl, *I*) ← Proc-Com(γ , *T*, **P**): used by existing group members to process control message *T* w.r.t. vector of proposals **P**, and reach new epoch; outputs updated group information GI (where GI = \perp if *T* is considered invalid) and a commit secret *I*;
 - (γ , GI, *I*) ← Join(γ , orig, W_{pub} , W_{priv} , isk): used by newly added group members to process welcome message W_{pub} and W_{priv} , generated via a commit operation by orig, and join a group using initial secret key isk; outputs group information GI (where GI = \perp if W_{pub} or W_{priv} are considered invalid) and a commit secret *I*.

A.2 Security

CGKA schemes must satisfy *correctness*, i.e., all group members output the same keys in every epoch. Furthermore, the keys must be private, and the CGKA scheme must satisfy post-compromise forward secrecy (PCFS).

Similarly to the SGM game, the CGKA game allows adversary \mathcal{A} to control the execution of and attack a single group. In particular, \mathcal{A} controls who creates the group, who is added and removed, who updates, who commits, etc. The attacker is also allowed to leak the state of any party (whether currently part of the group or not) at any time. The privacy of keys is captured by considering a challenge that outputs either the actual key output by the CGKA scheme or a

truly random one—which one is determined by an internal random bit *b*, which must be guessed by \mathcal{A} at the end of the game.

A.2.1 Bookkeeping. Similarly to SGM schemes, the CGKA security game keeps track of all the relevant execution data with the help of a so-called *history graph*, and the recorded data informs a *safety predicate* (pertaining to privacy) evaluated at the end of the game. The following paragraphs outline the differences in bookkeeping between the CGKA game and the SGM game.

History graphs. A history graph is again a directed tree whose nodes correspond to the group state in the various epochs of an execution. As with the SGM game, there are three types of nodes: v_{root} , v_{create} nodes, and commit nodes v. A node consists of the following values:

$$v = (vid, orig, pid),$$

where

- vid is the node's (unique) ID,
- orig is the party that caused the node's creation, i.e., orig is either the group creator or the committer,
- **pid** is the vector of (IDs of) proposals (see below) that were included in the commit.

The history graph is accessed by the security-game oracles via the "HG object" HG, which provides information via several methods (explained later as they are needed). The (ID of the) node corresponding to a party's current state is stored by the array V-Pt[ID].

Proposals. Similarly to HG and the history graph, the object Props keeps track of proposals. The information recorded about each proposal is a vector

$$p = (pid, vid, op, orig, data),$$

where

- pid is the proposal's (unique) ID,
- vid is the (ID of) the HG node corresponding to the epoch in which the proposal was created,
- $op \in \{add, rem, upd\}$ is the type of proposal,
- orig is the party that issued the proposal, and
- data is additional data.

Similarly to the HG object, Props will also export several useful methods (also explained later) to the oracles of the security game.

PKI bookkeeping. The PKI is significantly simpler in the CGKA game: it only handles so-called *init keys (IKs)*, each of which has an ID ikid. Arrays are maintained to map ikids to the corresponding public init key (IK-PK) and secret init key (IK-SK). Moreover, IK-St records the IKs stored by each party, and IK-Tr and IK-Lk keep track of "trash" IKs (old secret init keys not deleted) and leaked IKs.

Keys and Challenges. The array Key[vid] stores, for each vid the corresponding key, and (the Boolean arrays) Reveal[vid] and Chall[vid] store whether the key was revealed and challenged, respectively.

A.2.2 Initialization. The initialization of the CGKA game proceeds along very similar lines to that of the SGM game. It is depicted in Figure 9.

// General	$V\text{-}Lk \gets \emptyset$
$b \leftarrow \{0, 1\}$	$P\text{-}St[\cdot] \leftarrow \emptyset$
idCtr++	// Miscellaneous
$\forall ID : \gamma[ID] \leftarrow Init(ID)$	$Key[\cdot] \gets \emptyset$
// Communication	$Reveal[\cdot] \leftarrow \emptyset$
$CM[\cdot, \cdot] \leftarrow \epsilon$	$Chall[\cdot] \leftarrow \emptyset$
$PM[\cdot, \cdot] \leftarrow \epsilon$	$Del[\cdot] \leftarrow true$
$WM_{pub}[\cdot, \cdot] \leftarrow \epsilon$	$BR[\cdot] \leftarrow false$
$WM_{priv}[\cdot, \cdot] \leftarrow \epsilon$	// PKI
WM-Lk $\leftarrow \emptyset$	$IK\text{-}PK[\cdot] \leftarrow \epsilon$
// History Graph	$IK\text{-}SK[\cdot] \leftarrow \epsilon$
vid _{root} ← HG.init	$IK\operatorname{-St}[\cdot] \gets \emptyset$
$V\text{-}Pt[\cdot] \leftarrow vid_{root}$	$IK\text{-}Tr \gets \emptyset$
$V\text{-}Tr[\cdot] \leftarrow \emptyset$	$IK\text{-}Lk \gets \emptyset$

Figure 9: Initialization of the security game for CGKA schemes.

