


# Exploring How to Authenticate Application Messages in MLS: More Efficient, Post-Quantum, and Anonymous Blocklistable

Keitaro Hashimoto   
National Institute of Advanced  
Industrial Science and Technology  
(AIST)

Shuichi Katsumata   
AIST & PQShield

Guillermo Pascual-Perez   
Institute of Science and Technology Austria (ISTA)

## Abstract

The Message Layer Security (MLS) protocol has recently been standardized by the IETF. MLS is a scalable secure group messaging protocol expected to run more efficiently compared to the Signal protocol at scale, while offering a similar level of strong security. Even though MLS has undergone extensive examination by researchers, the majority of the works have focused on *confidentiality*.

In this work, we focus on the *authenticity* of the *application messages* exchanged in MLS. Currently, MLS authenticates every application message with an EdDSA signature and while manageable, the overhead is greatly amplified in the post-quantum setting as the NIST-recommended Dilithium signature results in a 40x increase in size. We view this as an invitation to explore new authentication modes that can be used instead. We start by taking a systematic view on how application messages are authenticated in MLS and categorize authenticity into four different security notions. We then propose several authentication modes, offering a range of different efficiency and security profiles. For instance, in one of our modes, COSMOS<sup>++</sup>, we replace signatures with one-time tokens and a MAC tag, offering roughly a 75x savings in the post-quantum communication overhead. While this comes at the cost of weakening security compared to the authentication mode used by MLS, the lower communication overhead seems to make it a worthwhile trade-off with security.

## 1 Introduction

### 1.1 Background

A *secure group messaging* (SGM) protocol allows a group of users to asynchronously communicate in an end-to-end encrypted fashion. The *Messaging Layer Security* (MLS) protocol [12, 20], a recently standardized SGM protocol by the IETF, is a proposal developed in a joint effort by academics and industry for a *scalable* SGM protocol supporting groups with tens of thousands of users. Similarly to the Signal protocol [36, 37, 46], considered the gold standard for two-user

SGMs, it offers a strong level of forward secrecy and post-compromise security, limiting the scope of device compromise. The draft versions of MLS are already running in production in Cisco’s Webex [65] and RingCentral [74], and other companies, including AWS, Cloudflare, and Google, are planning deployment.<sup>1</sup> Furthermore, with the recent adoption of the Digital Markets Act by the European Union, a standard like MLS is hoped to be a potential solution for the interoperability problem in secure messaging [59].

The security of MLS (and its variants) has undergone extensive examination by researchers during the standardization process, e.g., [4–8, 22, 29, 52, 53, 57, 78], and the protocol has been continuously updated leading up to 20 drafts in total<sup>2</sup> until the issuance of the RFC. The majority of works on MLS have focused on the *confidentiality* of the exchanged messages (or the shared group secret key). In contrast, relatively less attention has been directed towards the *authenticity* of messages, which is often viewed as a means to establish confidentiality.

In MLS, there are two types of messages being authenticated [12, Sec. 2]: *application* and *handshake* messages. While the former carry the actual payloads such as chat texts, the latter carry group operations affecting the group state (e.g., authenticating that user  $u$  added a new user  $v$  to the group). In this work, we revisit how MLS authenticates application messages motivated by the following two issues.

**Issue 1: Heavy Reliance on Signatures.** In MLS, every user  $u$  has a signature key pair  $(vk_u, sk_u)$  and signs the application message  $am$  for authentication. It further independently encrypts the message  $am$  and signature  $sig_u$  using a symmetric key encryption scheme whose key is derived from the group secret key to conceal the application message and its identity from the delivery server. The resulting (tuple of) ciphertext  $ct_u$  is then sent to the group. We call this mode of authentication *Enc-Sign mode*.<sup>3</sup> The recent work by Hashimoto et

<sup>1</sup><https://www.ietf.org/blog/mls-protocol-published/>

<sup>2</sup><https://datatracker.ietf.org/doc/rfc9420/>

<sup>3</sup>In contrast, handshake messages can be sent in *Sign mode*, where the user simply sends the pair  $(m, sig_u)$ .

al. [53] proposed adding an additional signature  $\overline{\text{sig}}_G$  on top of  $\text{ct}_u$ , using a signing key derived from the group secret key. This mode, called the *Sign-Enc-Sign mode*, is a simple but powerful enhancement of the Enc-Sign mode, allowing to anonymously block outsiders from injecting malicious messages to the group, similarly to Signal’s two-user Sealed Senders [60].

While adding signatures provides stronger authenticity guarantees, it comes with an increase in the communication and computational costs. This is currently manageable as MLS uses an EdDSA signature with an overhead of 64 B. However, this overhead is greatly amplified in the post-quantum setting. For instance, the NIST-recommended Dilithium signature is 2.4 KB, a 40x increase to EdDSA signatures. Given that a typical application message contains less than 100 B [47], the overhead has a noticeable effect. We thus view this as an invitation to explore alternative designs. We note that while handshake messages incur the same overhead when turning to post-quantum security, the effect is marginal as the size of the handshake message is larger, and the rate at which group operations are performed is less frequent compared to sending application messages.

**Issue 2: Lack of Formal Model for Authentication.** Compared to the comprehensive study of the confidentiality guarantees of MLS, authentication has drawn less attention. This lack of focus on authenticity may lead to unforeseen attacks on MLS that do not contradict confidentiality but still harm the protocol. As an illustrative example, the MLS is prone to abuse from malicious insiders (e.g., [8]). Notice that both Enc-Sign and Sign-Enc-Sign modes conceal the sender from the server. This allows a malicious insider to craft a malformed message and send it to the group. If the signature  $\text{sig}_u$  included in the ciphertext is malformed, even the group users cannot trace back the sender, meaning that a malicious sender can stealthily repeat the attack. While the users can reject these malformed messages, this can only happen *after* downloading them from the server and processing them. This opens the door for a malicious insider to mount a DoS attack on the group. A similar issue was pointed out by Tyagi et al. [75] for Signal’s two-user Sealed Senders [60], who experimentally verified that such an attack can easily drain a recipient’s battery in a short period of time.

A formal security model that comprehensively captures these properties allows us to better understand the strengths and limitations of a given authentication mode.

## 1.2 Our Contributions

In this work, we explore new approaches to authenticate application messages in MLS. Our contribution is explained below in more detail and an overview is provided in Tab. 1.

**Formal Model for Authentication.** In Sec. 2, we study how application messages are authenticated in MLS and systematically analyze the types of adversaries and threat models

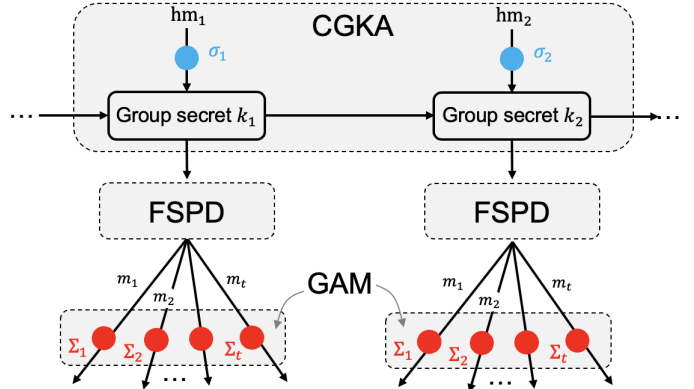


Figure 1: Relation between a CGKA, FSPD, and GAM protocols.  $hm$  denotes the handshake message used by the CGKA protocol.  $m_i$  denotes the output of the FSPD protocol; in MLS this is an *encryption* of the application message  $am$ . The blue and red circles indicate that the handshake and application messages are authenticated.

needed to be considered. More technically, the core of MLS can be regarded as a combination of two protocols: a *continuous group key agreement* (CGKA) and a *forward-secure payload delivery* (FSPD) protocol [5].<sup>4</sup> The former (resp. latter) handles handshake (resp. application) messages. In this paper, we formalize the authentication guarantees of the application messages handled by the FSPD protocol and introduce four different security notions: unforgeability, anonymity, anonymous blocklisting, and tracing soundness. To the best of our knowledge, this is the first work to put a focus on the authenticity of the application message; previous works on MLS studied the different types of CGKA protocol and focused on the confidentiality of the application message [4–8, 22, 29, 52, 53, 57, 78].

**Group Authenticated Messaging Protocol.** In Sec. 3, we propose the new notion of *group authenticated messaging* (GAM) protocol, allowing us to focus solely on the authenticity of the application messages while abstracting the confidentiality guarantees. More specifically, the FSPD protocol already entails the confidentiality of application messages and our GAM protocol can be viewed as adding authenticity guarantees to them. Fig. 1 gives an illustration on how a GAM protocol interacts with the CGKA and FSPD protocols. For instance, in MLS,  $m$  and  $\Sigma_1$  are the encryptions of the application message  $am$  and signature  $\text{sig}_u$  on  $am$ , respectively (i.e., Enc-Sign mode).

**New Authentication Modes.** In Secs. 4 and 5, we introduce five new GAM protocols: COSMOS, COSMAC, QUASAR, STARS, and GEMSTARS. All are based on generic building blocks such as one-way functions (OWFs), message authentication codes (MACs), and key encapsulation mechanisms (KEMs) that are instantiable from both classical and post-quantum assump-

<sup>4</sup>Alwen et al. [5] uses the term forward-secure group AEAD instead of FSPD.

Table 1: Comparison between different authentication modes for secure messaging protocols.  $N$  and  $T$  denote the size of the group and the number of messages each user sends. “Communication Cost Overhead per Msg per User” is defined as the sum of 1 (offline/online) upload cost and  $(N - 1)$  (offline/online) download cost for each user normalized by  $NT$ . For readability, we use the simplification  $(N + 1)/N \approx 1$ . sig, osig, and gsig denote a standard signature, a one-time signature, and a group signature, respectively. ovk denotes the verification key of a one-time signature. ct denotes a KEM ciphertext.  $\kappa$  denotes the security parameter, set to 128 bits.  $\checkmark^{(*)}$  denotes that it satisfies a weaker notion of unforgeability compared to  $\checkmark$  (see Sec. 2.3). “State Updates” comes with “-”, “local”, and “global”, where “-” means no state update is necessary (see Remark 3.2). COSMOS and COSMAC come with an optimized variants indicated by (+) and (++), whose respective total communication cost overheads and state updates are provided in parentheses.

Authentication Modes	Anon.	Unf.	Anonymous Blocklistable	Tracing Soundness	Comm. Cost Overhead per Msg per User	State Updates
Enc-Sign [12]	$\checkmark$	$\checkmark$	$\times$	$\times$	$ \text{sig} $	-
Sign-Enc-Sign [53]	$\checkmark$	$\checkmark$	$\checkmark$	$\times$	$2 \cdot  \text{sig} $	-
COSMOS(+,++) (Secs. 4)	$\times$	$\checkmark^{(*)}$	$\checkmark$	$\checkmark$	$3 \cdot \kappa (3 \cdot \kappa, (2 + \frac{3}{T}) \cdot \kappa)$	local (-, local)
COSMAC(+,++) (Secs. 4)	$\checkmark$	$\checkmark^{(*)}$	$\checkmark$	$\times$	$4 \cdot \kappa (4 \cdot \kappa, (3 + \frac{4}{T}) \cdot \kappa)$	local (-, local)
QUASAR (Sec. 5.1)	$\checkmark$	$\checkmark^{(*)}$	$\checkmark$	$\checkmark$	$6 \cdot \kappa + \frac{2 \cdot (\kappa +  \text{ct} )}{T}$	global
STARS (Sec. 5.2)	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$ \text{ovk}  + 2 \cdot  \text{osig}  + \frac{\kappa + 2 \cdot  \text{ct} }{T}$	global
GEMSTARS (Sec. 5.2)	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$ \text{sig}  +  \text{gsig} $	-

tions. Each mode fills a specific part of the design space with strengths and weaknesses, summarized in Tab. 1. In particular, COSMOS and COSMAC do not rely on signatures and the overhead (for their optimized variants) is merely 32 B and 48 B, respectively. This offers a roughly 75x savings in the post-quantum communication overhead compared to MLS, though at the cost of slightly weakening the unforgeability guarantee; we assume the malicious server does not collude with the malicious insider. See Sec. 2.3 for more detail. We believe this significantly lower communication overhead makes it a worthwhile trade-off with security.

**Efficiency Analysis.** In Sec. 7, we instantiate our proposed GAM protocols from both classical and post-quantum assumptions and compare their efficiency. For completeness, we also detail in Sec. 6 how to use each of our proposed GAM protocols inside MLS. While it is mostly a simple drop in, there are minor issues that require some explanation, since the syntax of GAM protocols intentionally leaves out some functionality provided by MLS, such as what is typically captured by the CGKA protocol (e.g., welcoming group members).

Lastly, we leave it as an important future work to analyze MLS in its entirety when using our notion of GAM protocol as a building block. While Alwen et al. [5] analyze MLS by composing the CGKA and FSPD protocols with a PRF-PRNG, the output of the FSPD protocol is explicitly signed (and not encrypted); that is, they assume the vanilla (unencrypted) GAM protocol used by MLS. Replacing this with a general GAM protocol and analyzing MLS is an interesting future work. We discuss further open problems in Sec. 8.

Other related work and preliminaries are deferred to Apps. A and B, respectively.

## 2 Setting: Authentication in SGM

This work focuses on secure group messaging (SGM) protocols, where group users share a unique common group secret key. In this section we use MLS as our primary example, but all of our constructions apply equally to most MLS variants (e.g., [2, 4, 6, 7, 52, 53, 57]) that rely on a group secret key to exchange messages.

Below, we give a brief background on how MLS authenticates application messages. We then take a close look at different security notions under the umbrella of authenticity and formally categorize them. Building on the systematization provided in this section, we introduce the concept of a *group authenticated messaging* (GAM) protocol in Sec. 3 and formally define the relevant security notions.

### 2.1 Secure Group Messaging and Our Goal

Following Alwen et al. [5], we view MLS as a combination of the CGKA and FSPD protocols (see Fig. 1 for illustration). At a high level, one can draw a parallel to hybrid encryption, where the heavy public key operations are handled by the CGKA protocol and the exchange of application messages is handled by the lightweight FSPD protocol.

In more detail, the CGKA protocol allows a group of users to agree on a continuous sequence of shared (symmetric) group secret keys. By regularly updating the group secret key (and user specific keys), strong notions of *forward secrecy* and *post-compromise security* [4–6, 8, 39] are guaranteed. The protocol is also responsible for handling group operations such as adding and removing users. *Handshake messages* is an umbrella term for the exchanged messages by the CGKA protocol, used to achieve the above objectives. A handshake

message, or an encryption of it, is signed using the user’s signing key to authenticate the sender. This plays an important role in guaranteeing the consistency and integrity of the group state. As can also be seen from Fig. 1, the authenticity of handshake messages is analyzed implicitly as a means to show confidentiality of the group secret keys; this is similar to standard (two-user) authenticated key exchange protocols where authenticity guarantees are implicit [15, 32].

The FSPD protocol then uses the established group secret key by the CGKA protocol to securely exchange *application messages*, containing various types of payload such as chat texts, images, and stamps. Compared to the CGKA protocol, the FSPD protocol is much simpler since there are no group operations (i.e., static groups) and the objective is only confidentiality with forward secrecy; authenticity is not a security requirement. In MLS, the output of the FSPD protocol — an encryption of the application message — is then signed using a signature scheme and encrypted (i.e., Enc-Sign mode), adding the necessary authenticity guarantee. This rather ad hoc way of adding authenticity seems to be justified by the simple nature of the FSPD protocol, and indeed most works on MLS mainly focus on the security of the CGKA protocol [4, 6–8, 22, 29, 52, 53, 57, 78]. To the best of our knowledge, Alwen et al. [5] is the only prior work to analyze MLS in its entirety. They do so by modularly combining the CGKA and FSPD protocols with a PRF-PRNG, assuming the output of the FSPD protocol is authenticated by a digital signature.

The goal of our paper is thus to put a spotlight onto the authentication of application messages or, to be more precise, the output of the FSPD protocol (see Fig. 1). We introduce a new primitive called group authenticated messaging (GAM) protocol and aim to more clearly and systematically explore alternative choices to the currently (implicitly) used GAM protocol by MLS, which is the Enc-Sign mode.

## 2.2 Environment

We first explain the environment in which MLS operates in. This involves introducing the relevant entities and outlining the network model under consideration.

### 2.2.1 Entities

*Group Users:* The set of users in a group. Depending on the considered security notion, the users are modeled to either be all honest or some malicious. For instance, we consider the latter case when modeling a security notion where a malicious insider (e.g., [8]) tries to impersonate an honest user.

*Server:* Any *asynchronous* messaging protocol requires a server to curate the messages between the group users. We consider two types of servers: honest and malicious. While servers are typically considered to be malicious by default in prior work, this is because the focus is mainly on the confidentiality of the CGKA protocol. For authentication, it makes

sense to consider honest servers as well. For example, the recent work by Hashimoto et al. [53] considers an honest server to anonymously block outsiders from injecting malicious messages to the group.

*Outsiders:* Any adversary that is not a group user or the server. For instance, a user of the secure messaging application not in the group.

### 2.2.2 Network Model

Due to asynchronicity, when group users exchange messages, they must upload and download these to and from the server. Depending on the anonymity guarantee we aim to achieve, there are two types of communication channels that can be used between the group users and the server.

*Non-Anonymous:* If the server is allowed to know the users in the group, then we assume a user-server authenticated channel is used. For instance, TLS or Noise [69] with user-side password-based authentication can be used.

*Anonymous:* If the group users are required to remain anonymous to the server, then we assume an authenticated anonymous channel such as TOR [42, 71] or a VPN is used.

We note dealing with authentication in the non-anonymous setting is trivial since the server can simply maintain the group list and explicitly authenticate the group users. In contrast, in the anonymous setting, such trivial solutions no longer exist and the issue of authentication becomes non-trivial. Indeed, prior works on *anonymous* secure messaging, e.g., [34, 53, 60, 75], overcome this by relying on some type of anonymous group authentication protocol.

## 2.3 Threat Model for Authentication

We now categorize authenticity into four different security notions: unforgeability, anonymity, anonymous blocklisting, and tracing soundness. This categorization of the application message is motivated by the security definitions used in well-studied anonymous authentication schemes, such as group signatures [13, 27, 35] and accountable ring signatures [18, 80].

Below, for each security notion, we explain *who* the adversary is, *what* the goal of the security notion is, and *why* we consider it. For simplicity, we leave outsider adversaries out of most security notions as they are strictly weaker than malicious servers and group users. The following security notions will be formalized in Sec. 3.3.

### Goal 1: Unforgeability.

**Adversary:** Malicious group users and/or a malicious server.

**Goal:** No adversary can forge a signature<sup>5</sup> of an honest user.

<sup>5</sup>Throughout this section, we use the term “signature” loosely and note that signatures are not the only way to authenticate. Using the terminology of our GAM protocol, this is more formally an “authentication token”.

This is the default notion that any secure messaging protocol must ensure. We can consider two levels of unforgeability: we call it *unforgeable* if the set of malicious group users and the malicious server can collude, and *non-colluding unforgeable* otherwise. The former guarantees that even a colluding malicious insider and server cannot forge a signature of an honest group user. In contrast, the latter restricts the adversary to be either the set of malicious group users or the malicious server; that is, unforgeability holds only if there is no collusion. While (standard) unforgeability is the more secure notion, sacrificing security against collusion of a malicious insider and server could be a reasonable compromise for better efficiency.

**Goal 2: Anonymity.**

**Adversary:** A malicious server.

**Goal:** The server cannot deanonymize and link the activity of the group users. E.g., the server cannot distinguish whether two uploaded messages came from the same user or from two different users.

For this security notion we must rely on an anonymous network model as, otherwise, communication will be linkable at the network level. We further assume all the group users to be honest, since a malicious user can always inform the server of who is in the group or who authored a message.

**Goal 3: Anonymous Blocklisting.**

**Adversary:** An outsider.

**Goal:** An honest server can block any outsider trying to upload messages on behalf of the group.

Observe that non-anonymous blocklisting is trivial to satisfy, since the server can perform access control by explicitly authenticating the group users. We therefore use the term “anonymous” to emphasize that the motivation of the server is to blocklist non-group users while preserving the anonymity of the users. The purpose of anonymous blocklisting is for the server to be able to prevent outsiders from launching a DoS attack on the group. Importantly, although group users can verify the authenticity of the messages by downloading them from the server, we require the server to directly reject invalid messages on behalf of the group. This is satisfied for example by Sealed Sender [60] used in the Signal protocol and the metadata-hiding MLS protocol by Hashimoto et al. [53].

**Goal 4: Tracing Soundness.**

**Adversary:** Malicious group users.

**Goal:** The set of honest group users can trace any (possibly maliciously crafted) signature back to a *unique* group user; if an honest user traces a signature back to a user  $u$  in the group, then all other honest users trace it back to the same user  $u$ .

Tracing soundness allows to keep the view of the honest users consistent. For instance, consider a malicious insider mounting a DoS attack against the group by spamming garbage application messages. With tracing soundness, the honest users can unanimously agree on who the malicious insider was and remove him from the group. One can draw a

parallel to anonymous blocklisting, that prevents such attacks from outsiders. Moreover, while similar, it is worth noting that tracing soundness is an orthogonal notion to unforgeability. Consider a malicious insider  $u$  that modifies the signature of an honest user  $v$  in such a way that for half of the group members it traces back to  $u$ , but for the other half traces back to  $v$ . While this does not contradict unforgeability, as the malicious user is effectively just “repurposing” somebody’s message, it clearly breaks the consistency of the group’s view.

## 2.4 Modeling Choices and Simplifications

Before introducing our GAM protocol in the next section, we clarify the modeling choices and simplifications we make.

*Trusted Setup.* The GAM protocol assumes the states of both the group users and the server are generated honestly by an initialization phase. This simplification is justified for protocols like MLS, since users are assumed to start the GAM protocol with the group secret key, derived from the CGKA protocol, already in their states.

*Static Groups.* The GAM protocol assumes static groups, following the way in which MLS’ FSPD protocol operates. Recall that in MLS a new FSPD protocol for a static group is initialized every time group membership changes, as this will trigger a new CGKA protocol epoch (see Fig. 1). More generally, though, we could consider a *continuous* GAM protocol where we do not need to reinitialize the protocol with every group change, similarly to a CGKA protocol. However, such a definition must be intertwined with that of the CGKA protocol responsible for group state updates, rendering the definition to be as complex as modeling MLS in its entirety. As a study investigating new security goals of authentication, we opt for making the security notions tractable and to improve the overall readability. Nonetheless, we explain in Sec. 6.2 with concrete examples on how each of our proposed GAM protocols can handle dynamic operations.

*Out-of-Order Messages.* In our work, we do not model authentication when messages arrive out-of-order. While this is arguably important for a comprehensive model, we highlight that, unlike confidentiality, lack of authentication does not harm the usability of the FSPD protocol. In the context of the MLS protocol, *immediate decryption* of the messages will still be maintained. The only difference between MLS is that we may lose *immediate authentication* when messages arrive out-of-order. Importantly, though security is lost while some messages are missing, assuming that every message eventually arrives, then out-of-order messages do not affect security. Instead, if some messages are permanently dropped, we can allow the recipients to fetch this missing authentication information, which they can do assuming the proper indexing of the messages required by out-of-order decryption. We note that in MLS [20, Section 5.2], whether messages

eventually arrive or not is controlled by the application that sets the policy.

### 3 Group Authenticated Messaging Protocol

We introduce *group authenticated messaging* (GAM) protocols and the associated security requirements.

#### 3.1 Definition

A GAM protocol is defined between a server and a group  $G$  of users. As explained above, there exists an initialization algorithm  $\text{Init}$  that prepares the initial state for the group users, possibly further preparing a secret key for the server. To send a message  $m$  (e.g., the output of a FSPD protocol), a user  $u \in G$  runs the  $\text{Send}$  algorithm, outputting a *group authentication token*  $\Sigma_G$ . A server verifies  $(m, \Sigma_G)$  using the  $\text{Verify}$  algorithm and prepares *user authentication tokens*  $(\sigma_i)_{i \in [N]}$ , where  $N = |G|$ . For example, in the context of the Sign mode in MLS (see Footnote 3),  $\Sigma_G$  is simply  $u$ 's signature and  $\sigma_i := \Sigma_G$ . To capture anonymity, we assume the server only knows the size of the group  $G^6$  and assume a bijective map  $\text{id}_x : G \rightarrow [N]$  is secretly known by the group users. Namely, a user  $u$  such that  $i = \text{id}_x(u)$  fetches  $\sigma_i$  from the server. It then runs the  $\text{Receive}$  algorithm to verify  $(m, \sigma_i)$  and traces the purported user  $v \in G$  that generated  $\sigma_i$ . It is worth highlighting that we make a distinction between a group authentication token  $\Sigma_G$  and a user authentication token  $\sigma_i$  to capture an optimization technique called *selective downloading* [7, 52]. This technique allows the server to sanitize the group authentication token  $\Sigma_G$  in a straightforward manner by delivering to each group user just the strictly necessary amount of data  $\sigma_i$ , while maintaining the same level of (dis)trust.

Finally, we endow a GAM protocol with an *offline-online* feature. In the offline phase, when the message is still unknown, a user can perform a possibly heavy state update, and share the update with the server and the group via the  $\text{UpdSend}$  algorithm. This algorithm is accompanied by algorithms  $\text{UpdVerify}$  and  $\text{UpdReceive}$  similarly to above. Once the message is known in the online phase, the user can send it using its updated state.<sup>7</sup> Formally, we have the following.

**Definition 3.1.** A GAM protocol for message space  $\mathcal{M}$  between a server  $S_v$  and a set of users in a group  $G$  consists of the following algorithms, where  $\text{id}_x : G \rightarrow [N]$  is a bijective function with  $N := |G|$ . Below, if an algorithm outputs  $\perp$ , we assume it reverts to the state before running the algorithm.

<sup>6</sup>While we could consider further hiding the size of the group to the server, we choose not to since it would resort in an inefficient padding strategy. This is the same level of anonymity satisfied by previous anonymous SGM protocols e.g., [34, 53].

<sup>7</sup>Naturally, protocols need not have such a differentiation and can simply only perform online state updates. This optimization allows us to improve the real-world usability of those protocols that do, as they can more evenly distribute their computation and communication over time.

$\text{Init}(1^k, G) \rightarrow (\text{pp}, \text{sk}_{S_v}, (\text{st}_u)_{u \in G})$  : On input the security parameter  $1^k$  and group information  $G \subset \{0, 1\}^*$ , it outputs public parameters  $\text{pp}$ , a secret key  $\text{sk}_{S_v}$  for the server  $S_v$ , and an initial state  $\text{st}_u$  for all users  $u \in G$ . We assume  $G \in \text{st}_u$ .

$\text{Send}(\text{st}_u, m) \rightarrow (\text{st}'_u, \Sigma_G)$  or  $\perp$  : On input a state  $\text{st}_u$  for user  $u \in G$  and a message  $m \in \mathcal{M}$ , user  $u$  outputs an updated state  $\text{st}'_u$  and a group authentication token  $\Sigma_G$ , or  $\perp$ .

$\text{Verify}(\text{pp}, \text{sk}_{S_v}, \Sigma_G, m) \rightarrow (\text{pp}', (\sigma_i)_{i \in [N]})$  or  $\perp$  : On input public parameters  $\text{pp}$ , a server secret key  $\text{sk}_{S_v}$ , a group authentication token  $\Sigma_G$ , and a message  $m \in \mathcal{M}$ , the server  $S_v$  outputs updated public parameters  $\text{pp}'$  and  $N$  user authentication tokens  $(\sigma_i)_{i \in [N]}$ , or  $\perp$ .

$\text{Receive}(\text{st}_u, \sigma, m) \rightarrow (\text{st}'_u, b \in \{\top, \perp\}, v \in G \cup \{\perp\})$  : On input a state  $\text{st}_u$  for user  $u \in G$ , a user authentication token  $\sigma$ , and a message  $m \in \mathcal{M}$ , user  $u$  outputs an updated state  $\text{st}'_u$ , a bit  $b$  indicating whether the token was valid ( $b = \top$ ) or invalid ( $b = \perp$ ), and a purported user  $v \in G \cup \{\perp\}$ , where  $v = \perp$  if tracing fails.

$\text{UpdSend}(\text{st}_u) \rightarrow (\text{st}'_u, \widehat{\Sigma}_G, \widehat{\text{ct}}_G)$  : On input a state  $\text{st}_u$  for user  $u \in G$ , user  $u$  outputs an updated state  $\text{st}'_u$ , a group update authentication token  $\widehat{\Sigma}_G$ , and group update information  $\widehat{\text{ct}}_G$ .

$\text{UpdVerify}(\text{pp}, \text{sk}_{S_v}, \widehat{\Sigma}_G, \widehat{\text{ct}}_G) \rightarrow (\text{pp}', (\widehat{\sigma}_i, \widehat{\text{ct}}_i)_{i \in [N]})$  or  $\perp$  : On input public parameters  $\text{pp}$ , a server secret key  $\text{sk}_{S_v}$ , a group update authentication token  $\widehat{\Sigma}_G$ , and group update information  $\widehat{\text{ct}}_G$ , the server  $S_v$  outputs updated public parameters  $\text{pp}'$  and a list of user update authentication tokens and user update information  $(\widehat{\sigma}_i, \widehat{\text{ct}}_i)_{i \in [N]}$ , or  $\perp$ .

$\text{UpdReceive}(\text{st}_u, \widehat{\sigma}, \widehat{\text{ct}}) \rightarrow (\text{st}'_u, b \in \{\top, \perp\}, v \in G \cup \{\perp\})$  : On input a state  $\text{st}_u$  for user  $u \in G$ , a user update authentication token  $\widehat{\sigma}$ , and user update information  $\widehat{\text{ct}}$ , user  $u$  outputs an updated state  $\text{st}'_u$ , a bit  $b$  indicating whether the token was valid ( $b = \top$ ) or invalid ( $b = \perp$ ), and a purported user  $v \in G \cup \{\perp\}$ , where  $v = \perp$  if tracing fails.

*Remark 3.2 (Local and Global State Updates).* For some protocols the user state may only allow signing up to  $T$  messages, and it may need to be updated before the user can sign again. There are two ways to perform state updates: *locally* and *globally*. In the former, a user regains the ability to send messages once it has updated its own state. In the latter, a user regains the ability to send messages only after every user in the group updates their states. Since global state updates are much more costly than local state updates, they are only useful if one state update allows to send a large number of messages  $T$ . Further, global updates can only guarantee security if users are online, a clear disadvantage over local updates. For the schemes presented in this paper there are no risks of a deadlock — i.e., a situation where a global state update cannot be completed and users are prevented to keep sending messages — as long as the users perform updates once coming online. However,

the general definition of global updates does not guarantee that such a deadlock does not occur.

### 3.2 Correctness

We define two types of signing correctness. One for signing messages and the other for signing updates. They stipulate that an honestly generated user authentication token is always valid and traceable.

**Definition 3.3 (Signing Correctness).** For any  $\kappa \in \mathbb{N}$ ,  $G \subset \{0, 1\}^*$ , and  $(pp, sk_{Sv}, (st_u)_{u \in G}) \in \text{Init}(1^\kappa, G)$ , if we execute  $(\text{Send}, \text{Verify}, \text{Receive}, \text{UpdSend}, \text{UpdVerify}, \text{UpdReceive})$  in an arbitrary but honest manner (i.e., we only run the algorithms on inputs that were output by another algorithm and run  $\text{Receive}$  (resp.  $\text{UpdReceive}$ ) for all users after running  $\text{Verify}$  (resp.  $\text{UpdVerify}$ )<sup>8</sup>, then we have the following, where  $pp'$  and  $(st'_u)_{u \in G}$  are arbitrary public parameters and states reachable from the initial  $pp$  and  $(st_u)_{u \in G}$ :

**Message Signing Correctness:** For any  $m \in \mathcal{M}$  and  $u \in G$ , if we execute  $(st''_u, \Sigma_G) \leftarrow \text{\$Send}(st'_u, m)$ , redefine  $st''_u := st''_u$ , and execute  $(pp'', (\sigma_i)_{i \in [N]}) \leftarrow \text{Verify}(pp', sk_{Sv}, \Sigma_G, m)$ ,  $(st''_v, b_v, u_v) \leftarrow \text{Receive}(st'_v, \sigma_{\text{idx}(v)}, m)$  for all  $v \in G$ , then conditioned on  $\Sigma_G \neq \perp$ , we have  $(b_v, u_v) = (T, u)$  (i.e., every user accepts and traces the message back to  $u$ ).