A.2.3 Oracles. All adversary oracles described below proceed according to the same pattern: (a) verifying the validity of the oracle call, (b) retrieving values needed for (c), (c) running the corresponding SGM algorithm, and (d) updating the bookkeeping. Validity checks (a) are described informally in the text below; a formal description is provided in Figure 14. Note that, most of the time, (b) and (c) are straight-forward and are not mentioned in the descriptions. To improve readability, lines (c) are highlighted.

A.2.4 PKI.. The spirit of the PKI (Figure 10) in the CGKA game differs somewhat from the SGM case. Init keys (IKs) are either honestly generated by a party, via oracle **gen-new-ik**, or registered by the attacker \mathcal{A} , via oracle **reg-ik**. Note that the CGKA game will allow the attacker to pick which IK a new group member is added with; hence, there is no association between IKs and identities.

Oracle **gen-new-ik**(ID) runs the IK-generation algorithm Gen-IK and stores the resulting key pair with a new ikid. The oracle also records in IK-St that ID now stores the IK with ID ikid. Oracle **reg-ik**(ipk) allows \mathcal{A} to register an IK public key ipk with the game (and get an ikid for it).

A.2.5 Main oracles. The main oracles of the CGKA game are split into two figures: oracles related to (i) group creation, proposals, and commits (Figure 11), and (ii) message processing, corruption, and challenges (Figure 13). The validity of all oracle calls is checked by the corresponding compatibility functions (Figure 14).

Group creation. The attacker can instruct a party ID to create a new group by calling the group-creation oracle (Figure 3), which works analogously to the corresponding oracle in the SGM game. The bookkeeping is updated as follows: A call is made to the HG.create(ID, skid) method. This causes HG to create a node

$$v = \begin{pmatrix} .vid \leftarrow idCtr++\\ .orig \leftarrow ID\\ .pid \leftarrow \bot \end{pmatrix}$$

// Instruct ID to generate a new initial key gen-new-ik(ID) $(\gamma[ID], (ipk, isk)) \leftarrow Gen-IK(\gamma[ID])$ $ikid \leftarrow idCtr++$ $IK-PK[ikid] \leftarrow ipk$ $IK-SK[ikid] \leftarrow isk$ IK-St[ID] +← ikid return (ikid, ipk) // Allows attacker to register IKs with game reg-ik(ipk) req ∄ ikid : IK-PK[ikid] = ipk ikid ← idCtr++ IK-Lk +← ikid $\mathsf{IK}\text{-}\mathsf{PK}[\mathsf{ikid}] \leftarrow \mathsf{ipk}$ return ikid Figure 10: PKI-related oracles of the security game for contin-

uous group key agreement schemes.

as a child of v_{root} and return vid = v.vid. Note that the CGKA groupcreation oracle also stores the key I output by Create in the array Key.

Proposal oracles. The oracles **prop-{add, rem, up}-user** (Figure 11) allow the attacker to instruct a party ID to issue add/remove/update proposals. These oracles work much like their counterparts in the SGM game. Bookkeeping is updated by calling Props.new(op, ID, data), which records proposal data

$$p = \begin{pmatrix} .pid & \leftarrow & idCtr \leftrightarrow \\ .vid & \leftarrow & V-Pt[ID] \\ .op & \leftarrow & op \\ .orig & \leftarrow & ID \\ .data & \leftarrow & data \end{pmatrix},$$

where op $\in \{ \mathsf{add}, \mathsf{rem}, \mathsf{upd} \}$ is the proposal type and where data stores

- (*add proposals*) data = (ID', ikid'), where ikid' is the (ID of the) init key ID' is to be added with;
- (remove proposals) data = ID';
- (update proposals) data = \perp .

Observe that one crucial difference between the CGKA and SGM games is that the CGKA game does not explicitly model the delivery of proposals. This is due to the fact that a CGKA is assumed to be run over authenticated channels, and hence there is no way for \mathcal{A} to inject malicious proposals, and, consequently, the proper delivery of proposals can be outsourced to the higher-level protocol.