**Update Signing Correctness:** For any  $u \in G$ , if we execute  $(st''_u, \widehat{\Sigma}_G, \widehat{ct}_G) \leftarrow \text{\$UpdSend}(st'_u)$ , redefine  $st''_u := st''_u$ , and execute  $(pp'', (\widehat{\sigma}_i, \widehat{ct}_i)_{i \in [N]}) \leftarrow \text{UpdVerify}(pp', sk_{Sv}, \widehat{\Sigma}_G, \widehat{ct}_G)$  followed by  $(st''_v, b_v, u_v) \leftarrow \text{UpdReceive}(st'_v, \widehat{\sigma}_{\text{idx}(v)}, \widehat{ct}_{\text{idx}(v)})$  for all  $v \in G$ , then conditioned on  $\widehat{\Sigma}_G \neq \perp$ , we have  $(b_v, u_v) = (T, u)$  (i.e., every user accepts and traces the update back to  $u$ ).

In some protocols, the states may occasionally need to be updated in order to regain the ability to send messages again. As explained in Remark 3.2, there are two ways a users can update their states. One is *local*, where it is sufficient that a user can simply update its state. The other one is *global*, where all the group users must update their states. Below, we define correctness of both state update modes.

**Definition 3.4 (State-Update Correctness).** Assume the same precondition as in Def. 3.3. Then, we have either of the following:

**Local State-Update Correctness:** For any  $m \in \mathcal{M}$  and user  $u \in G$ , if  $(st''_u, \Sigma_G) \leftarrow \text{\$Send}(st'_u, m)$  such that  $\Sigma_G = \perp$ , then if  $u$  executes  $(st''_u, \widehat{\Sigma}_G, \widehat{ct}_G) \leftarrow \text{\$UpdSend}(st'_u)$ , the server  $S_v$  executes  $(pp'', (\widehat{\sigma}_i, \widehat{ct}_i)_{i \in [N]}) \leftarrow \text{UpdVerify}(pp', sk_{Sv}, (\widehat{\Sigma}_G, \widehat{ct}_G))$ , and every user  $v \in G$

executes  $\text{UpdReceive}(st'_v, (\widehat{\sigma}_{\text{idx}(v)}, \widehat{ct}_{\text{idx}(v)}))$ , then the updated public parameters  $pp''$  and state  $st''_u$  allow user  $u$  to sign on  $m$ , that is,  $\Sigma_G \neq \perp$  and signing correctness holds (i.e., after user  $u$  sends a user update information  $\widehat{ct}_{\text{idx}(v)}$  to every user  $v$ , then user  $u$ 's state will be refreshed).

**State Update Correctness:** For any  $m \in \mathcal{M}$  and user  $u \in G$ , if  $(st''_u, \Sigma_G) \leftarrow \text{\$Send}(st'_u, m)$  such that  $\Sigma_G = \perp$ , then if all users  $v \in G$  execute  $(st''_v, \widehat{\Sigma}_G, \widehat{ct}_G) \leftarrow \text{\$UpdSend}(st'_v)$ , the server  $S_v$  executes  $\text{UpdVerify}$  for all  $(\widehat{\Sigma}_G, \widehat{ct}_G)_{v \in G}$  in an arbitrary order, and user  $u$  executes  $\text{UpdReceive}$  for all  $(\widehat{\sigma}_{\text{idx}(u)}, \widehat{ct}_{\text{idx}(u)})_{v \in G}$  output by  $\text{UpdVerify}$  in an arbitrary order, then the updated public parameters  $pp''$  and state  $st''_u$  allow user  $u$  to sign on  $m$ , that is,  $\Sigma_G \neq \perp$  and signing correctness holds (i.e., after every user  $v$  sends a user update information  $\widehat{ct}_{\text{idx}(u)}$  to user  $u$ , then user  $u$ 's state will be refreshed). Note that by symmetry, all users' state is refreshed.

### 3.3 Security

We formalize the threat models explained in Sec. 2.3: unforgeability, anonymity, anonymous blocklisting, and tracing soundness, via a security game defined in Fig. 2. The probability of the game outputting 1 against an efficient adversary must be negligible for every game except for anonymity. For anonymity, as it is a distinguishing game, the game must output 1 with probability negligibly close to  $\frac{1}{2}$ .

For every game, the adversary is given access to oracles  $\{O_{\text{Send}}, O_{\text{UpdSend}}\}$ , allowing it to invoke honest users to create group (update) authentication tokens. The adversary is further given access to either  $\{O_{\text{Receive}}, O_{\text{UpdReceive}}\}$  or  $\{O_{\text{GroupReceive}}, O_{\text{GroupUpdReceive}}\}$ . The former allows the adversary to directly invoke honest users to process (update) authentication tokens. This capture malicious server capabilities and is used by the unforgeability and anonymity games. In contrast, the latter only allows the adversary to query for *group* (update) authentication tokens. The oracle then individually invokes each honest users on the correctly processed (update) authentication tokens. Namely, this captures honest server behavior and models the fact that malicious users cannot directly send messages to group users. It is worth noting that, in this case, the authentication tokens created in  $\{O_{\text{Send}}, O_{\text{UpdSend}}\}$  are directly processed by  $\{O_{\text{GroupReceive}}, O_{\text{GroupUpdReceive}}\}$ , modeling the fact that the communication channel between an honest user and server is secure.

To aid readability, we highlight some features of the security game. We model two types of unforgeability by  $\text{Game}_{\mathcal{A}}^X$  with  $X \in \{\text{ncUnf}, \text{Unf}\}$ . In standard unforgeability, as the adversary models both a malicious user and server, it has unrestricted access to all oracles. In contrast, for non-colluding unforgeability, we have two case distinctions depending on whether the set of corrupted users  $\mathcal{C} = \emptyset$  or not.

<sup>8</sup>We impose the second condition to define correctness in a minimal yet well-defined manner. Without it, we must also include cases such as when only part of the users received a user update information.

In the former case the adversary is a malicious server, so the adversary is given the server secret key  $sk_{sv}$  and has access to  $\{O_{\text{Receive}}, O_{\text{UpdReceive}}\}$ . In the latter case the adversary is a set of malicious users, so the adversary is instead given the corrupted users' states and only has access to  $\{O_{\text{GroupReceive}}, O_{\text{GroupUpdReceive}}\}$ . For both types of unforgeability, an adversary wins if it can output a valid user (update) authentication token for an honest user that it has not seen before. For anonymity, we model a malicious server by giving the adversary the server secret key  $sk_{sv}$ . The adversary outputs two users and messages and the game creates the group authentication tokens for both users. To non-trivialize the game, we restrict the (group) authentication tokens to be valid. To perform this check, the adversary needs to further output a (possibly malformed) public parameter  $\overline{pp}$  so the game can run algorithm `Verify`. The adversary can further perform oracle queries under the restriction that it does not query the receive oracles on the challenge authentication tokens.

We now provide the formal definition of non-colluding and standard unforgeability, anonymity, anonymous blocklisting, and tracing soundness and some more intuitions on how to understand them.

**Unforgeability.** As already discussed above, we model standard and non-colluding unforgeability by giving the adversary access to different oracles and running it with different inputs (i.e., with or without server and user states). The adversary wins if it outputs a valid authentication token such that  $v_u \in \mathcal{H} \wedge (v_u, *, \bar{m}) \notin L_{\text{upd}}$  where  $\bar{m} \in \{m, \hat{ct}\}$  holds, where recall  $v_u$  is the traced user (cf. unforgeability game, lines 17 and 23). The former checks that the user  $v_u$  is not malicious; without this check a malicious user can trivially win unforgeability. The latter checks that the honest user  $v_u$  did not sign  $\bar{m}$ . Formally, we define unforgeability as follows.

**Definition 3.5 (Unforgeability).** We define  $\text{Game}_{\mathcal{A}}^{\text{ncUnf}}(1^\kappa)$  as in Fig. 2 for an adversary  $\mathcal{A}$ . We say a GAM protocol is no-colluding unforgeable if for any  $G \subset \{0, 1\}^*$ , injective function  $\text{idx} : G \rightarrow [N]$  with  $N = |G|$ , and any PPT adversary  $\mathcal{A}$ , we have

$$\text{Adv}_{\mathcal{A}}^{\text{ncUnf}}(1^\kappa) := \Pr[\text{Game}_{\mathcal{A}}^{\text{ncUnf}}(1^\kappa) = 1] = \text{negl}(\kappa).$$

We further say the scheme is (standard) unforgeable if the above holds for  $\text{Game}_{\mathcal{A}}^{\text{Unf}}(1^\kappa)$  as defined in Fig. 2.

**Anonymity.** As briefly explained above, the game checks if the group authentication token  $\Sigma_G$  and the individual authentication tokens  $(\sigma_i)_{i \in [N]}$  are valid. Without this check, an adversary may trivially break anonymity if the protocol requires state updates. Concretely, assume the adversary queries a user  $u_0$  to oracle  $O_{\text{Send}}$  until  $u_0$  can no longer sign without performing an update. At this point, if the adversary challenges user  $u_0$  and  $u_1$ , then it can trivially break anonymity as  $u_0$  cannot produce a group authentication token while  $u_1$  can.

Moreover, while we can easily define anonymity for updates, we chose not to do so as updates are sent far less often compared to messages, and we opted for simplicity of the security game. Formally, we define anonymity as follows.

**Definition 3.6 (Anonymity).** We define  $\text{Game}_{\mathcal{A}}^{\text{Anon}}(1^\kappa)$  as in Fig. 2 for an adversary  $\mathcal{A}$ . We say a GAM protocol is anonymous if for any  $G \subset \{0, 1\}^*$ , injective function  $\text{idx} : G \rightarrow [N]$  with  $N = |G|$ , and any PPT adversary  $\mathcal{A}$ , we have

$$\text{Adv}_{\mathcal{A}}^{\text{Anon}}(1^\kappa) := \left| \Pr[\text{Game}_{\mathcal{A}}^{\text{Anon}}(1^\kappa) = 1] - \frac{1}{2} \right| = \text{negl}(\kappa).$$

**Anonymous Blocklisting.** This game is quite intuitive as the adversary is an outsider. The adversary wins the game if it's able to output a valid group authentication token that nobody in the group created. Although similar, we note that anonymous blocklisting is different from unforgeability. To win anonymous blocklisting, the adversary is required to output a *group* authentication token  $\Sigma_G$  that verifies. This entails the fact that the server can check the validity of  $\Sigma_G$  and immediately block malformed group authentication tokens on behalf of the group users. In contrast, unforgeability does not capture this type of blocking by the server. Formally, we define anonymous blocklisting as follows.

**Definition 3.7 (Anonymous Blocklisting).** We define the security game  $\text{Game}_{\mathcal{A}}^{\text{AnonBlock}}(1^\kappa)$  as in Fig. 2 for an adversary  $\mathcal{A}$ . We say a GAM protocol is anonymous blocklistable if for any  $G \subset \{0, 1\}^*$ , injective function  $\text{idx} : G \rightarrow [N]$  with  $N = |G|$ , and any PPT adversary  $\mathcal{A}$ , we have

$$\text{Adv}_{\mathcal{A}}^{\text{AnonBlock}}(1^\kappa) := \Pr[\text{Game}_{\mathcal{A}}^{\text{AnonBlock}}(1^\kappa) = 1] = \text{negl}(\kappa).$$

**Tracing Soundness.** As discussed in Sec. 2.3, the adversary wins if it outputs a group authentication token for which the set  $L_{\text{tr}}$  of traced users by the honest users is not of the form  $L_{\text{tr}} = \{v\}$  for some group user  $v \in G$ . That is, if the group authentication token is valid, it must be traceable to some user in the group and this user must be unique among the honest users. Formally, we define tracing soundness as follows.

**Definition 3.8 (Tracing Soundness).** We define  $\text{Game}_{\mathcal{A}}^{\text{TraceSound}}(1^\kappa)$  as in Fig. 2 for an adversary  $\mathcal{A}$ . We say a GAM protocol is tracing sound if for any  $G \subset \{0, 1\}^*$ , injective function  $\text{idx} : G \rightarrow [N]$  with  $N = |G|$ , and any PPT adversary  $\mathcal{A}$ , we have

$$\text{Adv}_{\mathcal{A}}^{\text{TraceSound}}(1^\kappa) := \Pr[\text{Game}_{\mathcal{A}}^{\text{TraceSound}}(1^\kappa) = 1] = \text{negl}(\kappa).$$

*Remark 3.9 (Transparency of Server).* In any secure messaging protocol it may be important to have a transparent server, so as to limit the trust we put in it. In the context of a GAM protocol, notice that our current initialization algorithm `Init`



<p><b>Game<math>_{\mathcal{A}}^X(1^\kappa)</math> : <math>X \in \{\text{ncUnf}, \text{Unf}\}</math></b></p> <pre> 1: <math>C \leftarrow \mathcal{A}(1^\kappa)</math> 2: <math>\mathcal{H} := G \setminus C</math> 3: <math>L_{\text{msg}}, L_{\text{upd}} := \emptyset</math> / Book keeping 4: <math>(\text{pp}, \text{sk}_{Sv}, (\text{st}_u)_{u \in C}) \leftarrow \mathcal{A}^O(1^\kappa, G)</math> 5: <b>if</b> <math>X = \text{Unf}</math> <b>then</b> 6:   <math>(\text{label}, \text{obj}) \leftarrow \mathcal{A}^O(\text{pp}, \text{sk}_{Sv}, (\text{st}_u)_{u \in C})</math> 7: <b>else</b> / No collusion between malicious user and server 8:   <b>if</b> <math>C = \emptyset</math> / Honest users 9:     <math>(\text{label}, \text{obj}) \leftarrow \mathcal{A}^O(\text{pp}, \text{sk}_{Sv})</math> 10:  <b>else</b> / Honest server 11:    <math>(\text{label}, \text{obj}) \leftarrow \mathcal{A}^O(\text{pp}, (\text{st}_u)_{u \in C})</math> 12:  <b>if</b> <math>\text{label} = \text{msg}</math> <b>then</b> 13:    <b>parse</b> <math>(u, \sigma, m) \leftarrow \text{obj}</math> 14:    <b>req</b> <math>u \in \mathcal{H}</math> 15:    <math>(\text{st}'_u, b_v, v_u) \leftarrow \text{Receive}(\text{st}_u, \sigma, m)</math> 16:    <b>req</b> <math>b_v = \top</math> / Valid authentication token 17:    <math>b \leftarrow \llbracket v_u \in \mathcal{H} \wedge (v_u, *, m) \notin L_{\text{msg}} \rrbracket</math> 18:  <b>elseif</b> <math>\text{label} = \text{upd}</math> <b>then</b> 19:    <b>parse</b> <math>(u, \widehat{\sigma}, \widehat{\text{ct}}) \leftarrow \text{obj}</math> 20:    <b>req</b> <math>u \in \mathcal{H}</math> 21:    <math>(\text{st}'_u, b_v, v_u) \leftarrow \text{UpdReceive}(\text{st}_u, \widehat{\sigma}, \widehat{\text{ct}})</math> 22:    <b>req</b> <math>b_v = \top</math> / Valid authentication token 23:    <math>b \leftarrow \llbracket v_u \in \mathcal{H} \wedge (v_u, *, \widehat{\text{ct}}) \notin L_{\text{upd}} \rrbracket</math> 24:  <b>return</b> <math>b</math> </pre>	<p><b>Game<math>_{\mathcal{A}}^{\text{AnonBlock}}(1^\kappa)</math></b></p> <pre> 1: <math>\mathcal{H} := G</math> / No corrupt users 2: <math>L_{\text{msg}}, L_{\text{upd}} := \emptyset</math> / Book keeping 3: <math>(\text{pp}, \text{sk}_{Sv}, (\text{st}_u)_{u \in G}) \leftarrow \mathcal{A}^O(1^\kappa, G)</math> 4: <math>(\text{label}, \text{obj}) \leftarrow \mathcal{A}^O(\text{pp})</math> / Malicious outsiders 5: <b>if</b> <math>\text{label} = \text{msg}</math> <b>then</b> 6:   <b>parse</b> <math>(\Sigma_G, m) \leftarrow \text{obj}</math> 7:   <math>(\text{pp}', (\sigma_i)_{i \in [N]}) \leftarrow \text{Verify}(\text{pp}, \text{sk}_{Sv}, \Sigma_G, m)</math> 8:   <b>req</b> <math>\text{pp}' \neq \perp</math> / Require Verify to succeed 9:   / Sv accepts new non-member token 10:  <math>b \leftarrow \llbracket (*, \Sigma_G, *) \notin L_{\text{msg}} \rrbracket</math> 11:  <b>elseif</b> <math>\text{label} = \text{upd}</math> <b>then</b> 12:    <b>parse</b> <math>(\widehat{\Sigma}_G, \widehat{\text{ct}}_G) \leftarrow \text{obj}</math> 13:    <math>(\text{pp}', (\widehat{\sigma}_i, \widehat{\text{ct}}_i)_{i \in [N]})</math> 14:    <math>\leftarrow \text{UpdVerify}(\text{pp}, \text{sk}_{Sv}, \widehat{\Sigma}_G, \widehat{\text{ct}}_G)</math> 15:    <b>req</b> <math>\text{pp}' \neq \perp</math> / Require UpdVerify to succeed 16:    / Sv accepts new non-member token 17:    <math>b \leftarrow \llbracket (*, \widehat{\Sigma}_G, *) \notin L_{\text{upd}} \rrbracket</math> 18:  <b>return</b> <math>b</math> </pre> <p><b>Game<math>_{\mathcal{A}}^{\text{TraceSound}}(1^\kappa)</math></b></p> <pre> 1: <math>C \leftarrow \mathcal{A}(1^\kappa)</math> 2: <math>\mathcal{H} := G \setminus C</math> 3: <math>L_{\text{tr}} := \emptyset</math> / Book keeping 4: <math>(\text{pp}, \text{sk}_{Sv}, (\text{st}_u)_{u \in G}) \leftarrow \mathcal{A}^O(1^\kappa, G)</math> 5: <math>(\text{label}, \text{obj}) \leftarrow \mathcal{A}^O(\text{pp}, (\text{st}_u)_{u \in C})</math> 6: <b>if</b> <math>\text{label} = \text{msg}</math> <b>then</b> 7:   <b>parse</b> <math>(\Sigma_G, m) \leftarrow \text{obj}</math> 8:   <math>(\text{pp}', (\sigma_i)_{i \in [N]}) \leftarrow \text{Verify}(\text{pp}, \text{sk}_{Sv}, \Sigma_G, m)</math> 9:   <b>req</b> <math>\text{pp}' \neq \perp</math> / Require UpdVerify to succeed 10:  <b>foreach</b> <math>u \in \mathcal{H}</math> <b>do</b> 11:    <math>(\text{st}'_u, b_v, v_u) \leftarrow \text{Receive}(\text{st}_u, \sigma_{\text{idx}(u)}, m)</math> 12:    <b>if</b> <math>b_v = \top</math> <b>then</b> 13:      <math>L_{\text{tr}} \leftarrow L_{\text{tr}} \cup \{v_u\}</math> / If tracing fails, <math>v_u = \perp</math> 14:    <b>elseif</b> <math>\text{label} = \text{upd}</math> <b>then</b> 15:      <b>parse</b> <math>(\widehat{\Sigma}_G, \widehat{\text{ct}}_G) \leftarrow \text{obj}</math> 16:      <math>(\text{pp}', (\widehat{\sigma}_i, \widehat{\text{ct}}_i)_{i \in [N]})</math> 17:      <math>\leftarrow \text{UpdVerify}(\text{pp}, \text{sk}_{Sv}, \widehat{\Sigma}_G, \widehat{\text{ct}}_G)</math> 18:      <b>req</b> <math>\text{pp}' \neq \perp</math> / Require UpdVerify to succeed 19:      <b>foreach</b> <math>u \in \mathcal{H}</math> <b>do</b> 20:        <math>(\text{st}'_u, b_v, v_u) \leftarrow \text{UpdReceive}(\text{st}_u, \widehat{\sigma}_{\text{idx}(u)}, \widehat{\text{ct}}_{\text{idx}(u)})</math> 21:        <b>if</b> <math>b_v = \top</math> <b>then</b> 22:          <math>L_{\text{tr}} \leftarrow L_{\text{tr}} \cup \{v_u\}</math> / If tracing fails, <math>v_u = \perp</math> 23:        / Does not uniquely trace user 24:      <math>b \leftarrow \llbracket \exists v \in G : L_{\text{tr}} = \{v\} \rrbracket</math> 25:    <b>return</b> <math>b</math> </pre>	<p><b>Oracle <math>O_{\text{Send}}(u \in \mathcal{H}, m)</math></b></p> <pre> 1: <math>s_{\text{Rec}} := \emptyset</math> 2: <math>(\text{st}'_u, \Sigma_G) \leftarrow \mathcal{A}^O(\text{Send}(\text{st}_u, m))</math> 3: <math>L_{\text{msg}} \leftarrow L_{\text{msg}} \cup \{(u, \Sigma_G, m)\}</math> 4: / Server honestly processes group authentication 5: <b>if</b> <math>\mathcal{A}</math> has access to <math>O^*</math> <b>then</b> 6:   <math>s_{\text{Rec}} \leftarrow O_{\text{GroupReceive}}(\Sigma_G, m)</math> 7: <b>return</b> <math>(\Sigma_G, s_{\text{Rec}})</math> </pre> <p><b>Oracle <math>O_{\text{Receive}}(u \in \mathcal{H}, \sigma, m)</math></b></p> <pre> 1: <b>req</b> <math>\sigma \notin \text{Chall}_{\text{msg}}</math> / Only used by anonymity 2: <math>(\text{st}'_u, b_u, v_u) \leftarrow \text{Receive}(\text{st}_u, \sigma_{\text{idx}(u)}, m)</math> 3: <b>return</b> <math>(b_u, v_u)</math> </pre> <p><b>Oracle <math>O_{\text{GroupReceive}}(\Sigma_G, m)</math></b></p> <pre> 1: <math>(\text{pp}', (\sigma_i)_{i \in [N]}) \leftarrow \text{Verify}(\text{pp}, \text{sk}_{Sv}, \Sigma_G, m)</math> 2: <b>if</b> <math>\text{pp}' = \perp</math> <b>then return</b> <math>\perp</math> 3: <b>foreach</b> <math>u \in \mathcal{H}</math> <b>do</b> 4:   <math>(\text{st}'_u, b_u, v_u) \leftarrow \text{Receive}(\text{st}_u, \sigma_{\text{idx}(u)}, m)</math> 5: <b>return</b> <math>(b_u, v_u)_{u \in \mathcal{H}}</math> </pre> <p><b>Oracle <math>O_{\text{UpdSend}}(u \in \mathcal{H})</math></b></p> <pre> 1: <math>s_{\text{UpdRec}} := \emptyset</math> 2: <math>(\text{st}'_u, \widehat{\Sigma}_G, \widehat{\text{ct}}_G) \leftarrow \mathcal{A}^O(\text{UpdSend}(\text{st}_u))</math> 3: <math>L_{\text{upd}} \leftarrow L_{\text{upd}} \cup \{(u, \widehat{\Sigma}_G, \widehat{\text{ct}}_G)\}</math> 4: / Server honestly processes group authentication 5: <b>if</b> <math>\mathcal{A}</math> has access to <math>O^*</math> <b>then</b> 6:   <math>s_{\text{UpdRec}} \leftarrow O_{\text{GroupUpdReceive}}(\widehat{\Sigma}_G, \widehat{\text{ct}}_G)</math> 7: <b>return</b> <math>(\widehat{\Sigma}_G, \widehat{\text{ct}}_G, s_{\text{UpdRec}})</math> </pre> <p><b>Oracle <math>O_{\text{UpdReceive}}(u \in \mathcal{H}, \widehat{\sigma}, \widehat{\text{ct}})</math></b></p> <pre> 1: <math>(\text{st}'_u, b_u, v_u) \leftarrow \text{UpdReceive}(\text{st}_u, \widehat{\sigma}_{\text{idx}(u)}, \widehat{\text{ct}}_{\text{idx}(u)})</math> 2: <b>return</b> <math>(b_u, v_u)</math> </pre> <p><b>Oracle <math>O_{\text{GroupUpdReceive}}(\widehat{\Sigma}_G, \widehat{\text{ct}}_G)</math></b></p> <pre> 1: <math>(\text{pp}', (\widehat{\sigma}_i, \widehat{\text{ct}}_i)_{i \in [N]}) \leftarrow \text{UpdVerify}(\text{pp}, \text{sk}_{Sv}, \widehat{\Sigma}_G, \widehat{\text{ct}}_G)</math> 2: <b>if</b> <math>\text{pp}' = \perp</math> <b>then return</b> <math>\perp</math> 3: <b>foreach</b> <math>u \in \mathcal{H}</math> <b>do</b> 4:   <math>(\text{st}'_u, b_u, v_u) \leftarrow \text{UpdReceive}(\text{st}_u, \widehat{\sigma}_{\text{idx}(u)}, \widehat{\text{ct}}_{\text{idx}(u)})</math> 5:   <b>return</b> <math>(b_u, v_u)_{u \in \mathcal{H}}</math> </pre>
<p><b>Game<math>_{\mathcal{A}}^{\text{Anon}}(1^\kappa)</math></b></p> <pre> 1: <math>\mathcal{H} := G</math> / No corrupt users 2: <math>\text{Chall}_{\text{msg}} := \emptyset</math> 3: <math>\text{coin} \leftarrow \mathcal{A}(\{0, 1\})</math> 4: <math>(\text{pp}, \text{sk}_{Sv}, (\text{st}_u)_{u \in G}) \leftarrow \mathcal{A}^O(1^\kappa, G)</math> 5: <math>(\widehat{\text{pp}}, u_0, u_1, m_0, m_1) \leftarrow \mathcal{A}^O(\text{pp}, \text{sk}_{Sv})</math> 6: <b>foreach</b> <math>b \in \{0, 1\}</math> <b>do</b> 7:   <math>(\text{st}'_{u_b}, \Sigma_G^b) \leftarrow \mathcal{A}^O(\text{Send}(\text{st}_{u_b}, m_b \oplus \text{coin}))</math> 8:   <math>(\widehat{\text{pp}}', (\sigma_i^b)_{i \in [N]})</math> 9:   <math>\leftarrow \text{Verify}(\widehat{\text{pp}}, \text{sk}_{Sv}, \Sigma_G^b, m_b \oplus \text{coin})</math> 10:  / Require the authentication token to be valid 11:  <b>req</b> <math>\widehat{\text{pp}}' \neq \perp</math> 12:  <b>foreach</b> <math>u \in \mathcal{H}</math> <b>do</b> 13:    <math>(\text{st}'_u, b_u, v_u) \leftarrow \text{Receive}(\text{st}_u, \sigma_{\text{idx}(u)}^b, m_b \oplus \text{coin})</math> 14:    <b>req</b> <math>b_u \neq \perp</math> 15:    <math>\text{Chall}_{\text{msg}} \leftarrow \text{Chall}_{\text{msg}} \cup \{\sigma_i^b\}_{i \in [N]}</math> 16:    <math>\widehat{\text{pp}} \leftarrow \widehat{\text{pp}}'</math> 17:  <math>\widehat{\text{coin}} \leftarrow \mathcal{A}^O(\text{Chall}_{\text{msg}})</math> 18:  <b>return</b> <math>\llbracket \text{coin} = \widehat{\text{coin}} \rrbracket</math> </pre>		

Figure 2: Security games for (non-colluding) unforgeability, anonymity, anonymous blocklisting, and tracing soundness. We define a set of oracles  $O := \{O_{\text{Send}}, O_{\text{Receive}}, O_{\text{UpdSend}}, O_{\text{UpdReceive}}\}$  and  $O^* := \{O_{\text{Send}}, O_{\text{GroupReceive}}, O_{\text{UpdSend}}, O_{\text{GroupUpdReceive}}\}$ . We assume the game maintains the public parameter  $\text{pp}$  and (secret) user states  $\text{st}_u$ . Moreover, we assume the updated state  $\text{st}'_u$  is implicitly set as  $\text{st}_u$  and omit the substitution  $\text{st}'_u \leftarrow \text{st}_u$  for readability. When the condition in **req** does not hold, we assume the game outputs a random bit in the anonymity game and 0 in all other games. Lastly, for readability, we sometimes ignore creating the lists  $L_{\text{tr}}, L_{\text{msg}}, L_{\text{upd}}$  when they are not required by the game.

includes a server secret key  $sk_{sv}$ ; our security says nothing if  $sk_{sv}$  is *maliciously* generated.

One way to handle this issue of transparency of the server is to enforce that the server’s secret key  $sk_{sv}$  can be *deterministically* derived from any user state  $st_u$ . With such a restriction, any user can locally run the server’s algorithms `Verify` and `UpdVerify`, and potentially audit the server’s behavior. Indeed, all of the protocols proposed in this paper will have such a property as  $sk_{sv}$  is derived from the group secret key generated by the CGKA protocol.

## 4 COSMOS: Authentication with One-Time Tokens

In this section we propose a GAM protocol named COSMOS (Compact authenticated Secure Messaging with randomized One-time tokenS). When anonymity is not necessary, COSMOS is the most efficient and simplest protocol among all our proposed protocols. The additional total communication overhead is only  $3\kappa$  compared to a protocol where messages are sent without any authentication, where  $\kappa$  is the security parameter. Additionally, we show a simple method to bootstrap COSMOS to satisfy anonymity and anonymous blocklisting, which we name COSMAC. The added overhead to COSMOS is a single MAC tag. Lastly, we show how to optimize both protocols by batching sends and updates together.

### 4.1 Construction of COSMOS

The high level idea is as follows: each group user mints tokens  $(x_i, y_i) \in \{0, 1\}^\kappa \times \{0, 1\}^\kappa$  for  $i \in [T]$  such that  $y_i = \text{OWF}(x_i)$ ; stores the private tokens  $(x_i)_{i \in [T]}$  in its state; uploads the public tokens  $(y_i)_{i \in [T]}$  to the server in an *offline* phase; and ideally sends  $(x_i, m_i)$  to the server once the message  $m_i$  is defined in an *online* phase, where  $x_i$  acts as the authentication token, and delete  $x_i$  from its state. However, since  $x_i$  is not cryptographically tied to  $m_i$ , this is insecure. Thus, the user additionally MACs  $(x_i, m_i)$  using a MAC key only known among the group users. We highlight that such a MAC key can be generated from the *common* group secret key  $gsk$  maintained by the CGKA protocol.

More formally, after the initialization phase, each user  $u \in G$  and server maintain a list of public tokens  $\text{PubTOKEN} \in (\{0, 1\}^\kappa)^{NT}$ , where  $N = |G|$  and  $T$  is the number of messages a user can send before needing to update its state.  $\text{PubTOKEN}$  is a list such that, for each user  $u \in G$ ,  $\text{PubTOKEN}[u] \in (\{0, 1\}^\kappa)^T$  stores the  $T$  public tokens  $(y_i)_{i \in [T]}$  used by user  $u$ . User  $u$  also maintains a list of private tokens  $\text{PrivTOKEN}_u \in (\{0, 1\}^\kappa)^T$  storing the  $T$  private tokens  $(x_i)_{i \in [T]}$ .

To send a message  $m$ , user  $u$  retrieves an unused private token  $x$  from  $\text{PrivTOKEN}_u$ , along with the counter  $ctr \in [T]$  such that  $\text{PubTOKEN}[u][ctr] = \text{OWF}(x)$ , and sends  $(x, ctr, \Sigma_{MAC})$  as the group authentication token  $\Sigma_G$  to the

server, where  $\Sigma_{MAC}$  is a MAC tag using  $k_{MAC}$ . The server then checks if the token  $x$  is valid (i.e.,  $y_{ctr} = \text{OWF}(x)$ ) and relays  $(x, ctr, \Sigma_{MAC})$  as the user authentication token  $\sigma_i$  to all the users. Here,  $\Sigma_{MAC}$  does not need to (nor can it) be verified by the server. Now, since  $\text{PubTOKEN}$  and  $k_{MAC}$  is shared among the group, the users can verify the MAC tag and trace the user  $u$  that sent  $\sigma_i$ .

When only one private token  $x$  is left, user  $u$  performs a state update and mints new tokens. It generates a new batch of  $T$  tokens  $(x_i, y_i)_{i \in [T]}$  and uploads  $(y_i)_{i \in [T]}$  using the final token  $x$  along with a MAC tag. The server and users check that the newly minted tokens are from user  $u$  by validating  $x$  and update  $\text{PubTOKEN}[u] \leftarrow (y_i)_{i \in [T]}$ . Once user  $u$ ’s state is updated,  $u$  can send  $T$  messages again. Importantly, COSMOS is *locally* state-updatable since a user can start sending messages once they update their state. We provide the formal description of COSMOS in Fig. 3.

Lastly, COSMOS satisfies all the security notions except for anonymity: non-colluding unforgeability, (anonymous) blocklisting, and tracing soundness. At a high level, we argue non-colluding unforgeability by considering two cases: against a malicious server the authentication token is unforgeable as  $k_{MAC}$  is unknown. Importantly, the same authentication token  $(x, ctr, \Sigma_{MAC})$  cannot be reused by the malicious server since the users have already deleted the associated public token  $y$  when it receives  $x$  the second time. Otherwise, against a malicious user, it is unforgeable as the private token  $x$  is unknown. In the latter, we use the fact that an honest server correctly processes the private token sent from an honest user (i.e., delete it from the server), preventing a malicious user from replaying it. One can check that it is not *standard* unforgeable since if a malicious server and insider collude, both  $k_{MAC}$  and private tokens  $x$  will be known to the adversary, allowing for a trivial forgery. Moreover, we note that even though the MAC tag attached to the group authentication token cannot be verified by the server, and hence can be stealthily modified to a garbage MAC tag, this will not harm tracing soundness as we only use the private tokens for tracing. The formal security proof is deferred to App. C.1.

### 4.2 COSMAC: An Anonymous COSMOS with Anonymous Blocklisting

While COSMOS is efficient, it lacks anonymity. A server can link two tokens by looking at their corresponding locations in  $\text{PubTOKEN}$ . We present a simple method to transform COSMOS to have anonymity and anonymous blocklisting at an overhead of only one MAC tag. We name this GAM protocol COSMAC (COSMOS with MAC). Note that in exchange for anonymity, COSMAC loses tracing soundness.

The high level idea is for each user in the group to additionally derive a unique MAC key  $\overline{k_{MAC}}$  and a symmetric-key encryption (SKE) key  $k_{SKE}$  from  $gsk$  where, unlike in COSMOS,  $\overline{k_{MAC}}$  is uploaded to the server. When a group user uploads

<p><b>Init</b>(<math>1^K, G</math>)</p> <hr/> <pre> 1: <math>gsk \leftarrow \mathcal{S}\{0, 1\}^K</math> // Group secret key 2: <math>k_{MAC} \leftarrow PRF(gsk, 0)</math> // MAC key for group 3: // Prepare empty lists 4: <b>foreach</b> <math>u \in G</math> <b>do</b> 5:   <math>PubTOKEN_u[*] := \perp</math> 6:   <math>PrivTOKEN_u[*] := \perp</math> 7:   <b>foreach</b> <math>u \in G</math> <b>do</b> 8:     <math>(X_u, Y_u) \leftarrow *gen-auth-token(G, u)</math> 9:     <math>PrivTOKEN_u \leftarrow X_u</math> // <math>X_u \in (\{0, 1\}^K)^T</math> 10:    <b>foreach</b> <math>v \in G</math> <b>do</b> 11:      <math>PubTOKEN_v[u] \leftarrow Y_u</math> // <math>Y_u \in (\{0, 1\}^K)^T</math> 12:    <b>foreach</b> <math>u \in G</math> <b>do</b> 13:      <math>TOKEN_u := (PrivTOKEN_u, PubTOKEN_u)</math> 14:      <math>st_u \leftarrow (G, k_{MAC}, 1, TOKEN_u)</math> 15:      <math>DB[*] := \perp</math> // Prepare empty database for Sv 16:      <b>for</b> <math>u \in G</math> <b>do</b> 17:        <math>DB[u] \leftarrow Y_u</math> 18:      <math>pp \leftarrow DB</math> // <math>DB \in (\{0, 1\}^K)^{N \times T}</math> 19:      <b>return</b> <math>(pp, (st_u)_{u \in G})</math></pre>	<p><b>Send</b>(<math>st_u, m</math>)</p> <hr/> <pre> 1: <b>parse</b> <math>(G, k_{MAC}, ctr, TOKEN_u) \leftarrow st_u</math> 2: <b>if</b> <math>ctr \geq T - 1</math> <b>then return</b> <math>\perp</math> // Need to update tokens 3: <math>ctr' \leftarrow ctr + 1</math> 4: <math>\Sigma_G \leftarrow *attach-auth-token(st_u, m)</math> 5: <math>st'_u \leftarrow (G, k_{MAC}, ctr', TOKEN_u)</math> 6: <b>return</b> <math>(st'_u, \Sigma_G)</math></pre>
<p><b>UpdSend</b>(<math>st_u</math>)</p> <hr/> <pre> 1: <b>parse</b> <math>(G, k_{MAC}, ctr, TOKEN_u) \leftarrow st_u</math> 2: <b>parse</b> <math>(PrivTOKEN_u, PubTOKEN_u) \leftarrow TOKEN_u</math> 3: <math>(X_u, Y_u) \leftarrow *gen-auth-token(G, u)</math> 4: <math>PrivTOKEN_u \leftarrow X_u</math> 5: <math>PubTOKEN_u[u] \leftarrow Y_u</math> 6: <math>TOKEN_u \leftarrow (PrivTOKEN_u, PubTOKEN_u)</math> 7: // Refresh counter to 1 8: <math>st'_u \leftarrow (G, k_{MAC}, 1, TOKEN_u)</math> 9: <math>\hat{ct}_G \leftarrow (u, Y_u)</math> 10: <math>\hat{\Sigma}_G \leftarrow *attach-auth-token(st_u, \hat{ct}_G)</math> 11: <b>return</b> <math>(st'_u, \hat{\Sigma}_G, \hat{ct}_G)</math></pre>	<p><b>Verify</b>(<math>pp, \Sigma_G, m</math>)</p> <hr/> <pre> 1: <b>parse</b> <math>DB \leftarrow pp</math> 2: <b>try</b> <math>(pp', (\sigma_v)_{v \in G}) \leftarrow *verify-auth-token(pp, \Sigma_G)</math> 3: <b>return</b> <math>(pp', (\sigma_v)_{v \in G})</math></pre>
<p><b>UpdVerify</b>(<math>pp, \hat{\Sigma}_G, \hat{ct}_G</math>)</p> <hr/> <pre> 1: <b>parse</b> <math>DB \leftarrow pp</math> 2: <b>try</b> <math>(pp', (\hat{\sigma}_v)_{v \in G}) \leftarrow *verify-auth-token(pp, \hat{\Sigma}_G)</math> 3: <b>parse</b> <math>(u, Y_u) \leftarrow \hat{ct}_G</math> 4: <math>DB[u] \leftarrow Y_u</math> 5: <b>foreach</b> <math>v \in G</math> <b>do</b> 6:   <math>\hat{ct}_v \leftarrow (u, Y_u)</math> 7: <b>return</b> <math>(pp', (\hat{\sigma}_v, \hat{ct}_v)_{v \in G})</math></pre>	<p><b>Receive</b>(<math>st_u, \sigma, m</math>)</p> <hr/> <pre> 1: <b>try</b> <math>(st'_u, b, v) \leftarrow *trace-sender(st_u, \sigma, m)</math> 2: <b>return</b> <math>(st'_u, b, v)</math></pre>
<p><b>UpdReceive</b>(<math>st_u, \hat{\sigma}, \hat{ct}</math>)</p> <hr/> <pre> 1: <b>parse</b> <math>(G, k_{MAC}, ctr, TOKEN_u) \leftarrow st_u</math> 2: <b>parse</b> <math>(PrivTOKEN_u, PubTOKEN_u) \leftarrow TOKEN_u</math> 3: <b>try</b> <math>(st_u, b, v) \leftarrow *trace-sender(st_u, \hat{\sigma}, \hat{ct})</math> 4: <b>parse</b> <math>(v', Y) \leftarrow \hat{ct}</math> 5: <b>req</b> <math>v = v'</math> 6: <math>PubTOKEN_u[v] \leftarrow Y</math> // Update user v's tokens 7: <math>TOKEN_u \leftarrow (PrivTOKEN_u, PubTOKEN_u)</math> 8: <math>st'_u \leftarrow (G, k_{MAC}, ctr, TOKEN_u)</math> 9: <b>return</b> <math>(st'_u, b, v)</math></pre>	<p><b>Receive</b>(<math>st_u, \sigma, m</math>)</p> <hr/> <pre> 1: <b>try</b> <math>(st'_u, b, v) \leftarrow *trace-sender(st_u, \sigma, m)</math> 2: <b>return</b> <math>(st'_u, b, v)</math></pre>

Figure 3: COSMOS: A group authenticated messaging protocol with one-time tokens. The server  $S_v$  is assumed to implicitly apply the inverse of  $idx$ , i.e., it uses  $u \in G$  rather than  $i = idx(u) \in [N]$ . The helper algorithms used above are detailed in Fig. 4.

<p><b>Func *gen-auth-token(<math>G, u</math>)</b></p> <pre> 1: <math>X[*], Y[*] := \perp</math> 2: <b>foreach</b> <math>t \in [T]</math> <b>do</b> 3:   <math>X[t] \leftarrow \{0, 1\}^k</math> 4:   <math>Y[t] \leftarrow \text{OWF}(X[t])</math> 5: <b>return</b> <math>(X, Y)</math> </pre>	<p><b>Func *attach-auth-token(<math>st_u, m</math>)</b></p> <pre> 1: <b>parse</b> <math>(G, k_{\text{MAC}}, \text{ctr}, \text{TOKEN}_u) \leftarrow st_u</math> 2: <b>parse</b> <math>(\text{PrivTOKEN}_u, \text{PubTOKEN}_u) \leftarrow \text{TOKEN}_u</math> 3: <math>x_u^{(\text{ctr})} \leftarrow \text{PrivTOKEN}_u[\text{ctr}]</math> 4: <math>\text{PrivTOKEN}_u[\text{ctr}] \leftarrow \perp</math> 5: <math>\text{TOKEN}_u \leftarrow (\text{PrivTOKEN}_u, \text{PubTOKEN}_u)</math> 6: <math>\Sigma_{\text{MAC}} \leftarrow \text{MAC.TagGen}(k_{\text{MAC}}, (u, \text{ctr}, x_u^{(\text{ctr})}), m)</math> 7: <b>return</b> <math>(u, \text{ctr}, x_u^{(\text{ctr})}, \Sigma_{\text{MAC}})</math> </pre>
<p><b>Func *verify-auth-token(<math>pp, \Sigma_G</math>)</b></p> <pre> 1: <math>\text{DB} \leftarrow pp \quad / \text{DB} \in (\{0, 1\}^k)^{N \times T}</math> 2: <b>parse</b> <math>(u, \text{ctr}, x, \Sigma_{\text{MAC}}) \leftarrow \Sigma_G</math> 3: <b>if</b> <math>\text{DB}[u][\text{ctr}] \neq \text{OWF}(x)</math> <b>then</b> 4:   <b>return</b> <math>\perp</math> 5: <i>/ If check passes, set user authentication tokens and</i> 6: <i>/ delete entry from DB</i> 7: <b>foreach</b> <math>v \in G</math> <b>do</b> 8:   <math>\sigma_v \leftarrow (u, \text{ctr}, x, \Sigma_{\text{MAC}})</math> 9: <math>\text{DB}[u][\text{ctr}] \leftarrow \perp</math> 10: <math>pp' \leftarrow \text{DB}</math> 11: <b>return</b> <math>(pp', (\sigma_v)_{v \in G})</math> </pre>	<p><b>Func *trace-sender(<math>st_u, \sigma, m</math>)</b></p> <pre> 1: <b>parse</b> <math>(G, k_{\text{MAC}}, \text{ctr}, \text{TOKEN}_u) \leftarrow st_u</math> 2: <b>parse</b> <math>(\text{PrivTOKEN}_u, \text{PubTOKEN}_u) \leftarrow \text{TOKEN}_u</math> 3: <b>parse</b> <math>(v, \text{ctr}_v, x, \Sigma_{\text{MAC}}) \leftarrow \sigma</math> 4: <b>if</b> <math>\text{PubTOKEN}_u[v][\text{ctr}_v] \neq \text{OWF}(x)</math> 5:   <math>\vee \text{MAC.Verify}(k_{\text{MAC}}, (v, \text{ctr}_v, x, m), \Sigma_{\text{MAC}}) = \perp</math> <b>then</b> 6:   <b>return</b> <math>(st_u, \perp, \perp)</math> 7: <math>\text{PubTOKEN}_u[v][\text{ctr}_v] \leftarrow \perp</math> 8: <math>\text{TOKEN}_u \leftarrow (\text{PrivTOKEN}_u, \text{PubTOKEN}_u)</math> 9: <math>st'_u \leftarrow (G, k_{\text{MAC}}, \text{ctr}, \text{TOKEN}_u)</math> 10: <b>return</b> <math>(st'_u, \top, v)</math> </pre>

Figure 4: Helper functions used by COSMOS.

some content to the server, it runs the Send (resp. UpdSend) algorithm of COSMOS, encrypts the group (resp. update) authentication token using  $k_{\text{SKE}}$ , and MACs the ciphertext with  $\overline{k_{\text{MAC}}}$ . The server only accepts contents that have a valid tag under  $\overline{k_{\text{MAC}}}$ . A group user can verify the user authentication token by first decrypting the ciphertext using  $k_{\text{SKE}}$ , followed by the same check as COSMOS. The protocol is formally given in Fig. 5. The protocol consists of COSMOS and a wrapper protocol that encrypts and MACs the authentication tokens output by COSMOS. The main difference between COSMAC and COSMOS is highlighted with a box in Fig. 5.

Observe that the authentication tokens are now encrypted and the server no longer learns the identity of the user. This is how anonymity is achieved. Non-colluding unforgeability almost immediately follows from the non-colluding unforgeability of COSMOS. This is because from the users point of view, COSMAC and COSMOS are almost identical. The only difference is that COSMAC requires to first perform a decryption using  $k_{\text{SKE}}$ ; this does not make forging anymore easier for the adversary. Indeed, prove non-colluding unforgeability of COSMAC assuming the non-colluding unforgeability of COSMOS. Moreover, COSMAC satisfies anonymous blocklisting since an outsider without knowledge of  $\overline{k_{\text{MAC}}}$  cannot upload contents which the server will accept. Lastly, on the other hand, unlike COSMOS, a malicious user can now stealthily perform a DoS attack on the group since the server can only check the validity of the MAC tag and not the content. In particular, COSMAC loses tracing soundness, as the content, which can now be a malformed ciphertext, may not include the sender’s identity. We show in App. C.2 that COSMAC is non-colluding unforgeable, anonymous, and anonymous blocklistable.

### 4.3 Optimizations of COSMOS and COSMAC

We take advantage of the fact that COSMOS (and COSMAC) have an efficient *local* state update and apply two optimizations leading to COSMOS<sup>+</sup> and COSMOS<sup>++</sup> (and COSMAC<sup>+</sup> and COSMAC<sup>++</sup>, respectively). We focus on COSMOS as the case for COSMAC is almost identical.

**Removing Local Updates:** COSMOS<sup>+</sup>. Notice that local state-updates allow a user to execute the UpdSend algorithm as part of the Send algorithm. That is, users can send a message and perform an update *at the same time*. Concretely, this requires to maintain just one public token  $y$  per user  $u \in G$ . To send a message  $m$ ,  $u$  first mints a new token  $(x', y')$  and uploads both the message and public token  $(m, y')$  using the private token  $x$ , along with a MAC tag binding  $(m, x, y')$  together. The server and other group members replace  $y$  with  $y'$ . User  $u$  can repeat the process using the new private token  $x'$ . Effectively, the protocol now consists of only running an online phase, since the update is implicitly performed during a send. Compared to COSMOS, COSMOS<sup>+</sup> balances the throughput of the user without harming the total communication cost  $3 \cdot \kappa$ , while also reducing the storage cost of public tokens.

**Minimizing Communication Cost:** COSMOS<sup>++</sup>. This optimization reduces the communication cost of COSMOS<sup>+</sup> by  $1/3$  while keeping the local update of COSMOS.<sup>9</sup> The main idea is to make the private authentication token become the public token for the next message. As in COSMOS<sup>+</sup>, the server maintains a single public token  $y_{i,c}$  per user, where  $c \in \mathbb{N}$  will be the number of times user  $u$  ran UpdSend. As a result of running UpdSend the user will upload a public token  $y_{0,c} = \text{OWF}^T(x_{T,c})$ , i.e.,  $T^{\text{th}}$  invocation of OWF. This updated public token can be used to send  $T$  messages. To authenticate the  $i$ -th message  $m_i$ , user  $u$  simply sends token  $x_{i,c} = \text{OWF}^{T-i}(x_{T,c})$  along with a MAC tag on  $(m_i, x_{i,c})$ . The server and other group members update the public token to  $y_{i,c} := x_{i,c}$ . Since the public token generated in the offline phase is useful for sending  $T$  messages, the amortized cost of sending one message is  $(2 + \frac{1}{T}) \cdot \kappa$ . For a sufficiently large  $T$ , this reduces the communication cost of COSMOS<sup>+</sup> by  $1/3$ .

## 5 Anonymous and Tracing Sound GAMs

In this section we introduce three GAM protocols that simultaneously achieve anonymity and tracing soundness. These are the first authentication modes in the literature to do so. The first two protocols: QUASAR (Quick Authenticated Secure Anonymous messaging with Randomized one-time tokens) and STARS (Strongly-Authenticated anonymous messaging with Randomized one-time Signatures) satisfy these stronger authenticity guarantees at the cost of being only *global* as opposed to *local* state-updatable like COSMOS and COSMAC. Once a user  $u \in G$  exhausts its private tokens, it must wait till all other users perform an update before being able to send a message again. We discuss some ideas to mitigate the shortcoming of global state updates in App. D.3. Our third protocol GEMSTARS (Group Signature Modified STARS), eliminates updates altogether by relying on group signatures. Below we give intuitive overviews of the protocols, deferring the formal descriptions and security proofs to Apps. D and E.

### 5.1 QUASAR: Anonymous Authentication with Tokens

We first consider a non-anonymous variant of QUASAR and add anonymity later. Its core idea is to perform a relatively expensive *offline* phase (i.e., UpdSend) to make the *online* phase (i.e., Send) very cheap.

**Basic Idea.** Assume a group  $G = (u_i)_{i \in [N]}$ . Each user  $u_i$  mints tokens  $(x_{j \rightarrow i}^{(t)}, y_{j \rightarrow i}^{(t)}) \in \{0, 1\}^\kappa \times \{0, 1\}^\kappa$  for  $(j, t) \in [N] \times [T]$  such that  $y_{j \rightarrow i}^{(t)} = \text{OWF}(x_{j \rightarrow i}^{(t)})$ . Here,  $j \rightarrow i$  indicates that user

<sup>9</sup>This optimization was suggested to us by an anonymous reviewer; afterwards, another reviewer informed us that a similar idea is used in the S/KEY one-time password authentication protocol [48, 49].

<p><b>Init</b>(<math>1^k, G</math>)</p> <hr/> <pre> 1: <math>gsk \leftarrow \{0, 1\}^k</math> / Group secret key 2: <math>k_{MAC} \leftarrow PRF(gsk, 0)</math> / MAC key for group 3: <math>sk_{Sv} := k_{MAC} \leftarrow PRF(gsk, 1)</math> / MAC key for Sv and group 4: <math>k_{SKE} \leftarrow PRF(gsk, 2)</math> / SKE key for group 5: / Prepare empty lists 6: <b>foreach</b> <math>u \in G</math> <b>do</b> 7:   <math>PubTOKEN_u[*] := \perp</math> 8:   <math>PrivTOKEN_u[*] := \perp</math> 9: <b>foreach</b> <math>u \in G</math> <b>do</b> 10:  <math>(X_u, Y_u) \leftarrow *gen-auth-token(G, u)</math> 11:  <math>PrivTOKEN_u \leftarrow X_u</math> / <math>X_u \in (\{0, 1\}^k)^T</math> 12:  <b>foreach</b> <math>v \in G</math> <b>do</b> 13:    <math>PubTOKEN_v[u] \leftarrow Y_u</math> / <math>Y_u \in (\{0, 1\}^k)^T</math> 14: <b>foreach</b> <math>u \in G</math> <b>do</b> 15:  <math>TOKEN_u := (PrivTOKEN_u, PubTOKEN_u)</math> 16:  <math>st_u \leftarrow (G, k_{MAC}, \overline{k_{MAC}}, k_{SKE}, 1, TOKEN_u)</math> 17:  <math>pp := \perp</math> / No pp maintained by Sv 18: <b>return</b> <math>(pp, \overline{sk_{Sv}}, (st_u)_{u \in G})</math> </pre>	<p><b>Send</b>(<math>st_u, m</math>)</p> <hr/> <pre> 1: <b>parse</b> <math>(G, k_{MAC}, \overline{k_{MAC}}, k_{SKE}, ctr, TOKEN_u) \leftarrow st_u</math> 2: <b>if</b> <math>ctr \geq T - 1</math> <b>then return</b> <math>\perp</math> / Need to update tokens 3: <math>ctr' \leftarrow ctr + 1</math> 4: <math>\Sigma_G \leftarrow *attach-auth-token(st_u, m)</math> 5: <math>st'_u \leftarrow (G, k_{MAC}, \overline{k_{MAC}}, k_{SKE}, ctr', TOKEN_u)</math> 6: <b>return</b> <math>(st'_u, \Sigma_G)</math> </pre> <p><b>Verify</b>(<math>pp, \overline{sk_{Sv}}, \Sigma_G, m</math>)</p> <hr/> <pre> 1: <b>try</b> <math>(\perp, (\sigma_i)_{i \in [M]})</math>    <math>\leftarrow *verify-auth-token(\perp, \overline{sk_{Sv}}, \Sigma_G)</math> 2: <b>return</b> <math>(pp', (\sigma_i)_{i \in [M]})</math> </pre> <p><b>Receive</b>(<math>st_u, \sigma, m</math>)</p> <hr/> <pre> 1: <b>try</b> <math>(st'_u, b, v) \leftarrow *trace-sender(st_u, \sigma, m)</math> 2: <b>return</b> <math>(st'_u, b, v)</math> </pre>
<p><b>UpdSend</b>(<math>st_u</math>)</p> <hr/> <pre> 1: <b>parse</b> <math>(G, k_{MAC}, \overline{k_{MAC}}, k_{SKE}, ctr, TOKEN_u) \leftarrow st_u</math> 2: <b>parse</b> <math>(PrivTOKEN_u, PubTOKEN_u) \leftarrow TOKEN_u</math> 3: <math>(X_u, Y_u) \leftarrow *gen-auth-token(G, u)</math> 4: <math>\widehat{\Sigma}_G \leftarrow *attach-auth-token(st_u, Y_u)</math> 5: <math>PrivTOKEN_u \leftarrow X_u</math> 6: <math>PubTOKEN_u[u] \leftarrow Y_u</math> 7: <math>TOKEN_u \leftarrow (PrivTOKEN_u, PubTOKEN_u)</math> 8: / Refresh counter to 1 9: <math>st'_u \leftarrow (G, k_{MAC}, \overline{k_{MAC}}, k_{SKE}, 1, TOKEN_u)</math> 10: <math>\widehat{ct}_{SKE} \leftarrow Enc(k_{SKE}, (u, Y_u))</math> 11: <math>\widehat{ct}_G \leftarrow \widehat{ct}_{SKE}</math> 12: <b>return</b> <math>(st'_u, \widehat{\Sigma}_G, \widehat{ct}_G)</math> </pre>	<p><b>UpdVerify</b>(<math>pp, \overline{sk_{Sv}}, \widehat{\Sigma}_G, \widehat{ct}_G</math>)</p> <hr/> <pre> 1: <b>try</b> <math>(\perp, (\widehat{\sigma}_i)_{i \in [M]}) \leftarrow *verify-auth-token(\perp, \overline{sk_{Sv}}, \widehat{\Sigma}_G)</math> 2: <b>foreach</b> <math>i \in [M]</math> <b>do</b> 3:   <math>\widehat{ct}_i \leftarrow \widehat{ct}_G</math> 4: <b>return</b> <math>(pp', (\widehat{\sigma}_i, \widehat{ct}_i)_{i \in [M]})</math> </pre> <p><b>UpdReceive</b>(<math>st_u, \widehat{\sigma}, \widehat{ct}</math>)</p> <hr/> <pre> 1: <b>parse</b> <math>(G, k_{MAC}, \overline{k_{MAC}}, k_{SKE}, ctr, TOKEN_u) \leftarrow st_u</math> 2: <b>parse</b> <math>(PrivTOKEN_u, PubTOKEN_u) \leftarrow TOKEN_u</math> 3: <b>parse</b> <math>\widehat{ct}_{SKE} \leftarrow \widehat{ct}</math> 4: <math>(v', Y) \leftarrow SKE.Dec(k_{SKE}, \widehat{ct}_{SKE})</math> 5: <b>try</b> <math>(st_u, b, v) \leftarrow *trace-sender(st_u, \widehat{\sigma}, Y)</math> 6: <b>req</b> <math>v = v'</math> 7: <math>PubTOKEN_u[v] \leftarrow Y</math> / Update user v's tokens 8: <math>TOKEN_u \leftarrow (PrivTOKEN_u, PubTOKEN_u)</math> 9: <math>st'_u \leftarrow (G, k_{MAC}, \overline{k_{MAC}}, k_{SKE}, ctr, TOKEN_u)</math> 10: <b>return</b> <math>(st'_u, \top, v)</math> </pre>

Figure 5: COSMAC: An anonymous group authenticated messaging protocol with one-time tokens. The main differences between COSMOS is highlighted by a box. The helper algorithms used above are detailed in Fig. 6.

<p><b>Func</b> *gen-auth-token(<math>G, u</math>)</p> <pre> 1: <math>X[*], Y[*] := \perp</math> 2: <b>foreach</b> <math>t \in [T]</math> <b>do</b> 3:   <math>X[t] \leftarrow \{0, 1\}^k</math> 4:   <math>Y[t] \leftarrow \text{OWF}(X[t])</math> 5: <b>return</b> <math>(X, Y)</math> </pre>	<p><b>Func</b> *attach-auth-token(<math>st_u, m</math>)</p> <pre> 1: <b>parse</b> <math>\left( G, k_{\text{MAC}}, \overline{k_{\text{MAC}}}, k_{\text{SKE}}, \text{ctr}, \text{TOKEN}_u \right) \leftarrow st_u</math> 2: <b>parse</b> <math>(\text{PrivTOKEN}_u, \text{PubTOKEN}_u) \leftarrow \text{TOKEN}_u</math> 3: <math>x_u^{(\text{ctr})} \leftarrow \text{PrivTOKEN}_u[\text{ctr}]</math> 4: <math>\text{PrivTOKEN}_u[\text{ctr}] \leftarrow \perp</math> 5: <math>\text{TOKEN}_u \leftarrow (\text{PrivTOKEN}_u, \text{PubTOKEN}_u)</math> 6: <math>\Sigma_{\text{MAC}} \leftarrow \text{MAC.TagGen}(k_{\text{MAC}}, (u, \text{ctr}, x_u^{(\text{ctr})}, m))</math> 7: <math>\text{ctsKE} \leftarrow \text{SKE.Enc}(k_{\text{SKE}}, (u, \text{ctr}, x_u^{(\text{ctr})}, \Sigma_{\text{MAC}}))</math> 8: <math>\overline{\Sigma}_{\text{MAC}} \leftarrow \text{MAC.TagGen}(\overline{k_{\text{MAC}}}, \text{ctsKE})</math> 9: <b>return</b> <math>(\text{ctsKE}, \overline{\Sigma}_{\text{MAC}})</math> </pre>
<p><b>Func</b> *verify-auth-token(<math>pp = \perp, sk_{\text{Sv}}, \Sigma_G</math>)</p> <pre> 1: <b>parse</b> <math>(\text{ctsKE}, \overline{\Sigma}_{\text{MAC}}) \leftarrow \Sigma_G</math> 2: <b>if</b> <math>\text{MAC.Verify}(sk_{\text{Sv}}, \text{ctsKE}, \overline{\Sigma}_{\text{MAC}}) = \perp</math> <b>then</b> 3:   <b>return</b> <math>\perp</math> 4: / If check passes, set user authentication tokens 5: <b>foreach</b> <math>i \in [N]</math> <b>do</b> 6:   <math>\sigma_i \leftarrow \text{ctsKE}</math> 7: <b>return</b> <math>(\perp, (\sigma_i)_{i \in [N]})</math> </pre>	<p><b>Func</b> *trace-sender(<math>st_u, \sigma, m</math>)</p> <pre> 1: <b>parse</b> <math>\left( G, k_{\text{MAC}}, \overline{k_{\text{MAC}}}, k_{\text{SKE}}, \text{ctr}, \text{TOKEN}_u \right) \leftarrow st_u</math> 2: <b>parse</b> <math>(\text{PrivTOKEN}_u, \text{PubTOKEN}_u) \leftarrow \text{TOKEN}_u</math> 3: <b>parse</b> <math>\text{ctsKE} \leftarrow \sigma</math> 4: <math>(v, \text{ctr}_v, x, \Sigma_{\text{MAC}}) \leftarrow \text{SKE.Dec}(k_{\text{SKE}}, \text{ctsKE})</math> 5: <b>if</b> <math>\text{PubTOKEN}_u[v][\text{ctr}_v] \neq \text{OWF}(x)</math> 6:   <math>\vee \text{MAC.Verify}(k_{\text{MAC}}, (v, \text{ctr}_v, x, m), \Sigma_{\text{MAC}}) = \perp</math> <b>then</b> 7:   <b>return</b> <math>(st_u, \perp, \perp)</math> 8: <math>\text{PubTOKEN}_u[v][\text{ctr}_v] \leftarrow \perp</math> 9: <math>\text{TOKEN}_u \leftarrow (\text{PrivTOKEN}_u, \text{PubTOKEN}_u)</math> 10: <math>st'_u \leftarrow \left( G, k_{\text{MAC}}, \overline{k_{\text{MAC}}}, k_{\text{SKE}}, \text{ctr}, \text{TOKEN}_u \right)</math> 11: <b>return</b> <math>(st'_u, \top, v)</math> </pre>

Figure 6: Helper functions used by COSMAC. The main differences between COSMOS is highlighted by a box.

	$u_1$	$u_2$	$u_3$	$u_4$
$u_1$				
$u_2$	$x_{1 \rightarrow 2}^{(1)}$	$x_{2 \rightarrow 2}^{(2)}$	$x_{3 \rightarrow 2}^{(3)}$	$x_{4 \rightarrow 2}^{(4)}$
$u_3$		$x_{2 \rightarrow 3}^{(2)}$		
$u_4$		$x_{2 \rightarrow 4}^{(2)}$		

	$u_1$	$u_2$	$u_3$	$u_4$
$u_1$	$y_{1 \rightarrow 1}^{(1)}$			$y_{4 \rightarrow 1}^{(3)}$
$u_2$	$y_{1 \rightarrow 2}^{(1)}$	...	$y_{2 \rightarrow 2}^{(2)}$	...
$u_3$	$y_{1 \rightarrow 3}^{(1)}$		$y_{2 \rightarrow 3}^{(2)}$	$y_{4 \rightarrow 3}^{(3)}$
$u_4$	$y_{1 \rightarrow 4}^{(1)}$		$y_{2 \rightarrow 4}^{(2)}$	$y_{4 \rightarrow 4}^{(4)}$

Figure 7: A toy example of a non-anonymous variant of QUASAR.  $G = (u_i)_{i \in [4]}$  and  $T = 3$ . The upper box stores the private tokens  $(x_{i \rightarrow j}^{(t)})_{t \in [3]}$  for  $i, j \in [4]$ . The blue columns are tokens that user  $u_2$  uses to send messages and the red rows are tokens that user  $u_2$  minted. The bottom box stores the public tokens  $(y_{i \rightarrow j}^{(t)})_{t \in [3]}$  for  $i, j \in [4]$  held by the server.