Creating commits. The commit oracle (Figure 11) allows \mathcal{A} to instruct a party ID to execute a commit operations. It works analogously to its SGM-game counterpart. Bookkeeping calls HG.commit(ID, **pid**), which creates a new HG node

$$v = \begin{pmatrix} .vid & \leftarrow & idCtr ++ \\ .orig & \leftarrow & ID \\ .pid & \leftarrow & pid \end{pmatrix}$$

```
// ID creates new group
create-group(ID, ikid, r)
     req *compat-create(ID, ikid)
     isk \leftarrow IK-SK[ikid]
     (\gamma[ID], I) \leftarrow Create(\gamma[ID], isk; r)
     vid \leftarrow HG.create(ID, r)
     Key[vid] \leftarrow I
     \mathsf{BR}[\mathsf{vid}] \leftarrow (r \neq \bot)
     V-Pt[ID] \leftarrow vid
     return vid
// ID proposes to add ID'
prop-add-user(ID, ID', ikid')
     req *compat-prop(add, ID, ID', ikid')
     ipk' ← IK-PK[ikid']
     (\gamma[ID], P) \leftarrow Add(\gamma[ID], ID', ipk')
     pid ← Props.new(add, ID, (ID', ikid'))
     PM[pid] \leftarrow P
     return (pid, P)
// ID proposes to remove ID'
prop-rem-user(ID, ID')
     req *compat-prop(rem, ID, ID', ⊥)
     (\gamma[ID], P) \leftarrow \text{Remove}(\gamma[ID], ID')
     pid ← Props.new(rem, ID, ID')
     PM[pid] \leftarrow P
     return (pid, P)
// ID proposes to update
prop-up-user(ID, r)
     req \ ^{*}compat-prop(\mathsf{upd},\mathsf{ID},\bot,\bot)
     (\gamma[ID], P) \leftarrow Update(\gamma[ID]; r)
     pid \leftarrow Props.new(upd, ID, \perp)
     BR[vid] \leftarrow (r \neq \bot)
     \mathsf{PM}[\mathsf{pid}] \leftarrow P
     return (pid, P)
// ID commits proposals pid
commit(ID, pid, r)
     req *compat-commit(ID, pid)
     vid ← HG.commit(ID, pid)
     (\gamma[ID], I, W_{pub}, \mathbf{W}_{priv}, T) \leftarrow Commit(\gamma[ID], PM[pid]; r)
     Key[vid] \leftarrow I
     BR[vid] \leftarrow (r \neq \bot)
     for ID' \in HG.roster(V-Pt[ID])
      | CM[vid, ID'] \leftarrow T
     (ID_1, \ldots, ID_m) \leftarrow Props.addedIDs(pid)
     for i = 1, ..., m
          WM_{pub}[vid, ID_i] \leftarrow W_{pub}
          WM_{priv}[vid, ID_i] \leftarrow W_{priv}[ID_i]
     return (vid, T, W_{pub}, W_{priv})
Figure 11: Oracles for group creation, add, remove, update
```

Figure 11: Oracles for group creation, add, remove, update proposals and commits of the security game for continuous group key agreement schemes. The compatibility functions are described in the accompanying text; a formal description is provided in Figure 14.

```
// Process commit msg for ID and epoch vid
process(vid, ID)
    req *compat-process(vid, ID)
    T \leftarrow CM[vid, ID]
    P \leftarrow PM[HG[vid].pid]
    (\gamma[ID], GI, I) \leftarrow Proc-Com(\gamma[ID], T, P)
    if \negHG.checkGI(vid, GI) \lor Key[vid] \neq I
        win
    if ¬Del[ID]
        V-Tr[ID] \leftarrow V-Pt[ID]
    if HG.isRemoved(vid, ID)
        V-Pt[ID] \leftarrow vid_{root}
    else
        V-Pt[ID] \leftarrow vid
    i[\mathsf{ID}] \leftarrow 0
    P-St[ID] \leftarrow \emptyset
    return Gl
// Reveal the update secret of epoch vid
reveal(vid)
    req Key[vid] \neq \epsilon
    req \neg(Reveal[vid] \lor Chall[vid])
    Reveal[vid] \leftarrow true
    return Key[vid]
// Challenge the update secret of epoch vid
chall(vid)
    req Key[vid] \neq \epsilon
    req \neg(Reveal[vid] \lor Chall[vid])
    I_0 \leftarrow \text{Key[vid]}
    I_1 \leftarrow I
    Chall[vid] \leftarrow true
    return Ib
// Welcome msg. of epoch vid delivered to ID
dlv-WM(vid, ID)
    req *compat-dlv-WM(vid, ID)
    W_{\text{pub}} \leftarrow WM\text{-pub}[vid, ID]
    W_{priv} \leftarrow WM\text{-}priv[vid, ID]
    ikid \leftarrow HG.addedIK(vid, ID)
    isk \leftarrow IK-SK[ikid]
    let v s.t. v.vid = vid
    orig \leftarrow v.orig = vid
    (\gamma[ID], GI, I) \leftarrow Join(ID, orig, W_{pub}, W_{priv}, isk)
    if \negHG.checkGI(vid, GI) \lor Key[vid] \neq I
        win
    if ¬Del[ID]
     IK-Tr[ID] +← ikid
    IK-St[ID] -← ikid
    V-Pt[ID] \leftarrow vid
return Gl
Figure 12: Part one of oracles for message processing, cor-
```

Figure 12: Part one of oracles for message processing, corruption, and challenges of the security game for continuous group key agreement schemes. The compatibility functions are described in the accompanying text; a formal description is provided in Figure 14.