$u_i$  approves  $u_j$  to send a message to him. User  $u_i$  sends the private tokens  $(x_{j \rightarrow i}^{(t)})_{t \in [T]}$  to  $u_j$  by encrypting it with a public-key encryption (PKE) scheme using  $u_j$ 's public key. Moreover,  $u_i$  uploads the public tokens  $(y_{j \rightarrow i}^{(t)})_{(j,t) \in [N] \times [T]}$  to the server. A toy example is provided in Fig. 7. Throughout the protocol, each user  $u_i$  maintains two types of tokens: *sender* tokens  $(x_{i \rightarrow j}^{(t)})_{(j,t) \in [N] \times [T]}$  (blue box in Fig. 7) and *receiver* tokens  $(x_{i \rightarrow j}^{(t)})_{(j,t) \in [N] \times [T]}$  (red box in Fig. 7). To send the  $i^*$ -th ( $i^* < T$ ) message  $m$ ,  $u_i$  retrieves the sender tokens  $(x_{i \rightarrow j}^{(i^*)})_{j \in [N]}$  and uploads this as the group authentication token  $\Sigma_G$  along with  $m$ . Similarly to COSMOS and COSMAC,  $\Sigma_G$  will include  $N$  MAC tags for each message tuple  $(x_{i \rightarrow j}^{(i^*)}, m)$ , binding the tokens to  $m$ . The server checks that  $\Sigma_G$  maps to a specific column of public tokens in its database. If so, it parses  $\Sigma_G$ , sets the user authentication token  $\sigma_j$  for user  $u_j$  as  $x_{i \rightarrow j}^{(i^*)}$  and the  $j$ -th MAC tag. User  $u_j$  can verify and trace  $\sigma_j$  back to  $u_i$  by searching through its receiver tokens. Once the users exhaust their tokens, they perform an update by minting new one-time tokens and distributing them to the group as done above. Note that this is where we require global state updates: once the boxes in Fig. 7 become empty, every user has to update in order to fill them back again.

Informally, user traceability holds since the group authentication token  $\Sigma_G$  corresponds to a unique column in the database held by the server. All honest users can use this unique column index to trace the same sender.

**Amortized Efficiency.** To reduce the communication cost of the offline phase and make the dependence on the ciphertext size minimal (particularly important in the post-quantum regime), we replace the private tokens  $(x_{j \rightarrow i}^{(t)})_{t \in [T]}$  by one PRF seed  $seed_{j \rightarrow i}$ , allowing each user to locally derive the corresponding tokens. This reduces the number of ciphertexts by a factor of  $T$ , making the overhead in the offline communica-

tion cost to be  $2 \cdot \left(\frac{\text{ct}}{T} + \kappa\right)$  per message, where  $\text{ct}$  denotes a PKE ciphertext and  $\kappa$  is the bit-length of the public tokens.

**Adding Anonymity.** Fig. 7 illustrates how if, e.g.,  $u_2$  sends two messages using private tokens from the blue box, the server will be able to link these two messages together. Our final description of QUASAR fixes this by permuting the column indices using a permutation key derived from the group secret key. The server stills checks the group authentication tokens  $\Sigma_G$  with respect to a single column, and the users can map this randomly permuted column index to a unique sender.

Similarly to all the GAM protocols so far, QUASAR is non-colluding unforgeable since each users maintain a database of the public token and user pair. The added complexity only comes from adding tracing soundness while maintaining anonymity. Moreover, QUASAR is not standard unforgeable as a malicious server and malicious insider can collude to impersonate any honest user.

## 5.2 STARS and GEMSTARS

**STARS:** This is almost equivalent to QUASAR, except that it additionally achieves *standard* unforgeability by replacing the usage of one-time tokens (i.e., private and public tokens  $(x, y)$  such that  $\text{OWF}(x) = y$ ) with *one-time signatures* (OTS). Importantly, the usage must remain *one-time* as otherwise two messages sent with the same signing key becomes linkable.

**GEMSTARS:** This is essentially STARS without state updates. The main idea is to use a *group signature* (GS) [13, 27, 35].<sup>10</sup> Informally, a GS consists of three entity types: a group manager, group tracers (also referred to as *opener*), and users. A group manager is unique and handles the registration of a new user to the system. When a user  $u$  wishes to join the system, the group manager provides  $u$  with a certificate  $\text{cert}_u$  attesting to the fact that  $u$  is a valid user in the system. User  $u$  can specify a group tracer  $I$  and use  $\text{cert}_u$  to *anonymously* sign a message on behalf of all the users in the system. The signature can be publicly verifiable, but importantly, it remains anonymous to any entity (including the group manager) except to the specified group tracer  $I$ , who can *trace* the signature back to user  $u$ .

An initial attempt to build a secure GAM protocol from a GS is as follows. The server and users of a GAM protocol are mapped to the group manager and users of a GS, respectively. Since the group users in  $G$  should be the only users that can trace the signature, we then would like to map each group  $G$  in the GAM protocol to a group tracer  $I$ . We achieve this by noticing that each group  $G$  shares a common group secret key  $\text{gsk}$ . Thus, the group generates the keys  $(\text{gtvk}_I, \text{gtsk}_I)$  for the group tracer  $I$  by executing the key generation algorithm of the group tracer using randomness derived from  $\text{gsk}$ . When a user sends a message to the group  $G$ , it runs the signing

<sup>10</sup>Since we consider multiple group tracers, we can view it as an accountable ring signature as well [18, 80].



algorithm of GS by specifying  $g\text{tv}k_j$  as the group tracer’s key. The server then relays the message to the group  $G$  only if the signature verifies. The group users can locally run the tracing algorithm of  $I$  on the signature since they hold the group tracer’s secret key  $g\text{ts}k_j$ .

This GAM protocol inherits the anonymity and unforgeability of GS. Unfortunately, it does not achieve anonymous blocklisting nor tracing soundness. The issue is that *any* user in the system, even those outside of the group  $G$ , can sign on behalf of  $G$  — this stems from the mismatch in the scope of “groups” considered by the GAM protocol and GS. In particular, an adversary can mount the following Sybil attack on the group: The adversary creates fake users  $u'$  and obtains credentials  $\text{cert}_{u'}$  by joining the system. It then specifies  $g\text{tv}k_j$  as the group tracer’s key and signs on behalf of the group  $G$ . While traceability of GS allows the group users to trace back to this outsider  $u'$ , the server cannot block  $u'$  from uploading malicious contents to the group since it is a valid group signature.

We resolve this issue by proving group membership in a different layer, similarly to the metadata-hiding protocol of [53] and what we did with COSMAC. Using the group secret key  $\text{gsk}$ , the group further generates a unique signature key pair  $(\text{vk}, \text{sk})$  and uploads  $\text{vk}$  to the server. To send a message to the group, a group user generates the group signature as before, but further attaches a signature generated by  $\text{sk}$ . Similarly to COSMAC, this additional signature guarantees anonymous blocklisting, while GS guarantees that we can trace the signer to a group user. Lastly, since GS allows to sign an unbounded number of messages, GEMSTARS requires no updates.

Although GEMSTARS removes state updates while satisfying all the desired security, it is quite inefficient in the post-quantum setting due to the lack of efficient GS.

## 6 Running GAM Protocols on MLS

In this section, we explain how to integrate a GAM protocol into MLS. While this is fairly straightforward, some discussion is required since our GAM protocol assumes a trusted initialization algorithm  $\text{Init}$  that prepares the group user states. For the particular GAM protocol (i.e., Enc-Sign mode) currently used by MLS this is not an issue, since the initial user state (i.e., the group secret key and verification keys of the group users) is implicitly provided by the CGKA protocol. However, in general, this may not be the case.

Below, we explain this integration in two steps. We first consider the base case where a GAM protocol operates on an already established set of user states. We then consider how our *specific* GAM protocols can handle the  $\text{Init}$  algorithm without a trusted setup. Recall in MLS, a single user  $u$  initializes a group  $G = \{u\}$  and then dynamically adds users by sending welcome messages. We follow this approach and explain how the  $\text{Init}$  algorithm of a GAM protocol can be implemented through these dynamic group operations. As

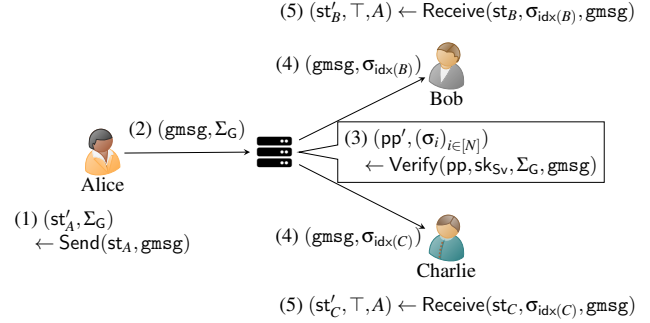


Figure 8: Using a GAM protocol on top of MLS’ FSPD protocol. When retrieving a message, a user  $u$  specifies its index  $\text{idx}(u)$  to the server; the server then returns the user authentication token  $\sigma_{\text{idx}(u)}$  along with  $\text{gmsg}$ .

explained in Sec. 2.4, a formal model for allowing dynamic groups in GAM protocols is left as an important future work.

### 6.1 Authentication in a Static Group

Assume the base case where the user states of the GAM protocol are already set up. Recall an application message of MLS’ FSPD protocol have the following format [12, Sec. 6]:

$$(\text{gid}, \text{epoch}_{\text{CGKA}}, \text{ct}_{\text{senderID}}, \text{ct}_{\text{Contents}}, \text{AuthData}) \quad (1)$$

$\text{gid}$  is the group identity;  $\text{epoch}_{\text{CGKA}}$  is a counter<sup>11</sup>;  $\text{senderID}$  is the sender’s identity;  $\text{Contents}$  stores the payload;  $\text{ct}_X$  is an SKE encryption of  $X$  under an SKE key derived from the group secret key; and  $\text{AuthData}$  is an authentication data field, including an encryption of the signature (i.e., Enc-Sign mode). Below, we denote  $\text{gmsg}$  as Eq. (1), excluding  $\text{AuthData}$ .

With this structure in mind, using a GAM protocol on top of MLS’ FSPD is straightforward. This is depicted in Fig. 8 (see also Fig. 1). To send a message  $\text{gmsg}$ , user  $u \in G$  runs the  $\text{Send}$  algorithm to create a group authentication token  $\Sigma_G$  and uploads  $(\text{gmsg}, \text{AuthData} := \Sigma_G)$  to the server. The server verifies  $\Sigma_G$  and prepares *user* authentication tokens  $(\sigma_i)_{i \in [N]}$ , where  $N = |G|$ . When a user  $v \in G$  with index  $i = \text{idx}(v)$  contacts the server, the server returns  $(\text{gmsg}, \text{AuthData}_i := \sigma_i)$ . User  $v$  can then verify and trace user  $u$  by running the  $\text{Receive}$  algorithm of the GAM protocol on  $(\text{gmsg}, \sigma_i)$ .

Lastly, in case the GAM protocol requires state updates (e.g., QUASAR and STARS), this can be simply run on top of FSPD by directly embedding the group update information  $\hat{\Sigma}_G$  into  $\text{Contents}$  or  $\text{ct}_{\text{Contents}}$  depending on the anonymity guarantee we require.

<sup>11</sup>Note that this  $\text{epoch}_{\text{CGKA}}$  is maintained by the CGKA protocol and differs from those used by QUASAR and STARS.

## 6.2 Authentication in a Dynamic Group

We now discuss how to add users to a preexisting GAM protocol using the welcome message functionality provided by MLS (or the CGKA protocol to be precise). Similarly to how the CGKA protocol is implemented, this procedure can be used to execute the Init algorithm. Since the process is inherently protocol specific, we explain them individually below.

**COSMOS.** This is simple since each user state is independent of the others. When an outsider joins the group via a welcome message, it mints a new token  $(x_0, y_0)$  and uploads the public token  $y_0$  to the server and group. The outsider can additionally authenticate  $y_0$  by including it in its key package maintained by the Delivery Service [20, Sec. 5] or by directly signing it to the group with its long-term key maintained by the Authentication Service [20, Sec. 4]. Once  $y_0$  is shared among the group, it can use  $x_0$  to start sending messages. Moreover, the outsider can fetch the current public tokens of the other users from the server. In case they want to be sure that these public tokens were indeed minted by the respective users, they can obtain the tokens directly from the senders. Note that the process is very similar to the Enc-Sign mode currently used by MLS’ FSPD protocol. The case for the optimized COSMOS is similar.

**COSMAC.** This is almost identical to COSMOS. The only difference is that the outsider, after processing the welcome message, recovers the current group secret key  $gsk$ . It then uses the  $gsk$  to derive the MAC key  $k_{MAC}$  and an SKE key  $k_{SKE}$ . Note that we can update  $k_{MAC}$  and  $k_{SKE}$  anytime the CGKA protocol performs a commit, which would allow some form of post-compromise security of the anonymity and anonymous blocklisting properties (see Sec. 8). Moreover, if we need sender anonymity of the welcome message, we can rely on existing anonymous *two-party* messaging protocols, such as Sealed Sender [60] or Orca [75], where the latter protocol also provides user traceability.

**QUASAR and STARS.** These are the most involved due to global state updates. Since the protocol flows of QUASAR and STARS are identical, we only focus on QUASAR. Looking at the toy example from Fig. 7, we cannot simply append columns corresponding to the new users since the server can always link public tokens associated to these appended columns. Thus, to ensure anonymity, the group users must all update their state and refresh the public tokens held by the server. In more detail and similarly to COSMOS, the outsider  $o$ , after having processed the welcome message, mints their one-time tokens, uploads the public tokens  $(y_{j \rightarrow o}^{(t)})_{(j,t) \in [M] \times [T]}$  to the server, and sends the PRF seeds  $seed_{j \rightarrow o}$  to each group user so that they can recover the corresponding private tokens. After every other user updates their states by further minting  $T$  extra public tokens for the new user  $o$ , the user can start sending messages.

An optimization of the above approach will have each

group member  $j$  send a new PRF seed  $seed_{o \rightarrow j}$  to  $o$ , but, for each other group member  $i$ , derive the corresponding seed for the new epoch from the seed  $seed_{j \rightarrow i}$  (respectively  $seed_{i \rightarrow j}$ ) for the previous epoch, and upload the corresponding public tokens  $y_{* \rightarrow j}^{(t)}$  to the server. This could be done by setting the new seed to be  $OWF(seed_{j \rightarrow i} || epoch || o)$  (respectively  $OWF(seed_{i \rightarrow j} || epoch || o)$ ). The advantage of such an approach is that each group member is only required to send a single ciphertext, and download nothing, as opposed to uploading and downloading  $N$  ciphertexts.

**GEMSTARS.** This is the only protocol that requires the MLS server to additionally run a group signature scheme. When a user joins the secure group messaging application, the server provides the user with a certificate (see Sec. 5.2 for details). Assuming this step has been finished, then joining a group is straightforward. This is because the only thing the outsider requires to run GEMSTARS is the groups tracer key and signature key, which are both derived from the group secret key  $gsk$ . Hence, following the same discussion in COSMAC, we can easily add new users to GEMSTARS.

Lastly, we note that removing users is straightforward for all of our GAM protocols. Since COSMOS is not anonymous, the server can simply maintain a list of group members. COSMAC and GEMSTARS natively supports removal at the CGKA layer since, once a commit occurs, the MAC key  $k_{MAC}$  is updated along with the group secret key  $gsk$ . Finally, for QUASAR, the users can simply remove the unused public tokens corresponding to the removed users from the server.

## 7 Bandwidth Efficiency Analysis

In this section we analyze the efficiency of our proposed GAM protocols and compare them with existing authentication modes. We are specifically interested in the bandwidth overhead incurred by each authentication mode compared to a messaging protocol where the application message (e.g., chat texts) is sent in the clear, i.e., without authentication.

### 7.1 Instantiation

We target the NIST Level I security<sup>12</sup> stating that breaking the protocol is no easier than key-recovery on a block cipher with a 128-bit key (e.g., AES-128). This provides a meaningful baseline to discuss post-quantum security and ignoring quantum attacks corresponds to a classical security level of 128 bits. The main cryptographic primitives used in our GAM protocols is summarized below. The sizes of the cryptographic artifacts used in our instantiations are shown in Tab. 2.

**OWF:** We use SHA-256 and truncate its output to 16 B.

<sup>12</sup>[https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/evaluation-criteria/security-\(evaluation-criteria\)](https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/evaluation-criteria/security-(evaluation-criteria))

Table 2: Instantiation of the building blocks for our GAM protocols. SIG and GS stand for signature schemes and group signatures. sig, osig, and gsig denote the signature a standard signature, an OTS, and a group signature, respectively. ovk denotes the verification key of an OTS. ek and ct denote a KEM encapsulation key and ciphertext, respectively. All sizes are given in bytes. † indicates the output of SHA-256 is truncated to 16 B. We use  $\kappa$  to denote the security parameter, set to 128 bits.

Primitives	Classic	Post-quantum	Related Auth. Modes
OWF	SHA-256: $\kappa = 16^\dagger$	SHA-256: $\kappa = 16^\dagger$	COSMOS, COSMAC, QUASAR
MAC	HMAC-SHA-256: $\kappa = 16^\dagger$	HMAC-SHA-256: $\kappa = 16^\dagger$	COSMAC
OTS	EdDSA: $ \text{ovk}  = 32,  \text{osig}  = 64$	WOTS <sup>+</sup> : $ \text{ovk}  = 1320,  \text{osig}  = 1072$	STARS
SIG	EdDSA: $ \text{sig}  = 64$	Dilithium: $ \text{sig}  = 2420$	Sign, Enc-Sign, Sign-Enc-Sign, GEMSTARS
GS	BBS: $ \text{gsig}  = 336$	LNP: $ \text{gsig}  = 9.2 \times 10^4$	GEMSTARS
KEM	Hashed ElGamal: $ \text{ek}  =  \text{ct}  = 32$	Kyber: $ \text{ek}  = 800,  \text{ct}  = 768$	QUASAR, STARS

**MAC and PRF:** We use HMAC [58] with SHA-256 and truncate its output to 16 B. Note that a deterministic MAC can be viewed as a PRF.

**Pseudo-random permutation:** We require a PRP to permute the set of  $NT$  tokens. We have the option of using either FastPRP [73] or the Thorp shuffle [67].

**One-time signature schemes:** For classical security, we use EdDSA [17]. For post-quantum security, we use WOTS<sup>+</sup> [54, 55] used as a building block of the NIST PQC standard signature SPHINCS+ [56]. We set the Winternitz parameter  $w = 16$ , and use SHA-256 as the underlying hash function.

**Signature schemes:** For classical security, we use EdDSA [17]. For post-quantum security, we use the NIST PQC standard Dilithium [61]. We did not consider Falcon [70], another PQC standard, since Dilithium is selected as the primary algorithm and NIST recommends it for most use cases.

**Group signatures:** For classical security, we use the pairing-based BBS scheme [26] with the BLS12-381 pairing-friendly curve. For post-quantum security, we use the lattice-based scheme proposed by Lyubashevsky, Nguyen, and Plançon (LNP) [62].

**KEM schemes:** For classical and post-quantum security, we use the Hashed ElGamal KEM [40] and the NIST PQC standard Kyber [72], respectively.

## 7.2 Efficiency

**Cost Metric.** Following [7, 53], analyzing the bandwidth efficiency of MLS (and its variants), we analyze our GAM protocol through three metrics: the upload and download cost, and the total cost. Each of these costs are further broken down into offline and online costs; online (resp. offline) cost is associated to the cost of uploading and downloading contents generated by Send and Verify (resp. UpdSend and UpdVerify).

We assume a user can send at most  $T$  messages once their states are updated. For protocols that require no updates, we can simply set  $T = 1$  as there are no offline costs. In more detail, we have the following, where the costs are defined per user in a group of size  $N$ .

*Total upload cost:* The cost of uploading  $T$  outputs of Send (online) and one output of UpdSend (offline).

*Total download cost:* The cost of downloading  $NT$  outputs of Send (online) and  $N$  outputs of UpdSend (offline).

*Total cost:* The sum of the upload and download costs.

The download cost is a factor  $N$  times larger than the upload cost since each user has  $N = |G|$  users to download from. Here, for COSMOS<sup>++13</sup>, COSMAC<sup>++</sup>, and GEMSTARS, we can slightly optimize the download cost by allowing the users to not download the messages they upload. In contrast, for QUASAR and STARS, the users *must* download what they uploaded to preserve anonymity. At a high level, looking at Fig. 7, if a user does not download what it uploaded, then the server can link the (permuted) columns together. Lastly, observe that even if the offline cost is larger compared to the online cost, it gets amortized by  $T$ : as  $T$  grows larger, the total online cost starts to dominate the total offline cost.

**Comparison of Communication Costs.** We analyze the communication cost of our proposed GAM protocols and compare them with existing authentication modes. The number of exchanged cryptographic elements in the GAM protocols is summarized in Tab. 3. We classify the GAM protocols into the following three categories and compare them.

- (1) Non-anonymous protocols: Sign mode and COSMOS<sup>++</sup>.
- (2) Anonymous protocols without tracing soundness: Enc-Sign mode, Sign-Enc-Sign mode, and COSMAC<sup>++</sup>.
- (3) Anonymous protocols with tracing soundness: QUASAR, STARS, and GEMSTARS.

Table 3: The number of total cryptographic elements exchanged in GAM protocols.  $N$  is the group size and  $T$  is the number of online messages per one offline update.  $\kappa$  denotes the security parameter. sig, osig, and gsig denote a standard signature, a one-time signature, and a group signature, respectively. ovk denotes the verification key of a one-time signature. ct denotes a KEM ciphertext.

Auth. Mode	Offline						Online							
	Upload			Download			Upload			Download				
	$\kappa$	ovk	ct	$\kappa$	ovk	ct	$\kappa$	osig	sig	gsig	$\kappa$	osig	sig	gsig
Enc-Sign (MLS)								$T$						$(N-1)T$
Sign-Enc-Sign [53]								$T$						$(N-1)T$
COSMOS <sup>++</sup>	3			$3(N-1)$			$2T$				$2(N-1)T$			
COSMAC <sup>++</sup>	4			$4(N-1)$			$3T$				$3(N-1)T$			
QUASAR	$2NT$		$N$	$2N$		$N$	$2NT$				$2NT$			
STARS		$NT$	$N$	$N$		$N$		$NT$				$NT$		
GEMSTARS								$T$	$T$					$(N-1)T$ $(N-1)T$

Table 4: Communication cost of each GAM protocols. The sizes are in bytes.  $N$  is the group size and  $T$  is the number of online messages per one offline update. In Tabs. 4a and 4b, we ignore the  $O(N)$  cost of the offline phase for readability. The column “Total” is normalized by  $NT$ , denoting the total overhead cost per message. The column “PQ?” is  $\times$  (resp.  $\checkmark$ ) for classical (resp. post-quantum) security. The mode “Sign” in Tab. 4a corresponds to the naïve approach of simply signing messages.

a: Non-anonymous GAM protocols

Auth. Mode	Online		Total	PQ?
	Upload	Download		
Sign	$64 \cdot T$	$64 \cdot (N-1)T$	64	$\times$
	$2420 \cdot T$	$2420 \cdot (N-1)T$	2420	$\checkmark$
COSMOS <sup>++</sup>	$32 \cdot T$	$32 \cdot (N-1)T$	$16 \cdot (2 + \frac{3}{T})$	$\checkmark$

b: Anonymous GAM protocols without tracing soundness

Auth. Mode	Online		Total	PQ?
	Upload	Download		
Enc-Sign (MLS)	$64 \cdot T$	$64 \cdot (N-1)T$	64	$\times$
	$2420 \cdot T$	$2420 \cdot (N-1)T$	2420	$\checkmark$
Sign-Enc-Sign [53]	$128 \cdot T$	$128 \cdot (N-1)T$	128	$\times$
	$4840 \cdot T$	$4840 \cdot (N-1)T$	4840	$\checkmark$
COSMAC <sup>++</sup>	$48 \cdot T$	$48 \cdot (N-1)T$	$16 \cdot (3 + \frac{4}{T})$	$\checkmark$

c: Anonymous GAM protocols with tracing soundness

Auth. Mode	Offline		Online		Total	PQ?
	Upload	Download	Upload	Download		
QUASAR	$32 \cdot NT + 32 \cdot N$	$64 \cdot N$	$64 \cdot NT$	$32 \cdot NT$	$(96 + \frac{96}{T}) \cdot \frac{N+1}{N}$	$\times$
	$32 \cdot NT + 768 \cdot N$	$800 \cdot N$	$32 \cdot NT$	$32 \cdot NT$	$(96 + \frac{1568}{T}) \cdot \frac{N+1}{N}$	$\checkmark$
STARS	$32 \cdot NT + 32 \cdot N$	$48 \cdot N$	$64 \cdot NT$	$64 \cdot NT$	$(160 + \frac{80}{T}) \cdot \frac{N+1}{N}$	$\times$
	$1320 \cdot NT + 768 \cdot N$	$784 \cdot N$	$1072 \cdot NT$	$1072 \cdot NT$	$(3464 + \frac{1552}{T}) \cdot \frac{N+1}{N}$	$\checkmark$
GEMSTARS			$400 \cdot T$	$400 \cdot (N-1)T$	400	$\times$
			$9.4 \times 10^4 \cdot T$	$9.4 \times 10^4 \cdot (N-1)T$	$9.4 \times 10^4$	$\checkmark$

*Category (1).* Tab. 4a compares the Sign mode (cf. Footnote 3) used in MLS and COSMOS<sup>++</sup>. Technically, the Sign mode is used exclusively by the CGKA protocol in MLS, and not used to authenticate the output of the FSPD protocol. Nonetheless, this mode was used to analyze MLS by Alwen et al. [5] and we view it as the vanilla non-anonymous GAM protocol. Considering that any cryptographic element added to satisfy authentication should be at least 128-bits, COSMOS<sup>++</sup> is near optimal. Compared to the Sign mode, the total post-quantum communication cost is a factor 75x smaller. On the other hand, Sign mode achieves *standard* unforgeability while COSMOS<sup>++</sup> does not. We believe COSMOS<sup>++</sup> offers a worthwhile tradeoff between efficiency and security.

<sup>13</sup>We only consider the most efficient version of COSMOS and COSMAC here.

*Category (2).* Tab. 4b compares the Enc-Sign mode used in MLS, Sign-Enc-Sign mode [53], and COSMAC<sup>++</sup>. Recall Enc-Sign mode does not achieve anonymous blocklisting while Sign-Enc-Sign and COSMAC<sup>++</sup> do (see Tab. 1). Out of the three protocols, COSMAC<sup>++</sup> has the lowest communication cost. Compared to the Enc-Sign mode used in MLS, the total post-quantum communication cost is a factor 50x smaller. As with COSMOS<sup>++</sup>, COSMAC<sup>++</sup> only achieves non-colluding unforgeability while Enc-Sign and Sign-Enc-Sign modes do.

*Category (3).* Tab. 4c compares QUASAR, STARS, and GEMSTARS. These are the only anonymous GAM protocols with tracing soundness. QUASAR and STARS have a variable total communication cost that becomes smaller as  $T$  (and  $N$ ) increases. This is because the KEM ciphertext encrypting the

PRF seed, exchanged during the offline phase, can be used to mint  $T$  tokens. Specifically, the cost of sending a large ciphertext is amortized by the number of messages  $T$  sent in the online phase. By setting  $T = 1000$ , the cost of sending a KEM ciphertext relative to the total cost is only 2 B or less per message, even in the post-quantum setting.

Out of the three protocols, QUASAR provides the most total-cost-efficient protocol. In fact, QUASAR is even comparable to COSMAC<sup>++</sup> that has no tracing soundness, e.g., it is 106 B when  $(N, T) = (10, 1000)$ . While larger than QUASAR, STARS also offers a relatively small total overhead, albeit more computationally expensive due to running an OTS. The benefit of using STARS over QUASAR is that it achieves *standard* unforgeability. Lastly, while both QUASAR and STARS have an  $O(N)$  online upload cost (i.e., maximum bandwidth consumption) per message, the concrete cost is only 16 KB for QUASAR even for a relatively large group of  $N = 1024$ . STARS uploads 64 KB and 1 MB of data in the classical and post-quantum settings, respectively. Lastly, recall one of the weaknesses of QUASAR and STARS are that they are only *globally* state-updatable. GEMSTARS removes updates altogether, with the cost of a larger total communication overhead; in the post-quantum setting, it is 94 KB.

## 8 Open Problems and Future Work

Other than those discussed in Sec. 2.4, we consider the following as interesting future work.

**FS and PCS.** Both forward secrecy (FS) and post-compromise security (PCS) are standard security notions in secure messaging. A natural question is then, *given the compromise of (one or more) users states to the adversary, what is the effect of this on the unforgeability, anonymity, anonymous blocklisting, and user traceability of past and future messages?* This opens interesting directions, both towards formalizing these notions and towards constructing authentication modes satisfying them.

Regarding FS, we first note that unforgeability, anonymous blocklisting, and user traceability are not relevant, as in our setting these notions are only concerned with the moment messages are processed by users.<sup>14</sup> Anonymity, on the other hand, is more interesting: can a state compromise allow an adversary to de-anonymize past messages from a user? The answer for both MLS and our proposals is “Yes”, at least in some cases. Indeed, messages in each MLS’ FSPD instantiation share an epoch and are thus all signed with the same signing key, and their sender identity and signature are encrypted with the shared secret in the epoch. In the latter, either key material is static (like in GEMSTARS), or anonymity relies on the key material that only gets rotated between CGKA

<sup>14</sup>The concept of *forward-secret signatures* [14], motivated by the will that the compromise of the current secret key does not enable an adversary to forge signatures pertaining to the past, is thus not relevant.

epochs, like the MAC key used in COSMAC, or the permutation key used in QUASAR or STARS. Thus, natural questions are: *how do we formalize “forward anonymity”?* and *can we design authentication modes that satisfy it?*

The matter regarding PCS for authentication is more involved, since all the security notions make sense in this setting, as indicated by the original work on PCS by Cohn-Gordon, Cremers, and Garratt [39]. For unforgeability, the work of Cremers, Hale, and Kohbrok [41] introduces the notion of PCS signatures, where key-pairs can be evolved to “heal” from a compromise. For anonymity, ideas from unlinkable sanitizable signatures [30, 45] could be useful. We leave concrete construction of a PCS GAM protocol as an interesting problem.

**Optimal Security with PQ Efficiency.** We provide several GAM protocols with different efficiency and security profiles, some of which offering much better post-quantum efficiency compared to the GAM protocol used in MLS, albeit weakening unforgeability. So far, the only GAM protocol satisfying optimal security (i.e., standard unforgeability, anonymity, anonymous blocklisting, tracing soundness) with no global state updates is GEMSTARS. However, this comes at a great cost as post-quantum group signatures are much more costly than signatures. We view it as an interesting open problem to find a GAM protocol achieving all the desirable properties while retaining efficiency.

## Acknowledgments

The authors thank the anonymous CCS and USENIX reviewers. Their reviews really helped us make our paper better. This research was partially supported by JST CREST JP-MJCR22M1, Japan and funded by the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No.665385.

## Ethical Considerations

Our work follows the ethical guidelines of the conference. It proposes improvements for a popular and widely studied IETF standard: MLS. Most of the improvements target efficiency, and the formalization of authenticity notions, while new, does not allow for the creation of non-previously-known primitives that could be seen as bringing controversial or negative consequences, i.e., anonymous (the only notion that could be seen as potentially having a negative outcome) messaging systems already exist and, in particular, anonymity on the FSPD layer is already an aim of MLS. All the other authenticity notions enabled by this work are hard to imagine to have a negative outcome if implemented in the real world. Finally, our work did not include any experiments with live systems and does not lead to negative results for already used or implemented systems.

## Open Science Policy

We do not have any artifacts (e.g., datasets, scripts, binaries) related to this paper.

## References

- [1] Joël Alwen, Benedikt Auerbach, Miguel Cueto Noval, Karen Klein, Guillermo Pascual-Perez, and Krzysztof Pietrzak. DeCAF: Decentralizable continuous group key agreement with fast healing. Cryptology ePrint Archive, Report 2022/559, 2022. <https://eprint.iacr.org/2022/559>.
- [2] Joël Alwen, Benedikt Auerbach, Miguel Cueto Noval, Karen Klein, Guillermo Pascual-Perez, Krzysztof Pietrzak, and Michael Walter. CoCoA: Concurrent continuous group key agreement. In Dunkelman and Dziembowski [44], pages 815–844.
- [3] Joël Alwen, Sandro Coretti, and Yevgeniy Dodis. The double ratchet: Security notions, proofs, and modularization for the Signal protocol. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 129–158. Springer, Heidelberg, May 2019.
- [4] Joël Alwen, Sandro Coretti, Yevgeniy Dodis, and Yiannis Tselekounis. Security analysis and improvements for the IETF MLS standard for group messaging. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part I*, volume 12170 of *LNCS*, pages 248–277. Springer, Heidelberg, August 2020.
- [5] Joël Alwen, Sandro Coretti, Yevgeniy Dodis, and Yiannis Tselekounis. Modular design of secure group messaging protocols and the security of MLS. In Vigna and Shi [76], pages 1463–1483.
- [6] Joël Alwen, Sandro Coretti, Daniel Jost, and Marta Mularczyk. Continuous group key agreement with active security. In Pass and Pietrzak [68], pages 261–290.
- [7] Joël Alwen, Dominik Hartmann, Eike Kiltz, and Marta Mularczyk. Server-aided continuous group key agreement. In Yin et al. [81], pages 69–82.
- [8] Joël Alwen, Daniel Jost, and Marta Mularczyk. On the insider security of MLS. In Dodis and Shrimpton [43], pages 34–68.
- [9] Michael Anastos, Benedikt Auerbach, Mirza Ahad Baig, Miguel Cueto Noval, Matthew Kwan, Guillermo Pascual-Perez, and Krzysztof Pietrzak. The cost of maintaining keys in dynamic groups with applications to multicast encryption and group messaging. In Elette Boyle and Mohammad Mahmoudy, editors, *Theory of Cryptography*, pages 413–443, Cham, 2025. Springer Nature Switzerland.
- [10] Benedikt Auerbach, Miguel Cueto Noval, Guillermo Pascual-Perez, and Krzysztof Pietrzak. On the Cost of Post-compromise Security in Concurrent Continuous Group-Key Agreement. In Guy Rothblum and Hoeteck Wee, editors, *Theory of Cryptography*, pages 271–300, Cham, 2023. Springer Nature Switzerland.
- [11] David Balbás, Daniel Collins, and Serge Vaudenay. Cryptographic administration for secure group messaging. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 1253–1270, Anaheim, CA, August 2023. USENIX Association.
- [12] Richard Barnes, Benjamin Beurdouche, Raphael Robert, Jon Millican, Emad Omara, and Katriel Cohn-Gordon. The Messaging Layer Security (MLS) Protocol. RFC 9420, July 2023.
- [13] Mihir Bellare, Daniele Micciancio, and Bogdan Warinschi. Foundations of group signatures: Formal definitions, simplified requirements, and a construction based on general assumptions. In Eli Biham, editor, *EUROCRYPT 2003*, volume 2656 of *LNCS*, pages 614–629. Springer, Heidelberg, May 2003.
- [14] Mihir Bellare and Sara K. Miner. A forward-secure digital signature scheme. In Michael J. Wiener, editor, *CRYPTO '99*, volume 1666 of *LNCS*, pages 431–448. Springer, Heidelberg, August 1999.
- [15] Mihir Bellare and Phillip Rogaway. Entity authentication and key distribution. In Douglas R. Stinson, editor, *CRYPTO '93*, volume 773 of *LNCS*, pages 232–249. Springer, Heidelberg, August 1994.
- [16] Mihir Bellare, Haixia Shi, and Chong Zhang. Foundations of group signatures: The case of dynamic groups. In Alfred Menezes, editor, *CT-RSA 2005*, volume 3376 of *LNCS*, pages 136–153. Springer, Heidelberg, February 2005.
- [17] Daniel J. Bernstein, Niels Duif, Tanja Lange, Peter Schwabe, and Bo-Yin Yang. High-speed high-security signatures. In Bart Preneel and Tsuyoshi Takagi, editors, *CHES 2011*, volume 6917 of *LNCS*, pages 124–142. Springer, Heidelberg, September / October 2011.
- [18] Ward Beullens, Samuel Dobson, Shuichi Katsumata, Yi-Fu Lai, and Federico Pintore. Group signatures and more from isogenies and lattices: Generic, simple, and efficient. In Dunkelman and Dziembowski [44], pages 95–126.

- [19] Benjamin Beurdouche, Eric Rescorla, Emad Omara, Srinivas Inguva, and Alan Duric. The Messaging Layer Security (MLS) Architecture. Internet-Draft draft-ietf-mls-architecture-10, Internet Engineering Task Force, December 2022. Work in Progress.
- [20] Benjamin Beurdouche, Eric Rescorla, Emad Omara, Srinivas Inguva, and Alan Duric. The Messaging Layer Security (MLS) Architecture. Internet-Draft draft-ietf-mls-architecture-15, Internet Engineering Task Force, August 2024. Work in Progress.
- [21] Karthikeyan Bhargavan, Richard Barnes, and Eric Rescorla. TreeKEM: Asynchronous Decentralized Key Management for Large Dynamic Groups A protocol proposal for Messaging Layer Security (MLS). Research report, Inria Paris, May 2018.
- [22] Karthikeyan Bhargavan, Benjamin Beurdouche, and Prasad Naldurg. Formal Models and Verified Protocols for Group Messaging: Attacks and Proofs for IETF MLS. Research report, Inria Paris, December 2019.
- [23] Alexander Bienstock, Yevgeniy Dodis, Sanjam Garg, Garrison Grogan, Mohammad Hajiabadi, and Paul Rösler. On the worst-case inefficiency of CGKA. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022, Part II*, volume 13748 of *LNCS*, pages 213–243. Springer, Heidelberg, November 2022.
- [24] Alexander Bienstock, Yevgeniy Dodis, and Paul Rösler. On the price of concurrency in group ratcheting protocols. In Pass and Pietrzak [68], pages 198–228.
- [25] Alexander Bienstock, Jaiden Fairoze, Sanjam Garg, Pratyay Mukherjee, and Srinivasan Raghuraman. A more complete analysis of the Signal double ratchet algorithm. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 784–813. Springer, Heidelberg, August 2022.
- [26] Dan Boneh, Xavier Boyen, and Hovav Shacham. Short group signatures. In Matthew Franklin, editor, *CRYPTO 2004*, volume 3152 of *LNCS*, pages 41–55. Springer, Heidelberg, August 2004.
- [27] Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Essam Ghadafi, and Jens Groth. Foundations of fully dynamic group signatures. In Mark Manulis, Ahmad-Reza Sadeghi, and Steve Schneider, editors, *ACNS 16*, volume 9696 of *LNCS*, pages 117–136. Springer, Heidelberg, June 2016.
- [28] Jacqueline Brendel, Rune Fiedler, Felix Günther, Christian Janson, and Douglas Stebila. Post-quantum asynchronous deniable key exchange and the signal handshake. In Goichiro Hanaoka, Junji Shikata, and Yohei Watanabe, editors, *PKC 2022, Part II*, volume 13178 of *LNCS*, pages 3–34. Springer, Heidelberg, March 2022.
- [29] Chris Brzuska, Eric Cornelissen, and Konrad Kohbrok. Security analysis of the MLS key derivation. In *2022 IEEE Symposium on Security and Privacy*, pages 2535–2553. IEEE Computer Society Press, May 2022.
- [30] Christina Brzuska, Marc Fischlin, Anja Lehmann, and Dominique Schröder. Unlinkability of sanitizable signatures. In Phong Q. Nguyen and David Pointcheval, editors, *PKC 2010*, volume 6056 of *LNCS*, pages 444–461. Springer, Heidelberg, May 2010.
- [31] Ran Canetti, Palak Jain, Marika Swanberg, and Mayank Varia. Universally composable end-to-end secure messaging. In Dodis and Shrimpton [43], pages 3–33.
- [32] Ran Canetti and Hugo Krawczyk. Security analysis of IKE’s signature-based key-exchange protocol. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 143–161. Springer, Heidelberg, August 2002. <https://eprint.iacr.org/2002/120/>.
- [33] Melissa Chase, Sarah Meiklejohn, and Greg Zaverucha. Algebraic MACs and keyed-verification anonymous credentials. In Gail-Joon Ahn, Moti Yung, and Ninghui Li, editors, *ACM CCS 2014*, pages 1205–1216. ACM Press, November 2014.
- [34] Melissa Chase, Trevor Perrin, and Greg Zaverucha. The signal private group system and anonymous credentials supporting efficient verifiable encryption. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020*, pages 1445–1459. ACM Press, November 2020.
- [35] David Chaum and Eugène van Heyst. Group signatures. In Donald W. Davies, editor, *EUROCRYPT’91*, volume 547 of *LNCS*, pages 257–265. Springer, Heidelberg, April 1991.
- [36] Katriel Cohn-Gordon, Cas Cremers, Benjamin Dowling, Luke Garratt, and Douglas Stebila. A formal security analysis of the signal messaging protocol. In *IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 451–466, 2017.
- [37] Katriel Cohn-Gordon, Cas Cremers, Benjamin Dowling, Luke Garratt, and Douglas Stebila. A formal security analysis of the signal messaging protocol. *Journal of Cryptology*, 33(4):1914–1983, October 2020.
- [38] Katriel Cohn-Gordon, Cas Cremers, Luke Garratt, Jon Millican, and Kevin Milner. On ends-to-ends encryption: Asynchronous group messaging with strong security guarantees. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018*, pages 1802–1819. ACM Press, October 2018.

- [39] Katriel Cohn-Gordon, Cas J. F. Cremers, and Luke Garratt. On post-compromise security. In Michael Hicks and Boris Köpf, editors, *CSF 2016 Computer Security Foundations Symposium*, pages 164–178. IEEE Computer Society Press, 2016.
- [40] Ronald Cramer and Victor Shoup. Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. *SIAM Journal on Computing*, 33(1):167–226, 2003.
- [41] Cas Cremers, Britta Hale, and Konrad Kohbrok. The complexities of healing in secure group messaging: Why cross-group effects matter. In Michael Bailey and Rachel Greenstadt, editors, *USENIX Security 2021*, pages 1847–1864. USENIX Association, August 2021.
- [42] Roger Dingledine, Nick Mathewson, and Paul F. Syverson. Tor: The second-generation onion router. In Matt Blaze, editor, *USENIX Security 2004*, pages 303–320. USENIX Association, August 2004.
- [43] Yevgeniy Dodis and Thomas Shrimpton, editors. *CRYPTO 2022, Part II*, volume 13508 of *LNCS*. Springer, Heidelberg, August 2022.
- [44] Orr Dunkelman and Stefan Dziembowski, editors. *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*. Springer, Heidelberg, May / June 2022.
- [45] Nils Fleischhacker, Johannes Krupp, Giulio Malavolta, Jonas Schneider, Dominique Schröder, and Mark Simkin. Efficient unlinkable sanitizable signatures from signatures with re-randomizable keys. In Chen-Mou Cheng, Kai-Min Chung, Giuseppe Persiano, and Bo-Yin Yang, editors, *PKC 2016, Part I*, volume 9614 of *LNCS*, pages 301–330. Springer, Heidelberg, March 2016.
- [46] Signal Foundation. Signal protocol: Technical documentation. (Accessed on 04/25/2023).
- [47] GreenNet. Understanding file sizes. <https://www.greenet.org.uk/support/understanding-file-sizes>. (Accessed on 04/25/2023).
- [48] Neil M. Haller. The S/KEY one-time password system. In *Proceedings of the ISOC Symposium on Network and Distributed System Security*, February 1994.
- [49] Neil M. Haller. The S/KEY One-Time Password System. RFC 1760, February 1995.
- [50] Keitaro Hashimoto, Shuichi Katsumata, Kris Kwiatkowski, and Thomas Prest. An efficient and generic construction for signal’s handshake (X3DH): Post-quantum, state leakage secure, and deniable. In Juan Garay, editor, *PKC 2021, Part II*, volume 12711 of *LNCS*, pages 410–440. Springer, Heidelberg, May 2021.
- [51] Keitaro Hashimoto, Shuichi Katsumata, Kris Kwiatkowski, and Thomas Prest. An efficient and generic construction for signal’s handshake (X3DH): Post-quantum, state leakage secure, and deniable. *Journal of Cryptology*, 35(3):17, July 2022.
- [52] Keitaro Hashimoto, Shuichi Katsumata, Eamonn Postlethwaite, Thomas Prest, and Bas Westerbaan. A concrete treatment of efficient continuous group key agreement via multi-recipient PKEs. In Vigna and Shi [76], pages 1441–1462.
- [53] Keitaro Hashimoto, Shuichi Katsumata, and Thomas Prest. How to hide MetaData in MLS-like secure group messaging: Simple, modular, and post-quantum. In Yin et al. [81], pages 1399–1412.
- [54] Andreas Hülsing. W-OTS+ - shorter signatures for hash-based signature schemes. In Amr Youssef, Abderrahmane Nitaj, and Aboul Ella Hassanien, editors, *AFRICACRYPT 13*, volume 7918 of *LNCS*, pages 173–188. Springer, Heidelberg, June 2013.
- [55] Andreas Hülsing. WOTS+ – shorter signatures for hash-based signature schemes. Cryptology ePrint Archive, Report 2017/965, 2017. <https://eprint.iacr.org/2017/965>.
- [56] Andreas Hülsing, Daniel J. Bernstein, Christoph Dobraunig, Maria Eichlseder, Scott Fluhrer, Stefan-Lukas Gazdag, Panos Kampanakis, Stefan Kolbl, Tanja Lange, Martin M Lauridsen, Florian Mendel, Ruben Niederhagen, Christian Rechberger, Joost Rijneveld, Peter Schwabe, Jean-Philippe Aumasson, Bas Westerbaan, and Ward Beullens. SPHINCS+. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>.
- [57] Karen Klein, Guillermo Pascual-Perez, Michael Walter, Chethan Kamath, Margarita Capretto, Miguel Cueto, Iliia Markov, Michelle Yeo, Joël Alwen, and Krzysztof Pietrzak. Keep the dirt: Tainted TreeKEM, adaptively and actively secure continuous group key agreement. In *2021 IEEE Symposium on Security and Privacy*, pages 268–284. IEEE Computer Society Press, May 2021.
- [58] Hugo Krawczyk, Mihir Bellare, and Ran Canetti. HMAC: Keyed-hashing for message authentication. IETF Internet Request for Comments 2104, February 1997.



- [59] Julia Len, Esha Ghosh, Paul Grubbs, and Paul Rösler. Interoperability in end-to-end encrypted messaging. Cryptology ePrint Archive, Report 2023/386, 2023. <https://eprint.iacr.org/2023/386>.
- [60] Joshua Lund. Technology preview: Sealed sender for signal. <https://signal.org/blog/sealed-sender/>, October 2018. (Accessed on 04/19/2023).
- [61] Vadim Lyubashevsky, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Peter Schwabe, Gregor Seiler, Damien Stehlé, and Shi Bai. CRYSTALS-DILITHIUM. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>.
- [62] Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Maxime Plançon. Lattice-based zero-knowledge proofs and applications: Shorter, simpler, and more general. In Dodis and Shrimpton [43], pages 71–101.
- [63] Moxie Marlinspike and Trevor Perrin. The double ratchet algorithm, November 2016. <https://signal.org/docs/specifications/doubleratchet/>.
- [64] Moxie Marlinspike and Trevor Perrin. The X3DH key agreement protocol, November 2016. <https://signal.org/docs/specifications/x3dh/>.
- [65] Jaroslav Martan. Webex meetings security. Technical report, September 2021. [https://www.cisco.com/c/dam/m/cs\\_cz/training-events/webinars/tech-club-webinars/webex-meetings-security.pdf](https://www.cisco.com/c/dam/m/cs_cz/training-events/webinars/tech-club-webinars/webex-meetings-security.pdf).
- [66] Ian Martiny, Gabriel Kaptchuk, Adam J. Aviv, Daniel S. Roche, and Eric Wustrow. Improving signal’s sealed sender. In *NDSS 2021*. The Internet Society, February 2021.
- [67] Ben Morris, Phillip Rogaway, and Till Stegers. Deterministic encryption with the Thorp shuffle. *Journal of Cryptology*, 31(2):521–536, April 2018.
- [68] Rafael Pass and Krzysztof Pietrzak, editors. *TCC 2020, Part II*, volume 12551 of *LNCS*. Springer, Heidelberg, November 2020.
- [69] Trevor Perrin. The noise protocol framework, 2018. (Accessed on 04/23/2023).
- [70] Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. FALCON. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>.
- [71] Guardian Project. Orbot: Proxy with tor. (Accessed on 04/19/2023).
- [72] Peter Schwabe, Roberto Avanzi, Joppe Bos, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, John M. Schanck, Gregor Seiler, Damien Stehlé, and Jintai Ding. CRYSTALS-KYBER. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>.
- [73] Emil Stefanov and Elaine Shi. FastPRP: Fast pseudorandom permutations for small domains. Cryptology ePrint Archive, Report 2012/254, 2012. <https://eprint.iacr.org/2012/254>.
- [74] Roman Tobe. Ringcentral expands end-to-end encryption to phone and messaging. <https://www.ringcentral.com/us/en/blog/ringcentral-update-e2ee-2022/>, 2022. (Accessed on 04/25/2023).
- [75] Nirvan Tyagi, Julia Len, Ian Miers, and Thomas Ristenpart. Orca: Blocklisting in sender-anonymous messaging. In Kevin R. B. Butler and Kurt Thomas, editors, *USENIX Security 2022*, pages 2299–2316. USENIX Association, August 2022.
- [76] Giovanni Vigna and Elaine Shi, editors. *ACM CCS 2021*. ACM Press, November 2021.
- [77] Théophile Wallez, Jonathan Protzenko, Benjamin Beurdouche, and Karthikeyan Bhargavan. TreeSync: Authenticated group management for messaging layer security. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 1217–1233, Anaheim, CA, August 2023. USENIX Association.
- [78] Matthew Weidner. Group messaging for secure asynchronous collaboration. Mphil dissertation, University of Cambridge, Cambridge, UK, 2019.
- [79] Matthew Weidner, Martin Kleppmann, Daniel Hugenroth, and Alastair R. Beresford. Key agreement for decentralized secure group messaging with strong security guarantees. In Vigna and Shi [76], pages 2024–2045.
- [80] Shouhuai Xu and Moti Yung. Accountable ring signatures: A smart card approach. In *Smart Card Research and Advanced Applications VI*, pages 271–286. Springer, 2004.
- [81] Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors. *ACM CCS 2022*. ACM Press, November 2022.

## A Other Related Work

**Secure Two-Party Messaging.** A well-known secure two-party messaging protocols is the Signal protocol, which comprises two sub-protocols: X3DH protocol [64] and Double Ratchet protocol [63]. The full security of Signal protocol was analyzed by Cohn-Gordon et al. [36, 37]. Subsequently, these two sub-protocols were studied separately. The X3DH protocol, which is used to establish the initial secret key required to start a secure conversation, was analyzed by [28, 50, 51]. With respect to the Double Ratchet protocol, Alwen et al. [3] abstracted it as Continuous Key Agreement (CKA) (CGKA in two-party setting) and studied the security of Signal’s Double Ratchet protocol in this abstraction. Recently, Bienstock et al. [25] and Canetti et al. [31] concurrently formalized an ideal functionality for the Signal protocol and proved its security using the Universally Composable (UC) framework.

**Secure Group Messaging.** Signal supports a group of more than two users by sending the same message to the users on the group with each of two-user messaging channels (i.e., pair-wise channels). It becomes less efficient as the group grows. To address this limitation, the Messaging Layer Security (MLS) protocol was developed in an IETF working group and is ready to be published as an RFC. MLS consists of two sub-protocols: Continuous Group Key Agreement (CGKA) protocol and Forward Secure Payload Delivery (FSPD) protocol. The former establishes a group secret key and continuously updates it, while the latter delivers payloads (a.k.a. application messages [12]) securely using the group secret key.

MLS uses TreeKEM [21] as a CGKA protocol, which is based on the Asynchronous Ratcheting Trees proposed by Cohn-Gordon et al. [38]. The security of TreeKEM as a stand-alone protocol was analyzed with game-based [4], simulation-based (UC framework) [8], and machine-checkable [22] security definitions. Alwen et al. [5] proposed a game-based security definitions for MLS, and analyzed MLS in its entirety. Their paper is the first to formally cast group messaging as the composition of CGKA and FSPD (termed FS-GAEAD in their model). They also identify a third component primitive, which takes the form of a hash function with particular properties. Bienstock et al. [23] proved a lower bound on the communication cost of any group messaging protocol, as well showing no optimal (with respect to distributions of group operations) protocol exists. Anastos et al. [9] prove a lower bound on the communication cost of replacing a set of users in CGKA, in particular showing MLS is optimal in this regard among protocols built in a black-box way from a standard set of primitives.

In addition to studying the standard TreeKEM in MLS, constructing new CGKA protocols is an active research topic. Re-randomized TreeKEM [4] and TreeKEM with active security [6] was proposed to improve forward security. Tainted TreeKEM [57] introduced an alternative approach to remov-

ing users which can be more efficient in certain settings, such as those with group administrators. Hashimoto et al. [52] and Alwen et al. [7] constructed communication-efficient CGKA protocols by allowing receivers to selectively download uploaded contents. Hashimoto et al. [52] proposed a CGKA protocol based on multi-recipient PKEs called Chained CmpKE, which is designed to reduce the total communication cost. Alwen et al. [7] proposed a server-aided CGKA protocol by extending MLS’s TreeKEM, which is designed to reduce the upload and download costs per message. Weidner et al. [79] proposed a decentralized CGKA (DCGKA) protocol to realize secure group messaging for decentralized networks that have no central authority. An issue with CGKA protocols is how to handle concurrently uploaded key updates. To address this, CGKA protocols supporting concurrent key updates and exhibiting different trade-offs have been proposed [1, 2, 24, 78]. Bienstock et al. [24] also showed a lower bound on the communication costs of achieving PCS in two rounds of concurrent communication. This lower bound was extended by Auerbach et al. [10] to capture the more general setting of achieving PCS in a (potentially) higher number of rounds.

**Authentication for Group States.** In secure group messaging, group members share a common *group state* that includes information such as the list of group members and their cryptographic keys. When new members join the group, they must verify that the received group state is synchronized across all current members. If not, the group suffers from insider and outsider attacks e.g., double join attacks [22]. Signal uses an anonymous group management protocol called *private groups* [34]. In this protocol, the group membership list is encrypted with a group secret key and stored on the server. When members want to modify the list, the server anonymously authenticates them to ensure that they have the permission to do so. Further, a recent paper [11] introduces the notion of *administrated CGKA*, which enables a set of administrators to cryptographically manage group membership.

In MLS, new members receive the group state as a welcome message from an existing member. The group state includes the membership of the group and the structure of the public key tree for TreeKEM. MLS uses an authentication mechanism called *tree-signing* to allow new members to agree on and authenticate the group’s membership and public key tree structure, without which powerful insider attacks could be performed [8]. Wallez et al. [77] formalized this authentication mechanism as *TreeSync* and provided a machine-checked formal specification for it.

It is worth noting that these works are orthogonal to the authentication we consider. Whereas they consider the integrity of *group states* and the authentication of values therein, we consider the authenticity of *conversation messages*.

**Anonymous Secure Messaging.** In the two-party setting, Signal provides sender-anonymous messaging through the Sealed Sender protocol [60]. Each user distributes an *access key* to their friends and registers it with the server in advance.

The server authenticates senders by checking whether the sender holds the access key of the intended receiver. Martiny et al. [66] demonstrated statistical analysis attacks that break sender anonymity in Sealed Sender, and suggested countermeasures to address their attacks. Tyagi et al. [75] pointed out that the access key distribution during the setup in Sealed Sender is not anonymous, and users could suffer from a DoS attack due to the lack of user traceability. To address these issues, they proposed a new sender-anonymous messaging protocol called Orca. Their protocol ensures that all communication, even in the setup phase, is anonymous and provides user traceability.

In the group setting, Signal already supports anonymous group messaging by using the sender-anonymous two-party protocol mentioned above as the underlying two-party messaging protocol. MLS also considers sender-anonymous group messaging according to its specification [12, Sec. 6.3.2.] (see also [19, Sec. 7.1.2]). Hashimoto et al. [53] studied the anonymity that MLS provides and defined a new security models that covers anonymity of senders and receivers against the server. They proposed a *metadata-hiding* CGKA protocol satisfying their security definition. Their protocol offers an implicit anonymous blocklisting property that allows the server to block access from outside the group, similar to Signal’s Sealed Sender.

**Group Signatures.** Group signatures [13, 27, 35] can be used to achieve anonymous authentication with user traceability. They allow users to anonymously sign messages on behalf of a group, while allowing a special entity called the *group tracer* to trace the signature back to the user. In the context of secure messaging, we view the server as the group manager and verifier, and users as signers and group tracers. The server manages a group of all users in the system. Users join this group by setting up the account, and receiving a certificate from the server. With this set up in place, a sender can generate a signature with its certificate, and designate the intended receiver as tracer. The server can then verify the signature and, if verification is successful, deliver it to the intended receiver, who can then trace it back to the sender with its tracing key. In the context of group messaging, the set of users in a group is regarded as one instantiation of the group tracer, thus allowing all members to trace the sender. Tyagi et al. [75] and this work propose anonymous (group) messaging protocols following this framework.

To use group signatures in secure messaging as explained above, they must support dynamic groups [16] and multiple group tracers [75]. It is worth noting that the same key is used for both issuing certificates and verifying signature in this setting, since the server acts as both group manager and verifier. This property is called *keyed-verification* [33, 34, 75] in the literature, and it could be useful to construct efficient schemes. Tyagi et al. [75] constructed an efficient group signature satisfying these properties based on the Diffie-Hellman assumption.

## B Preliminary: Cryptographic Tools

In this section, we define all the standard cryptographic tools used in our work.

**One-Way Function.** Let  $\text{OWF} : \mathcal{D} \rightarrow \mathcal{R}$  be an efficient function family with domain  $\mathcal{D}$  and finite range  $\mathcal{R}$ . We define a one-way function as follows.

**Definition B.1 (One-Way Function).** We say  $\text{OWF}$  is a one-way function if for all PPT adversary  $\mathcal{A}$ , we have

$$\Pr[x \leftarrow \mathcal{D}, x' \leftarrow \mathcal{A}(1^\kappa, \text{OWF}(x)) : x = x'] \leq \text{negl}(\kappa).$$

**Lemma B.2.** Let  $n, m$  be integers such that  $n, m \geq \kappa$ . Then, a hash function  $H : \{0, 1\}^n \rightarrow \{0, 1\}^m$  instantiated as a random oracle is a OWF.

*Proof.* Let  $\mathcal{A}$  be a PPT adversary that makes at most  $Q = \text{poly}(\kappa)$  queries to the random oracle. First, the probability that  $\mathcal{A}$  queries  $x$  is bounded by  $Q/2^n$ . Moreover, if it queries  $z \neq x$ , then the probability that  $H(z) = H(x)$  is  $Q/2^m$ . Hence, the probability that  $\mathcal{A}$  wins is upper bounded by  $Q(1/2^n + 1/2^m) \leq Q/2^{\kappa-1} = \text{negl}(\kappa)$ .  $\square$

**Pseudorandom Function.** Let  $\text{PRF} : \mathcal{K} \times \mathcal{D} \rightarrow \mathcal{R}$  be an efficient function family with key space  $\mathcal{K}$ , domain  $\mathcal{D}$  and finite range  $\mathcal{R}$ . We define a pseudorandom function as follows.

**Definition B.3 (Pseudorandom Function).** We say  $\text{PRF}$  is a pseudo-random function if for all PPT adversary  $\mathcal{A}$ , we have

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

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

**Pseudorandom Permutation.** Let  $\text{PRP} : \mathcal{K} \times \mathcal{R} \rightarrow \mathcal{R}$  be an efficient function family of one-to-one functions from  $\mathcal{R}$  to  $\mathcal{R}$  with key space  $\mathcal{K}$  for which  $\text{PRP}^{-1}$  is also efficiently computable given the first input (i.e., key). We define a pseudorandom permutation as follows.

**Definition B.4 (Pseudorandom Permutation).** We say  $\text{PRP}$  is a pseudo-random permutation if for all PPT adversary  $\mathcal{A}$ , we have

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

where  $\mathcal{R}\mathcal{P}$  is the set of all permutations over  $\mathcal{R}$ , and  $\mathcal{P}(\cdot)$  (resp.  $\mathcal{P}^{-1}(\cdot)$ ) is defined as  $\text{PRP}(K, \cdot)$  (resp.  $\text{PRP}^{-1}(K, \cdot)$ ) if  $b = 0$ , and  $\text{RP}(\cdot)$  (resp.  $\text{RP}^{-1}(\cdot)$ ) otherwise.

**Secret Key Encryption.** We provide the standard notion of secret key encryption (SKE).

**Definition B.5 (Secret-Key Encryption).** A secret-key encryption (SKE) over key space  $\mathcal{K}$  and message space  $\mathcal{M}$  consists of the following two algorithms:

$\text{Enc}(k, m) \rightarrow \text{ct}$ : On input a secret key  $k \in \mathcal{K}$  and a message  $m \in \mathcal{M}$ , it outputs a ciphertext  $\text{ct}$ .

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

**Definition B.6 (Correctness).** An SKE is correct if  $\Pr[\text{Dec}(k, \text{Enc}(k, m)) = m] = 1$  holds for all  $m \in \mathcal{M}$  and  $k \in \mathcal{K}$ .

**Definition B.7 (IND-CCA).** An SKE is IND-CCA secure if for all PPT adversary  $\mathcal{A}$ , we have

$$\left| \Pr \left[ b = b' : \begin{array}{l} (b, k) \leftarrow \mathcal{S}\{0, 1\} \times \mathcal{K}, \\ b' \leftarrow \mathcal{S}A^{\mathcal{C}(\cdot, \cdot), \mathcal{D}(\cdot)}(\text{ct}^*) \end{array} \right] - \frac{1}{2} \right| \leq \text{negl}(\kappa),$$

where  $\mathcal{C}(m_0, m_1)$  outputs  $(\text{ct}_0, \text{ct}_1)$ , where  $\text{ct}_i \leftarrow \mathcal{S}\text{Enc}(k, m_{i \oplus b})$  for  $i \in \{0, 1\}$ , and  $\mathcal{D}(\text{ct})$  returns  $\text{Dec}(k, \text{ct})$  conditioned on  $\text{ct}$  not being an output of  $\mathcal{C}(\cdot, \cdot)$ .

**Message Authentication Code.** We provide the standard notion of (deterministic) message authentication codes (MAC).

**Definition B.8 (MAC).** A (deterministic) message authentication code MAC over key space  $\mathcal{K}$  and message space  $\mathcal{M}$  consists of the following algorithms:

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

$\text{Verify}(k, \text{tag}, m) \rightarrow \perp/\top$ : On input a key  $k$ , a tag  $\text{tag}$ , and a message  $m$ , it (deterministically) outputs  $\top$  or  $\perp$ .

Since the  $\text{TagGen}$  algorithm is deterministic, we can simply define  $\text{Verify}$  to run  $\text{TagGen}$  on  $(k, m)$  and check if the generated  $\text{tag}'$  is identical to the provided  $\text{tag}$ .

**Definition B.9 (Correctness).** A MAC is correct if for all keys  $k \in \mathcal{K}$  and all messages  $m \in \mathcal{M}$ ,

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

**Definition B.10 (sEUF-CMA).** A MAC is sEUF-CMA secure if for all PPT adversary  $\mathcal{A}$ , we have

$$\Pr \left[ \begin{array}{l} \text{Verify}(k, m^*, \text{tag}^*) = \top \\ \wedge (m^*, \text{tag}^*) \notin L^* \end{array} : \begin{array}{l} k \leftarrow \mathcal{S}\mathcal{K}; \\ (m^*, \text{tag}^*) \leftarrow \mathcal{S}\mathcal{A}^{\mathcal{T}(\cdot)}(1^\kappa) \end{array} \right] \leq \text{negl}(\kappa)$$

where  $\mathcal{T}$  is the MAC oracle which on input  $m$  returns  $\text{TagGen}(k, m)$ , and  $L^*$  is the set of pairs of message and tag generated by the MAC oracle.

**(One Time) Signature** We provide the standard notion of a signature scheme.

**Definition B.11 (Signature Scheme).** A signature scheme SIG over a message space  $\mathcal{M}$  consists of the following algorithms:

$\text{KeyGen}(1^\kappa) \rightarrow (\text{vk}, \text{sk})$ : On input the security parameter  $1^\kappa$ , it outputs a verification and signing key pair  $(\text{vk}, \text{sk})$ .

$\text{Sign}(\text{sk}, m) \rightarrow \sigma$ : On input a signing key  $\text{sk}$  and a message  $m \in \mathcal{M}$ , it outputs a signature  $\sigma$ .