// Corrupt ID

corr(ID)

// Safety for privacy

*priv-safe

return \forall vid : Chall[vid] \implies *CGKA-priv(vid)

Figure 13: Part two of oracles for message processing, corruption, and challenges of the security game for continuous group key agreement schemes. The compatibility functions are described in the accompanying text; a formal description is provided in Figure 14.

as a child of ID's current epoch V-Pt[ID] and returns vid = v.vid. Note that the key output by Commit is stored in array Key. Furthermore, observe that the algorithm outputs (one) public and several private welcome messages (one for each newly added party). The private welcome messages are not returned to \mathcal{A} and are delivered securely to their intended recipient—the idea being that these messages are encrypted by the higher-level application.

Process control messages. The oracles **process** and **dlv-WM** (Figure 13) allow the attacker to deliver commit messages (to existing group members) resp. welcome messages (to new group members). The oracles works much like their SGM counterparts, except that all the information required by Proc-Com resp. Join is supplied by the game. The oracles also check that the key *I* output by Proc-Com resp. Join and matches the one stored in array Key during the corresponding call to **commit**.

Key-reveal and challenge oracles. For each epoch, the attacker \mathcal{A} gets to either see the actual key output by the protocol or a challenge by calling **reveal** or **chall** (Figure 13), respectively. Oracle **chall** outputs either the real key or a completely random one, depending on the secret internal bit *b* chosen at the onset of the game.

Corruption oracles. The oracles related to corruption in the CGKA game (Figure 13) are:

- **corr**(ID): leaks the state of ID to \mathcal{A} ;
- no-del(ID): instructs ID to stop deleting old values.

A.2.6 Safety. At the end of the execution of the CGKA security game, the procedure ***priv-safe** ensures that the attacker has only challenged in epochs that are considered secure by the (generic) safety predicate ***CGKA-priv**. If the condition is not satisfied, the attacker loses the game.

A.2.7 Advantage. Let $\Pi = {}^{*}CGKA$ -priv be the generic safety predicate used in the CGKA definition. The attacker \mathcal{A} is parameterized by it's running time, *t*, and the number of challenge queries, *q*, and

referred to as (t, q)-attacker. The advantage of \mathcal{A} against a CGKA scheme K w.r.t. to predicate Π is denoted by $\operatorname{Adv}_{\operatorname{CGKA, II}}^{\mathsf{K}}(\mathcal{A})$.

Definition A.1. A CGKA scheme K is (t, q, ε) -secure w.r.t. predicate II, if for all (t, q)-attackers,

$$\operatorname{Adv}_{\operatorname{CGKA,\Pi}}^{\operatorname{K}}(\mathcal{A}) \leq \varepsilon$$
.

```
*compat-create(ID, ikid)
     chk ikid ∈ IK-St[ID]
     return V-Pt[ID] = vid<sub>root</sub>
*compat-prop(op, ID, ID', ikid)
     \mathsf{vid} \leftarrow \mathsf{V}\text{-}\mathsf{Pt}[\mathsf{ID}]
     chk vid \neq vid<sub>root</sub>
     G \leftarrow \text{HG.roster(vid)}
     select op
           case add do
                chk ∃ipk' : IK-PK[ikid'] = ipk'
                chk \mathsf{ID}' \notin G
           case rem do
            chk \mathsf{ID}' \in G
     return true
*compat-commit(ID, pid)
     vid \leftarrow V-Pt[ID]
     chk vid \neq vid<sub>root</sub>
     G \leftarrow HG.roster(vid)
     for pid ∈ pid
           p \leftarrow \text{Props[pid]}
           chk p.vid = vid
           \mathbf{G} \leftarrow {}^*\mathbf{app}\text{-}\mathbf{prop}(\mathbf{G},p)
           \mathbf{chk} \mathrel{\mathsf{G}} \neq \bot
     chk ID ∉ G
     return true
*compat-process(vid, ID)
     chk HG.isChild(V-Pt[ID], vid)
     return true
*app-prop(G, p)
     req p.orig \in G
     select p.op
           case add do
                (ID', \cdot) \leftarrow p.data
                 req ID′ ∉ G
                G \leftrightarrow ID'
           case rem do
                ID' \leftarrow p.data
                 req ID' \in G
                 G \rightarrow ID'
     return G
*compat-dlv-WM(vid, ID)
     chk V-Pt[ID] = vid<sub>root</sub>
     ikid \leftarrow HG.addedIK(vid, ID)
     chk ikid \neq \bot
            \land ikid \in IK-St[ID]
```

return true

Figure 14: Compatibility oracles of the security game for continuous group key agreement schemes.