$\text{Verify}(\text{vk}, \sigma, m) \rightarrow \perp/\top$ : On input a key  $k$ , a signature  $\sigma$ , and a message  $m$ , it (deterministically) outputs  $\top$  or  $\perp$ .

**Definition B.12 (Correctness).** A signature scheme SIG is correct if for all  $\kappa \in \mathbb{N}$ ,  $m \in \mathcal{M}$ , and  $(\text{vk}, \text{sk}) \in \text{KeyGen}(1^\kappa)$ , we have

$$\Pr[\text{Verify}(\text{vk}, \sigma, m) = \top : \sigma \leftarrow \mathcal{S}\text{Sign}(\text{sk}, m)] = 1 - \text{negl}(\kappa).$$

**Definition B.13 (EUF-CMA).** A signature scheme SIG is EUF-CMA secure if for all PPT adversary  $\mathcal{A}$ , we have

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

where  $\mathcal{S}$  is the signing oracle which on input  $m$  returns  $\sigma \leftarrow \mathcal{S}\text{Sign}(\text{sk}, m)$ , and  $L^*$  is the set of messages queried to the signing oracle. We say SIG is one-time secure if  $|L^*| = 1$ .

**Key Encapsulation Mechanism** We provide the standard notion of a key encapsulation mechanism (KEM).

**Definition B.14 (Key Encapsulation Mechanism).** A key encapsulation mechanism (KEM) with key space  $\mathcal{K}$  consists of the following PPT algorithms:

$\text{KeyGen}(1^\kappa) \rightarrow (\text{ek}, \text{dk})$ : On input a security parameter  $1^\kappa$ , it outputs a pair of encryption and decryption keys  $(\text{ek}, \text{dk})$ .

$\text{Enc}(\text{ek}) \rightarrow \text{ct}$ : On input an encryption key  $\text{ek}$ , it outputs a key  $k \in \mathcal{K}$  and a ciphertext  $\text{ct}$ .

$\text{Dec}(\text{dk}, \text{ct}) \rightarrow m/\perp$ : On input a decryption key  $\text{dk}$  and a ciphertext  $\text{ct}$ , it outputs a key  $k \in \mathcal{K} \cup \{\perp\}$ .

**Definition B.15 ((1 -  $\delta$ )-Correctness).** A KEM is  $(1 - \delta)$ -correct if for all  $\kappa \in \mathbb{N}$ , we have

$$(1 - \delta) \leq \Pr \left[ \begin{array}{l} (\text{ek}, \text{dk}) \leftarrow \mathcal{S}\text{KeyGen}(1^\kappa), \\ (k, \text{ct}) \leftarrow \mathcal{S}\text{Enc}(\text{ek}) \end{array} : \text{Dec}(\text{dk}, \text{ct}) = k \right].$$

**Definition B.16 (IND-CCA Security).** A KEM is IND-CCA secure if for all PPT adversary  $\mathcal{A}$  we have

$$\left| \Pr \left[ \begin{array}{l} (b = b') : \begin{array}{l} (\text{ek}, \text{dk}) \leftarrow \mathcal{S}\text{KeyGen}(1^\kappa), \\ (b, k_0^*) \leftarrow \mathcal{S}\{0, 1\} \times \mathcal{K}, \\ (k_1^*, \text{ct}^*) \leftarrow \mathcal{S}\text{Enc}(\text{ek}), \\ b' \leftarrow \mathcal{S}A^{\mathcal{D}(\cdot)}(\text{ek}, \text{ct}^*, k_b) \end{array} \right] - \frac{1}{2} \right| \leq \text{negl}(\kappa),$$

where  $\mathcal{D}(\text{ct})$  outputs  $\text{Dec}(\text{dk}, \text{ct})$  conditioned on  $\text{ct} \neq \text{ct}^*$ .

## C Security Proofs for COSMOS and COSMAC

In this section, we provide the security proofs for COSMOS and COSMAC.

### C.1 Security Proof of COSMOS

We prove that COSMOS is correct and secure. Here, note that COSMOS is not anonymous since the list of public tokens PubTOKEN is indexed by the users  $u \in G$  and when user  $u$  updates its tokens  $(x_i, y_i)_{i \in [T]}$ , the list of public tokens PubTOKEN[ $u$ ] is updated to  $(y_i)_{i \in [T]}$ . Any two messages with private tokens  $x$  and  $x'$  such that  $\text{OWF}(x), \text{OWF}(x') \in (y_i)_{i \in [T]}$  can thus be linked together. Also, recall that if a GAM protocol is not anonymous, then blocklisting becomes trivial (see Secs. 2.2.2 and 2.3).

Formally, we have the following theorem.

**Theorem C.1.** *The GAM protocol COSMOS in Fig. 3 is signing correct, local state update correct, non-colluding unforgeable, and tracing sound assuming OWF is one-way, PRF is pseudorandom, and MAC is sEUF-CMA.*

*Proof.* Correctness of signing and local state-updates is clear from construction. In the following, we prove COSMOS is non-colluding unforgeable and tracing sound in Lems. C.2 and C.3, respectively.  $\square$

#### C.1.1 Proof of Lem. C.2

**Lemma C.2.** *COSMOS is non-colluding unforgeable assuming OWF is one-way, PRF is pseudorandom, and MAC is sEUF-CMA.*

*Proof.* There are two cases we must consider for non-colluding unforgeability:  $C \neq \emptyset$  (i.e., the server is honest and  $\mathcal{A}$  has access to  $O^*$ ) and  $C = \emptyset$  (i.e., the server is malicious and  $\mathcal{A}$  has access to  $O$ ).

We first consider the first case. Let us assume  $\mathcal{A}$  outputs  $(\text{label}, \text{obj}) = (\text{msg}, (v, \sigma, m))$  for  $v \in \mathcal{H}$ ; the other case when  $\text{label} = \text{upd}$  is proven identically. Since it is a valid adversary, the authentication token  $\sigma$  is valid and traces back to some honest user  $u^* \in \mathcal{H}$ . That is,  $\sigma$  is the form  $(u^*, \text{ctr}, x, \Sigma_{\text{MAC}})$  and  $(u^*, *, m) \notin L_{\text{msg}}$ . Now notice that a group authentication token of the form  $\Sigma_G = (u^*, \text{ctr}, x, *)$  could not have been output by  $O_{\text{Send}}$  since when  $\mathcal{A}$  has access to  $O^*$ ,  $\Sigma_G$  is immediately processed by  $O_{\text{GroupReceive}}$ . In particular, the honest server and group users delete the public token  $y = \text{OWF}(x)$  from DB and will not verify the second time it is used. Thus, for  $\mathcal{A}$  to forge  $u^*$  it must output an unknown private token  $x$  for a public token  $y = \text{OWF}(x)$  without querying  $O_{\text{Send}}$  on  $u^*$ . Due to the one-wayness of OWF, this can be done with at most negligible probability.

We next consider the second case. Similarly as before, let us assume  $\mathcal{A}$  outputs  $(\text{label}, \text{obj}) = (\text{msg}, (v, \sigma, m))$  for  $v \in \mathcal{H}$

and traces to some honest user  $u^* \in \mathcal{H}$ . Moreover, due to the winning condition, we have  $\sigma = (u^*, \text{ctr}, x, \Sigma_{\text{MAC}})$  with  $\text{MAC.Verify}(k_{\text{MAC}}, (u^*, \text{ctr}, x, m), \Sigma_{\text{MAC}}) = \top$  and  $(u^*, *, m) \notin L_{\text{msg}}$ . However, since the malicious server does not know the MAC key  $k_{\text{MAC}}$ , this can occur with at most negligible probability assuming sEUF-CMA security of the MAC. Concretely, if  $(u^*, *, m) \notin L_{\text{msg}}$ , then the game will have never signed the tuple  $(u^*, \text{ctr}, x, m)$ . Therefore, if an adversary outputs a MAC tag  $\Sigma_{\text{MAC}}$  for such a tuple, it breaks sEUF-CMA security. Here, to simulate the sEUF-CMA security game, we rely on the pseudorandomness of the PRF to sample the MAC key  $k_{\text{MAC}}$  directly without running PRF.  $\square$

It is worth mentioning that in the above proof, we crucially rely on the fact that either the set of group users or the server is honest. When considering *standard* unforgeability, a malicious user and server can easily collude to break unforgeability as secrecy of both the OWF inputs and MAC key  $k_{\text{MAC}}$  are known to the adversary.

#### C.1.2 Proof of Lem. C.3

**Lemma C.3.** *COSMOS is (unconditionally) tracing sound.*

*Proof.* Assume the adversary  $\mathcal{A}$  outputs  $(\text{label}, \text{obj}) = (\text{msg}, (\Sigma_G, m))$ . The game then runs  $(\text{pp}', (\sigma_i)_{i \in [N]}) \leftarrow \text{Verify}(\text{pp}, \Sigma_G, m)$ . Since  $\text{pp}'$  verifies,  $\Sigma_G = (u, \text{ctr}, x, \Sigma_{\text{MAC}})$  for some  $u \in G$ ,  $\text{ctr} \in [T]$ , and  $\text{DB}[u][\text{ctr}] = \text{OWF}(x)$ , where DB is the database stored in the public parameter  $\text{pp}$ . Then, a user authentication token for user  $v \in G \cap \mathcal{H}$  is set as  $\sigma_v = (u, \text{ctr}, x, \Sigma_{\text{MAC}})$ . Next, notice that all  $v \in G \cap \mathcal{H}$  include the same PubTOKEN in their state as those included in DB. Then, by the description of  $*_{\text{trace-sender}}$  in Fig. 4, all honest users  $v$  either all reject the authentication token because  $\Sigma_{\text{MAC}}$  is invalid or uniquely traces  $u$  given  $\sigma_v$ . The case when  $\mathcal{A}$  outputs  $(\text{label}, \text{obj}) = (\text{upd}, (\hat{\Sigma}_G, \hat{\text{ctr}}_G))$  is proven identically. This completes the proof of the lemma.  $\square$

#### C.1.3 Security of the Optimizations

Lifting the security proofs of COSMOS from App. C to the optimized schemes COSMOS<sup>+</sup> and COSMOS<sup>++</sup> are straightforward. Indeed, it is clear from the similarity of the construction that COSMOS<sup>+</sup> will satisfy the same security guarantees. When it comes to COSMOS<sup>++</sup>, we note that tracing soundness immediately follows from the original proof in Lemma C.3. The proof for non-colluding unforgeability in Lemma C.2 shows that, in order to forge, an adversary must either break sEUF-CMA security of the MAC or find a pre-image of the employed OWF. Relying on the same argument and noting that any revealed private token becomes a pre-image challenge for the adversary, it follows that COSMOS<sup>++</sup> is non-colluding unforgeable as well.

## C.2 Security Proof of COSMAC

We show in the following theorem that COSMAC is non-colluding unforgeable, anonymous, and anonymous blocklistable. We note that COSMAC is not tracing sound since any malicious user can upload a valid MAC tag along with a ciphertext that does not decrypt to any valid authentication tokens of COSMOS. The optimized schemes COSMAC<sup>+</sup> and COSMAC<sup>++</sup> can be proven almost identically as per the discussion in App. C.1.3.

**Theorem C.4.** *The GAM protocol COSMAC in Fig. 5 is signing correct, local state update correct, non-colluding unforgeable, anonymous, and anonymous blocklistable assuming OWF is one-way, PRF is pseudorandom, MAC is sEUF-CMA secure, and SKE is IND-CCA secure.*

*Proof.* Correctness of signing and local state-updates are inherited from COSMOS, additionally relying on the fact that the MAC and SKE are correct. In the following we prove COSMAC is non-colluding unforgeable, anonymous, and anonymous blocklistable in Lems. C.5 to C.7, respectively. Proving these lemmas completes the proof of the theorem.  $\square$

### C.2.1 Proof of Lem. C.5

**Lemma C.5.** *COSMAC is non-colluding unforgeable assuming OWF is one-way, PRF is pseudorandom, and MAC is sEUF-CMA.*

*Proof.* An adversary  $\mathcal{A}$  that breaks the non-colluding unforgeability of COSMAC can be used to construct an adversary  $\mathcal{B}$  that breaks the non-colluding unforgeability of COSMOS. Concretely,  $\mathcal{B}$  generates  $\overline{k_{\text{MAC}}}$  and  $k_{\text{SKE}}$  randomly over their respective domains without invoking the PRF; this is indistinguishable assuming the pseudorandomness of PRF. To simulate  $O_{\text{Send}}$  and  $O_{\text{UpdSend}}$  to  $\mathcal{A}$ ,  $\mathcal{B}$  calls its own oracle; encrypts its output using  $k_{\text{SKE}}$ ; and signs it using  $\overline{k_{\text{MAC}}}$  as specified by COSMAC. To simulate  $O_{\text{Receive}}$ ,  $O_{\text{UpdReceive}}$ ,  $O_{\text{GroupReceive}}$ , and  $O_{\text{GroupUpdReceive}}$  to  $\mathcal{A}$ ,  $\mathcal{B}$  checks if the attached tag and content verifies and decrypts under  $\overline{k_{\text{MAC}}}$  and  $k_{\text{SKE}}$ , respectively. If so, it queries its own oracle on the decrypted message. Finally, if  $\mathcal{A}$  outputs a valid forgery,  $\mathcal{B}$  removes the tag, decrypts it using  $k_{\text{SKE}}$ , and submits it as its own forgery. Thus, assuming that COSMOS is non-colluding unforgeable (which is established in Lem. C.2), COSMAC is non-colluding unforgeable.  $\square$

### C.2.2 Proof of Lem. C.6

**Lemma C.6.** *COSMAC is anonymous assuming PRF is pseudorandom and SKE is IND-CCA secure.*

*Proof.* We first use the pseudorandomness of the PRF to modify the security game so that  $k_{\text{MAC}}$ ,  $\overline{k_{\text{MAC}}}$ , and  $k_{\text{SKE}}$  are sampled independently. We show that an adversary  $\mathcal{A}$  that breaks anonymity of this modified security game can be used

to construct an adversary  $\mathcal{B}$  that breaks IND-CCA security of the SKE.

Concretely,  $\mathcal{B}$  samples  $sk_{\text{SV}} = \overline{k_{\text{MAC}}}$  and provides it to  $\mathcal{A}$ . It then prepares an empty list  $L_{\text{ct}}$ . When  $\mathcal{A}$  queries  $O_{\text{Send}}$  and  $O_{\text{UpdSend}}$ ,  $\mathcal{B}$  runs COSMAC except that it queries  $(\overline{m}, \overline{m})$  to its encryption oracle  $C(\cdot, \cdot)$  where  $\overline{m} = (u, \text{ctr}, x_u^{(\text{ctr})}, \Sigma_{\text{MAC}})$ , rather than generating it itself. On receiving  $(\text{ct}_{\text{SKE},0}, \text{ct}_{\text{SKE},1})$  from the oracle,  $\mathcal{B}$  uses  $\text{ct}_{\text{SKE},0}$  and  $\overline{k_{\text{MAC}}}$  to create the authentication tokens. It then updates  $L_{\text{ct}} \leftarrow L_{\text{ct}} \cup (m, \text{ct}_{\text{SKE},0})$ , where  $m$  is the message for which  $\text{attach-auth-token}$  is invoked when  $\mathcal{A}$  queries  $O_{\text{Send}}$  and  $O_{\text{UpdSend}}$ . This simulates COSMAC perfectly.

When  $\mathcal{A}$  makes a challenge query  $(u_0, u_1, m_0, m_1)$ ,  $\mathcal{B}$  runs COSMAC except that it queries  $((u_0, \text{ctr}_0, x_{u_0}^{(\text{ctr}_0)}, \Sigma_{\text{MAC},0}), (u_1, \text{ctr}_1, x_{u_1}^{(\text{ctr}_1)}, \Sigma_{\text{MAC},1}))$  to its encryption oracle  $C(\cdot, \cdot)$ . Here, note that we ignore the public parameter  $\overline{pp}$  from  $\mathcal{A}$ 's output as it is empty for COSMAC. On receiving  $(\text{ct}_{\text{SKE},0}^*, \text{ct}_{\text{SKE},1}^*)$  from the oracle,  $\mathcal{B}$  generates MAC tags  $(\overline{\Sigma}_{\text{MAC},0}^*, \overline{\Sigma}_{\text{MAC},1}^*)$  using  $\overline{k_{\text{MAC}}}$ , respectively, and then outputs  $(((\text{ct}_{\text{SKE},0}^*, \overline{\Sigma}_{\text{MAC},0}^*), m_0), ((\text{ct}_{\text{SKE},1}^*, \overline{\Sigma}_{\text{MAC},1}^*), m_1))$ . Finally,  $\mathcal{B}$  executes the Verify and Receive algorithms to check the resulting (group) authentication tokens are valid. Here, we note that  $\mathcal{B}$  can ignore decrypting the ciphertext  $(\text{ct}_{\text{SKE},0}^*, \text{ct}_{\text{SKE},1}^*)$  in the Receive algorithm assuming correctness of the SKE. It can be checked that the challenge bit of the IND-CCA security game is consistent with the challenge bit coin implicitly sampled by  $\mathcal{B}$ . Hence, this simulates COSMAC perfectly.

Lastly, when  $\mathcal{A}$  queries  $O_{\text{Receive}}$  and  $O_{\text{UpdReceive}}$  on an authentication token  $\sigma$ ,  $\mathcal{B}$  checks  $\sigma = \text{ct}'_{\text{SKE}} \notin \text{Chall}_{\text{msg}}$ , where recall  $\text{Chall}_{\text{msg}}$  is the  $2N$  authentication tokens generated by Verify above. If the ciphertext  $\text{ct}'_{\text{SKE}}$  satisfies  $(m', \text{ct}'_{\text{SKE}}) \in L_{\text{ct}}$  for some  $m'$ , it runs COSMAC except that it uses  $m'$ . Otherwise, if  $\text{ct}'_{\text{SKE}} \notin L_{\text{ct}} \cup \{\text{ct}_{\text{SKE},0}^*, \text{ct}_{\text{SKE},1}^*\}$ , it runs COSMAC except that it decrypts  $\text{ct}'_{\text{SKE}}$  using its decryption oracle  $\mathcal{D}(\cdot)$ . Finally, we never have the case  $\text{ct}'_{\text{SKE}} \in \{\text{ct}_{\text{SKE},0}^*, \text{ct}_{\text{SKE},1}^*\}$  due to the condition  $\sigma \notin \text{Chall}_{\text{msg}}$ . Thus, this simulates COSMAC perfectly.

At the end of the game, when  $\mathcal{A}$  outputs a guess  $\widehat{\text{coin}}$ ,  $\mathcal{B}$  outputs this as its guess. Since  $\mathcal{B}$  perfectly simulates the game to  $\mathcal{A}$ , if  $\mathcal{A}$  breaks anonymity of COSMAC, then  $\mathcal{B}$  breaks IND-CCA security of SKE. This completes the proof.  $\square$

### C.2.3 Proof of Lem. C.7

**Lemma C.7.** *COSMAC is anonymous blocklistable assuming PRF is pseudorandom and MAC is sEUF-CMA secure.*

*Proof.* We first use the pseudorandomness of the PRF to modify the security game so that  $k_{\text{MAC}}$ ,  $\overline{k_{\text{MAC}}}$ , and  $k_{\text{SKE}}$  are sampled independently. We show that an adversary  $\mathcal{A}$  that breaks anonymous blocklisting of this modified security game can be used to construct an adversary  $\mathcal{B}$  that breaks sEUF-CMA security of the MAC.

Concretely,  $\mathcal{B}$  samples everything as in the anonymous blocklistable security game except for implicitly setting  $sk_{Sv} = \overline{k_{MAC}}$  as the MAC key used by the sEUF-CMA security game. It then proceeds as follows. When  $\mathcal{A}$  queries  $O_{Send}$  and  $O_{UpdSend}$ ,  $\mathcal{B}$  runs COSMAC except that it queries the MAC oracle  $TG(\cdot)$  on  $ct_{SKE}$ , rather than generating the MAC tag on its own. This perfectly simulates the view to  $\mathcal{A}$  since MAC is perfectly correct.

When  $\mathcal{A}$  queries  $O_{GroupReceive}$  and  $O_{GroupUpdReceive}$ , it queries its MAC oracle on  $ct_{SKE}$  attached to  $\mathcal{A}$ 's query. If the returned MAC tag is identical to the MAC tag  $\overline{\Sigma}_{MAC}$  attached to  $\mathcal{A}$ 's query, it completes COSMAC. Since MAC is deterministic, this perfectly simulates the view to  $\mathcal{A}$ . (Note that we can rely on sEUF-CMA security instead if we assume a randomized MAC). Here, without loss of generality, we can assume  $ct_{SKE}$  was never queried to the MAC oracle by  $\mathcal{B}$ . Otherwise,  $\mathcal{A}$  can use the pair  $(ct_{SKE}, \overline{\Sigma}_{MAC})$  to win the anonymous blocklisting security game.

Finally, when  $\mathcal{A}$  outputs a forgery  $(\Sigma_G, m)$  or  $(\widehat{\Sigma}_G, \widehat{ct}_G)$ , it parses the authentication token as  $(ct_{SKE}^*, \overline{\Sigma}_{MAC}^*)$  and submits it as its forgery. Due to the winning condition of the security game, the tuple  $(ct_{SKE}^*, \overline{\Sigma}_{MAC}^*)$  was never output by  $O_{Send}$  or  $O_{UpdSend}$ . Moreover, as discussed above, we can assume this tuple was never queried to  $O_{GroupReceive}$  or  $O_{GroupUpdReceive}$ . This implies that  $\mathcal{B}$  never queried to its MAC oracle. Thus the tuple  $(ct_{SKE}^*, \overline{\Sigma}_{MAC}^*)$  forms a valid forgery against the sEUF-CMA security game as desired. This completes the proof.  $\square$

## D More Details on QUASAR

In this section, we provide the omitted details from Sec. 5.1.

### D.1 Formal Description of QUASAR

We provide the formal description of QUASAR in Fig. 9.

### D.2 Security Proof of QUASAR

The following establishes the correctness and security of QUASAR.

**Theorem D.1.** *The GAM protocol QUASAR in Fig. 9 is signing correct, global state-update correct, non-colluding unforgeable, anonymous, anonymous blocklistable, and tracing sound assuming OWF is one-way, PRF and PRP are pseudorandom, MAC is sEUF-CMA secure, and KEM is IND-CCA secure.*

*Proof.* Correctness of signing follows from the construction. In the following, we prove QUASAR is global state-update correct, non-colluding unforgeable, anonymous, anonymous blocklistable, and tracing sound in Lems. D.2 to D.6, respectively. Proving these lemmas completes the proof of the theorem.  $\square$

#### D.2.1 Proof of Lem. D.2

**Lemma D.2.** *QUASAR is global state-update correct.*

*Proof.* Assume user  $u \in G$  used up all of its tokens in epoch, i.e.,  $\Sigma_G = \perp$ . After each user (including  $u$ ) run UpdSend; the server runs UpdVerify on all group update authentication tokens; and the users run UpdReceive on all user update authentication tokens, each user  $v$ 's  $SendSEED_v[\text{epoch} + 1]$  will be updated and hence the epoch in their states will be incremented by one and the counter  $ctr$  will be refreshed to one. Therefore, every user (including  $u$ ) will be able to run Send after ever users update. This shows global state-update correctness.  $\square$

#### D.2.2 Proof of Lem. D.3

**Lemma D.3.** *QUASAR is non-colluding unforgeable assuming OWF is one-way, PRF is pseudorandom, MAC is sEUF-CMA secure, and KEM is IND-CCA secure.*

*Proof.* We first focus on the case the adversary  $\mathcal{A}$  outputs  $(\text{label}, \text{obj}) = (\text{msg}, (v, \sigma, m))$  for  $v \in \mathcal{H}$ ; the other case when  $\text{label} = \text{upd}$  is proven identically. Since it is a valid adversary, the authentication token  $\sigma$  is valid and traces back to some honest user  $u^* \in \mathcal{H}$ . That is,  $\sigma$  is the form  $(\text{epoch}^*, \widetilde{id}_{\text{send}}, \widetilde{x}_{u^* \rightarrow v}^{(t)}, \Sigma_{MAC})$  where  $\widetilde{x}_{u^* \rightarrow v}^{(t)} = \text{PRF}(\text{seed}_{u^* \rightarrow v}, \text{epoch}^* || t || v)$ ,  $(\text{seed}_{u^* \rightarrow v}, t) = \text{ReceiveSEED}_v[\text{epoch}^*][u^*]$ , and  $\text{MAC.Verify}(k_{MAC}, (\widetilde{id}_{\text{send}}, \widetilde{x}_{u^* \rightarrow v}^{(t)}, m), \Sigma_{MAC}) = \top$ . We also have  $(u^*, *, m) \notin L_{\text{msg}}$ .

Recall there are two cases we must consider for non-colluding unforgeability:  $C \neq \emptyset$  (i.e., the server is honest and  $\mathcal{A}$  has access to  $O^*$ ) and  $C = \emptyset$  (i.e., the server is malicious and  $\mathcal{A}$  has access to  $O$ ). Below, our goal is to argue that no  $\mathcal{A}$  can guess  $x$  in the first case and no  $\mathcal{A}$  can create  $\Sigma_{MAC}$  in the second case. Towards this goal, we first invoke IND-CCA security of the KEM to argue that  $\text{seed}_{u^* \rightarrow v}$  for  $u^*, v \in \mathcal{H}$  is distributed uniformly random over  $\{0, 1\}^k$  from the view of  $\mathcal{A}$ .

Game 1: This is the real non-colluding unforgeability game.

Game 2: In this game, when  $\mathcal{A}$  invokes  $O_{UpdSend}(v)$  on  $\text{epoch}^* - 1$ , it stores a random  $\text{seed}_{u^* \rightarrow v}$  in  $\text{ReceiveSEED}_v[\text{epoch}^*][u^*]$ . Note that in the previous game, it generated a user update information  $(\text{seed}_{u^* \rightarrow v}, ct_{u^*}) \leftarrow \text{sKEM.Enc}(ek_{u^*})$  and stored  $\text{seed}_{u^* \rightarrow v}$  in  $\text{ReceiveSEED}_v[\text{epoch}^*][u^*]$ . Moreover, in this game, when user  $u^*$  is invoked on UpdReceive with user update information  $ct_{u^*}$ , it simply uses  $\text{seed}_{u^* \rightarrow v}$  rather than decrypting it by  $\text{KEM.Dec}(dk_{u^*}, ct_{u^*})$ . Otherwise, the game proceeds identically to the previous game.

<p><b>Init(<math>1^K, G</math>)</b></p> <ol style="list-style-type: none"> <li>1: <math>\text{gsk} \leftarrow \mathcal{S}\{0, 1\}^K</math> / Group secret key</li> <li>2: <math>k_{\text{MAC}} \leftarrow \text{PRF}(\text{gsk}, 0)</math> / MAC key for group</li> <li>3: <math>K \leftarrow \text{PRF}(\text{gsk}, 1)</math> / PRF key for group</li> <li>4: / Prepare empty list and public key for users</li> <li>5: <b>foreach</b> <math>u \in G</math> <b>do</b></li> <li>6:   <math>\text{SendSEED}_u[*], \text{ReceiveSEED}_u[*] := \perp</math></li> <li>7:   <math>(\text{ek}_u, \text{dk}_u) \leftarrow \text{KEM.KeyGen}(1^K)</math></li> <li>8:   <math>\text{EK} \leftarrow (u, \text{ek}_u)_{u \in G}</math> / <math>\text{EK}[u] = \text{ek}_u</math></li> <li>9:   <math>\text{epoch} := 0</math></li> <li>10: <math>k_{\text{epoch}} \leftarrow \text{PRF}(K, \text{epoch})</math> / Permutation keys for tokens</li> <li>11: <b>foreach</b> <math>u \in G</math> <b>do</b></li> <li>12:   <math>(\text{ReceiveSEED}_u, Y_u, (\text{ct}_v)_{v \in G})</math>            <math>\leftarrow * \text{gen-auth-token}(G, \text{EK}, \text{epoch}, u, \text{ReceiveSEED}_u)</math></li> <li>13:   <b>foreach</b> <math>v \in G</math> <b>do</b></li> <li>14:     <math>(\text{seed}_{v \rightarrow u}, 1) \leftarrow \text{ReceiveSEED}_u[\text{epoch}][v]</math></li> <li>15:     <math>\text{SendSEED}_v[\text{epoch}][u] \leftarrow \text{seed}_{v \rightarrow u}</math></li> <li>16:   <b>foreach</b> <math>u \in G</math> <b>do</b></li> <li>17:     <math>\text{SEED}_u := (\text{SendSEED}_u, \text{ReceiveSEED}_u)</math></li> <li>18:     <math>\text{st}_u \leftarrow (G, k_{\text{MAC}}, K, \text{EK}, \text{dk}_u, \text{epoch}, 1, \text{SEED}_u)</math></li> <li>19:   <math>\text{DB}[*] := \perp</math> / Prepare empty database for Sv</li> <li>20:   <b>for</b> <math>(u, j) \in G \times [NT]</math> <b>do</b></li> <li>21:     <math>\tilde{j} \leftarrow \text{PRP}(k_{\text{epoch}}, j)</math> / Reorder public tokens and store it in DB</li> <li>22:     <math>\text{DB}[\text{epoch}][\text{id}_X(u)][\tilde{j}] \leftarrow Y_u[j]</math></li> <li>23:   <math>\text{pp} \leftarrow \text{DB}</math> / <math>\text{DB}[0] \in (\{0, 1\}^*)^{N \times NT}</math></li> <li>24:   <b>return</b> <math>(\text{pp}, (\text{st}_u)_{u \in G})</math></li> </ol>	<p><b>Send(<math>\text{st}_u, m</math>)</b></p> <ol style="list-style-type: none"> <li>1: <b>parse</b> <math>(G, k_{\text{MAC}}, K, \text{EK}, \text{dk}_u, \text{epoch}, \text{ctr}, \text{SEED}_u) \leftarrow \text{st}_u</math></li> <li>2: <b>if</b> <math>\text{ctr} \geq T - 1</math> <b>then return</b> <math>\perp</math> / Need to update tokens</li> <li>3: <math>\text{ctr}' \leftarrow \text{ctr} + 1</math></li> <li>4: <math>\Sigma_G \leftarrow (\text{epoch}, * \text{attach-auth-token}(\text{st}_u, \text{msg} :: m))</math></li> <li>5: <math>\text{st}'_u \leftarrow (G, k_{\text{MAC}}, K, \text{EK}, \text{dk}_u, \text{epoch}, \text{ctr}', \text{SEED}_u)</math></li> <li>6: <b>return</b> <math>(\text{st}'_u, \Sigma_G)</math></li> </ol> <p><b>Verify(<math>\text{pp}, \Sigma_G, m</math>)</b></p> <ol style="list-style-type: none"> <li>1: <b>parse</b> <math>\text{DB} \leftarrow \text{pp}</math></li> <li>2: <b>try</b> <math>(\text{pp}', (\sigma_i)_{i \in [N]}) \leftarrow * \text{verify-auth-token}(\text{pp}, \Sigma_G)</math></li> <li>3: <b>return</b> <math>(\text{pp}', (\sigma_i)_{i \in [N]})</math></li> </ol> <p><b>Receive(<math>\text{st}_u, \sigma, m</math>)</b></p> <ol style="list-style-type: none"> <li>1: <b>try</b> <math>(\text{st}'_u, b, v) \leftarrow * \text{trace-sender}(\text{st}_u, \sigma)</math></li> <li>2: <b>return</b> <math>(\text{st}'_u, b, v)</math></li> </ol>
<p><b>UpdSend(<math>\text{st}_u</math>)</b></p> <ol style="list-style-type: none"> <li>1: <b>parse</b> <math>(G, k_{\text{MAC}}, K, \text{EK}, \text{dk}_u, \text{epoch}, \text{ctr}, \text{SEED}_u) \leftarrow \text{st}_u</math></li> <li>2: / Cannot update again if others haven't</li> <li>3: <b>if</b> <math>\text{ctr} = T</math> <b>then return</b> <math>\perp</math></li> <li>4: <b>parse</b> <math>(\text{SendSEED}_u, \text{ReceiveSEED}_u) \leftarrow \text{SEED}_u</math></li> <li>5: <math>\text{epoch}' \leftarrow \text{epoch} + 1</math> / Move to new epoch</li> <li>6: <math>k_{\text{epoch}' } \leftarrow \text{PRF}(K, \text{epoch}')</math></li> <li>7: / Create seeds/tokens for new epoch</li> <li>8: <math>(\text{ReceiveSEED}_u, Y_u, (\text{ct}_v)_{v \in G})</math>            <math>\leftarrow * \text{gen-auth-token}(G, \text{EK}, \text{epoch}', u, \text{ReceiveSEED}_u)</math></li> <li>9: <math>\tilde{Y}_u[*] := \perp</math> / Reorder public tokens to be uploaded to Sv</li> <li>10: <b>for</b> <math>j \in [NT]</math> <b>do</b></li> <li>11:   <math>\tilde{j} \leftarrow \text{PRP}(k_{\text{epoch}'}, j)</math></li> <li>12:   <math>\tilde{Y}_u[\tilde{j}] \leftarrow Y_u[j]</math></li> <li>13:   <math>\hat{\text{ct}}_G \leftarrow (\text{id}_X(u), \tilde{Y}_u, (\text{ct}_v)_{v \in G})</math></li> <li>14:   <math>\hat{\Sigma}_G \leftarrow (\text{epoch}, * \text{attach-auth-token}(\text{st}_u, \text{upd} :: (\text{ct}_v)_{v \in G}))</math></li> <li>15:   <math>\text{SendSEED}_u[\text{epoch}] \leftarrow \perp</math> / Remove seeds from previous epoch</li> <li>16: / Update own seed/private tokens for next epoch</li> <li>17: <math>\text{SendSEED}_u[\text{epoch}'][u] \leftarrow \text{ReceiveSEED}_u[\text{epoch}'][u]</math></li> <li>18: <math>\text{SEED}_u \leftarrow (\text{SendSEED}_u, \text{ReceiveSEED}_u)</math></li> <li>19: / Max out counter to T but do not increment epoch yet</li> <li>20: <math>\text{st}'_u \leftarrow (G, k_{\text{MAC}}, K, \text{EK}, \text{dk}_u, \text{epoch}, T, \text{SEED}_u)</math></li> <li>21: <b>return</b> <math>(\text{st}'_u, \hat{\Sigma}_G, \hat{\text{ct}}_G)</math></li> </ol>	<p><b>UpdVerify(<math>\text{pp}, \hat{\Sigma}_G, \hat{\text{ct}}_G</math>)</b></p> <ol style="list-style-type: none"> <li>1: <b>parse</b> <math>\text{DB} \leftarrow \text{pp}</math></li> <li>2: <b>try</b> <math>(\text{pp}', (\hat{\sigma}_i)_{i \in [N]}) \leftarrow * \text{verify-auth-token}(\text{pp}, \text{upd} :: (\hat{\Sigma}_G, \hat{\text{ct}}_G))</math></li> <li>3: <b>parse</b> <math>(\text{epoch}, \hat{\text{id}}_{\text{send}}, (x_i, \Sigma_{\text{MAC}, i})_{i \in [N]}) \leftarrow \hat{\Sigma}_G</math></li> <li>4: <b>parse</b> <math>(\text{id}_{\text{receive}}, \tilde{Y}, (\text{ct}_i)_{i \in [N]}) \leftarrow \hat{\text{ct}}_G</math></li> <li>5: <b>req</b> <math>\text{DB}[\text{epoch} + 1][\text{id}_{\text{receive}}] = \perp</math> / No public token for <math>\text{id}_{\text{receive}}</math> in <math>\text{epoch} + 1</math></li> <li>6: <math>\text{DB}[\text{epoch} + 1][\text{id}_{\text{receive}}] \leftarrow \tilde{Y}</math> / Update DB for next epoch</li> <li>7: <b>foreach</b> <math>i \in [N]</math> <b>do</b></li> <li>8:   <math>\hat{\text{ct}}_i \leftarrow \text{ct}_i</math></li> <li>9: <b>return</b> <math>(\text{pp}', (\hat{\sigma}_i, \hat{\text{ct}}_i)_{i \in [N]})</math></li> </ol> <p><b>UpdReceive(<math>\text{st}_u, \hat{\sigma}, \hat{\text{ct}}</math>)</b></p> <ol style="list-style-type: none"> <li>1: <b>parse</b> <math>(G, k_{\text{MAC}}, K, \text{EK}, \text{dk}_u, \text{epoch}, \text{ctr}, \text{SEED}_u) \leftarrow \text{st}_u</math></li> <li>2: <b>parse</b> <math>(\text{SendSEED}_u, \text{ReceiveSEED}_u) \leftarrow \text{SEED}_u</math></li> <li>3: <b>try</b> <math>(\text{st}'_u, b, v) \leftarrow * \text{trace-sender}(\text{st}_u, \hat{\sigma}, \hat{\text{ct}})</math></li> <li>4: <b>if</b> <math>v = u</math> <b>then return</b> <math>(\text{st}_u, T, u)</math> / <math>\text{SendSEED}_u</math> is already updated with <math>\text{seed}_{u \rightarrow u}</math></li> <li>5: <b>req</b> <math>\text{SendSEED}_u[\text{epoch} + 1][v] = \perp</math> / <math>v</math> hasn't updated yet in <math>\text{epoch} + 1</math></li> <li>6: <math>\text{seed}_{u \rightarrow v} \leftarrow \text{KEM.Dec}(\text{dk}_u, \hat{\text{ct}})</math> / Seed used by <math>u</math> to send message to <math>v</math> in <math>\text{epoch} + 1</math></li> <li>7: <math>\text{SendSEED}_u[\text{epoch} + 1][v] \leftarrow \text{seed}_{u \rightarrow v}</math></li> <li>8: <math>\text{SEED}_u \leftarrow (\text{SendSEED}_u, \text{ReceiveSEED}_u)</math></li> <li>9: <b>if</b> <math>\forall w \in G, \text{SendSEED}_u[\text{epoch} + 1][w] \neq \perp</math> <b>then</b></li> <li>10:   / Increment epoch and refresh counter to 1 if everybody updated</li> <li>11:   <math>\text{st}'_u \leftarrow (G, k_{\text{MAC}}, K, \text{EK}, \text{dk}_u, \text{epoch} + 1, 1, \text{SEED}_u)</math></li> <li>12: <b>else</b> <math>\text{st}'_u \leftarrow (G, k_{\text{MAC}}, K, \text{EK}, \text{dk}_u, \text{epoch}, \text{ctr}, \text{SEED}_u)</math></li> <li>13: <b>return</b> <math>(\text{st}'_u, T, v)</math></li> </ol>

Figure 9: QUASAR: An anonymous group authenticated messaging protocol with tracing soundness.  $\text{label} :: \text{obj}$  denotes that  $\text{obj}$  has a type label, where  $\text{label}$  is a special string used nowhere else. The helper algorithms used above are detailed in Fig. 10



Func *gen-auth-token( $G, EK, \text{epoch}, u, \text{ReceiveSEED}$ )	Func *attach-auth-token( $\text{st}_u, \text{label} :: \text{obj}$ )
<pre> 1: req ReceiveSEED[epoch] = <math>\perp</math> 2: foreach <math>v \in G</math> do 3:   (seed<math>_{v \rightarrow u}, \text{ct}_v</math>) <math>\leftarrow</math> <math>\\$</math>KEM.Enc(EK[v]) 4:   ReceiveSEED[epoch][v] <math>\leftarrow</math> (seed<math>_{v \rightarrow u}, 1</math>) 5:   foreach <math>t \in [T]</math> do 6:     // <math>t</math>-th private token <math>v</math> sends to <math>u</math> in epoch 7:     <math>x_{v \rightarrow u}^{(t)} \leftarrow</math> PRF(seed<math>_{v \rightarrow u}, \text{epoch}    t    u</math>) 8:     <math>y_{v \rightarrow u}^{(t)} \leftarrow</math> OWF(<math>x_{v \rightarrow u}^{(t)}</math>) // Corresponding <math>t</math>-th public token 9:   // Arrange all public tokens in order of users <math>G = (u_1, \dots, u_N)</math> 10:  <math>Y \leftarrow (y_{u_1 \rightarrow u}^{(1)}, \dots, y_{u_1 \rightarrow u}^{(T)},</math> 11:    <math>\dots, y_{u_N \rightarrow u}^{(1)}, \dots, y_{u_N \rightarrow u}^{(T)}) \in (\{0, 1\}^k)^{NT}</math> 12:  return (ReceiveSEED, <math>Y, (\text{ct}_v)_{v \in G}</math>) </pre>	<pre> 1: parse (<math>G, k_{\text{MAC}}, K, EK, dk_u, \text{epoch}, \text{ctr}, \text{SEED}_u</math>) <math>\leftarrow</math> <math>\text{st}_u</math> 2: parse (SendSEED<math>_u, \text{ReceiveSEED}_u</math>) <math>\leftarrow</math> SEED<math>_u</math> 3: <math>k_{\text{epoch}} \leftarrow</math> PRF(<math>K, \text{epoch}</math>) 4: foreach <math>v \in G</math> do 5:   seed<math>_{u \rightarrow v} \leftarrow</math> SendSEED<math>_u[\text{epoch}][v]</math> 6:   <math>x_{u \rightarrow v}^{(\text{ctr})} \leftarrow</math> PRF(seed<math>_{u \rightarrow v}, \text{epoch}    \text{ctr}    v</math>) // <math>\text{ctr}</math>-th private token <math>u</math> sends to <math>v</math> 7:   if label = msg then 8:     parse <math>m \leftarrow</math> obj 9:     <math>\Sigma_{\text{MAC}, v} \leftarrow</math> <math>\\$</math>MAC.TagGen(<math>k_{\text{MAC}}, (x_{u \rightarrow v}^{(\text{ctr})}, m)</math>) 10:  elseif label = obj then 11:    parse (<math>\text{ct}_v</math>)<math>_{v \in G} \leftarrow</math> obj 12:    // Sign <math>v</math>'s ciphertext using otsk from <math>v</math> 13:    <math>\Sigma_{\text{MAC}, v} \leftarrow</math> <math>\\$</math>MAC.TagGen(<math>k_{\text{MAC}}, (x_{u \rightarrow v}^{(\text{ctr})}, \text{ct}_v)</math>) 14:  // Compute permuted index of the corresponding public token in DB 15:  id<math>_{\text{send}} \leftarrow</math> ctr + (idx(<math>u</math>) - 1) <math>\cdot</math> <math>T</math> 16:  <math>\tilde{\text{id}}_{\text{send}} \leftarrow</math> PRP(<math>k_{\text{epoch}}, \text{id}_{\text{send}}</math>) // id<math>_{\text{send}}, \tilde{\text{id}}_{\text{send}} \in [NT]</math> 17:  return (<math>\tilde{\text{id}}_{\text{send}}, (x_{u \rightarrow v}^{(\text{ctr})}, \Sigma_{\text{MAC}, v})_{v \in G}</math>) </pre>
<pre> Func *verify-auth-token(<math>\text{pp}, \Sigma_G</math>) 1: DB <math>\leftarrow</math> pp // <math>Y \in (\{0, 1\}^k)^{N \times NT}</math> 2: parse (epoch, <math>\tilde{\text{id}}_{\text{send}}, (x_i, \Sigma_{\text{MAC}, i})_{i \in [N]}</math>) <math>\leftarrow</math> <math>\Sigma_G</math> 3: foreach <math>i \in [N]</math> do 4:   if DB[epoch][<math>\tilde{\text{id}}_{\text{send}}</math>][<math>i</math>] <math>\neq</math> OWF(<math>x_i</math>) then 5:     return <math>\perp</math> 6: // If check passes, set user authentication tokens and 7: // delete column to update DB 8: foreach <math>i \in [N]</math> do 9:   <math>\sigma_i \leftarrow</math> (epoch, <math>\tilde{\text{id}}_{\text{send}}, x_i, \Sigma_{\text{MAC}, i}</math>) 10:  DB[epoch][<math>\tilde{\text{id}}_{\text{send}}</math>][<math>i</math>] <math>\leftarrow</math> <math>\perp</math> 11:  <math>\text{pp}' \leftarrow</math> DB 12:  return (<math>\text{pp}', (\sigma_i)_{i \in [N]}</math>) </pre>	<pre> Func *trace-sender(<math>\text{st}_u, \sigma, m</math>) 1: parse (<math>G, k_{\text{MAC}}, K, EK, dk_u, \text{epoch}, \text{ctr}, \text{SEED}_u</math>) <math>\leftarrow</math> <math>\text{st}_u</math> 2: parse (SendSEED<math>_u, \text{ReceiveSEED}_u</math>) <math>\leftarrow</math> SEED<math>_u</math> 3: parse (epoch', <math>\tilde{\text{id}}_{\text{send}}, x, \Sigma_{\text{MAC}}</math>) <math>\leftarrow</math> <math>\sigma</math> 4: req epoch = epoch' 5: <math>k_{\text{epoch}} \leftarrow</math> PRF(<math>K, \text{epoch}</math>) 6: id<math>_{\text{send}} \leftarrow</math> PRP<math>^{-1}</math>(<math>k_{\text{epoch}}, \tilde{\text{id}}_{\text{send}}</math>) 7: <math>i \leftarrow \lfloor \text{id}_{\text{send}} / T \rfloor + 1</math> // Index of the user who sent the message 8: <math>t \leftarrow \text{id}_{\text{send}} - \lfloor \text{id}_{\text{send}} / T \rfloor \cdot T</math> // <math>t</math>-th token from user in epoch 9: <math>v \leftarrow</math> <math>i</math>-th user in <math>G</math> // <math>v \in G</math> and id<math>_{\text{send}} = t + (\text{idx}(v) - 1) \cdot T</math> 10: (seed<math>_{v \rightarrow u}, \text{ctr}</math>) <math>\leftarrow</math> ReceiveSEED<math>_u[\text{epoch}][v]</math> 11: req ctr = <math>t</math> // Require that token wasn't used yet 12: <math>x_{v \rightarrow u}^{(t)} \leftarrow</math> PRF(seed<math>_{v \rightarrow u}, \text{epoch}    t    u</math>) 13: if <math>x_{v \rightarrow u}^{(t)} \neq x</math> 14:   <math>\vee</math> MAC.Verify(<math>k_{\text{MAC}}, (\tilde{\text{id}}_{\text{send}}, x_{v \rightarrow u}^{(\text{ctr})}, m), \Sigma_{\text{MAC}})</math>) = <math>\perp</math> then 15:   return (<math>\text{st}_u, \perp, \perp</math>) 16: ReceiveSEED<math>_u[\text{epoch}][v] \leftarrow</math> (seed<math>_{v \rightarrow u}, \text{ctr} + 1</math>) // Mark <math>t</math>-th private token to be used 17: SEED<math>_u \leftarrow</math> (SendSEED<math>_u, \text{ReceiveSEED}_u</math>) 18: <math>\text{st}'_u \leftarrow</math> (<math>G, k_{\text{MAC}}, K, EK, dk_u, \text{epoch}, \text{ctr}, \text{SEED}_u</math>) 19: return (<math>\text{st}'_u, T, v</math>) </pre>

Figure 10: Helper functions used by QUASAR.

These two games are indistinguishable assuming the IND-CCA security of the KEM. Concretely, an adversary  $\mathcal{B}$  against the IND-CCA security game is given  $(ek^*, ct^*, seed^*)$  as its challenge. It then embeds  $ek^*$  into  $ek_{u^*}$  and simulates the non-colluding unforgeability security game to  $\mathcal{A}$ . When  $\mathcal{A}$  invokes  $O_{\text{UpdSend}}(v)$  on epoch $^*$ , it sets the user update information  $ct_{u^*}$  as  $ct^*$  and stores  $seed^*$  in  $\text{ReceiveSEED}_v[\text{epoch}^*][u^*]$ . Moreover, when user  $u^*$  is invoked on  $\text{UpdReceive}$  with user update information  $ct'_{u^*}$ , it queries its decryption oracle if  $ct'_{u^*} \neq ct^*$ . Otherwise, it simply sets the decapsulated value as  $seed^*$  without running the decapsulation algorithm. Assuming the  $(1 - \delta)$ -correctness,  $\mathcal{B}$  correctly simulates Game 1 (resp. Game 2) when the challenge bit is 1 (resp. 0) as desired.

We are now ready to invoke the pseudorandomness of the PRF.

**Game 3:** In this game, when  $\mathcal{A}$  invokes  $O_{\text{UpdSend}}(v)$  on epoch $^* - 1$ , it samples  $T$  random private tokens for  $u^*$ :  $\bar{x}_{u^* \rightarrow v}^{(t)} \leftarrow \{0, 1\}^k$  for  $t \in [T]$  (cf. line 7 of `*gen-auth-token` in Fig. 10). Note that in the previous game, it generated them by using  $seed_{u^* \rightarrow v}$  and the PRF. Moreover, in this game, when user  $u^*$  runs  $\text{Send}$  or  $\text{UpdSend}$  with  $\text{SendSEED}_{u^*}[\text{epoch}][w] = seed_{u^* \rightarrow v}$  for some  $w \in G$ , it uses  $\left(\bar{x}_{u^* \rightarrow v}^{(t)}\right)_{t \in [T]}$  when epoch = epoch $^*$  rather than generating them via the PRF. Otherwise, if epoch  $\neq$  epoch $^*$ , it samples  $T$  random fresh private tokens for this epoch and uses them. Here, note that we do not necessarily have  $(w, \text{epoch}) = (u, \text{epoch}^*)$  since a malicious user  $w$  can reuse the user update information  $ct_{u^*}$  that  $v$  sent to  $u^*$  at some later epoch  $>$  epoch $^*$ .

Finally, when `*trace-sender`( $st_u, \sigma, m$ ) with  $\sigma = (\text{epoch}^*, \tilde{id}_{\text{send}}, x, \Sigma_{\text{MAC}})$  is invoked it checks if  $\tilde{id}_{\text{send}}$  is the permuted index of user  $u^*$ 's  $\text{ctr}$ -th token (cf. lines 7 to 9 of `*trace-sender` in Fig. 10), where  $(\text{ReceiveSEED}_v[\text{epoch}^*][u^*] = (seed_{u^* \rightarrow v}, \text{ctr}))$ . If not, it proceeds as in Game 2. Otherwise, it checks if  $x = \bar{x}_{u^* \rightarrow v}^{(\text{ctr})}$  rather than using the PRF. Otherwise, the game proceeds identically to the previous game.

Assuming the pseudorandomness of the PRF, it is clear that Game 2 and Game 3 are indistinguishable.

We now make two case distinctions based on  $C \neq \emptyset$  or not. Let us consider the first case where the adversary is the set of malicious sender and the server is honest. In Game 3, notice the private tokens  $\left(\bar{x}_{u^* \rightarrow v}^{(t)}\right)_{t \in [T]}$  in epoch $^*$  are distributed uniformly random over  $(\{0, 1\}^k)^T$  and shared only between the honest users  $u^*$  and  $v$ . Due to the one-wayness of OWF, the corresponding public tokens  $\left(y_{u^* \rightarrow v}^{(t)} = \text{OWF}(\bar{x}_{u^* \rightarrow v}^{(t)})\right)_{t \in [T]}$  do not leak the private tokens in any meaningful way. Moreover, since each private tokens are tied to epoch $^*$  and a counter  $t \in [T]$ ,  $\bar{x}_{u^* \rightarrow v}^{(t)}$  is only revealed when  $u^*$  is invoked on the  $t$ -th

$\text{Send}$  or  $\text{UpdSend}$  in epoch $^*$ . As explained above,  $u^*$  may reveal the same private token  $\bar{x}_{u^* \rightarrow v}^{(t)}$  multiple times in epoch $^*$  on the  $t$ -th invocation of  $\text{Send}$  or  $\text{UpdSend}$  since a malicious user  $w$  can set  $\bar{x}_{u^* \rightarrow v}^{(t)} = \bar{x}_{u^* \rightarrow v}^{(t)}$  by replaying the ciphertext  $ct_{u^*}$ . However, as long as the adversary  $\mathcal{A}$  cannot reuse the private token for a different  $t' \neq t$  or epoch $' \neq$  epoch $^*$ , this does not constitute in an attack. Indeed, since we assume an honest server, the private tokens sent by  $v$  for a particular pair (epoch,  $t$ ) is processed correctly, and in particular, the adversary cannot reuse them in a different (epoch $'$ ,  $t'$ )  $\neq$  (epoch,  $t$ ). Thus, we conclude that  $\mathcal{A}$  cannot output a private token that hasn't been used by the honest user  $u^*$ .

Let us consider the second case where the adversary is the malicious server but all users are honest. The hardness almost immediately follows from assuming sEUF-CMA security of the MAC as the malicious server does not know the MAC key  $k_{\text{MAC}}$ . Concretely, if  $(u^*, *, m) \notin L_{\text{msg}}$ , then user  $u^*$  will have never signed the tuple  $(\bar{x}_{u^* \rightarrow v}^{(t)}, m)$ . Moreover, since  $\bar{x}_{u^* \rightarrow v}^{(t)}$  is now sampled uniformly random over  $\{0, 1\}^k$ , the probability that another user  $w \neq u^*$  signed the tuple is negligible. Therefore, we conclude that if an adversary outputs a MAC tag  $\Sigma_{\text{MAC}}$  for such a tuple, it breaks sEUF-CMA security. Here, to simulate the sEUF-CMA security game, we rely on the pseudorandomness of the PRF to sample the MAC key  $k_{\text{MAC}}$  directly without running PRF.

This completes the proof of the lemma.  $\square$

## D.2.3 Proof of Lem. D.4

**Lemma D.4.** *QUASAR is anonymous assuming assuming PRF and PRP are pseudorandom and KEM is IND-CCA secure.*

*Proof.* Before we get into the proof, we make some useful observations. Due to global state-update correctness, if  $u_0$  and  $u_1$  can both run  $\text{Send}$ , the generated global authentication tokens  $\Sigma_G^0$  and  $\Sigma_G^1$  must contain the same epoch $^*$ . Moreover, due to  $\mathcal{A}$ 's winning condition that all the users verify, every user in  $G$  must have been invoked on  $O_{\text{UpdSend}}$  at epoch $^* - 1$  as otherwise, they will not be able to verify  $u_0$  and  $u_1$ 's user authentication token at epoch $^*$ . With this in mind, we perform several game hybrids, where anonymity trivially holds in the final game. The first two game transition is almost identical to those done in the proof of non-colluding unforgeability.

Game 1: This is the real anonymity game.

Game 2: In this game, for any  $u \in G$ , when  $\mathcal{A}$  invokes  $O_{\text{UpdSend}}(u)$  on epoch $^* - 1$ , it stores a random  $seed_{v \rightarrow u}$  in  $\text{ReceiveSEED}_u[\text{epoch}^*][v]$  for all  $v \in G$ . Note that in the previous game, it generated a user update information  $(seed_{v \rightarrow u}, ct_v) \leftarrow \text{sKEM.Enc}(ek_v)$  and stored  $seed_{v \rightarrow u}$  in  $\text{ReceiveSEED}_u[\text{epoch}^*][v]$ . Moreover, in this game, when user  $v$  is invoked on  $\text{UpdReceive}$  with user update information  $ct_v$ , it simply uses  $seed_{v \rightarrow u}$  rather

than decrypting it by running  $\text{KEM.Dec}(\text{dk}_v, \text{ct}_v)$ . Otherwise, the game proceeds identically to the previous game.

Following the same argument made to move from Game 1 to Game 2 in Lem. D.3, the two games are indistinguishable assuming the IND-CCA security of the KEM. We next invoke the pseudorandomness of the PRF.

**Game 3:** In this game, for any  $u \in G$ , when  $\mathcal{A}$  invokes  $O_{\text{UpdSend}}(u)$  on epoch<sup>\*</sup> - 1, it samples  $T$  random private tokens for every  $v \in G$ :  $\tilde{x}_{v \rightarrow u}^{(t)} \leftarrow \{0, 1\}^k$  for  $t \in [T]$  (cf. line 7 of `*gen-auth-token` in Fig. 10). Note that in the previous game, it generated them by using  $\overline{\text{seed}}_{v \rightarrow u}$  and the PRF. Moreover, in this game, when any user  $v$  runs `Send` or `UpdSend` with  $\text{SendSEED}_v[\text{epoch}][w] = \overline{\text{seed}}_{v \rightarrow u}$ , it uses  $\left(\tilde{x}_{v \rightarrow u}^{(t)}\right)_{t \in [T]}$  when epoch = epoch<sup>\*</sup> rather than generating them via the PRF. Otherwise, if epoch  $\neq$  epoch<sup>\*</sup>, it samples  $T$  random fresh private tokens for this epoch and uses them. Here, note that we do not necessarily have  $(w, \text{epoch}) = (u, \text{epoch}^*)$  since the malicious server may reuse the user update information  $\text{ct}_v$  that  $u$  sent to  $v$  by modifying user  $w$ 's user update information — this is due to group update authentication tokens not being cryptographically tied to the group update information.

Finally, when `*trace-sender`( $\text{st}_u, \sigma, m$ ) with  $\sigma = (\text{epoch}^*, \tilde{\text{id}}_{\text{send}}, x, \Sigma_{\text{MAC}})$  is invoked (as part of `Receive` or `UpdReceive`) it checks if  $\tilde{\text{id}}_{\text{send}}$  is the permuted index of user  $v$ 's ctr-th token (cf. lines 7 to 9 of `*trace-sender` in Fig. 10), where  $(\text{ReceiveSEED}_u[\text{epoch}^*][v] = (\overline{\text{seed}}_{v \rightarrow u}, \text{ctr}))$ . If not, it proceeds as in Game 2. Otherwise, it checks if  $x = \tilde{x}_{v \rightarrow u}^{(\text{ctr})}$  rather than using the PRF. Otherwise, the game proceeds identically to the previous game.

Assuming the pseudorandomness of the PRF, it is clear that Game 2 and Game 3 are indistinguishable. At this point, the private tokens  $\left(\tilde{x}_{v \rightarrow u}^{(t)}\right)_{(t,v,u) \in [T] \times G^2}$  generated via  $O_{\text{UpdSend}}$  in epoch<sup>\*</sup> are independently distributed. We next argue that the challenge users  $u_b \in G$ ,  $b \in \{0, 1\}$  must also receive these private tokens  $\left(\tilde{x}_{u_b \rightarrow v}^{(t)}\right)_{(t,v) \in [T] \times G}$ . Namely, we consider the following game.

**Game 4:** In this game, the game aborts if there exists a user  $v \in G$  such that the following holds for either  $b \in \{0, 1\}$ :

$$\text{ReceiveSEED}_v[\text{epoch}^*][u_b] \neq \text{SendSEED}_{u_b}[\text{epoch}^*][v]. \quad (2)$$

Put differently, the game aborts if the user update information  $\text{ct}_v$  from  $v$  to  $u_b$  on epoch<sup>\*</sup> was modified.

We argue that this modification has no impact on  $\mathcal{A}$ 's advantage. Assume if the adversary  $\mathcal{A}$  won under this condition. Due to the winning condition of the anonymity game, the group authentication token generated by  $u_0$  and  $u_1$  must be accepted by user  $v$ . As we established in Game 3, the private tokens  $\left(\tilde{x}_{u_b \rightarrow v}^{(t)}\right)_{(t,b) \in [T] \times \{0,1\}}$  stored in  $v$ 's state at epoch<sup>\*</sup> are distributed uniformly random over  $\{0, 1\}^k$  from the view of  $\mathcal{A}$ . Hence, if Equation (2) holds, then the probability that  $v$  accepts  $u_b$ 's private token is at most  $1/2^k$ . Thus, the advantage of  $\mathcal{A}$  cannot change with all but a negligible probability.

At this point, we have that the private tokens included in the group authentication tokens of  $u_0$  and  $u_1$  are uniformly distributed. We finally invoke the pseudorandomness of the PRF one last time along with that of the PRP.

**Game 5:** In this game, the permutation key  $k_{\text{epoch}}$  at each epoch is no longer sampled. Instead, it samples a uniformly random permutation over  $[NT]$  at each epoch and uses that instead. Otherwise, it is identical to the previous game.

Since the group secret key  $\text{gsk}$  is hidden to the adversary, we can first invoke the pseudorandomness of the PRF to argue that the permutation key at each epoch is uniform random and independent. Then, we invoke the pseudorandomness of the PRP to replace the PRP by a truly random permutation.

By Game 5,  $\tilde{\text{id}}_{\text{seed},0}$  and  $\tilde{\text{id}}_{\text{seed},1}$  included in the group authentication tokens  $\Sigma_G^0$  and  $\Sigma_G^1$ , respectively, no longer leak the identity of  $u_0$  and  $u_1$ . Thus combining everything, we conclude that  $\Sigma_G^0$  and  $\Sigma_G^1$  no longer leak the identity of  $u_0$  and  $u_1$ . This completes the proof of the lemma.  $\square$

#### D.2.4 Proof of Lem. D.5

**Lemma D.5.** *QUASAR is anonymous blocklistable assuming OWF is one-way, PRF is pseudorandom, and KEM is IND-CCA secure.*

*Proof.* The proof follows almost directly from the proof of non-colluding unforgeability (cf. Lem. D.3). In fact, the proof is much simpler since unlike in the non-colluding unforgeability game, we can assume without of generality that the adversary  $\mathcal{A}$  never queries  $O_{\text{Receive}}$  or  $O_{\text{UpdReceive}}$  such that `Verify` or `UpdVerify` verify, respectively. This is because such a query can be used to directly win the anonymous blocklistable game since the private tokens can only be used once. (Concretely, we only require IND-CPA security of the KEM). Then, following an almost exact proof, we can argue that the only information leakage of the private tokens corresponding to the public tokens stored inside DB maintained by the server are only from the public parameter  $\text{pp} = \text{DB}$ . Hence, under the one-wayness of OWF, no  $\mathcal{A}$  can output a group (update) authentication token for which the server verifies.  $\square$

### D.2.5 Proof of Lem. D.6

**Lemma D.6.** QUASAR is (unconditionally) tracing sound.

*Proof.* We first focus on the case the adversary  $\mathcal{A}$  outputs  $(\text{label}, \text{obj}) = (\text{msg}, (\Sigma_G, m))$ . The game then runs  $(\text{pp}', (\sigma_i)_{i \in [N]}) \leftarrow \text{Verify}(\text{pp}, \Sigma_G, m)$ . Since  $\text{pp}' \neq \perp$ ,  $\Sigma_G = \left( \text{epoch}, \tilde{\text{id}}_{\text{send}}, (x_i, \Sigma_{\text{MAC}, i}) \right)_{i \in [N]}$  for some  $\tilde{\text{id}}_{\text{send}} \in [N]$ ,  $\Sigma_{\text{MAC}, i}$  is a valid MAC tag under the message tuple  $(\text{id}_{\text{send}}, x_i, m)$ , and  $\text{DB}[\text{epoch}][i][\tilde{\text{id}}_{\text{send}}] = \text{OWF}(x_i)$  for all  $i \in [N]$ , where DB is the database stored in the public parameter pp. Then, a user authentication token for user  $u \in G \cap \mathcal{H}$  is set as  $\sigma_u = (\text{epoch}, \tilde{\text{id}}_{\text{send}}, x_{\text{id}_x(u)}, \Sigma_{\text{MAC}, \text{id}_x(u)})$ . Noticing that each honest users maintain the same group secret key gsk and  $\sigma_u$  contains the same  $(\text{epoch}, \tilde{\text{id}}_{\text{send}})$ , if  $(\text{st}'_u, b_u, v_u) \leftarrow \text{Receive}(\text{st}_u, \sigma_u, m)$  and  $b_u = \top$ , then  $v_u$  must be the same user for all honest user (cf. lines 7 to 9 of \*trace-sender in Fig. 10). Finally, the case when  $\mathcal{A}$  outputs  $(\text{label}, \text{obj}) = (\text{upd}, (\widehat{\Sigma}_G, \widehat{\text{ct}}_G))$  is proven identically. This completes the proof of the lemma.  $\square$

## D.3 Alternatives to Global State Updates

Lastly, we discuss below some ideas to mitigate the shortcoming of global state updates discussed in Sec. 5.1.

**Unbalancing the Number of Tokens.** Depending on the group, some users may have a higher frequency of communication than others. In such scenarios, if we allocate all the users the same number of tokens  $T$ , some may take significantly longer to exhaust their budget than others. As a result, those who have already consumed their tokens may face prolonged waiting periods before everyone updates their state.