B FORWARD-SECURE GROUP AEAD

B.1 Syntax

- v ← Init(k_e, n, ID): takes as input a key k_e, the group size n, as well as a party ID ID and generates the initial state.
- (v', e) ← Send(v, a, m) generates ciphertext e encrypting plaintext m and authenticating associated data a;
- (v', S, i, m) ← Rcv(v, a, e): decrypts ciphertext e to plaintext m and verifies associated data a; also outputs a pair (S, i) consisting of sender ID S and message index i.

B.2 Security

B.2.1 Main oracles. The oracles of the FS-GAEAD game are depicted in Figure 15.

Initialization. At the onset of the FS-GAEAD security game, a random bit *b* and a uniformly random key k_e are chosen, and the initialization algorithm is run for all group members in *G* (which is specified by an argument to the initialization procedure). The security game also initializes a message counter *i*[ID] for each party. Additionally, the game maintains the following variables, which work analogously to their counterparts in the SGM game.

- a set Chall of pairs (S, *i*) recording that the *i*th message sent by S was a challenge,
- a trash array AM-Tr[ID] of pairs (S, i) keeping track of nondeleted key material by ID,
- a set AM-Lk of elements (ID, AM-Rcvd, AM-Tr[ID]), where AM-Tr[ID] keeps track of the messages whose key material is leaked via corruption of ID, and AM-Rcvd are the messages that have been already received at the time of corruption.
- a Boolean array Del[ID] keeping track of which parties are deleting old values, and
- an array AM[S, *i*, R] of triples (*a*, *m*, *e*) consisting of associated data (AD), plaintext, and ciphertext.

The recorded data informs a safety predicate ***priv-safe** evaluated at the end of the game to determine whether a given execution was legal.

Sending messages and challenges. The oracle **send**(S, a, m) allows the attacker to have party S send AD a and message m (to all other group members). The oracle runs algorithm Send, which produces a ciphertext e. The triple (a, m, e) is recorded in array AM. Oracle **chall** works similarly, except that it takes two (equal-length) messages as input and passes one of them to Send; which one is chosen is determined by the secret random bit b chosen initially. Furthermore, the pair (S, i[S]) is recorded as being a challenge.

Delivering and injecting ciphertexts. Oracle **dlv-AM**(S, *i*, R) allows \mathcal{A} to have the ciphertext corresponding to the *i*th message sent by S delivered to R. The ciphertext and the corresponding AD are fed to algorithm Rcv, which must correctly decrypt the ciphertext and identify sender S and message number *i*. The pair (S, *i*) is set as "received" by R by setting AM[S, *i*, R] \leftarrow received which indicates that the corresponding key material should now have been deleted by R; in case R does not delete old values (i.e., Del[R] = false), the pair is added to AM-Tr.

// Initialize group

```
init(G)
      b \leftarrow \{0, 1\}
      k_e \leftarrow \mathcal{K}
      for ID \in G
             v[\mathsf{ID}] \leftarrow \mathsf{Init}(k_e, |G|, \mathsf{ID})
            i[ID] \leftarrow 0
      \mathsf{Chall} \gets \emptyset
      \mathsf{AM}\text{-}\mathsf{Tr}[\cdot] \gets \emptyset
      \mathsf{AM}\text{-}\mathsf{Lk} \gets \emptyset
      \mathsf{AM}[\cdot,\cdot,\cdot] \leftarrow \epsilon
      \text{Del}[\cdot] \leftarrow \epsilon
// S sends m with AD a
send(S, a, m)
      (v[S], e) \leftarrow \text{Send}(v[S], a, m)
      i[S] ++
      for R \in G \setminus \{S\}
        | AM[S, i[S], R] \leftarrow (a, m, e)
      return e
// Corruption of ID
corr(ID)
      AM-Rcvd \leftarrow \{(S, i) \mid AM[S, i, ID] = received\}
      AM-Lk +← (ID, AM-Rcvd, AM-Tr[ID])
      return v[ID]
// Message delivery
dlv-AM(S, i, R)
      req AM[S, i, R] \neq {\epsilon, received}
      (a, m, e) \leftarrow AM[S, i, R]
      (v[R], S', i', m') \leftarrow \operatorname{Rev}(v[R], a, e)
      if (S', i', m') \neq (S, i, m)
             win
      if ¬Del[R]
            AM-Tr[R] + \leftarrow (S, i)
      AM[S, i, R] \leftarrow received
// stop deletions for ID
no-del(ID)
      Del[ID] \leftarrow false
// Message injection
inj-AM(a', e', R)
      req \forall S, i, i : AM[S, i, R] \neq (a', i, e')
      (v[S], S', i', m') \leftarrow \operatorname{Rev}(v[R], a, e)
      if m' \neq \bot
             if (*AM-sec(S', i')) \lor AM[S', i', R] = received
      | | win
AM[S', i', R] \leftarrow received
      return (S', i', m')
// S challs m_0, m_1
chall(S, a, m_0, m_1)
      reg |m_0| = |m_1|
      (v[S], e) \leftarrow \text{Send}(v[S], a, m_b)
      i[S]++
      for R \in G \setminus \{S\}
           \mathsf{AM}[\mathsf{S}, i[\mathsf{S}], \mathsf{R}] \leftarrow (a, m, e)
      Chall + \leftarrow (S, i)
      return e
// Safety predicate
*priv-safe
      if \forall (S, i) : (S, i) \in \text{Chall} \implies {}^{*}FS\text{-sec}(S, i)
        return true
       return false
```

Figure 15: The main oracles of the FS-GAEAD security game.

Oracle **inj-AM** can be used by \mathcal{A} to inject any non-honestly generated AD/ciphertext pair. Algorithm Rcv must reject all such pairs unless they are compromised, which is the case if \mathcal{A} has learned the corresponding key material via state compromise. Whether or not this is the case is determined by the safety helper function ***FS-sec**, which is discussed below. Irrespective of whether a compromise has occurred, Rcv is required to detect and prevent replays.

Corruption. By calling oracle **corr**(ID) the attacker can learn the current state of party ID. The game adds the triple (ID, AM-Rcvd, AM-Tr[ID]) to the set AM-Lk to indicate that ID's state is now compromised, but the messages recorded in AM-Rcvd are supposed to remain secure (because the corresponding key material must have been deleted); the game also adds potential trash, AM-Tr[ID], stored by ID to the same set.

B.2.2 Privacy-related safety. At the end of the execution of the FS-GAEAD security game, the procedure ***priv-safe** ensures that the attacker has only challenged messages that are considered secure by the (generic) safety predicate ***FS-sec**. If the condition is not satisfied, the attacker loses the game.

B.2.3 Advantage. Let Π = ***FS-sec** be the generic safety predicate used in the FS-GAEAD definition. The attacker \mathcal{A} is parameterized by it's running time *t* and the number of challenge queries *q*, referred to as (*t*, *q*)-attacker. The advantage of \mathcal{A} against an FS-GAEAD scheme F w.r.t. to predicate Π is denoted by $\operatorname{Adv}_{fen}^{FS}(\mathcal{A})$.

Definition B.1. An FS-GAEAD scheme F is (t, q, ε) -secure w.r.t. predicate II, if for all (t, q)-attackers,

$$\operatorname{Adv}_{\operatorname{FS-GAEAD},\Pi}^{\mathsf{F}}(\mathcal{A}) \leq \varepsilon$$
.

C PRF-PRNGS

C.1 Syntax

A PRF-PNRG PP is an algorithm $(\sigma', R) \leftarrow PP(\sigma, I, C)$: it takes the current state σ , absorbs input *I* along with context information *C*, and produces a new state σ' as well as an output string *R*.

C.2 Security

A PRF-PRNG must satisfy PCFS (cf. Section 3.2) and be resilient to splitting-attacks. Therefore, the security game for PRF-PRNGs (cf. Figure 16) follows the same history-graph approach as the definitions of SGM and CKGA. However, since the game only consists of the state of the PRF-PRNG and there are no parties, it suffices to keep track of a much smaller amount information:

- Nodes of the history graph only consist of the vid.
- For every node vid,
 - the value σ [vid] stores the corresponding state of the PRF-PRNG,
 - the value *R*[vid] stores the corresponding output of the PRF-PRNG, and
 - the value BI[vid] is a flag indicating whether the input *I* absorbed to reach the state σ [vid] is known to the attacker.
- The set V-Lk records the vids for which the PRF-PRNG state is leaked to the attacker.