QUASAR can be easily modified to a protocol where each user is allocated different number of tokens  $(T_u)_{u \in G}$ . In fact, the users can adaptively modify the number of tokens per epoch  $T_u^{\text{epoch}}$ , where one epoch corresponds to one global state updates. This is possible because, in QUASAR, an unlimited number of tokens can be minted from the seed using the PRF. As before, the tokens are then shuffled by using a PRP, whose keys are updated each epoch. While the server can observe the number of total tokens  $\sum_{u \in G} T_u^{\text{epoch}}$  fluctuating between epochs, due to anonymity, it cannot deduce how many tokens are allocated to each user. From the server's perspective, it could be a group where everybody is actively talking (i.e.,  $T_u \approx T_v$  for any  $u, v \in G$ ) or a group where only one user is allocated most of the tokens (i.e.,  $T_u \gg T_v$  for any  $v \in G \setminus \{u\}$ ).

**Relaxing the Global State Updates Restriction.** While unbalancing the number of tokens, explained in App. D.3, mitigates one aspect of the shortcoming of global state updates, it does not solve the leading issue: once a user exhausts its tokens, it must wait till all the users update to be able to send messages again. An approach to address this is to let the users

who exhausted its tokens in epoch\* - 1 to send messages in epoch\* with the limited updates it received. More specifically, assuming a subset of the group  $\overline{G} \subset G$  performed an update for epoch\*, any  $u \in \overline{G}$  can start sending a message using the user authentication tokens received from all the users in  $v \in \overline{G}$ . Visually, the rows in Fig. 7 corresponding to users in  $\overline{G}$  are filled with tokens while the other rows remain empty. In the extreme case when  $u$  is the only user that performed an update for epoch\*, then  $\overline{G} = \{u\}$ . It is worth noting that this approach is possible with QUASAR but not something that any GAM protocol can do.

The upside of this approach is that users can locally update their state via UpdSend and can start sending messages again without waiting for all the other users to perform an update. Any user  $v \in G \setminus \overline{G}$  can at any point perform an update for epoch\* to join  $\overline{G}$ . Moreover, users  $v \in \overline{G}$  can trace back any user authentication token exchanged in epoch\* to a specific user in  $\overline{G}$ , thus achieving a limited scope of tracing soundness.

The downside of this approach is that the anonymity set of the sender is limited to the size of the subgroup  $\overline{G}$ . This is because only the users in  $\overline{G}$  are capable of sending a message with a token from epoch\*. In the aforementioned extreme case, when  $\overline{G} = \{u\}$ , then  $u$  is the only user with tokens from epoch\*. As a result, while the server does not learn who  $u$  is, it can link together any messages  $u$  sent in epoch\*. Another downside is that there are no tracing soundness guarantees for the users  $w \in G$  outside the subgroup  $\overline{G}$  as they have not minted any private tokens for epoch\*. Tracing soundness is restored only once  $w$  updates.

**Combining QUASAR with COSMAC<sup>+</sup>.** When  $|\overline{G}| \ll |G|$ , we have seen that the anonymity for the users in  $\overline{G}$  and the tracing soundness of  $G \setminus \overline{G}$  is weakened. In this case, it may be more appropriate to use COSMAC<sup>+</sup> instead of QUASAR, given that COSMAC<sup>+</sup> is anonymous, more efficient but lacks tracing soundness. To balance these trade-offs, we propose a hybrid protocol that primarily employs QUASAR, but switches to COSMAC<sup>+</sup> when a predetermined number of users have not yet performed an update for the next epoch\*. This approach combines the benefits of both QUASAR and COSMAC<sup>+</sup>.

## E More Details on STARS

In this section, we provide the omitted details from Sec. 5.2.

### E.1 Formal Description of STARS

We provide the formal description of STARS. The main difference between QUASAR is that STARS replaces one-time tokens and MAC tags with one-time signatures (OTS). Specifically, the exchanged messages are signed by the OTS, effectively making the GAM protocol *standard* unforgeable. The difference is highlighted with a box in Fig. 11.

<p><b>Init(<math>1^K, G</math>)</b></p> <ol style="list-style-type: none"> <li>1: <math>\text{gsk} \leftarrow \mathcal{S}\{0, 1\}^K</math> / Group secret key</li> <li>2: / Prepare empty list and public key for users</li> <li>3: <b>foreach</b> <math>u \in G</math> <b>do</b></li> <li>4:   <math>\text{SendSEED}_u[*], \text{ReceiveSEED}_u[*] := \perp</math></li> <li>5:   <math>(\text{ek}_u, \text{dk}_u) \leftarrow \text{KEM.KeyGen}(1^K)</math></li> <li>6:   <math>\text{EK} \leftarrow (u, \text{ek}_u)_{u \in G}</math> / <math>\text{EK}[u] = \text{ek}_u</math></li> <li>7:   <math>\text{epoch} := 0</math></li> <li>8:   <math>k_{\text{epoch}} \leftarrow \text{PRF}(\text{gsk}, \text{epoch})</math> / Permutation keys for OTS otkvs</li> <li>9:   <b>foreach</b> <math>u \in G</math> <b>do</b></li> <li>10:     <math>(\text{ReceiveSEED}_u, Y_u, (\text{ct}_v)_{v \in G})</math>              <math>\leftarrow * \text{gen-auth-token}(G, \text{EK}, \text{epoch}, u, \text{ReceiveSEED}_u)</math></li> <li>11:     <b>foreach</b> <math>v \in G</math> <b>do</b></li> <li>12:       <math>(\text{seed}_{v \rightarrow u}, 1) \leftarrow \text{ReceiveSEED}_u[\text{epoch}][v]</math></li> <li>13:       <math>\text{SendSEED}_v[\text{epoch}][u] \leftarrow \text{seed}_{v \rightarrow u}</math></li> <li>14:     <b>foreach</b> <math>u \in G</math> <b>do</b></li> <li>15:       <math>\text{SEED}_u := (\text{SendSEED}_u, \text{ReceiveSEED}_u)</math></li> <li>16:       <math>\text{st}_u \leftarrow (G, \text{gsk}, \text{EK}, \text{dk}_u, \text{epoch}, 1, \text{SEED}_u)</math></li> <li>17:     <math>\text{DB}[*] := \perp</math> / Prepare empty database for Sv</li> <li>18:     <b>for</b> <math>(u, j) \in G \times [NT]</math> <b>do</b></li> <li>19:       <math>\tilde{j} \leftarrow \text{PRP}(k_{\text{epoch}}, j)</math> / Reorder OTS otkvs and store it in DB</li> <li>20:       <math>\text{DB}[\text{epoch}][\text{idx}(u)][\tilde{j}] \leftarrow Y_u[j]</math></li> <li>21:     <math>\text{pp} \leftarrow \text{DB}</math> / <math>\text{DB}[0] \in \{0, 1\}^{K \times N \times NT}</math></li> <li>22:     <b>return</b> <math>(\text{pp}, (\text{st}_u)_{u \in G})</math></li> </ol>	<p><b>Send(<math>\text{st}_u, m</math>)</b></p> <ol style="list-style-type: none"> <li>1: <b>parse</b> <math>(G, \text{gsk}, \text{EK}, \text{dk}_u, \text{epoch}, \text{ctr}, \text{SEED}_u) \leftarrow \text{st}_u</math></li> <li>2: <b>if</b> <math>\text{ctr} \geq T - 1</math> <b>then return</b> <math>\perp</math> / Need to update tokens</li> <li>3: <math>\text{ctr}' \leftarrow \text{ctr} + 1</math></li> <li>4: <math>\Sigma_G \leftarrow (\text{epoch}, * \text{attach-auth-token}(\text{st}_u, \text{msg} :: m))</math></li> <li>5: <math>\text{st}'_u \leftarrow (G, \text{gsk}, \text{EK}, \text{dk}_u, \text{epoch}, \text{ctr}', \text{SEED}_u)</math></li> <li>6: <b>return</b> <math>(\text{st}'_u, \Sigma_G)</math></li> </ol> <p><b>Verify(<math>\text{pp}, \Sigma_G, m</math>)</b></p> <ol style="list-style-type: none"> <li>1: <b>parse</b> <math>\text{DB} \leftarrow \text{pp}</math></li> <li>2: <b>try</b> <math>(\text{pp}', (\sigma_i)_{i \in [N]})</math>              <math>\leftarrow * \text{verify-auth-token}(\text{pp}, \boxed{\text{msg} :: (\Sigma_G, m)})</math></li> <li>3: <b>return</b> <math>(\text{pp}', (\sigma_i)_{i \in [N]})</math></li> </ol> <p><b>Receive(<math>\text{st}_u, \sigma, m</math>)</b></p> <ol style="list-style-type: none"> <li>1: <b>try</b> <math>(\text{st}'_u, b, v) \leftarrow * \text{trace-sender}(\text{st}_u, \sigma, m)</math></li> <li>2: <b>return</b> <math>(\text{st}'_u, b, v)</math></li> </ol>
<p><b>UpdSend(<math>\text{st}_u</math>)</b></p> <ol style="list-style-type: none"> <li>1: <b>parse</b> <math>(G, \text{gsk}, \text{EK}, \text{dk}_u, \text{epoch}, \text{ctr}, \text{SEED}_u) \leftarrow \text{st}_u</math></li> <li>2: / Cannot update again if others haven't</li> <li>3: <b>if</b> <math>\text{ctr} = T</math> <b>then return</b> <math>\perp</math></li> <li>4: <b>parse</b> <math>(\text{SendSEED}_u, \text{ReceiveSEED}_u) \leftarrow \text{SEED}_u</math></li> <li>5: <math>\text{epoch}' \leftarrow \text{epoch} + 1</math></li> <li>6: <math>k_{\text{epoch}' } \leftarrow \text{PRF}(\text{gsk}, \text{epoch}' )</math></li> <li>7: / Create seeds/OTS keys for new epoch</li> <li>8: <math>(\text{ReceiveSEED}_u, Y_u, (\text{ct}_v)_{v \in G})</math>              <math>\leftarrow * \text{gen-auth-token}(G, \text{EK}, \text{epoch}', u, \text{ReceiveSEED}_u)</math></li> <li>9: <math>\tilde{Y}_u[*] := \perp</math> / Reorder OTS otkvs to be uploaded to Sv</li> <li>10: <b>for</b> <math>j \in [NT]</math> <b>do</b></li> <li>11:   <math>\tilde{j} \leftarrow \text{PRP}(k_{\text{epoch}' }, j)</math></li> <li>12:   <math>\tilde{Y}_u[\tilde{j}] \leftarrow Y_u[j]</math></li> <li>13:   <math>\hat{\text{ct}}_G \leftarrow (\text{idx}(u), \tilde{Y}_u, (\text{ct}_v)_{v \in G})</math></li> <li>14:   <math>\hat{\Sigma}_G \leftarrow (\text{epoch}, * \text{attach-auth-token}(\text{st}_u, \text{upd} :: (\hat{\Sigma}_G, \hat{\text{ct}}_G)))</math></li> <li>15:   <math>\text{SendSEED}_u[\text{epoch}] \leftarrow \perp</math> / Remove seeds from previous epoch</li> <li>16: / Update own seed/OTS otkvs for next epoch</li> <li>17: <math>\text{SendSEED}_u[\text{epoch}'][u] \leftarrow \text{ReceiveSEED}_u[\text{epoch}'][u]</math></li> <li>18: <math>\text{SEED}_u \leftarrow (\text{SendSEED}_u, \text{ReceiveSEED}_u)</math></li> <li>19: / Max out counter to T but do not increment epoch yet</li> <li>20: <math>\text{st}'_u \leftarrow (G, \text{gsk}, \text{EK}, \text{dk}_u, \text{epoch}, T, \text{SEED}_u)</math></li> <li>21: <b>return</b> <math>(\text{st}'_u, \hat{\Sigma}_G, \hat{\text{ct}}_G)</math></li> </ol>	<p><b>UpdVerify(<math>\text{pp}, \hat{\Sigma}_G, \hat{\text{ct}}_G</math>)</b></p> <ol style="list-style-type: none"> <li>1: <b>parse</b> <math>\text{DB} \leftarrow \text{pp}</math></li> <li>2: <b>try</b> <math>(\text{pp}', (\hat{\sigma}_i)_{i \in [N]}) \leftarrow * \text{verify-auth-token}(\text{pp}, \boxed{\text{upd} :: (\hat{\Sigma}_G, \hat{\text{ct}}_G)})</math></li> <li>3: <b>parse</b> <math>(\text{epoch}, \tilde{\text{id}}_{\text{send}}, (x_1, \dots, x_N)) \leftarrow \hat{\Sigma}_G</math></li> <li>4: <b>parse</b> <math>(\text{id}_{\text{receive}}, \tilde{Y}, (\text{ct}_i)_{i \in [N]}) \leftarrow \hat{\text{ct}}_G</math></li> <li>5: <b>req</b> <math>\text{DB}[\text{epoch} + 1][\text{id}_{\text{receive}}] = \perp</math> / No OTS otkvs for <math>\text{id}_{\text{receive}}</math> in epoch + 1</li> <li>6: <math>\text{DB}[\text{epoch} + 1][\text{id}_{\text{receive}}] \leftarrow \tilde{Y}</math> / Update DB for next epoch</li> <li>7: <b>foreach</b> <math>i \in [N]</math> <b>do</b></li> <li>8:   <math>\hat{\text{ct}}_i \leftarrow \text{ct}_i</math></li> <li>9: <b>return</b> <math>(\text{pp}', (\hat{\sigma}_i, \hat{\text{ct}}_i)_{i \in [N]})</math></li> </ol> <p><b>UpdReceive(<math>\text{st}_u, \hat{\sigma}, \hat{\text{ct}}</math>)</b></p> <ol style="list-style-type: none"> <li>1: <b>parse</b> <math>(G, \text{gsk}, \text{EK}, \text{dk}_u, \text{epoch}, \text{ctr}, \text{SEED}_u) \leftarrow \text{st}_u</math></li> <li>2: <b>parse</b> <math>(\text{SendSEED}_u, \text{ReceiveSEED}_u) \leftarrow \text{SEED}_u</math></li> <li>3: <b>try</b> <math>(\text{st}_u, b, v) \leftarrow * \text{trace-sender}(\text{st}_u, \hat{\sigma}, \hat{\text{ct}})</math></li> <li>4: <b>if</b> <math>v = u</math> <b>then return</b> <math>(\text{st}_u, \top, u)</math> / <math>\text{SendSEED}_u</math> is already updated with <math>\text{seed}_{u \rightarrow u}</math></li> <li>5: <b>req</b> <math>\text{SendSEED}_u[\text{epoch} + 1][v] = \perp</math> / v hasn't updated yet in epoch + 1</li> <li>6: <math>\text{seed}_{u \rightarrow v} \leftarrow \text{KEM.Dec}(\text{dk}_u, \hat{\text{ct}})</math> / Seed used by u to send message to v in epoch + 1</li> <li>7: <math>\text{SendSEED}_u[\text{epoch} + 1][v] \leftarrow \text{seed}_{u \rightarrow v}</math></li> <li>8: <math>\text{SEED}_u \leftarrow (\text{SendSEED}_u, \text{ReceiveSEED}_u)</math></li> <li>9: <b>if</b> <math>\forall w \in G, \text{SendSEED}_u[\text{epoch} + 1][w] \neq \perp</math> <b>then</b></li> <li>10:   / Increment epoch and refresh counter to 1 if everybody updated</li> <li>11:   <math>\text{st}'_u \leftarrow (G, \text{gsk}, \text{EK}, \text{dk}_u, \text{epoch} + 1, 1, \text{SEED}_u)</math></li> <li>12: <b>else</b> <math>\text{st}'_u \leftarrow (G, \text{gsk}, \text{EK}, \text{dk}_u, \text{epoch}, \text{ctr}, \text{SEED}_u)</math></li> <li>13: <b>return</b> <math>(\text{st}'_u, \top, v)</math></li> </ol>

Figure 11: STARS: An anonymous group authenticated messaging protocol with standard unforgeability. The main differences between QUASAR is highlighted by a box. label :: obj denotes that obj has a type label, where label is a special string used nowhere else. The helper algorithms used above are detailed in Fig. 12.

<pre> Func *gen-auth-token(G, EK, epoch, u, ReceiveSEED) 1: req ReceiveSEED[epoch] = ⊥ 2: foreach v ∈ G do 3:   (seed<sub>v→u</sub>, ct<sub>v</sub>) ← \$KEM.Enc(EK[v]) 4:   ReceiveSEED[epoch][v] ← (seed<sub>v→u</sub>, 1) 5:   foreach t ∈ [T] do 6:     / t-th OTS key v uses to send a message to u in epoch 7:     r<sub>v→u</sub><sup>(t)</sup> ← PRF(seed<sub>v→u</sub>, epoch    t    u) 8:     (otvk, otsk) ← OTS.KeyGen(1<sup>k</sup>, r<sub>v→u</sub><sup>(t)</sup>) 9:     y<sub>v→u</sub><sup>(t)</sup> := otvk 10: / Arrange all OTS otvks in order of users G = (u<sub>1</sub>, ..., u<sub>N</sub>) 11: Y ← (y<sub>u<sub>1</sub>→u</sub><sup>(1)</sup>, ..., y<sub>u<sub>1</sub>→u</sub><sup>(T)</sup>, 12:       ..., y<sub>u<sub>N</sub>→u</sub><sup>(1)</sup>, ..., y<sub>u<sub>N</sub>→u</sub><sup>(T)</sup>) ∈ ({0, 1}<sup>k</sup>)<sup>NT</sup> 13: return (ReceiveSEED, Y, (ct<sub>v</sub>)<sub>v∈G</sub>) </pre>	<pre> Func *attach-auth-token(st<sub>u</sub>, label :: obj) 1: parse (G, gsk, EK, dk<sub>u</sub>, epoch, ctr, SEED<sub>u</sub>) ← st<sub>u</sub> 2: parse (SendSEED<sub>u</sub>, ReceiveSEED<sub>u</sub>) ← SEED<sub>u</sub> 3: foreach v ∈ G do 4:   seed<sub>u→v</sub> ← SendSEED<sub>u</sub>[epoch][v] 5:   / Derive ctr-th OTS key u uses to send a message to v in epoch 6:   r<sub>u→v</sub><sup>(ctr)</sup> ← PRF(seed<sub>u→v</sub>, epoch    ctr    v) 7:   (otvk, otsk) ← OTS.KeyGen(1<sup>k</sup>, r<sub>u→v</sub><sup>(ctr)</sup>) 8:   if label = msg then 9:     parse m ← obj 10:    x<sub>u→v</sub><sup>(ctr)</sup> := sig ← OTS.Sign(otsk, m) 11:   elseif label = obj then 12:     parse (ct<sub>v</sub>)<sub>v∈G</sub> ← obj 13:     / Sign v's ciphertext using otsk from v 14:     x<sub>u→v</sub><sup>(ctr)</sup> := sig ← OTS.Sign(otsk, ct<sub>v</sub>) 15:   k<sub>epoch</sub> ← PRF(gsk, epoch) 16:   / Compute permuted index of the corresponding OTS otvk in DB 17:   id<sub>send</sub> ← ctr + (idx(u) - 1) · T 18:   id̃<sub>send</sub> ← PRP(k<sub>epoch</sub>, id<sub>send</sub>) / id<sub>send</sub>, id̃<sub>send</sub> ∈ [NT] 19:   return (id̃<sub>send</sub>, (x<sub>u→v</sub><sup>(ctr)</sup>)<sub>v∈G</sub>) </pre>
<pre> Func *verify-auth-token(pp, label :: obj) 1: DB ← pp 2: Y ← DB[epoch] / Y ∈ ({0, 1}<sup>k</sup>)<sup>N×NT</sup> 3: if label = msg then 4:   parse ((epoch, id̃<sub>send</sub>, (x<sub>1</sub>, ..., x<sub>N</sub>)), m) ← obj 5:   foreach i ∈ [N] do 6:     if OTS.Verify(Y[i][id̃<sub>send</sub>], m, x<sub>i</sub>) = ⊥ then 7:       return ⊥ 8:   elseif label = upd then 9:     parse (Σ̂<sub>G</sub>, ct̂<sub>G</sub>) ← obj 10:    parse (epoch, id̃<sub>send</sub>, (x<sub>1</sub>, ..., x<sub>N</sub>)) ← Σ̂<sub>G</sub> 11:    parse (id<sub>receive</sub>, Ỹ, (ct<sub>i</sub>)<sub>i∈G</sub>) ← ct̂<sub>G</sub> 12:    foreach i ∈ [N] do 13:      if OTS.Verify(Y[i][id̃<sub>send</sub>], ct<sub>i</sub>, x<sub>i</sub>) = ⊥ then 14:        return ⊥ 15:   / If check passes, set user authentication tokens and 16:   / delete column from Y 17:   foreach i ∈ [N] do 18:     σ<sub>i</sub> ← (epoch, id̃<sub>send</sub>, x<sub>i</sub>) 19:     Y[i][id̃<sub>send</sub>] ← ⊥ 20:   DB[epoch] ← Y / Update DB 21:   pp' ← DB 22:   return (pp', (σ<sub>i</sub>)<sub>i∈[N]</sub>) </pre>	<pre> Func *trace-sender(st<sub>u</sub>, σ, m) 1: parse (G, gsk, EK, dk<sub>u</sub>, epoch, ctr, SEED<sub>u</sub>) ← st<sub>u</sub> 2: parse (SendSEED<sub>u</sub>, ReceiveSEED<sub>u</sub>) ← SEED<sub>u</sub> 3: parse (epoch', id̃<sub>send</sub>, x) ← σ 4: req epoch = epoch' 5: k<sub>epoch</sub> ← PRF(gsk, epoch) 6: id<sub>send</sub> ← PRP<sup>-1</sup>(k<sub>epoch</sub>, id̃<sub>send</sub>) 7: i ← ⌊id<sub>send</sub>/T⌋ + 1 / Index of the user who sent the message 8: t ← id<sub>send</sub> - ⌊id<sub>send</sub>/T⌋ · T / t-th OTS otvk from user in epoch 9: v ← i-th user in G / v ∈ G and id<sub>send</sub> = t + (idx(v) - 1) · T 10: (seed<sub>v→u</sub>, ctr) ← ReceiveSEED<sub>u</sub>[epoch][v] 11: req ctr = t / Require that token wasn't used yet 12: r<sub>v→u</sub><sup>(ctr)</sup> ← PRF(seed<sub>v→u</sub>, epoch    ctr    u) 13: (otvk, otsk) ← OTS.KeyGen(1<sup>k</sup>, r<sub>v→u</sub><sup>(ctr)</sup>) 14: if OTS.Verify(otvk, m, x) = ⊥ then return (st<sub>u</sub>, ⊥, ⊥) 15: / Mark t-th OTS otsk to be used 16: ReceiveSEED<sub>u</sub>[epoch][v] ← (seed<sub>v→u</sub>, ctr + 1) 17: SEED<sub>u</sub> ← (SendSEED<sub>u</sub>, ReceiveSEED<sub>u</sub>) 18: st'<sub>u</sub> ← (G, gsk, EK, dk<sub>u</sub>, epoch, ctr, SEED<sub>u</sub>) 19: return (st'<sub>u</sub>, T, v) </pre>

Figure 12: Helper functions used by STARS. The main differences between QUASAR is highlighted by a box.

## E.2 Security Proof of STARS

We prove that STARS is correct and secure. The proof follows almost identically to the proof given for QUASAR.

**Theorem E.1.** *The GAM protocol STARS in Fig. 11 is signing correct, global state-update correct, standard unforgeable, anonymous, anonymous blocklistable, and tracing sound assuming PRF and PRP are pseudorandom, OTS is EUF-CMA, KEM is IND-CCA secure.*

*Proof.* Correctness of signing follows from construction. Moreover, we omit the proof of global state-update correctness and tracing soundness as they follow almost identically to the proof given for QUASAR (cf. Apps. D.2.1 and D.2.5). In the following, we prove STARS is unforgeable, anonymous, and anonymous blocklistable in Lems. E.2 to E.4, respectively. Proving these lemmas completes the proof of the theorem.  $\square$

### E.2.1 Proof of Lem. E.2

**Lemma E.2.** *STARS is unforgeable assuming PRF is pseudorandom, OTS is EUF-CMA, and KEM is IND-CCA secure.*

*Proof.* The proof is identical up to Game 2 of the non-colluding unforgeability proof given for QUASAR (cf. App. D.2.2). We pick up the proof from Game 3. Below, recall we focus on two honest users  $u$  and  $v$  (that may possibly be  $u = v$ ) and assume the adversary  $\mathcal{A}^*$  forges a group (update) authentication on epoch\*.

Game 3: In this game, when  $\mathcal{A}$  invokes  $O_{\text{UpdSend}}(u)$  on epoch\* - 1, it samples  $T$  fresh randomness for each  $w \in G$  used to run algorithm  $\text{OTS.KeyGen}: \bar{r}_{w \rightarrow u}^{(t)} \leftarrow \mathcal{S}\{0, 1\}^*$  for  $t \in [T]$  (cf. line 7 of \*gen-auth-token in Fig. 10). Note that in the previous game, it generated them by running  $\text{PRF}(\overline{\text{seed}}_{w \rightarrow u}, \text{epoch}^* || t || u)$ . Moreover, in this game, when user  $v$  runs  $\text{Send}$  or  $\text{UpdSend}$  with  $\text{SendSEED}_v[\text{epoch}][w] = \overline{\text{seed}}_{v \rightarrow u}$ , it uses  $\left(\bar{r}_{w \rightarrow u}^{(t)}\right)_{t \in [T]}$  when epoch = epoch\* rather than generating them via the PRF. Otherwise, if epoch  $\neq$  epoch\*, it samples  $T$  random fresh randomness for the pair  $(w, \text{epoch})$  and uses them. Here, note that we do not necessarily have  $(w, \text{epoch}) = (u, \text{epoch}^*)$  since a malicious user  $w$  can reuse the user update information  $\text{ct}_v$ , that  $u$  sent to  $v$  at some later epoch.

Finally, when  $\text{*trace-sender}(\text{st}_u, \sigma)$  with  $\sigma = (\text{epoch}^*, \tilde{\text{id}}_{\text{send}}, x, m)$  is invoked (as part of  $\text{Receive}$  or  $\text{UpdReceive}$ ) it checks if  $\tilde{\text{id}}_{\text{send}}$  is the permuted index of user  $v$ 's ctr-th token (cf. lines 7 to 9 of \*trace-sender in Fig. 10), where  $(\text{ReceiveSEED}_u[\text{epoch}^*][v] = (\overline{\text{seed}}_{v \rightarrow u}, \text{ctr}))$ . If not, it proceeds as in Game 2. Otherwise, it generates the OTS keys  $(\text{otvk}, \text{otsk})$  using randomness  $\bar{r}_{v \rightarrow u}^{(\text{ctr})}$  and check if

the signature message pair  $(x = \text{sig}, m)$  verifies under  $\text{otvk}$ . Otherwise, the game proceeds identically to the previous game.

Assuming the pseudorandomness of the PRF, it is clear that Game 2 and Game 3 are indistinguishable.

We are now ready to invoke the EUF-CMA security of OTS to show standard unforgeability. In Game 3, notice the OTS signing keys  $\left(\text{otsk}_{v \rightarrow u}^{(t)} := \bar{x}_{v \rightarrow u}^{(t)}\right)_{t \in [T]}$  in epoch\* are distributed uniformly random over the signing key space and shared only between the honest users  $u$  and  $v$ . The only information on these signing keys are provided to the adversary  $\mathcal{A}$  through the corresponding public OTS verification keys  $\left(\text{otvk}_{v \rightarrow u}^{(t)} := y_{v \rightarrow u}^{(t)}\right)_{t \in [T]}$ . Moreover, since each OTS signing keys are generated by fresh randomness  $\bar{r}_{v \rightarrow u}^{(t)}$ ,  $\text{otsk}_{v \rightarrow u}^{(t)}$  is only used when  $v$  is invoked on the  $t$ -th  $\text{Send}$  or  $\text{UpdSend}$  in epoch\* to prepare a signature for  $u$ . Note that this is why we modified QUASAR to include the sender of  $\text{seed}_{v \rightarrow u}$  to derive the randomness  $\bar{r}_{v \rightarrow u}^{(t)}$  (cf. line 7 of \*gen-auth-token in Fig. 10). Without this modification, a malicious user  $w$  that sets  $\text{SendSEED}_v[\text{epoch}^*][w] = \text{seed}_{v \rightarrow u}$  can trick  $v$  to sign multiple times using  $\text{otsk}$  since  $\bar{r}_{v \rightarrow u}^{(t)}$  will only depend on epoch\* and counter  $t$ .

Due to the winning condition, if  $\mathcal{A}$  outputs a valid forgery that traces to  $v$  in epoch\*, it must be on a message that  $v$  hasn't signed yet. This is exactly the winning condition of the EUF-CMA security game of OTS. Moreover, an adversary against the EUF-CMA security can simulate Game 3 to  $\mathcal{A}$  by embedding its challenge into one of its  $\text{otsk}_{v \rightarrow u}^{(t)}$  since each OTS signing key is only used once. Finally, the case when  $\mathcal{A}$  outputs  $(\text{label}, \text{obj}) = (\text{upd}, (\bar{\Sigma}_G, \hat{\text{ct}}_G))$  is proven identically. This completes the proof of the lemma.  $\square$

### E.2.2 Proof of Lem. E.3

**Lemma E.3.** *STARS is anonymous assuming PRF and PRP are pseudorandom, OTS is EUF-CMA, and KEM is IND-CCA secure.*

*Proof.* As the proof is a straightforward modification of the anonymity proof of QUASAR (cf. App. D.2.3), we only sketch the proof. The only modification is how we argue indistinguishability of Game 3 and Game 4 in Lem. D.4. In QUASAR, we used the entropy of the private token to argue that these two games are statistically indistinguishable. In STARS, we use the EUF-CMA security of the OTS. In particular, if the challenge users  $u_0$  or  $u_1$  generate a user authentication token for which user  $v$  accepts under condition Equation (2), then that token can be used to win the EUF-CMA security game. This completes the proof of the lemma.  $\square$

### E.2.3 Proof of Lem. E.4

**Lemma E.4.** *STARS is anonymous blocklistable assuming PRF is pseudorandom, OTS is EUF-CMA, and KEM is IND-CCA secure.*

*Proof.* The proof follows almost directly from the proof of unforgeability (cf. Lem. E.2), combined with the discussion provided in the proof of anonymous blocklistability of QUASAR (cf. App. D.2.4).  $\square$

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Background . . . . .	1
1.2	Our Contributions . . . . .	2
<b>2</b>	<b>Setting: Authentication in SGM</b>	<b>3</b>
2.1	Secure Group Messaging and Our Goal . . . . .	3
2.2	Environment . . . . .	4
2.3	Threat Model for Authentication . . . . .	4
2.4	Modeling Choices and Simplifications . . . . .	5
<b>3</b>	<b>Group Authenticated Messaging Protocol</b>	<b>6</b>
3.1	Definition . . . . .	6
3.2	Correctness . . . . .	7
3.3	Security . . . . .	7
<b>4</b>	<b>COSMOS: Authentication with One-Time Tokens</b>	<b>10</b>
4.1	Construction of COSMOS . . . . .	10
4.2	COSMAC: An Anonymous COSMOS with Anonymous Blocklisting . . . . .	10
4.3	Optimizations of COSMOS and COSMAC . . . . .	13
<b>5</b>	<b>Anonymous and Tracing Sound GAMs</b>	<b>13</b>
5.1	QUASAR: Anonymous Authentication with To- kens . . . . .	13
5.2	STARS and GEMSTARS . . . . .	16
<b>6</b>	<b>Running GAM Protocols on MLS</b>	<b>17</b>
6.1	Authentication in a Static Group . . . . .	17
6.2	Authentication in a Dynamic Group . . . . .	18
<b>7</b>	<b>Bandwidth Efficiency Analysis</b>	<b>18</b>
7.1	Instantiation . . . . .	18
7.2	Efficiency . . . . .	19
<b>8</b>	<b>Open Problems and Future Work</b>	<b>21</b>
<b>A</b>	<b>Other Related Work</b>	<b>26</b>
<b>B</b>	<b>Preliminary: Cryptographic Tools</b>	<b>27</b>
<b>C</b>	<b>Security Proofs for COSMOS and COSMAC</b>	<b>29</b>
C.1	Security Proof of COSMOS . . . . .	29
C.2	Security Proof of COSMAC . . . . .	30
<b>D</b>	<b>More Details on QUASAR</b>	<b>31</b>
D.1	Formal Description of QUASAR . . . . .	31
D.2	Security Proof of QUASAR . . . . .	31
D.3	Alternatives to Global State Updates . . . . .	36
<b>E</b>	<b>More Details on STARS</b>	<b>36</b>
E.1	Formal Description of STARS . . . . .	36
E.2	Security Proof of STARS . . . . .	39