The root node vid_{root} of the history graph corresponds to the initial state of the PRF-PRNG, which is assumed to be the all-zero string. The attacker \mathcal{A} has the following capabilities:

• He may create a new child state of any node vid by calling oracle **process** and specifying an input/context pair (*I*, *C*); of course, only one such call per triple (vid, *I*, *C*) is allowed. If

init

```
b \leftarrow_{\mathsf{R}} \{0, 1\}
     vid_{root} \leftarrow HG.init
      \sigma[vid_{root}] \leftarrow 0
      \mathsf{V}\text{-}\mathsf{L}\mathsf{k} \gets \emptyset
      R[\cdot] \leftarrow \emptyset
      Reveal[\cdot] \leftarrow \emptyset
      Chall[\cdot] \leftarrow \emptyset
     BI[\cdot] \leftarrow \emptyset
corr (vid)
     V-Lk +← vid
     return \sigma[vid]
process(vid, I, C)
     req \nexists child of vid for (I, C)
     vid' \leftarrow HG.create(vid, I)
      \mathsf{BI}[\mathsf{vid'}] \leftarrow (I \neq \bot)
      (\sigma[vid'], R[vid']) \leftarrow PP(\sigma[vid], I, C)
     return vid'
reveal(vid)
     req R[vid] \neq \varepsilon
      req \neg(Reveal[vid] \lor Chall[vid])
      Reveal[vid] \leftarrow true
     return R[vid]
chall(vid)
      req R[vid] \neq \varepsilon
      req \neg(Reveal[vid] \lor Chall[vid])
      R_0 \leftarrow R[vid]
      R_1 \leftarrow \mathcal{R}
      Chall[vid] \leftarrow true
     return R<sub>h</sub>
```

safe

```
return ∀vid : Chall[vid] ⇒ *PP-secure(vid)
Figure 16: PRF-PRNG security game.
```

the call is made with $I \neq \bot$, the game samples *I* it randomly. The value BI[vid] is set accordingly.

- He may reveal or challenge outputs *R*[vid] corresponding to arbitrary nodes vid ≠ vid_{root} by calling the corresponding oracles **reveal** and **chall**, respectively. The flags Reveal[vid] resp. store Chall[vid] whether a reveal resp. a challenge has been requested for vid.
- Finally, A can also leak the state σ[vid] for any epoch vid ≠ vid_{root} using oracle corr.

As per usual, at the end of the game, the oracle **safe** ensures that \mathcal{A} does not win the game trivially; **safe** uses a generic safety predicate ***PP-secure**.

C.2.1 Advantage. Let Π = ***PP-secure** be the generic safety predicate used in the PRF-PRNG definition. The attacker \mathcal{A} is parameterized by it's running time *t*, referred to as *t*-attacker. The advantage of \mathcal{A} against a PP scheme PRF-PRNG w.r.t. to predicate Π is denoted by Adv^{PP}_{PRF-PRNG,\Pi}(\mathcal{A}).

Definition C.1. A PRF-PRNG scheme PP is (t, ε) -secure w.r.t. predicate Π , if for all *t*-attackers,

 $\operatorname{Adv}_{\mathsf{PRF}-\mathsf{PRNG}\Pi}^{\mathsf{PP}}(\mathcal{A}) \leq \varepsilon$.

D SECURE GROUP MESSAGING

In this section we define compatibility helpers for SGM.

*compat-create(ID, skid, wkid) *compat-dlv-CM(ID, vid) chk V-Pt[ID] = vid_{root} chk HG.isChild(V-Pt[ID], vid) chk SK-ID[skid] = ID **chk** HG[vid].pid ⊆ P-St[ID] chk WK-SK[wkid] = skid return true **chk** skid ∈ SK-St[ID] *compat-inj-CM(ID, T') $vid \leftarrow V-Pt[ID]$ **chk** wkid ∈ WK-St[ID] **chk** vid \neq vid_{root} return true *compat-prop(op, ID, ID', skid) **chk** \forall vid' \in HG.children(vid) : $T' \neq CM[vid', ID]$ vid ← V-Pt[ID] **chk** vid ≠ vid_{root} return true *compat-dlv-CM(ID, vid) $G \leftarrow \text{HG.roster(vid)}$ chk HG.isChild(V-Pt[ID], vid) select op $\mathbf{chk}\;\mathsf{HG}[\mathsf{vid}].\mathsf{pid}\subseteq\mathsf{P}\text{-}\mathsf{St}[\mathsf{ID}]$ case add do return true chk $CL-KB[ID, ID'] \neq \emptyset$ *compat-inj-CM(ID, T') **chk** $\mathsf{ID}' \notin G$ vid \leftarrow V-Pt[ID] **chk** vid \neq vid_{root} case rem do $chk |D' \in G$ **chk** \forall vid' \in HG.children(vid) : $T' \neq \mathsf{CM}[\mathsf{vid}',\mathsf{ID}]$ case upd do **chk** skid ∈ SK-St[ID] return true return true *compat-dlv-WM(ID, vid) *compat-dlv-PM(ID, pid) chk V-Pt[ID] = vid_{root} chk Props[pid].vid = V-Pt[ID] wkid \leftarrow HG.addedWK(ID, vid) return true **chk** wkid $\neq \bot$ ∧ wkid ∈ WK-St[ID] *compat-inj-PM(ID, P') $\mathsf{vid} \gets \mathsf{V}\text{-}\mathsf{Pt}[\mathsf{ID}]$ return true **chk** vid ≠ vid_{root} *compat-inj-WM(ID, W') | vid \leftarrow V-Pt[ID] chk 7 pid : Props[pid].vid = vid $\land P' = PM[pid]$ $\mathbf{chk} \ \mathsf{vid} = \mathsf{vid}_{\mathsf{root}}$ chk ∀vid' : $W' \neq CM[vid', ID]$ return true *compat-commit(ID, pid) return true $vid \leftarrow V-Pt[ID]$ *compat-send(S) chk V-Pt[ID] ≠ vid_{root} chk vid ≠ vid_{root} chk pid \subseteq P-St[ID] return true $G \leftarrow HG.roster(vid)$ *compat-chall(S, m_0, m_1) **chk** V-Pt[ID] \neq vid_{root} for pid ∈ pid **chk** $|m_0| = |m_1|$ $G \leftarrow *app-prop(G, pid)$ **chk** G $\neq \bot$ return true *compat-dlv-AM(vid, S, i, R) return true $chk vid \in V-St[R]$ *app-prop(G, pid) **chk** AM[vid, S, *i*, R] ∉ $p \leftarrow \text{Props[pid]}$ {\epsilon, received} req p.orig \in G return true select op *compat-inj-AM(a, e, R) case add do **chk** \forall S, vid, *i*, *m* $(\mathsf{ID}',\cdot,\cdot) \gets p.\mathsf{data}$ $AM[vid, S, i, R] \neq (a, m, e)$ req ID' ∉ G return true G +← ID' case rem do $ID' \leftarrow p.data$ $\mathbf{req}\;\mathsf{ID}'\in G$ -← ID' G

return GFigure 17: Compatibility oracles of the security game for secure group-messaging schemes.

E SGM CONSTRUCTION

In this section we formally define the helper functions of our SGM construction.

```
// Returns ids, wpks, of new parties
*added(P)
     ID[\cdot] \leftarrow \varepsilon
     wpk[\cdot] \leftarrow \varepsilon
     i \leftarrow 1 for P \in P
           if P = ("add", \cdot, \cdot, (ID_a, kb'), \cdot)
                 (wpk, \cdot) \leftarrow kb'
                 ID[i] \leftarrow ID_a
                 wpk[i] \leftarrow wpk
     return (ID, wpk)
    Returns spks after applying props
11
*new-spks(epid, P)
     spk \leftarrow s.Ep-SPK[epid]
     for P \in P
           if P = ("add", \cdot, \cdot, (ID_a, kb'), \cdot)
                (\cdot, spk) \leftarrow kb'
                 spk[ID_a] \leftarrow spk
           \mathbf{if} \ P = ("\mathsf{rem}", \cdot, \cdot, \mathsf{ID}_r, \cdot)
               spk[ID_r] \leftarrow \varepsilon
           if P = ("upd", \cdot, ID_u, spk, \cdot)
            | spk[ID<sub>u</sub>] \leftarrow spk
     return spk
    Returns position of ID in roster
11
*roster-pos(ID, G)
     [\mathsf{ID}_1,\ldots,\mathsf{ID}_n] \leftarrow \mathsf{G}
     for i \in [1, n]
           if ID_i = ID
            | return i
// Returns Proposal Info
*get-propInfo(P')
     if P' = (\text{"add"}, \cdot, \text{ID}_s, (\text{ID}_a, \text{kb'}), \bar{P})
           op = add
           orig = ID_s
           (wpk, spk) \leftarrow kb'; data = (ID_a, wpk, spk)
     if P' = ("rem", \cdot, ID_s, ID_r, \overline{P})
           op = rem
           orig = ID_s
           data = ID_r
     if P' = ("upd", \cdot, ID_s, spk, \bar{P})
           op = upd
           orig = ID_s
           data = spk
     return (op, orig, data)
     Figure 18: The SGM Construction : Helper Algorithm.
```

Helpers. *added receives a vector of proposals **P** and returns the ids and welcome public keys of newly added members. *new-spks receives an epoch id and vector of proposals, (epid, **P**), and returns the public verification keys for the ids affected by the proposals in **P**. *roster-pos returns the position of ID in G, and *get-propInfo receives a proposal P' and returns the proposal information, namely the operation, op, the proposal origin, orig, and the data, data, where if op = add, data holds the id, ID_a, welcome key material, wpk, and the signature verification key spk, of the newly added member. If op = rem, data holds the id of the removed party ID_r, and if op = upd, data holds the updated verification key spk.